

Iris-WasmFX: Modular Reasoning for Wasm Stack Switching

MAXIME LEGOUPIL, Aarhus University, Denmark

MATHIAS PEDERSON, Aarhus University, Denmark

LARS BIRKEDAL, Aarhus University, Denmark

SAM LINDLEY, University of Edinburgh, United Kingdom

JEAN PICHON-PHARABOD, Aarhus University, Denmark

WebAssembly, abbreviated Wasm, is a low-level portable bytecode. WasmFX is a proposed extension of Wasm with primitives for explicitly manipulating execution stacks as continuations. By exposing an interface in the style of effect handlers, WasmFX enables non-local control flow features to be compiled in a modular way: one handcrafts a library that directly implements such features in WasmFX, and compilation then merely calls into the library. Alas, code involving non-local control flow is notoriously challenging, and so this proposal raises the questions of the soundness of the language extension, and of the correctness of such handcrafted libraries. In this paper, we first describe WasmFXCert, a mechanisation of WasmFX in Rocq, and prove the expected type soundness result. We then develop Iris-WasmFX, a program logic to reason about WebAssembly programs that use effect handlers, and illustrate its application to two key use cases of effect handlers: a coroutine library, and a generator. Together, these validate the design of stack switching, and provide a modular framework for verifying future effect-based libraries.

1 Introduction

Wasm [Haas et al. 2017; Rossberg 2019, 2023] is a low-level portable bytecode extensively used both on the client- and server-side. As such, it is the compilation target for code that involves non-local control flow, including features such as promises, async/await, generators, and more. Compiling such features to current versions of Wasm requires a whole-program transformation such as the one used by Binaryen’s Asyncify [Zakai 2019]. However, such transformations suffer from three problems. First, such transformations introduce complexity. Second, they incur a significant overhead, in particular in code size. And third, they are not compositional: transformed code can only be composed with transformed code, not with off-the-shelf Wasm code.

To enable simpler, more efficient, and compositional compilation to Wasm, WasmFX [Phipps-Costin et al. 2023] extends Wasm with non-local control flow primitives based on Plotkin and Pretnar’s effect handlers [Plotkin and Pretnar 2009, 2013]. WasmFX is in the process of being standardised under the name *stack switching* [McCabe et al. 2025] as it provides a means for switching stacks. Previously the design was also referred to as *typed continuations* [Hillerström 2021; Rossberg 2020; Rossberg et al. 2020], as it is founded on a notion of continuations that follow the stack typing discipline of Wasm. In this paper we stick with *WasmFX* for conciseness. We will sometimes refer to the official standardisation documents for WasmFX [McCabe et al. 2025] as the *WasmFX proposal*. WasmFX already has implementations and industrial uses, but two important aspects remain to be investigated. First, although the static and dynamic semantics of Wasm are given in the proposal, the typing rules for runtime *administrative* instructions and a type soundness proof are missing. Second, the expected use case of WasmFX is to combine compiled code with hand-optimised libraries efficiently implementing non-local control flow features of source languages, raising the question of how to prove correctness of such libraries.

Authors’ Contact Information: Maxime Legoupil, Aarhus University, Aarhus, Denmark, maxime@cs.au.dk; Mathias Pederson, Aarhus University, Aarhus, Denmark, mp@cs.au.dk; Lars Birkedal, Aarhus University, Aarhus, Denmark, birkedal@cs.au.dk; Sam Lindley, University of Edinburgh, Edinburgh, United Kingdom, mymail@mail.com; Jean Pichon-Pharabod, Aarhus University, Aarhus, Denmark, jean.pichon@cs.au.dk.

Contributions. In this paper, we make two main contributions to the formal study of WasmFX: WasmFXCert, a mechanisation of the operational semantics and type system of WasmFX; and Iris-WasmFX, a program logic that enables modular reasoning for programs using WasmFX.

First, in order to formalise the semantics of WasmFX, we introduce WasmFXCert, a mechanisation of WasmFX in the Rocq proof assistant, built atop the existing WasmCert [Watt et al. 2021] formalisation of plain Wasm. WasmFXCert defines the full operational semantics and syntactic type system of WasmFX, including formally defined typing of administrative instructions, which existing specifications did not cover beyond an informal description [Phipps-Costin et al. 2023, §3.4]. WasmFXCert also provides the first result on the type safety of WasmFX, in the form of a mechanised proof which covers administrative instructions — thus validating the basic consistency of the WasmFX proposal.

Second, because effect handlers introduce non-local control flow, they famously hinder modular reasoning at the level of the operational semantics [Dijkstra 1968; Madore 2001]. In particular, capturing and storing of continuations make it challenging [Timany and Birkedal 2019] to reason about a program module-by-module, as doing so means that control does not flow in a well-bracketed way — and therefore not in a module-bracketed way — anymore. To address this difficulty, we introduce Iris-WasmFX, the first modular reasoning method for WasmFX programs. Iris-WasmFX is a program logic for WasmFX defined in the Iris separation logic framework [Jung et al. 2018] by extending the existing Iris-Wasm program logic [Rao et al. 2023] for plain WebAssembly, and takes inspiration from Hazel [de Vilhena 2022; de Vilhena and Pottier 2021, 2023], an Iris-based program logic for an ML-style calculus with effect handlers.

In this paper, we outline the difficulties that we faced when defining Iris-WasmFX, and highlight the challenges of doing machine-checked proofs of properties of industrial-size languages. Our mechanisation and program logic are crucial tools for gaining a deep understanding of the behaviour of WasmFX programs and establishing confidence in desired properties of specific programs. Libraries whose behaviour is critical for safety can be verified manually, and users of WasmFX can establish a high-degree of assurance in the soundness of its design. Together, WasmFXCert and Iris-WasmFX validate the design of WasmFX, and illustrate how to achieve modular reasoning for the prototypical kind of libraries that WasmFX was designed to support.

Plan. We first present our running example of a typical use-case of WasmFX, which we use throughout the paper to demonstrate how to program, specify, and verify with WasmFX (§2). We then describe the design of WasmFX and our formalisation, WasmFXCert, (§3), and establish type safety. On top of this foundation, we describe our program logic, Iris-WasmFX (§4), and apply it to two case studies (§5). Finally, we discuss related work (§6), and conclude (§7).

2 A Canonical Use Case of WasmFX: a Coroutine Library

In the remainder of this paper, we illustrate the design of WasmFX, our mechanisation WasmFXCert, and our program logic Iris-WasmFX, on a canonical use case of WasmFX: a small cooperative coroutine library module. The coroutine module exposes an interface comprising two functions. Library function `$par` takes two function references as arguments; these define two coroutines whose execution is interleaved. Library function `$yield` is used to signal that control should be handed over to the other coroutine, thereby cooperatively interleaving the execution of the coroutines. Behind the scenes, `$par` implements a scheduler that starts running one function, and if that function stops by invoking `$yield`, then starts running the other function, and so on.

This example illustrates the kind of code that WasmFX makes possible. Without it, this kind of functionality would require a somewhat complex compilation of the client code, likely incurring some code size inflation. Instead, with WasmFX, a concise definition can be given to `$par` and

`$yield`, as we show below. Thanks to Wasm’s modularity, from the point of view of a client, the functions `$par` and `$yield` can be called like normal functions, and the code of the client can thus be written without using the new instructions in the WasmFX semantics.

A Simple Specification of Coroutines. A natural specification for this coroutine module in the style of Concurrent Separation Logic (CSL) [Brookes and O’Hearn 2016] is centred around the notion of an invariant: if both client functions $\$f_i$ (for i in $\{1, 2\}$) being run ensure that they establish the invariant I before calling `$yield`, then they can rely on the fact that this same invariant is restored when control returns to them — that is, when `$yield` returns. Assuming this, if each function $\$f_i$, starting from precondition P_i , ensures postcondition Q_i at termination, with access to invariant I at the start and relinquishing it at the end, then running both functions via `$par` starting with P_1, P_2 and I ensures that at termination, Q_1, Q_2 and the invariant hold. We give a more detailed intuition for the shape of the specification by way of Hoare triples for the `$yield` and `$par` functions. The Hoare triple for `$par` uses nested Hoare triples [Schwinghammer et al. 2009]: it requires in its precondition corresponding specifications for $\$f_1$ and $\$f_2$ in the form of Hoare triples:

$$\forall P_1, P_2, Q_1, Q_2, I. \\ \{I\}[\text{invoke } \text{addryield}] \{I\} * \\ \left\{ \begin{array}{l} P_1 * P_2 * I * \\ \{P_1 * I\}[\text{invoke } \text{addrf}_1] \{Q_1 * I\} * \\ \{P_2 * I\}[\text{invoke } \text{addrf}_2] \{Q_2 * I\} \end{array} \right\} \left[\begin{array}{l} \text{func.ref } \$f_1; \\ \text{func.ref } \$f_2; \\ \text{invoke } \text{addrpar} \end{array} \right] \{Q_1 * Q_2 * I\}$$

where `addryield`, `addrpar`, `addrf1` and `addrf2` are the addresses of the closures for `$yield`, `$par`, $\$f_1$ and $\$f_2$. We show how to make this informal specification precise in Appendix E.

Modularity. Importantly, this specification can be proved for the coroutine module by itself, without access to the code of the client functions $\$f_i$, which can themselves be verified against this specification, without reference to the code of the coroutine module.

Implementation. This coroutine module can be implemented using WasmFX: `$yield` suspends the current program, saving the current continuation, and `$par` resumes the continuations of $\$f_1$ and $\$f_2$, round-robin, using the corresponding WasmFX instructions. We describe this simple implementation of coroutines in Section 3, and give its code in Figure 3.

Library Client. A simple example use of this coroutine library is a client implementing a classic ‘message passing’ concurrency example [Sewell et al. 2010, 8–1][Maranget et al. 2012, §3]. We implement this in Wasm (the code is in the appendix in Figure 13) in the form of a client module that uses the coroutine module to concurrently run a ‘writer’ function $\$f_1$ that writes 42 to a ‘data’ global $\$x$, and then 1 to a ‘flag’ global $\$y$ to signal that the data is ready in $\$x$, with a ‘reader’ function $\$f_2$ that reads the ‘flag’ $\$y$ and then reads the ‘data’ in $\$x$, with both functions calling `$yield` in between their two accesses to globals. Using this specification of the coroutine library, the goal is to prove that if the reader sees that the flag is set, then it should see 42 as the data. In the rest of this paper, we build up to making the above specification precise and prove it correct.

3 WasmFXCert: Formal Operational Semantics for WasmFX

In this section, we introduce WasmFX and our mechanisation of it in the Rocq proof assistant, WasmFXCert. Phipps-Costin et al. [2023, §2–3] give a more detailed introduction to WasmFX.

3.1 WasmFX, Informally

Design of WasmFX. WasmFX is effectively a variant of *effect handlers*, a mechanism for non-local control flow that can be thought of as resumable exceptions: when an effect is performed, execution

```

148 1 (loop $l
149 2   (block $b
150 3     ...                ;; prepare function to run
151 4     resume ... (on $t $b) ;; run
152 5     ...                ;; handle normal termination
153 6     br $l)             ;; resume the loop, circumventing the code for yield below
154 7     ...                ;; handle yield, changing function to run
155 8     br $l)             ;; resume the loop
156
157
158
159

```

Fig. 1. High-level structure of the implementation of `$par` (full code in Figure 3)

is halted, and control is handed over to the innermost *handler* for that effect, thereby defining a delimited continuation [Felleisen 1988]. If it was intended for programmers to routinely write code directly in Wasm, it would make most sense to define a specific, user-friendly mechanism. But since Wasm is designed as a compilation target, it makes most sense to offer a flexible mechanism that can be used to implement different non-local control flow mechanisms, even if it is at the cost of being less human-readable. Therefore, the design of Wasm’s effect handlers in WasmFX is somewhat unusual, and we illustrate it on our motivating example. Phipps-Costin et al. [2023, §4] show how to implement other types of non-local control flow using the instructions of WasmFX.

WasmFX Instructions. The core of WasmFX consists of two instructions: *suspend*, which stops the current execution context, and reifies it as a *continuation*, a new kind of Wasm value with a corresponding new Wasm type; and *resume*, which starts executing such a continuation — thereby switching the stack, hence the name of the Wasm proposal. In our coroutine example, `$yield` can be implemented by a *suspend*, interrupting the execution of the current function, and `$par` can be implemented by *resume*-ing the other function, in a loop, as sketched in Figure 1. We now focus on the *resume* instruction on line 4.

Block-Based Effect Handling. The main unusual feature of WasmFX is that *suspend* suspends to a code block specified in the *resume*, as opposed to some more user-friendly mechanism. As such, WasmFX continuations are *delimited* continuations [Felleisen 1988]. In our coroutine example, this means that the code just after the *resume* (line 5, within the block) handles normal termination, and the code after the block (line 7) handles the `$yield`. Because execution is fall-through, normal termination (in block `$b`) has to jump over the handling of the yield (lines 7–8, after block `$b`). Confusingly, in Wasm, *loop* only declares a block to loop to: actual looping only takes place given an explicit `br $l` (lines 6 and 8).

Tags. To compose different effects modularly, *suspend* is annotated with a tag (called an *effect operation* in the effect handler literature). Each tag specifies what types of values (the *payload* of the effect) are to be passed along with it. Then *resume* can specify, for each tag, to which block it should be suspended, using the *on* keyword. In our coroutine example, we only need one tag to (ask `$par` to) pass over control to the other function: `$yield` then just suspends with that tag.

3.2 Operational Semantics

Pedigree. The WasmFX proposal builds upon Wasm 3.0. However, it only makes use of a few features not present in Wasm 1.0: function references, and exception handling [Ahn and Titze 2022]. Our mechanisation, WasmFXCert, takes advantage of this, and is built on top of WasmCert 1.0 [Watt et al. 2021], to which we have manually added the features of Wasm 3.0 that WasmFX

uses. This is because Iris-Wasm [Rao et al. 2023], which we used as a departure point to define our program logic (see §4), is defined for Wasm 1.0; defining it atop Wasm 3.0 would have required dealing with the numerous features of Wasm 3.0 that are orthogonal to WasmFX.

In the rest of this section, we give an introduction to the semantics of Wasm. A reader familiar with Wasm can safely skip to §3.3, where we describe the new instructions in WasmFX. A reader familiar with WasmFX can skip to §3.4, where we describe type safety.

Syntax. We give the syntax for WasmFX in Figure 2, and illustrate how a Wasm module is structured with our implementation of a coroutine module in Figure 3. To avoid confusion with the symbol $*$ for the separating conjunction, we depart from the Wasm standard and use an ‘s’ suffix to indicate a list; e.g. bs indicates a list of b elements. The coroutine module in Figure 3 defines a type (line 2), a continuation type (line 3), a tag (line 4), and the two library functions $\$yield$ (lines 5–6) and $\$par$ (lines 7–39).

Frame and Store. The state of a Wasm program is separated in two parts: the *frame* and the *store*. The store S contains all objects (e.g. global variables, functions, memories, as well as, in WasmFX, tags and continuations) declared by all modules instantiated thus far, making imports and exports easy to deal with. The frame F contains information about the environment of the function currently being run. It consists of two components: the values of the function’s local variables and the *instance*, a dictionary that translates local indices into addresses in the store. This indirection is crucial to obtain Wasm’s local state encapsulation property, but increases the complexity of defining a mechanisation and program logic.

Reductions. Reductions are of the form $(S, F, es) \hookrightarrow (S', F', es')$, where S and F are the aforementioned store and frame, and es is a list of *administrative instructions*, a superset of basic instructions that includes instructions used to represent a program mid-execution; WasmFX introduces a number of administrative instructions, and we describe them in the rest of this subsection.

Control Flow. To express the semantics of loop and block, Wasm defines an administrative instruction $\text{label}_n\{es_{cont}\} es \text{ end}$, that both loop and block reduce to, where es represents the body of the loop or block; es_{cont} represents the code that should be run in case of a break (i.e. the entire loop again in the case of a loop, or nothing in case of a block), and n is the number of values that should be on the stack in case of a break. Wasm defines evaluation by the composition of head reduction rules with the following evaluation-under-context rule:

$$\text{REDUCE-LABEL} \quad \frac{(S, F, es) \hookrightarrow (S', F', es')}{(S, F, L[es]) \hookrightarrow (S', F', L[es'])}$$

where evaluation contexts L are as per Figure 2 (they are written lh_i in Iris-Wasm [Rao et al. 2023], and B^i in the standard [Rossberg 2023].) Taking $L = [\text{label}_n\{es_{cont}\} [_] \text{ end}]$, we get that after a loop or block has reduced to a `label`, the code inside the label can run. If that code terminates on a value, the final value exits the label and execution continues with the instructions that follow the label. If however a `br` instruction (for *break* or *branch*) is encountered, the appropriate amount of values is taken from the stack and code es_{cont} is run. In WasmFX, as we will detail in §3.3, the `resume` instruction also reduces to a `br $b` when an effect is performed, meaning that the code after the instruction corresponds to the code that will be run if no effect is performed, whereas the code es_{cont} of the `label` instruction labelled `$b` will run if an effect is performed.

Function Calls. Wasm defines several mechanisms for function calls, and an administrative instruction `invoke addr` that all the primitive call instructions reduce to. The simplest calling instruction is `call $f`. The frame is used to determine at what address `addr` in the store to find the

(numerical value type) $nt ::= i32 \mid i64 \mid f32 \mid f64$ (value type) $t ::= nt \mid rt$
 (reference value type) $rt ::= \boxed{ft.func \mid exn \mid ft.cont}$ (function type) $ft ::= ts \rightarrow ts$
 (continuation clause) $cc ::= on\ i\ i' \mid on\ i\ switch$ (value) $v ::= nt.const\ c \mid rv$
 (immediate) $i, addr, max ::= \mathbb{N}$ (i for local indices, $addr$ for store addresses)
 (reference value) $rv ::= rt.null \mid ref.func\ addr \mid ref.exn\ addr\ addr' \mid ref.cont\ addr$
 (basic instructions) $b ::= nt.const\ c \mid t.add \mid other\ stackops \mid local.\{get/set\}\ i \mid$
 $global.\{get/set\}\ i \mid t.load\ flags \mid t.store\ flags \mid memory.size \mid$
 $memory.grow \mid block\ ft\ bs \mid loop\ ft\ bs \mid if\ ft\ bs\ bs' \mid br\ i \mid br_if\ i \mid$
 $br_table\ is \mid call\ i \mid call_indirect\ i \mid return \mid rt.null \mid$
 $ref.is_null \mid ref.function\ i \mid ref.call\ i \mid$
 exception handling instructions $\mid resume\ i\ ccs \mid$
 $suspend\ i \mid switch\ i\ i' \mid cont.bind\ i\ i' \mid resume_throw\ i\ i'\ ccs$
 (translated continuation clause) $dcc ::= on\ addr\ i \mid on\ addr\ switch$
 (administrative instruction) $e ::= b \mid trap \mid invoke\ addr \mid label_n\{es'\}\ es\ end \mid$
 $frame_n\{F\}\ es\ end \mid call_host\ ft\ i\ vs \mid ref.func\ addr \mid$
 exception handling instructions $\mid ref.cont\ addr \mid$
 $prompt_{ts}\{dcs\}\ es \mid suspend.addr\ \boxed{vs}\ addr$
 $switch.addr\ \boxed{vs}\ addr\ ft\ addr'$
 (functions) $func ::= func\ i\ ts\ bs$ (tables) $tab ::= tab\ min\ max$
 (memories) $mem ::= mem\ min\ max$ (globals) $glob ::= glob\ mutable\ t\ e_{init}$
 (import descriptions) $importdesc ::= func_i\ i \mid tab_i\ i\ i' \mid mem_i\ i \mid glob_i\ mutable^? t \mid tag_i\ i$
 (export descriptions) $exportdesc ::= func_e\ i \mid tab_e\ i \mid mem_e\ i \mid glob_e\ i \mid tag_e\ i$
 (modules) $m ::= \{types : fts, tags : fts, funcs : funcs, globs : globs, mems : mems, \dots\}$
 (module instance) $inst ::= \{types : fts, funcs : addrs, globs : addrs, mems : addrs,$
 $tabs : addrs, tags : addrs\}$
 (frame) $F ::= \{locs : vs, inst : inst\}$
 (evaluation context) $L ::= \boxed{vs ++ [_] ++ es} \mid vs ++ label_n\{es'\}\ L\ end ++ es \mid$
 $vs ++ prompt_{ts}\{dcs\}\ L\ end ++ es$
 (handler context) $H ::= vs ++ [_] ++ es \mid vs ++ label_n\{es'\}\ H\ end ++ es \mid$
 $vs ++ frame_n\{F\}\ H\ end ++ es \mid$
 $vs ++ prompt_{ts}\{dcs\}\ H\ end ++ es$
 (function instances) $finst ::= \{(inst; ts); bs\}_{ft}^{NativeCl} \mid \{hidx\}_{ft}^{HostCl}$
 (continuation instance) $cinst ::= cont\ \boxed{ft}\ H \mid dagger\ \boxed{ft}$
 (store) $S ::= \{funcs : finsts, globs : ginsts, mems : minsts, tabs : tinsts,$
 $tags : fts, exns : einsts, conts : cinsts\}$
 (the meaning of some of the fields in the store is orthogonal to stack switching and is omitted)

Fig. 2. The AST of the WasmFXCert language. Black parts are in WasmCert 1.0 (plain Wasm 1.0). Indigo parts are reference instructions from Wasm 3.0. Orange parts are the additions of exception handling, that are also needed for WasmFX. Magenta parts belong to WasmFX proper. A full AST can be found in appendix, Figures 7 and 8. Places where WasmFXCert departs from WasmFX are shown boxed.

```

295 1 (module ;; coroutine
296 2   (type $func (func))
297 3   (type $cont (cont $func))
298 4   (tag $t)
299 5   (func $yield                                ;; takes no arguments
300 6     (suspend $t))                             ;; and just suspends
301 7   (func $par
302 8     (param $f1 (ref $func))                   ;; takes two functions as arguments
303 9     (param $f2 (ref $func))
304 10    (local $current (ref $cont))               ;; local variables:
305 11    (local $next (ref $cont))
306 12    (local $next_is_done i32)
307 13    i32.const 0                               ;; $next_is_done := false
308 14    local.set $next_is_done
309 15    local.get $f1                             ;; $current := cont(f1)
310 16    cont.new $cont
311 17    local.set $current
312 18    local.get $f2                             ;; $next := cont(f2)
313 19    cont.new $cont
314 20    local.set $next
315 21    (loop $l                                  ;; while true do
316 22      (block $b (result (ref $cont))
317 23        local.get $current                     ;; resume continuation $current
318 24        resume $cont (on $t $b)
319 25        local.get $next_is_done               ;; if it terminates, and $next_is_done
320 26        br_if 2                               ;; return
321 27        i32.const 1                           ;; $next_is_done := true
322 28        local.set $next_is_done
323 29        local.get $next                       ;; $current := $next
324 30        local.set $current
325 31        br $l)                                ;; loop
326 32      local.set $current                       ;; else (if effect is triggered) $current := cont
327 33      local.get $next_is_done                 ;; if $next_is_done
328 34      br_if $l                                ;; loop
329 35      local.get $current                       ;; swap $current $next
330 36      local.get $next
331 37      local.set $current
332 38      local.set $next
333 39      br $l)                                ;; loop
334 40    (export "yield" (func $yield))             ;; make $yield available
335 41    (export "par" (func $par))                 ;; and $par
336
337
338
339

```

Fig. 3. Implementation of our simple coroutine module using WasmFX

closure for $\$f$, and call $\$f$ reduces to invoke *addr*. A module can also define a function table and use *call_indirect* to call functions from it. In Wasm 3.0, the *ref.function* $\$f$ instruction creates a reference to an existing function, which can be called using the *ref.call* instruction.

Closures. The closures in the store can either be defined by the embedding host language (see §3.5, in which case control is handed over to the host language, which in WasmCert is rendered by reducing to a special `call_host` administrative instruction [Rao et al. 2023, §2.2]); or it can be defined in Wasm, in which case `invoke addr` reduces to $(\text{frame}_{|ts_2|}\{F'\} \text{ block } ([\rightarrow ts_2] \text{ bs end})$ where *bs* is the body of the closure, *ts₂* is the function’s return type, and *F'* is the frame in which this body is meant to be executed. This frame *F'* contains the function’s arguments as local variables, and the instance of the module in which the function was defined, hence ensuring the code of the function can only access the elements of the store it is intended to.

The $\text{frame}_n\{F\} \text{ es end}$ administrative instruction (sometimes called `local` [Haas et al. 2017; Rao et al. 2023]) represents the execution of a function call, and is responsible for ensuring encapsulation. Its rule is similar to `REDUCE-LABEL`, but updates the frame in its premise:

$$\begin{array}{c} \text{REDUCE-FRAME} \\ \hline (S, F, \text{es}) \hookrightarrow (S', F', \text{es}') \\ \hline (S, F_0, [\text{frame}_n\{F\} \text{ es end}]) \hookrightarrow (S', F_0, [\text{frame}_n\{F'\} \text{ es' end}]) \end{array}$$

A return instruction is defined similarly to `br` and allows to exit a function by jumping out of the closest frame instruction. Since an `invoke` reduces to the body wrapped in both a frame and a block, the programmer can alternatively use a `br` instruction instead of a return, for example to take advantage of the `br_if` instruction, as is done in the code of `$par` at line 26.

3.3 New Instructions in WasmFX

We now introduce formally the new WasmFX instructions sketched in §3.1. WasmFX defines a new type, *ft.cont*, for continuations of type *ft*. All new continuations are created from a function reference using the `cont.new` instruction. The new continuation is placed in the store at an address *addr*, and `ref.cont addr` is returned:

$$\begin{array}{c} \text{REDUCE-CONTNEW} \\ \hline F.\text{inst.types}[i] = \text{ft} \quad S' = \{S \text{ with } \text{conts} \mathrel{+=} \text{cont } \text{ft } ([_]\mathrel{++} [\text{ref.func } \text{addr}; \text{ref.call } \text{ft}])\} \\ \hline (S, F, [\text{ref.func } \text{addr}; \text{cont.new } i]) \hookrightarrow (S', F, [\text{ref.cont } |S.\text{conts}|]) \end{array}$$

Behaviour of resume and prompt. The `resume` instruction expects a value of the form `ref.cont addr` on the stack, and runs the continuation present in the store at address *addr*. The instruction takes a list of *clauses* as an immediate argument, describing the intended behaviour in case an effect is triggered. This means that in WasmFX, there are not two distinct instructions for running a continuation and installing a handler: both operations are dealt with by the `resume` instruction. This means the usual distinction between shallow and deep handlers is meaningless.

Similarly to `label` and `frame`, WasmFX defines an administrative instruction `prompt` to represent a continuation being run. Just like `resume`, `prompt` is decorated with a list of *clauses* that inform the desired behaviour in case of a suspend or a switch. One small difference is that the `resume` instruction’s clauses reference the local tag names as written by the programmer, whereas the clauses in the `prompt` instruction has the explicit store addresses of these tags.

To account for the fact that in WasmFX, continuations are one-shot [Phipps-Costin et al. 2023, §2.4], upon resuming, the store is updated with a dagger token to indicate the continuation has already been resumed, and that therefore resuming it again would cause a trap:

$$\begin{array}{c} \text{REDUCE-RESUME} \\ \hline F.\text{inst.types}[i] = \text{ft} \quad \text{ft} = ts_1 \rightarrow ts_2 \quad |vs| = |ts_1| \quad S.\text{conts}[\text{addr}] = \text{cont } \text{ft } H \\ H[vs] = \text{es} \quad F \vdash \text{ccs} \rightarrow \text{dccc} \quad S' = \{S \text{ with } \text{conts}[\text{addr}] = \text{dagger } \text{ft}\} \\ \hline (S, F, vs \mathrel{++} [\text{ref.cont } \text{addr}; \text{resume } i \text{ ccs}]) \hookrightarrow (S', F, [\text{prompt}_{ts_2}\{\text{dccc}\} \text{ es end}]) \end{array}$$

The *ccs* clauses in the resume are translated into explicit *dccs* clauses in the prompt instruction: on *i i'* and on *i switch* are translated to on *addr i'* and on *addr switch* if *addr = F.inst.tags[i]*. We write $F \vdash ccs \rightarrow dccs$ to describe this translation. WasmFX extends execution contexts *L* to also include prompt contexts, so execution carries out inside of a prompt exactly like it does under a label. Likewise, if this execution terminates in a value, WasmFX defines a reduction rule making the value exit the prompt and continuing execution with the instructions following the prompt: this constitutes the ‘normal return’ case of the effect handling.

Behaviour of suspend. The suspend instruction takes a tag *\$t* as an immediate argument. Informally, it takes the required amount of values from the stack to constitute the effect’s payload, creates a delimited continuation out of the current state of the stack, and yields control to the innermost enclosing prompt that has a clause handling tag *\$t*. This clause has the form on *\$t \$l*: this means that the handler intends for a break to label *\$l* to occur when the effect identified by tag *\$t* is triggered; hence the right-hand-side of reduction rule EWP-SUSPEND below is a break.

Translated suspend. One caveat is that the suspend instruction’s immediate argument is a tag name *\$t* and the prompt instructions use tag *addresses* (into the store) in their clauses. The WasmFX proposal originally attempted to do this translation implicitly but ended up using the wrong function instance to do so. This bug in the proposal was exposed by the mechanisation and has now been fixed. The reference interpreter avoided this by using a smaller-step operational semantics.

Our proposed solution in WasmFXCert is to make this translation from tag names to tag addresses explicit: we distinguish a basic instruction suspend with a local index for the tag, and an administrative instruction suspend.addr, with the actual address into the store. This solution has now also been adopted by the WasmFX proposal, albeit with a slight difference in naming, as the proposal prefers to broaden the definition of the suspend instruction by allowing it to alternatively take an address as an argument. This means that in the WasmFX proposal, suspend *i* is translated to suspend *addr* instead of a new suspend.addr instruction as we do in WasmFXCert (see rule REDUCE-SUSPEND-TRANSLATE below). This is in order to maintain consistency with other places in the standard and limit the proliferation of administrative instructions. WasmFXCert uses two distinct instructions for clarity, but both solutions have equivalent semantics.

In WasmFXCert, we find it convenient to also let this translation take the payload of the effect from the stack and place it into the suspend.addr instruction itself. This has no effect on the behaviour of the execution, but gives the language a property that will be crucial when defining the Iris-WasmFX program logic (see Figure 5): inspecting an expression (without knowledge of the store) is enough to know whether the expression is an effect, and if so what the effect’s payload is. Without placing the payload into the instruction itself, it is impossible to know how many values to take from the stack without knowing the type signature of the tag, which is in the store and not visible in the expression itself. We explain in §4.2 why this is important.

Thus we introduce the following rule for explicit translation of the suspend instruction; the standard has now introduced the same rule, but without relocating the payload:

$$\frac{\text{REDUCE-SUSPEND-TRANSLATE} \quad F.\text{inst.tags}[i] = \text{addr} \quad S.\text{tags}[\text{addr}] = ts_1 \rightarrow ts_2 \quad |vs| = |ts_1|}{(S, F, vs \mathrel{++} [\text{suspend } i]) \hookrightarrow (S, F, [\text{suspend.addr } vs \text{ addr}])}$$

The Rule for Suspension. The suspend.addr instruction itself is stuck and needs to be placed under an appropriate context to reduce. When reducing $H[\text{suspend.addr } vs \text{ addr}]$, we look up for the innermost prompt in *H* which has a clause for tag address *addr*: this allows us to decompose the expression into $H_1[\text{prompt}_{ts}\{dccs\} H_2[\text{suspend.addr } vs \text{ addr}] \text{ end}]$ where H_2 is capture-avoiding; we can use the REDUCE-LABEL and REDUCE-FRAME rules to focus reasoning inside of H_1 , and then

apply the rule REDUCE-SUSPEND rule below. We use notation $H_{addr}^{sus}[es]$ to denote capture-avoiding plugging (i.e. no prompts with a suspend clause handling $addr$ in H) of es into context H .

$$\begin{array}{c} \text{REDUCE-SUSPEND} \\ \frac{S.\text{tags}[addr] = ts_1 \rightarrow ts_2 \quad \text{first_suspend_occurrence}(addr, dcs) = \text{on } addr \ i \quad H_{addr}^{sus}[\text{suspend.addr } vs \ addr] = es \quad S' = \{S \text{ with conts } ++= \text{cont } (ts_2 \rightarrow ts) \ H\}}{(S, F, [\text{prompt}_{ts}\{dcs\} \ es \ \text{end}]) \hookrightarrow (S', F, vs ++ [\text{ref.cont } |S.\text{conts}|; \text{br } i])} \end{array}$$

Behaviour of switch. The switch instruction optimises the common combination of a suspend followed by immediately resuming a different continuation. It expects a continuation H on the stack, creates a continuation out of the current state of the stack, and yields control to H . Just like for suspend, in WasmFXCert, we explicitly translate the local index by distinguishing two instructions switch and switch.addr. We write $H_{addr}^{sw}[es]$ to denote capture-avoiding (i.e. no prompts with a switch clause handling $addr$ in H) plugging of expression es in context H .

$$\begin{array}{c} \text{REDUCE-SWITCH-TRANSLATE} \\ \frac{F.\text{inst.types}[i] = ft \quad ft = ts_1 \rightarrow ts_2 \quad F.\text{inst.tags}[i'] = addr' \quad |vs| + 1 = |ts_1|}{(S, F, vs ++ [\text{ref.cont } addr; \text{switch } i \ i']) \hookrightarrow (S, F, [\text{switch.addr } vs \ addr \ ft \ addr'])} \end{array}$$

REDUCE-SWITCH

$$\begin{array}{c} \text{REDUCE-SWITCH} \\ \frac{(on \ \text{addr}' \ \text{switch}) \in dcs \quad ft = (ts_1 ++ ft'.\text{cont}) \rightarrow ts_2 \quad S.\text{conts}[addr] = \text{cont } (ts'_1 \rightarrow ts'_2) \ H' \quad H_{addr'}^{sw}[\text{switch.addr } vs \ addr \ ft \ addr'] = es \quad S' = \{S \text{ with conts } ++= \text{cont } ft' \ H\} \quad S'' = \{S' \text{ with conts } [addr] = \text{dagger } (ts'_1 \rightarrow ts'_2)\}}{(S, F, [\text{prompt}_{ts}\{dcs\} \ es \ \text{end}]) \hookrightarrow (S'', F, [\text{prompt}_{ts}\{dcs\} \ (H' [vs ++ \text{ref.cont } |S.\text{conts}|]) \ \text{end}])} \end{array}$$

Behaviour of cont.bind and resume_throw. Finally, WasmFX defines two more instructions for ease of programming: cont.bind partially applies a continuation, and resume_throw resumes a continuation but plugs in an exception-throwing instruction, as a way to cleanly discard a continuation that is no longer useful. These instructions are not crucial to understanding WasmFX, so we omit further mention of them, but we model them in WasmFXCert, and give them rules in our program logic Iris-WasmFX, which can be seen in our Rocq development.

Departures from the WasmFX Proposal in WasmFXCert. We make a few minor departures from the official WasmFX semantics. These are shown boxed in Figure 2:

- Different versions of the standard and proposals use equivalent but different grammars for reference types and contexts. We pick one that is readable for the features we use.
- As discussed, we place the effect payload in suspend.addr and switch.addr.
- We add a type annotation on the prompt instruction and on the continuations in the store. Those apply only to administrative instructions, not source instructions, so this does not affect typechecking of source code. We added these annotations because identifying the type of these expressions is likely to be useful for future users of Iris-WasmFX when conducting deeper analyses, like proofs by logical relation.

3.4 Syntactic Typing and Type Safety

To prove type safety for WasmFX, we need to extend the type system of Phipps-Costin et al. [2023] to define typing rules for administrative instructions. In doing so, we identify a few minor defects in their definitions (see Appendix C). We then prove the following theorem in Rocq (see Appendix C or the Rocq development for precise definitions):

THEOREM 3.1 (TYPE SAFETY). *If the store S typechecks and $S, C(S, F) \vdash es: [] \rightarrow ts$ and $(S, F, es) \hookrightarrow^* (S', F', es')$, then either es' reduces, or it is a constant value, or it is stuck on a host call, an unhandled suspend, switch, or exception throw.*

3.5 Instantiation

The process of preparing the modules for running is called *instantiation*: the code is typechecked, the imports are fetched, the objects defined by the module itself are added to the store, and the module's exports are prepared for subsequent imports by other modules.

This process cannot be performed by Wasm itself but instead is done by an embedding *host language*. The Iris-Wasm program logic [Rao et al. 2023] comes with a custom-designed host language that can instantiate Wasm modules and do a few other Wasm operations like calling Wasm functions or reading Wasm global variables. For example, to run our example from §2, the host program (shown in Theorem 5.1) consists of three steps: instantiating the coroutine module, instantiating the client module, and calling the `$main` function of the client module (see top of Figure 13 in appendix for how to run these three steps in the reference interpreter of WasmFX). This host language is mechanised in Rocq and has its own program logic; for simplicity, it does introduce its own effects.

4 Iris-WasmFX: Modular Reasoning for WasmFX Programs

In this section, we introduce Iris-WasmFX, our program logic for WasmFX.

4.1 Iris and Iris-Wasm

Our program logic is defined in Iris [Jung et al. 2018], a framework for higher-order separation logic which has already been used to reason about Wasm with Iris-Wasm [Rao et al. 2023] and Memory-Safe Wasm (MSWasm) [Michael et al. 2023] with Iris-MSWasm [Legoupil et al. 2024], among others. Once instantiated with our operational semantics, Iris provides us with proof rules that can be used to reason about the execution of programs, and we can derive instruction-specific rules to allow for mostly syntax-oriented verification proofs.

The central construct in our logic, used to give specifications, is our *extended weakest precondition*

$$\text{ewp } es; F \langle \Psi \rangle \{w \mid F, \Phi(w, F)\}$$

We define it by adding Wasm-specific components (see §4.2 and §4.4) to Hazel's extended weakest precondition [de Vilhena 2022; de Vilhena and Pottier 2021, 2023], which is itself built on the (plain) *weakest precondition* statement from the Iris logic, $\text{wp } es \{w, \Phi(w)\}$.

This weakest precondition statement can be read as 'the expression es runs safely, and if it terminates on a value w , then *postcondition* Φ holds of w '. For readability, we also write specifications under the form of a *Hoare triple* $\{P\} es \{w, \Phi(w)\}$, which means that as long as the *precondition* P holds before execution, es runs safely and if it terminates on a value w , then Φ holds of w . This Hoare triple is defined in terms of the weakest precondition, as $\Box(P \multimap \text{wp } es \{w, \Phi(w)\})$, where \multimap is a separating implication, and the persistent modality \Box means that this separating implication can be used any number of times, but cannot rely on non-persistent knowledge when it is proven.

Resources. In the above definition, P and $\Phi(w)$ are separation logic predicates that describe individual pieces of the program state. These can be thought of as (the conjunction of) *resources*, each representing *ownership* of part of the state. In Iris-Wasm [Rao et al. 2023], these represent individual pieces of the Wasm store, like function closures ($\text{addr} \xrightarrow{\text{wf}} \text{cl}$), globals ($\text{addr} \xrightarrow{\text{wg}} \text{gval}$), etc. The arrow label differentiates the various kinds of resources. To accommodate WasmFX, Iris-WasmFX defines two new resources: one for tags, and one for continuations.

Tag Resources. Tag resource $addr \xrightarrow{\text{wtag}} ft$ represents ownership of a tag. We often use fractional resources $addr \xrightarrow{\text{wtag}}_q ft$ ($q \in \mathbb{Q}$), as an easy way to share the resource between parts of a program.

Continuation Resources. The second resource we introduce represents exclusive ownership of a safe continuation and its type: $addr \xrightarrow{\text{wcont}} (ft, scont)$. We define a safe continuation as being either a dead continuation (which arise because WasmFX continuations are one-shot), or a *safe context*, which is either the context $[_] \mathrel{++} [\text{ref.func } \$f; \text{ref.call } ft]$ for some $\$f$ and ft , or a context of the form $[\text{frame}_n\{F\} H \text{ end}]$ (with H a handler context as defined in Figure 2). This means that in our program logic, we enforce that continuations always start off as a function call. If H is a safe context, then once plugged with any expression es , the resulting expression $H[es]$ reduces independently of the current frame (i.e. (S, F_1, es) reduces to (S', F'_1, es') if and only if (S, F_2, es) reduces to (S', F'_2, es')), since a `ref.call` reduces to a call no matter the frame, and reducing the body of a frame is always done with the frame F specified in the frame instruction itself rather than the frame F_1 or F_2 used when reducing the frame instruction. We show why this property is crucial for soundness when we introduce the proof rule for resume in §4.3.

Proof Rules. The Iris-Wasm program logic [Rao et al. 2023] for Wasm 1.0 defines proof rules for all Wasm instructions using the plain weakest precondition statement, most of which we inherit almost as-is. For example, the Iris-Wasm rule for updating a global variable is:

$$\frac{\text{WP-GLOBAL-SET} \quad \xrightarrow{\text{Fr}} F * F.\text{inst.globs}[\$g] = addr * addr \xrightarrow{\text{wg}} v_1}{\text{wp } [v_2; \text{global.set } \$g] \left\{ w, w = \text{immV } [] * \xrightarrow{\text{Fr}} F * addr \xrightarrow{\text{wg}} v_2 \right\}} \quad \text{in plain Iris-Wasm}$$

where `immV` is the constructor for *logical values* of Iris-Wasm:

Logical Values. In order to enable modular reasoning, Iris-Wasm defines several kinds of *logical values*, all representing expressions that do not reduce:

- Immediate values `immV vs` represent lists of Wasm values.
- Trap values `trapV` represent a trap failure value.
- Breaking values `brV i vh` represent `br i` expressions in a context vh syntactically enforced (by dependent typing, which we omit here) to be too shallow to break out from.
- Return values `retV sh` represent return expression in a context sh without a frame.
- Host call values `callhostV vs k tf llh` represent calls to host closures.

The reason that Iris-Wasm defines logical values more broadly than Wasm values is that it makes reasoning modular by making it possible to have a standard *bind rule* [Jung et al. 2018, Fig 13], which decomposes reasoning about a compound expression into reasoning about its parts in turn: reasoning about es in context L , that is, $L[es]$, can be decomposed into reasoning first about es itself, and then about the result w of es in context L . In Iris-Wasm, this rule looks like this:

$$\frac{\text{WP-BIND-LABEL} \quad \text{wp } es \{w, \text{wp } L[w] \{v, \Phi(v)\}\}}{\text{wp } L[es] \{v, \Phi(v)\}} \quad \text{in plain Iris-Wasm}$$

This rule mirrors reduction rule `REDUCE-LABEL`, and makes it possible to focus reasoning on a specific instruction so that one can, say, apply the proof rule specific to that instruction. Thanks to `brV i vh` being a value, one can focus on code that will get stuck: once it is reduced to the value $w = \text{brV } i \text{ } vh$, the postcondition must hold of w , which means there remains to prove a weakest precondition on $L[vh[\text{br } i]]$. If L contains the block targeted by `br i`, then one can apply Iris-Wasm's `WP-BR` rule to continue reasoning about the execution of the code after the jump. Otherwise, the expression is again a value `brV i (L ◦ vh)`, and one must now prove that Φ holds of that value.

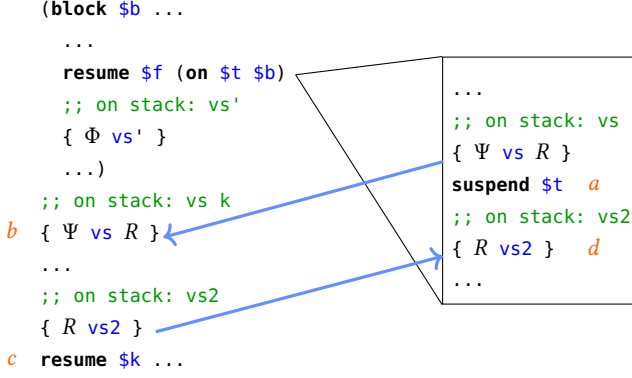


Fig. 4. The protocol Ψ relates the payload vs of the suspend (from a to b) to the resumption condition R : the condition under which the continuation can be resumed (from c to d).

This rule is sound in Iris-WasmFX *provided one limits* the evaluation context L to only contain label and not prompt, since non-local control flow breaks the bind rule [Timany and Birkedal 2019]. Binding into a prompt requires a more intricate rule, as we show in §4.3 and Appendix B.

Frame. Iris-Wasm features a resource $\xrightarrow{\text{Fr}} F$ that symbolises ownership of the frame. Instead, our custom weakest precondition considers the frame as part of the expression rather than part of the state. As this is both orthogonal to our main contribution and can be used intuitively, we describe it in detail in Appendix B, and merely use our new approach in the rest of the paper.

4.2 Hazel

Motivation. When dealing with effects, the standard weakest precondition definition of Iris does not provide a satisfying level of modularity. The problem is caused by the suspend instruction, which translates to `suspend.addr` as per rule `REDUCE-SUSPEND-TRANSLATE`, and is then stuck: it is neither a value, nor can it take a step *by itself*, only in the context of a prompt, as per rule `REDUCE-SUSPEND`. Contrary to `br`, the enclosing prompt handling the effect might be in a different function or even a different module, and hence making suspend a logical value as is done for `br` in Iris-Wasm would require knowledge of the entire handling environment and break modularity.

In our running example, this would mean that to reason about the client functions that invoke `$yield`, and thereby suspend, one would have to consider the code of `$par`, breaking modularity.

Extended Weakest Precondition. The approach of Hazel [de Vilhena 2022; de Vilhena and Pottier 2021, 2023] is to extend the weakest precondition with a new component that locally accounts for non-local effects. Hazel’s *extended weakest precondition* $\text{ewp } es \langle \Psi \rangle \{w, \Phi(w)\}$ means not only (as usual) that es runs safely, and that if it terminates on a value w , then postcondition Φ holds of w , but also (as new in Hazel) that if es performs an effect, then that effect *follows* protocol Ψ . That protocol relates the effect’s payload (some values vs) with a resumption condition [Liu et al. 2023] (a predicate on values) describing the desired behaviour of the program when the suspended continuation is later resumed, as illustrated in Figure 4. As such, the type of a protocol is $vs \rightarrow (vs \rightarrow \text{iProp}) \rightarrow \text{iProp}$, where iProp is the type of Iris propositions. $\Psi \vee R$ is thus a proposition that should hold whenever an effect is performed (site a in Figure 4) with payload v , and $R(w)$ is expected to hold when the continuation is later resumed (site c in Figure 4) with argument w . This resumption condition must

be established when suspending (site *a* in Figure 4) and can be relied upon when handling (site *b* in Figure 4). We give a few examples below of what protocols can look like in practice.

Logical Effects. In Iris-WasmFX, suspend is not a logical value (in fact, we retain the same definition of logical value as Iris-Wasm), but a *logical effect*, defined to be of the form suspendE, switchE or throwE. For instance, effect suspendE *vs addr H* represents expression $H[\text{suspend.addr } vs \text{ addr}]$ where H is syntactically enforced to not contain a prompt that captures the tag with address *addr* (otherwise, the expression would not be stuck). Similarly, switchE and throwE represent a switch.addr or a throw (from the exception handling suite) under a sufficiently shallow context.

Our Extended Weakest Precondition. In Iris-WasmFX, each tag gets its own protocol, so our extended weakest precondition uses a *metaprotocol* Ψ instead of a simple protocol. A metaprotocol is defined as the combination of three maps: the first is used for suspend effects, on which we focus now, and reused for switch effects; the second only for switch effects, which we return to in §4.4; and the last for exception handling. The first map in a metaprotocol maps tag addresses to protocols. In this section, when we write $\Psi \text{ addr}$, we mean application of the first map in metaprotocol Ψ .

Hence, our custom extended weakest-precondition is of the form

$$\text{ewp } es ; F \langle \Psi \rangle \{ w F', \Phi(w, F') \}$$

where the Wasm frame F comes alongside the expression *es*, the metaprotocol Ψ is given to specify the intended behaviour in case of an effect being performed, and the postcondition Φ is a predicate on final values w and final frames F' .

Practical Protocols. While in principle a protocol can be any predicate at all, de Vilhena and Pottier [2021] define a specific constructor inspired by session types [Honda et al. 1998] that is sufficient for many practical examples: $!x(v_1)\langle P \rangle ?y(v_2)\langle Q \rangle$. Informally, this says “the suspender sends payload v_1 and resources P to the handler; if execution is ever resumed, it expects to receive return value v_2 and resources Q ”. Formally, this protocol applied to values *vs* and predicate Φ is

$$\exists x, (vs = v_1 * P * \forall y, (Q \multimap \Phi(v_2)))$$

where x , which can be used to describe the sent values, can be mentioned in v_1 , P , v_2 and Q ; and y , which can be used to describe the received values, can be mentioned in v_2 and Q .

Another useful protocol is the always false ‘bottom’ protocol \perp , which symbolises an effect that should not be performed. In practice, we often use at top-level a metaprotocol where all tags map to the bottom protocol, and only start introducing non-bottom protocols when reasoning about continuations run by way of a resume. As we will see in the next section, the reasoning rule for resume lets the verifier use a different protocol in the premises than in the conclusion.

Running Example Protocol. For our coroutine example, the tag does not expect any values (and this is enforced by typing), but the resumption condition is that the invariant has been restored. Concretely, the metaprotocol Ψ_0 is the bottom protocol \perp for all tags, except the tag address corresponding to local index $\$t$, for which the protocol is $!() \{I\} ?() \{I\}$.

Definition. We define the extended weakest precondition for Iris-WasmFX formally in Figure 5. We give an informal reading before explaining the more technical elements.

- When *es* is a value, we check (as usual) that the postcondition holds.
- When *es* is a suspend effect, that is, a suspend.addr with some values *vs* and a tag address *addr*, under a (capture-avoiding) context H , we mandate that the protocol $\Psi(\text{addr})$ must be followed. If that protocol is of the form $!x(v)\langle P \rangle ?y(w)\langle Q \rangle$, this means there must exist an x such that the payload $\text{immV } vs$ is equal to v and we have the resources P ; and we

$$\begin{aligned}
& \text{ewp } es; F \langle \Psi \rangle \{w F, \Phi(w, F)\} = \text{match to_val}(es) \text{ with} \\
\text{value} & \quad | \text{Some } w \implies \models \Phi(w, F) \\
& \quad | \text{None} \implies \text{match to_eff}(es) \text{ with} \\
\text{effect: suspend} & \quad | \text{Some (suspendE vs addr H)} \implies \\
& \quad \quad \uparrow (\Psi \text{ addr})(\text{immV vs}) (\lambda w, \triangleright \text{ewp } H[w]; F \langle \Psi \rangle \{w F, \Phi(w, F)\}) \\
\text{switch} & \quad | \text{Some (switchE vs kaddr ft taddr H)} \implies \text{see §4.4} \\
\text{neither} & \quad | \text{None} \implies \forall S, \text{state_interp}(S) \Rightarrow \\
& \quad (\exists S' F' es', (S, F, es) \hookrightarrow (S', F', es')) * \\
& \quad \forall S' F' es', (S, F, es) \hookrightarrow (S', F', es') \Rightarrow \\
& \quad \triangleright (\text{state_interp}(S') * \text{ewp } es'; F' \langle \Psi \rangle \{w F, \Phi(w, F)\})
\end{aligned}$$

Fig. 5. The extended weakest precondition of Iris-WasmFX

must guarantee that as long as we are provided with a y such that we are given return value w and the resources Q , we can continue execution safely, i.e. we have a weakest precondition statement for $H[w]$. This is the point that motivated moving the payload into the `suspend.addr` instruction, as per §3.3: inspecting the expression without knowledge of the store needs to be enough to identify that it is an effect and what its payload is.

- We explain the case where es is a switch effect in §4.4
- When es is neither a value nor an effect, we require (as usual) that for any *valid* store S , configuration (S, F, es) is reducible, and for any configuration (S', F', es') that it reduces to, S' is still valid, and the extended weakest precondition holds of (F', es') .

Formal Notation. In Figure 5, the update modality \Rightarrow means that Iris ghost state can be updated – a reader unfamiliar with Iris can think of it as a regular implication. The later modality $\triangleright P$ means that P holds after one execution step; this is crucial to avoid cyclicity in the recursive definition, but can be ignored by readers unfamiliar with Iris. Functions `to_val` and `state_interp` are language-specific functions present in the language-agnostic definition of a weakest precondition in Iris. The first attempts to transform an expression into a value and returns `None` if it cannot, and the second describes what a *valid* store is. We use mostly unchanged definitions from Iris-Wasm. Third, the function `to_eff` is a language-specific function in Hazel’s definition of the extended weakest precondition. It attempts to transform an expression into an effect and returns `None` if the expression is not an effect. We provide our own definition for it in the context of WasmFX. Finally, the upwards closure of the protocol $\uparrow \Psi$ allows strengthening the resumption condition ($\uparrow \Psi \vee \Phi$ is defined as $\exists \Phi', \Psi \vee \Phi' * (\forall w, \Phi'(w) \multimap \Phi(w))$), which helps in practice but can be ignored for simplicity [de Vilhena and Pottier 2021, §3.4].

4.3 Proof Rules

In this section, we show the reasoning rules for the new instructions of WasmFX. We focus our presentation on the most crucial rules; see our Rocq development for a full view of the logic, including failure rules, exception handling, the `cont.bind` and `resume_throw` instructions, etc.

Building Continuations. The proof rule for `cont.new` is shown in Figure 6. It mirrors its reduction rule: using the frame to translate the function type parameter of the `cont.new` function, the rule allows reasoning about creating a continuation from a function reference. This rule is the only one that creates a new continuation resource from scratch.

EWP-CONTNEW

$$\frac{F.\text{inst.types}[i] = ft}{\text{ewp} [\text{ref.func } addr; \text{cont.new } i] ; F \langle \Psi \rangle \left\{ \begin{array}{l} \exists kaddr, w = \text{immV} [\text{ref.cont } kaddr] * F' = F * \\ w F', \quad kaddr \xrightarrow{\text{wcont}} \\ (ft, [_]) ++ [\text{ref.func } addr; \text{ref.call } ft] \end{array} \right\}}$$

EWP-SUSPEND

$$\frac{\begin{array}{l} |vs| = |ts_1| * addr \xrightarrow{\text{wtag}}_q (ts_1 \rightarrow ts_2) * F.\text{inst.tags}[i] = addr * \\ \triangleright (addr \xrightarrow{\text{wtag}}_q (ts_1 \rightarrow ts_2) \multimap \uparrow (\Psi \text{ addr}) (\text{immV } vs) (\lambda v, \triangleright \Phi v F)) \end{array}}{\text{ewp } vs ++ [\text{suspend } i] ; F \langle \Psi \rangle \{v F, \Phi(v, F)\}}$$

EWP-RESUME

$$\frac{\begin{array}{l} (1) \quad F.\text{inst.types}[i] = ts_1 \rightarrow ts_2 * |vs| = |ts_1| * F \vdash ccs \rightarrow dcs * \\ (2) \quad (\forall addr, addr \notin dcs \implies \Psi(addr) = \Psi'(addr)) * addr \xrightarrow{\text{wcont}} (ts_1 \rightarrow ts_2, H) * \\ (3) \quad (\forall F, \neg \Phi(\text{trapV}, F)) * \neg \Phi'(\text{trapV}) * \triangleright \text{ewp } H[vs] ; \emptyset \langle \Psi \rangle \{w F, \Phi(w, F)\} * \\ (4) \quad \triangleright (\forall w, \Phi(w, \emptyset) \multimap \text{ewp} [\text{prompt}_{ts_2} \{dcs\} \text{ w end}] ; \emptyset \langle \Psi' \rangle \{v _, \Phi'(v)\}) * \\ (5) \quad \triangleright (\forall dcc \in dcs, \langle \Psi \rangle \{ \Phi \} dcc; ts_2 \langle \Psi' \rangle \{ \lambda v _, \Phi'(v) \}) \end{array}}{\text{ewp } vs ++ [\text{ref.cont } addr; \text{resume } i \text{ ccs}] ; F \langle \Psi' \rangle \{v F', \Phi'(v) * F' = F\}}$$

Fig. 6. The reasoning rule for the WasmFX instructions. See §4.3 and §4.4 for the definition of (5)

Suspend. The rule for suspend is shown in Figure 6. It merely appeals to the protocol corresponding to the tag. The last premise indicates that if given back the tag resource, the payload vs should follow the specified protocol, and if the continuation is ever resumed, postcondition Φ will hold. To understand what this means in practice, let us consider the case where protocol $\Psi(addr)$ has form $!x(v)\{P\} ?y(w)\{Q\}$. In that case, we must show that $\exists x, vs = v * P * \forall y, Q \multimap \Phi(w, F)$, i.e. that there exists an x such that payload vs is equal to v , we have resource P and for all y , having Q is enough to guarantee the postcondition $\Phi(w, F)$. Note that the frame F in $\Phi(v, F)$ is the same as the one in the conclusion of the rule: when the continuation is resumed and control flow comes back to the current point, we are back in the same environment and should thus use the same frame F as before the suspend.

Running Example. We can now describe our proof of `$yield`: we want to show

$$\forall F I q, I \multimap \text{addyield} \xrightarrow{\text{wf}} cl_{\text{yield}} \multimap \text{addrt} \xrightarrow{\text{wtag}}_q ([_] \rightarrow [_]) \multimap$$

$$\text{ewp} [\text{invoke } \text{addyield}] ; F \langle \Psi_0 \rangle \left\{ w F', \quad \begin{array}{l} w = \text{immV } [_] * F = F' * I * \\ \text{addyield} \xrightarrow{\text{wf}} cl_{\text{yield}} * \text{addrt} \xrightarrow{\text{wtag}}_q ([_] \rightarrow [_]) \end{array} \right\}$$

where $\text{addyield} = F.\text{inst.funcs}[\$yield]$, $\text{addrt} = F.\text{inst.tags}[\$t]$, and cl_{yield} is the closure of `$yield`. We phrase this specification using our extended weakest precondition, with metaprotocol Ψ_0 . We begin by applying Iris-Wasm rules to unfold the function invocation and bind onto the code of the function. This requires using the function closure resources, and gives it back. We show this in detail in Appendix B. Next, we reason about the code of the function itself, which consists of a single suspend for which rule EWP-SUSPEND requires the tag resource, which we have, and it remains to show that when given back the tag resource, protocol $\Psi_0(\$t)$ is followed. This protocol requires us, for the sending, to show that the payload is the empty list, which is straightforward, and to provide I , which is the precondition; and for the receiving, to show that the postcondition for $\text{immV } [_]$ and F can be established from I . Since we hold both resources, it is easy to conclude.

Resume. We show the rule for the resume instruction in Figure 6. The premises can be divided into three parts. First, we want the code being resumed to be safe to run. This is rendered with an extended weakest precondition on line (3). The H in that premise is tied by the continuation resource relinquished (to symbolise that the continuations are one-shot) on line (2) to the address $addr$ that resume took as an argument from the stack, and per line (1) the right number of values vs are taken from the stack. Note that the metaprotocol used on line (3) can differ from the one in the conclusion (though, as per line (2), only for tags handled by the resume): this is because the effects handled by the clauses of the resume can now be triggered in the resumed code.

Second, line (4) covers the ‘normal return’ case, when the continuation terminates on a value w . We know that in that case w satisfies postcondition Φ , and must show that this is enough to safely plug w into the prompt environment the resumed continuation was running in.

Third, line (5) covers the effect case: for each clause in the list of clauses specified by the resume, we must show that the effect handled in the clause is safe to run. We capture this by defining a *clause triple* which each clause has to respect. We explain how this clause triple is defined for a suspend effect now, and we come back to the switch effect in §4.4.

Clause Triple. When a suspend is performed, reduction rule REDUCE-SUSPEND mandates that a continuation is created out of the state of the stack, and the handler reduces to a br instruction with the label specified by the clause from the resume instruction, with the effect’s payload and the newly created continuation on the stack. Hence, we wish to require that in that case, we hold an extended weakest precondition statement for $vs \mathrel{++} [\text{ref.cont } kaddr; \text{br } ilab]$, where $kaddr$ is the address of the newly created continuation. Little is known about the form of that continuation, except that the suspend instruction will have followed a protocol from metaprotocol Ψ . Accordingly, the clause triple for a suspend (which closely follows that of Hazel [de Vilhena and Pottier 2021, Fig. 7]) requires us to prove the extended weakest precondition statement on the br using a continuation resource for the new continuation, and knowledge that the effect was performed in a way that follows the protocol Ψ $taddr$. Since the continuation H is universally quantified, the only way to reason about resuming the continuation again is by knowing that the protocol is followed:

$$\begin{aligned} &\langle \Psi \rangle \{ \Phi \} \text{ on } taddr \text{ } ilab; \text{ } ts \langle \Psi' \rangle \{ \Phi' \} = \\ &\quad \exists ts_1 \text{ } ts_2 \text{ } q, \text{ } taddr \xrightarrow{\text{wtag}}_q (ts_1 \rightarrow ts_2) * \\ &\quad \forall vs \text{ } kaddr \text{ } H, \text{ } kaddr \xrightarrow{\text{wcont}} (ts_2 \rightarrow ts, H) \multimap taddr \xrightarrow{\text{wtag}}_q (ts_1 \rightarrow ts_2) \multimap \\ &\quad \uparrow (\Psi \text{ } taddr) (\text{immV } vs) (\lambda w, \triangleright \text{ewp } H[w]; \emptyset \langle \Psi \rangle \{ w \text{ } F, \Phi(w, F) \}) \multimap \\ &\quad \text{ewp } vs \mathrel{++} [\text{ref.cont } kaddr; \text{br } ilab]; \emptyset \langle \Psi' \rangle \{ w \text{ } F, \Phi'(w, F) \} \end{aligned}$$

Frame Agnosticity. The last two lines of the definition of the clause triple refer to an empty frame. According to the operational semantics, the frame to be used in this position is the frame as it would be at the suspension site (for the last line) and re-resumption site (for the penultimate line), but when scrutinising a resume instruction, the frame at those later times is unknown. Fortunately, the choice of the frame to use at these points is irrelevant since continuations always reduce in a frame-agnostic way, as mentioned in §4.1. Therefore, we can choose whatever frame we want. (This does not work for exceptions, so we define much stronger rules for the exception handling suite and the resume_throw instruction. We leave it as a future work to refine these rules, perhaps using adjoint separation logic [Wagner et al. 2025].) We find it convenient to use the empty frame everywhere (including lines (3–4) in EWP-RESUME), so that all the frames match.

To link the frame F in the conclusion of EWP-RESUME with the empty frame in premise (3), we show that if es reduces frame-agnostically and is safe to run under the empty frame, then it is safe to run with any frame, ends with an unchanged frame, and the same post-condition.

Proving the Resume Rule. To prove soundness of the EWP-RESUME rule above, we unfold the operational semantics one step, which reduces resume to a prompt instruction. As discussed, binding into a prompt using the simple EWP-LABEL rule would be unsound, hence we define a bind rule specific to prompt, as we discuss in Appendix D. The premises of this rule are similar to those in EWP-RESUME: a premise for the body, a premise for the ‘normal return’ case, and a premise for the handled effects using clause triples.

4.4 Reasoning About the Switch Instruction

Given that the operational semantics for the switch instruction is intuitively just a suspension immediately followed by the resumption of a different continuation, one might expect that the insertion of switch into the logic would be straightforward and that the corresponding rule would merely be a combination of the other rules. However, this is not the case, because we want to reason at a different site. When reasoning about separate resume and suspend, what executes after the suspend is static: it is determined by the code around the resume and by the values it gets passed from the suspend, which is described by the protocol. On the other hand, with a switch, what executes next is dynamic: it is determined by the argument to the switch, which we want, for modularity, to be able to reason about at the site of switch, not at the site of the resume.

To address this, we use an extra component in our metaprotocols. For tags used for switch, the metaprotocol specifies not only a protocol, but also a predicate on continuations that the switch instruction is meant to switch to. We write $\Psi^1 \text{ taddr}$ and $\Psi^2 \text{ taddr}$ respectively for the protocol and for the predicate. Using this, we now complete the definition of our extended weakest precondition from Figure 5 with the case for switch:

$$\begin{aligned} & | \text{Some} (\text{switchE } vs \text{ kaddr } ft \text{ taddr } H) \implies \exists H' \, ts_1 \, ts_2 \, ft' \, ts \, q, \\ & \quad \text{taddr} \xrightarrow{\text{wtag}}_q [] \rightarrow ts * \text{kaddr} \xrightarrow{\text{wcont}} (ft, H') * ft' = ts_1 \rightarrow ts * \\ & \quad ft = (ts_1 ++ [ft'.\text{cont}]) \rightarrow ts_2 * \Psi^2 \text{ taddr } H' * \\ & \quad (\text{taddr} \xrightarrow{\text{wtag}}_q [] \rightarrow ts \multimap \uparrow (\Psi^1 \text{ taddr}) \text{ vs } (\lambda w, \triangleright \text{ewp } H[w]; F \langle \Psi \rangle \{ \Phi \})) \end{aligned}$$

In order to switch, we need to give up a tag resource and a continuation resource corresponding to the continuation we are switching to; we need to know that the types have the required form; and that the metaprotocol is being followed, meaning that the continuation satisfied the required predicate, and when given back the tag resource, the payload follows the protocol. The continuation resource is not given back, because WasmFX continuations are one-shot.

The proof rule for the switch instruction itself simply mirrors the definition above.

EWP-SWITCH

$$\begin{aligned} & ft = ts_1 \rightarrow ts * F.\text{inst.types}[i] = ft' * ft' = (ts_1 ++ [ft.\text{cont}]) \rightarrow ts_2 * F.\text{inst.tags}[i'] = \text{taddr} * \\ & | vs | = | ts_1 | * \text{taddr} \xrightarrow{\text{wtag}}_q [] \rightarrow ts * \text{kaddr} \xrightarrow{\text{wcont}} (ft', H) * \Psi^2 \text{ taddr } H * \\ & \quad \triangleright (\text{taddr} \xrightarrow{\text{wtag}}_q [] \rightarrow ts \multimap \uparrow (\Psi^1 \text{ taddr}) \text{ vs } \Phi) \\ & \hline & \text{ewp } vs ++ [\text{ref.cont } \text{kaddr}; \text{switch } i \, i']; F \langle \Psi \rangle \{ \Phi \} \end{aligned}$$

Recall that the ‘effect’ premise of the EWP-RESUME rule uses the clause triple to require that for all clauses, if an effect is triggered targeting that clause, continuing execution is safe. We can now give the following definition of a clause triple for a switch clause:

$$\begin{aligned} & \langle \Psi \rangle \{ \Phi \} \text{ on } \text{taddr } \text{switch}; ts \langle \Psi' \rangle \{ \Phi' \} = \\ & \quad \exists q, \text{taddr} \xrightarrow{\text{wtag}}_q [] \rightarrow ts * \\ & \quad \square \forall vs \text{ kaddr } H \, H' \, ts_1, \text{kaddr} \xrightarrow{\text{wcont}} (ts_1 \rightarrow ts, H) \multimap \Psi^2 \text{ taddr } H' \multimap \\ & \quad \quad \uparrow (\Psi^1 \text{ taddr}) (\text{immV } vs) (\lambda w, \triangleright \text{ewp } H[w]; \emptyset \langle \Psi \rangle \{ w \, F, \Phi(w, F) \}) \multimap \\ & \quad \quad \text{ewp } H' [vs ++ [\text{ref.cont } \text{kaddr}]]; \emptyset \langle \Psi \rangle \{ w \, F, \Phi(w, F) \} \end{aligned}$$

In the case of a `switch`, the operational semantics mandates that execution continues by running the continuation given to `switch` as an argument with the payload taken from the stack and a newly created continuation corresponding to the current environment. Hence we must establish an extended weakest precondition for running that new continuation, when given an appropriate payload and continuation. The \square ensures duplicability, because the new continuation is run under the same clauses and hence we need the clause triple to hold again for the new continuation.

We have used these rules to reason about simple examples using `switch`, these can be found in our Rocq development. We note that realistic programs using the `switch` instructions typically require recursive types, which we do not support. We leave this as future work, as it is orthogonal to supporting the WasmFX operations themselves.

4.5 Soundness and Adequacy

All proof rules showed in this paper have been proved sound in the Rocq proof assistant; that is, we show that holding all premises is sufficient to obtain the extended weakest precondition statement in the conclusion, as defined in Figure 5. Our Rocq development also includes proof rules for every instruction of WasmFX.

THEOREM 4.1 (SOUNDNESS). *All proof rules displayed in this paper have been proven sound with respect to WasmFXCert in the Rocq proof assistant.*

We also prove *adequacy*, which allows us to extract, from each Iris-WasmFX proof of a pure extended weakest precondition, a theorem phrased purely in terms of the operational semantics, making no mention of Iris. The proof is similar to that of plain Iris [Jung et al. 2018, §6.4].

THEOREM 4.2 (ADEQUACY). *If $\text{state_interp}(S) * \text{ewp es}; F \langle \perp \rangle \{w F, \Phi(w, F)\}$ for some pure predicate Φ (meaning Φ does not contain any resourceful Iris propositions), then for any reduction trace $(S, F, \text{es}) \hookrightarrow^* (S', F', \text{vs})$ (where vs are values), it holds that $\Phi(\text{vs}, F')$.*

5 Case Studies

In this section, we briefly consider two canonical features with which we exemplify the kind of modular reasoning that Iris-WasmFX enables. Coroutines illustrate switching of stacks, where the tag is internal to the library, and passes through the client. Generators illustrate passing of values with tags, where the tag is used as the interface between the library and client to *send* values from the generator to the client. Full details are given in Appendix E.

Coroutines. We now precisely state a modular specification for our coroutine library using the protocol of Section 4.2, and reason about the client module with respect to that specification, as desired, and use adequacy to derive a statement in terms of the operational semantics, without reference to Iris:

THEOREM 5.1. *If Cor is the coroutine module whose code is shown in Figure 3, and Cl is the message-passing client module, and host program `[instantiate Cor; instantiate Cl; invoke addrmain]` terminates on a value w , then w is of the form $w = \text{immV } [\text{vala}; \text{valb}]$, and if vala is 1, then valb is 42.*

Generators. Phipps-Costin et al. [2023] use effects to implement a generator of the naturals: 0, 1, 2, ... The generator consists of an infinite loop which suspends with the value of the current counter, and then increments it. Every time a client resumes the generator, they get the next natural number. They use this generator to implement a function `$sum_until`, which adds all naturals from 0 to a given number. Again, we give a modular proof: we first verify the generator, giving it a generic specification, and then verify the client in terms of that specification, to show:

THEOREM 5.2. *For any i32 n , if Gen is the generator module containing the generator of the naturals and the `$sum_until` function, and if the host program `[instantiate Gen; invoke addrsum_until n]` terminates on a value w , then $w = \text{immV} \left[\sum_{i=0}^n i \right]$ (in i32 arithmetic).*

6 Related Work

Iris-WasmFX is, to our knowledge, the first program logic for WasmFX, and as such pulls together several strands of research:

Formalising WasmFX. There is an ongoing project [Liang et al. 2025] to model WasmFX in SpecTec [Youn et al. 2024], a domain-specific language designed specifically to write the semantics of Wasm with tooling to export theorem prover definitions. SpecTec has been officially adopted for authoring the Wasm specification [Rossberg 2025] (from 3.0 on), and could eventually replace the hand-written definitions of WasmFXCert. However, this would require more mature theorem prover output, and would in any case not replace the proofs of §3.4. Moreover, it raises the question of how to handle the modifications we make to the language definition for the sake of compatibility with the theorem prover and program logic, as described in §3.3.

Program Logics for Effect Handlers. Iris-WasmFX applies the ideas of Hazel, a program logic for an ML-style calculus with effect handlers, to a full-fledged programming language, Wasm. Osiris [Daby-Seesaram et al. 2023; Seassau et al. 2025a,b] is a program logic for OLang, a substantial fragment of another full-fledged programming language that features effect handlers, namely sequential OCaml 5.3. Osiris has a sibling logic, Horus, for pure expressions, that makes reasoning simpler for them, and is compatible with Osiris; we have not explored this angle for Wasm, as it is much more imperative, but the idea could nevertheless prove useful. The two projects differ in the challenges they face: one of the main sources of complexity of OLang is the loose evaluation order of OCaml, which Wasm was designed to avoid; on the other hand, the functional style of OLang integrates with effect handlers in a more human-readable way than in WasmFX. They also differ in their implementations: OLang is defined by elaboration into a monad, whereas WasmFXCert follows the usual operational semantics style, as used in the Wasm standard. The combined existence of Osiris and Iris-WasmFX opens up the possibility of formally relating programs in OCaml and Wasm, be it by verified compilation, translation validation, etc.

Integration with JavaScript. The extensive use of non-local control flow in JavaScript/ECMAScript is another motivation for extending Wasm with WasmFX. Khayam et al. [2022] formalise parts of ECMAScript and describe how, even though they do not model generators and asynchronous function definitions, they identified corner cases of the language semantics to do with manipulation of the evaluation context, showcasing the challenges raised by non-local control flow, and how precise definitions and formal verification can be helpful in this context.

JavaScript Promise Integration (JSPI) [McCabe et al. 2024] is the ‘reverse’ of WasmFX: it allows plain Wasm, not written to be asynchronous, to work with host code (in JavaScript) that uses promises. WasmFXCert is the first building block to consider the semantics of Wasm running in a host language with its own non-local control flow features, and to study their interaction. In particular, part of the contract between JavaScript and Wasm is that Wasm is only meant to suspend or resume the execution of JavaScript via promises. Verifying this is outside of the scope of this paper, but we lay the foundations to do so.

Compilation to WasmFX. RichWasm [Fitzgibbons et al. 2024] is a richly typed language based on Wasm, designed as a target for typed compilation enforcing strong memory safety guarantees. Extending RichWasm to cover WasmFX, as modelled by WasmFXCert, could offer a useful staging post in proving correctness of the compilation of non-local control features into WasmFX.

7 Conclusion

By developing WasmFXCert and Iris-WasmFX, we have validated the design of the WasmFX proposal in time to feed back into its development. In addition, with Iris-WasmFX, we have laid the foundations for verifying effect-based libraries.

Acknowledgments

Ohad Kammar, for helpful discussions. Andreas Rossberg, for comments on a draft of this paper. Amal Ahmed, for ... JP: AUFF LB: Villum???

References

- Heejin Ahn and Ben Titser. 2022. *Exception Handling Proposal for WebAssembly*. Technical Report. <https://webassembly.github.io/exception-handling/>
- Stephen Brookes and Peter W. O’Hearn. 2016. Concurrent separation logic. *ACM SIGLOG News* 3, 3 (2016), 47–65. doi:10.1145/2984450.2984457
- Arnaud Daby-Seesaram, François Pottier, and Armaël Guéneau. 2023. Osiris: an Iris-based program logic for OCaml. OCaml’23: Caml Users and Developers Workshop 2023. <https://icfp23.sigplan.org/details/ocaml-2023-papers/7/Osiris-an-Iris-based-program-logic-for-OCaml>.
- Paulo de Vilhena. 2022. *Proof of Programs with Effect Handlers. (Preuve de Programmes avec Effect Handlers)*. Ph. D. Dissertation. Paris Cité University, France. <https://tel.archives-ouvertes.fr/tel-03891381>
- Paulo Emilio de Vilhena and François Pottier. 2021. A separation logic for effect handlers. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–28. doi:10.1145/3434314
- Paulo Emilio de Vilhena and François Pottier. 2023. Verifying an Effect-Handler-Based Define-By-Run Reverse-Mode AD Library. *Log. Methods Comput. Sci.* 19, 4 (2023). doi:10.46298/LMCS-19(4:5)2023
- Edsger W. Dijkstra. 1968. Letters to the editor: go to statement considered harmful. *Commun. ACM* 11, 3 (1968), 147–148. doi:10.1145/362929.362947
- Matthias Felleisen. 1988. The Theory and Practice of First-Class Prompts. In *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages, San Diego, California, USA, January 10-13, 1988*, Jeanne Ferrante and Peter Mager (Eds.). ACM Press, 180–190. doi:10.1145/73560.73576
- Jean-Christophe Filliâtre and Mário Pereira. 2016. A Modular Way to Reason About Iteration. In *NASA Formal Methods - 8th International Symposium, NFM 2016, Minneapolis, MN, USA, June 7-9, 2016, Proceedings (Lecture Notes in Computer Science, Vol. 9690)*, Sanjai Rayadurgam and Oksana Tkachuk (Eds.). Springer, 322–336. doi:10.1007/978-3-319-40648-0_24
- Michael Fitzgibbons, Zoe Paraskevopoulou, Noble Mushtak, Michelle Thalakkottur, Jose Sulaiman Manzur, and Amal Ahmed. 2024. RichWasm: Bringing Safe, Fine-Grained, Shared-Memory Interoperability Down to WebAssembly. *Proc. ACM Program. Lang.* 8, PLDI (2024), 1656–1679. doi:10.1145/3656444
- Andreas Haas, Andreas Rossberg, Derek L. Schuff, Ben L. Titser, Michael Holman, Dan Gohman, Luke Wagner, Alon Zakai, and J. F. Bastien. 2017. Bringing the web up to speed with WebAssembly. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017*, Albert Cohen and Martin T. Vechev (Eds.). ACM, 185–200. doi:10.1145/3062341.3062363
- Daniel Hillerström. 2021. Typed Continuations: A Basis for Stack-switching in Wasm. <https://raw.githubusercontent.com/WebAssembly/meetings/main/stack/2021/presentations/2021-9-20-TypedContinuations.pdf>
- Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *Programming Languages and Systems - ESOP’98, 7th European Symposium on Programming, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS’98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings (Lecture Notes in Computer Science, Vol. 1381)*, Chris Hankin (Ed.). Springer, 122–138. doi:10.1007/BFB0053567
- Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-order ghost state. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*. 256–269. doi:10.1145/2951913.2951943
- 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. doi:10.1017/S0956796818000151
- Adam Khayam, Louis Noizet, and Alan Schmitt. 2022. A Faithful Description of ECMAScript Algorithms. In *PPDP 2022: 24th International Symposium on Principles and Practice of Declarative Programming, Tbilisi, Georgia, September 20 - 22, 2022*. ACM, 8:1–8:14. doi:10.1145/3551357.3551381

- Maxime Legoupil, June Rousseau, Aina Linn Georges, Jean Pichon-Pharabod, and Lars Birkedal. 2024. Iris-MSWasm: Elucidating and Mechanising the Security Invariants of Memory-Safe WebAssembly. *Proc. ACM Program. Lang.* 8, OOPSLA2 (2024), 304–332. doi:10.1145/3689722
- Yalun Liang, Sam Lindley, and Andreas Rossberg. 2025. Experience Report: Stack Switching in Wasm SpecTec. presented at WAW: the WebAssembly Workshop <https://effect-handlers.org/static/papers/sss-waw-2025.pdf>.
- Zongyuan Liu, Sergei Stepanenko, Jean Pichon-Pharabod, Amin Timany, Aslan Askarov, and Lars Birkedal. 2023. VMSL: A Separation Logic for Mechanised Robust Safety of Virtual Machines Communicating above FF-A. *Proc. ACM Program. Lang.* 7, PLDI (2023), 1438–1462. doi:10.1145/3591279
- David Madore. 2001. The Unlambda Programming Language. <http://www.madore.org/~david/programs/unlambda/>.
- Luc Maranget, Susmit Sarkar, and Peter Sewell. 2012. A Tutorial Introduction to the ARM and POWER Relaxed Memory Models. <http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test7.pdf>.
- Francis McCabe, Sam Lindley, and WebAssembly Community Group. 2025. Stack Switching WebAssembly proposal. <https://github.com/WebAssembly/stack-switching/blob/main/proposals/stack-switching/Explainer.md>.
- Francis McCabe, Thibaud Michaud, Ilya Rezvov, and Brendan Dahl. 2024. Introducing the WebAssembly JavaScript Promise Integration API. <https://v8.dev/blog/jsapi>
- Alexandra E. Michael, Anitha Gollamudi, Jay Bosamiya, Evan Johnson, Aidan Denlinger, Craig Disselkoen, Conrad Watt, Bryan Parno, Marco Patrignani, Marco Vassena, and Deian Stefan. 2023. MSWasm: Soundly Enforcing Memory-Safe Execution of Unsafe Code. *Proc. ACM Program. Lang.* 7, POPL (2023), 425–454. doi:10.1145/3571208
- Luna Phipps-Costin, Andreas Rossberg, Arjun Guha, Daan Leijen, Daniel Hillerström, K. C. Sivaramakrishnan, Matija Pretnar, and Sam Lindley. 2023. Continuing WebAssembly with Effect Handlers. *Proc. ACM Program. Lang.* 7, OOPSLA2 (2023), 460–485. doi:10.1145/3622814
- Gordon D. Plotkin and Matija Pretnar. 2009. Handlers of Algebraic Effects. In *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22–29, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5502)*, Giuseppe Castagna (Ed.). Springer, 80–94. doi:10.1007/978-3-642-00590-9_7
- Gordon D. Plotkin and Matija Pretnar. 2013. Handling Algebraic Effects. *Log. Methods Comput. Sci.* 9, 4 (2013). doi:10.2168/LMCS-9(4:23)2013
- François Pottier. 2017. Verifying a hash table and its iterators in higher-order separation logic. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs, CPP 2017, Paris, France, January 16–17, 2017*, Yves Bertot and Viktor Vafeiadis (Eds.). ACM, 3–16. doi:10.1145/3018610.3018624
- Xiaojia Rao, Aina Linn Georges, Maxime Legoupil, Conrad Watt, Jean Pichon-Pharabod, Philippa Gardner, and Lars Birkedal. 2023. Iris-Wasm: Robust and Modular Verification of WebAssembly Programs. *Proc. ACM Program. Lang.* 7, PLDI (2023), 1096–1120. doi:10.1145/3591265
- Andreas Rossberg. 2019. WebAssembly (Release 1.0). <https://webassembly.github.io/spec/>. Accessed 2020-01-01.
- Andreas Rossberg. 2020. Issue #1359: Typed continuations to model stacks. <https://github.com/WebAssembly/design/issues/1359>
- Andreas Rossberg. 2023. WebAssembly (Release 2.0). <https://webassembly.github.io/spec/>. Accessed 2023-20-02.
- Andreas Rossberg. 2025. SpecTec has been adopted. <https://webassembly.org/news/2025-03-27-spectec/>
- Andreas Rossberg, Daan Leijen, Daniel Hillerström, KC Sivaramakrishnan, Matija Pretnar Sam Lindley, and Stephen Dolan. 2020. Stacks and Continuations for Wasm — Idea Sketch. <https://github.com/WebAssembly/meetings/blob/master/main/2020/presentations/2020-02-rossberg-continuations.pdf> Accessed Jul 4 2020.
- Jan Schwinghammer, Lars Birkedal, Bernhard Reus, and Hongseok Yang. 2009. Nested Hoare Triples and Frame Rules for Higher-Order Store. In *Computer Science Logic, 23rd international Workshop, CSL 2009, 18th Annual Conference of the EACSL, Coimbra, Portugal, September 7–11, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5771)*, Erich Grädel and Reinhard Kahle (Eds.). Springer, 440–454. doi:10.1007/978-3-642-04027-6_32
- Remy Seassau, Irene Yoon, Jean-Marie Madiot, and François Pottier. 2025a. Formal Semantics & Program Logics for a Fragment of OCaml. In *ICFP*.
- Remy Seassau, Irene Yoon, Jean-Marie Madiot, and François Pottier. 2025b. Osiris: Towards Formal Semantics and Reasoning for OCaml.
- Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. 2010. x86-TSO: a rigorous and usable programmer’s model for x86 multiprocessors. *Commun. ACM* 53, 7 (2010), 89–97. doi:10.1145/1785414.1785443
- Amin Timany and Lars Birkedal. 2019. Mechanised relational verification of concurrent programs with continuations. *Proc. ACM Program. Lang.* 3, ICFP (2019), 105:1–105:28. doi:10.1145/3341709
- Andrew Wagner, Zachary Eisbach, and Amal Ahmed. 2025. An Adjoint Separation Logic for the Wasm Call Stack. presentation at Discussions at Dagstuhl Seminar 25241: Utilising and Scaling the WebAssembly Semantics. <https://www.andrewwagner.io/assets/slides/adj-wasm-dagstuhl.pdf>

- Conrad Watt, Xiaojia Rao, Jean Pichon-Pharabod, Martin Bodin, and Philippa Gardner. 2021. Two Mechanisations of WebAssembly 1.0. In *Formal Methods - 24th International Symposium, FM 2021, Virtual Event, November 20-26, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 13047)*, Marieke Huisman, Corina S. Pasareanu, and Naijun Zhan (Eds.). Springer, 61–79. doi:10.1007/978-3-030-90870-6_4
- Dongjun Youn, Wonho Shin, Jaehyun Lee, Sukyoung Ryu, Joachim Breitner, Philippa Gardner, Sam Lindley, Matija Pretnar, Xiaojia Rao, Conrad Watt, and Andreas Rossberg. 2024. Bringing the WebAssembly Standard up to Speed with SpecTec. *Proc. ACM Program. Lang.* 8, PLDI (2024), 1559–1584. doi:10.1145/3656440
- Alon Zakai. 2019. Pause and Resume WebAssembly with Binaryen’s Asyncify. <https://kriipken.github.io/blog/wasm/2019/07/16/asyncify.html>.

A WasmFXCert AST

```

1128 (numerical value type) nt ::= i32 | i64 | f32 | f64
1129
1130 (reference value type) rt ::= ft.func | exn | ft.cont
1131
1132 (value type) t ::= nt | rt (function type) ft ::= ts → ts
1133
1134 (immediate) i, min, max ::=  $\mathbb{N}$  (i for local indices)
1135
1136 (exception clause) ec ::= catch i i' | catch_ref i i' | catch_all i | catch_all_ref i
1137 (continuation clause) cc ::= on i i' | on i switch
1138
1139 (basic instructions) b ::= nt.const c | t.add | other stackops | local.{get/set} i |
1140 global.{get/set} i | t.load flags | t.store flags | memory.size |
1141 memory.grow | block ft bs | loop ft bs | if ft bs bs' | br i | br_if i |
1142 br_table is | call i | call_indirect i | return | rt.null |
1143 ref.is_null | ref.function i | ref.call i | try_table i ecs bs |
1144 throw i | throw_ref | cont.new i | resume i ccs |
1145 suspend i | switch i i' | cont.bind i i' | resume_throw i i' ccs
1146
1147 (functions) func ::= func i ts bs (tables) tab ::= tab min max
1148 (memories) mem ::= mem min max (globals) glob ::= glob mutable t einit
1149 (elem segments) elem ::= elem i bsoff is (data segments) data ::= data i bsoff bytes
1150
1151 (import descriptions) importdesc ::= funci i | tabi min max | memi min max |
1152 globi mutable? t | tagi i
1153 (imports) import ::= import string string importdesc
1154 (export descriptions) exportdesc ::= funce i | tabe i | meme i | globe i | tage i
1155 (exports) export ::= export string exportdesc
1156 (start) start ::= Some i | None
1157
1158 (modules) m ::= { types : fts, tags : fts, funcs : funcs, globs : globs, mems : mems, tabs : tabs,
1159 data : datas, elem : elems, imports : imports, exports : exports, start : start }
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176

```

Fig. 7. The AST for structures necessary to evaluate code in WasmFXCert. Black parts are borrowed directly from WasmCert (plain Wasm 1.0). Parts in indigo pertain to the references from Wasm 3.0, parts in orange pertain to exception handling, and parts in magenta are new in WasmFX. Places where WasmFXCert departs from WasmFX are shown boxed.

```

1177             (immediate)  $i, addr, max ::= \mathbb{N}$ 
1178             ( $i$  refers to local indices, while  $addr$  refers to addresses in the store)
1179
1180 (reference value)  $rv ::= rt.null \mid ref.func \textcolor{blue}{addr} \mid ref.exn \textcolor{orange}{addr} \textcolor{orange}{addr'} \mid ref.cont \textcolor{magenta}{addr}$ 
1181 (value)  $v ::= nt.const \textcolor{blue}{c} \mid rv$ 
1182 (module instance)  $inst ::= \{types : fts, \text{ funcs : } addr, \text{ globs : } addr, \text{ mems : } addr,$ 
1183              $\text{ tabs : } addr, \text{ tags : } \textcolor{orange}{addrs} \}$ 
1184 (frame)  $F ::= \{locs : vs, \text{ inst : } inst \}$ 
1185
1186 (translated exception clause)  $dec ::= catch \textcolor{orange}{addr} \textcolor{orange}{i} \mid catch\_ref \textcolor{orange}{addr} \textcolor{orange}{i} \mid$ 
1187              $catch\_all \textcolor{orange}{i} \mid catch\_all\_ref \textcolor{orange}{i}$ 
1188 (translated continuation clause)  $dcc ::= on \textcolor{magenta}{addr} \textcolor{magenta}{i} \mid on \textcolor{magenta}{addr} \textcolor{magenta}{switch}$ 
1189 (administrative instruction)  $e ::= b \mid trap \mid invoke \textcolor{blue}{addr} \mid label_n\{es'\} \textcolor{blue}{es} \text{ end} \mid$ 
1190              $frame_n\{F\} \textcolor{blue}{es} \text{ end} \mid call\_host \textcolor{blue}{Vft} \textcolor{blue}{i} \textcolor{blue}{vs} \mid ref.func \textcolor{blue}{addr} \mid$ 
1191              $ref.exn \textcolor{orange}{addr} \textcolor{orange}{addr'} \mid handler\{decs\} \textcolor{blue}{es} \text{ end} \mid$ 
1192              $throw\_ref.addr \boxed{vs} \textcolor{orange}{addr} \textcolor{orange}{addr'} \mid ref.cont \textcolor{magenta}{addr} \mid$ 
1193              $prompt_{ts}\{dcs\} \textcolor{magenta}{es} \mid suspend.addr \boxed{vs} \textcolor{magenta}{addr}$ 
1194              $switch.addr \textcolor{magenta}{vs} \boxed{addr} \textcolor{magenta}{ft} \textcolor{magenta}{addr'}$ 
1195
1196 (evaluation context)  $L ::= \boxed{vs \text{ ++ } [] \text{ ++ } es} \mid vs \text{ ++ } label_n\{es'\} \textcolor{blue}{L} \text{ end ++ } es \mid$ 
1197              $vs \text{ ++ } handler\{decs\} \textcolor{blue}{L} \text{ end ++ } es \mid$ 
1198              $vs \text{ ++ } prompt_{ts}\{dcs\} \textcolor{magenta}{L} \text{ end ++ } es$ 
1199 (handler context)  $H ::= vs \text{ ++ } [] \text{ ++ } es \mid vs \text{ ++ } label_n\{es'\} \textcolor{blue}{H} \text{ end ++ } es \mid$ 
1200              $vs \text{ ++ } frame_n\{F\} \textcolor{blue}{H} \text{ end ++ } es \mid$ 
1201              $vs \text{ ++ } handler\{decs\} \textcolor{blue}{H} \text{ end ++ } es \mid$ 
1202              $vs \text{ ++ } prompt_{ts}\{dcs\} \textcolor{magenta}{H} \text{ end ++ } es$ 
1203
1204 (function instances)  $finst ::= \{(inst; ts); bs\}_{ft}^{NativeCl} \mid \{hidx\}_{ft}^{HostCl}$ 
1205 (table instances)  $tinst ::= \{elem : is, \text{ max : } max^?\}$ 
1206 (memory instance)  $minst ::= \{data : bytes, \text{ max : } max^?\}$ 
1207 (global instance)  $ginst ::= \{mut : mutable^?, \text{ value : } v\}$ 
1208 (exception instance)  $einst ::= \{etag : \textcolor{orange}{addr}, \text{ efields : } vs\}$ 
1209 (continuation instance)  $cinst ::= cont \boxed{ft} \textcolor{magenta}{H} \mid dagger \boxed{ft}$ 
1210 (store)  $S ::= \{\text{funcs : } finsts, \text{ globs : } ginsts, \text{ mems : } minst, \text{ tabs : } tinst,$ 
1211              $\text{ tags : } fts, \text{ exns : } einst, \text{ conts : } cinst \}$ 
1212
1213 (typing context)  $C ::= \{ \text{ types : } fts, \text{ func : } fts, \text{ global : } gts, \text{ table : } tts, \text{ memory : } mts,$ 
1214              $\text{ local : } ts, \text{ label : } (ts)s, \text{ return : } ts^?, \text{ exns : } \textcolor{orange}{addrs}, \text{ tags : } fts \}$ 
1215
1216 (the meaning of some of the fields in the store is orthogonal to stack switching and is omitted)

```

Fig. 8. The AST for structures necessary to evaluate code in WasmFXCert. Black parts are borrowed directly from WasmCert (plain Wasm 1.0). Parts in **indigo** pertain to the references from Wasm 3.0, parts in **orange** pertain to exception handling, and parts in **magenta** are new in WasmFX. Places where WasmFXCert departs from WasmFX are shown **boxed**.

B Frame

Another novelty we bring in Iris-WasmFX with respect to the existing Iris-Wasm program logic is the way we deal with the Wasm frame. In Iris-Wasm, aside from resources corresponding to different parts of the store, there is a resource $\xrightarrow{\text{Fr}} F$ that symbolises ownership of the entire frame. For technical reasons,¹ Iris-Wasm mandates that the frame resource must be present in the proof environment every time a reduction step is taken, cluttering the proof rules and leading to the confusing situation where the frame resource appears in proof rules for instructions that in no way interact with the frame.

Instead, our custom weakest precondition considers the frame as part of the expression rather than part of the state. Thus, our first modification to Iris-Wasm's weakest precondition statement is that ours has the shape $\text{wp } es ; F \{w \ F, \Phi(w, F)\}$ where the expression es sits alongside the frame F , and the postcondition is a predicate on both the final logical value w and the final frame F .

We illustrate this new weakest precondition statement by showing the bind rule for the frame instruction. Having the frame baked into the weakest precondition statement makes it much easier to update the frame in order to run the body of the frame in the correct environment:

$$\frac{\text{WP-BIND-FRAME} \quad \text{wp } es ; F \{w \ F', \text{wp } [\text{frame}_n\{F'\} \ w \ \text{end}] ; F_0 \{v \ F_1, \Phi(v, F_1)\}\}}{\text{wp } [\text{frame}_n\{F\} \ es \ \text{end}] ; F_0 \{v \ F_1, \Phi(v, F_1)\}}$$

i.e. when reasoning under frame F_0 , if we reach a frame instruction (representing a function call being executed), when reasoning about the code inside the frame, we use the frame F specified by the frame instruction. Since the frame contains the local variables and the instance (which translates local indices into addresses into the store), this change of frames is what guarantees local state encapsulation. The bind rule above specifies that once the body es has reduced to a value w and the 'inner' frame F has become a frame F' , we can then place that value and that frame back in a frame instruction and resume reasoning from there.

Example. Let us showcase how the two bind rules above are used and how Iris-WasmFX deals with the frame, by going through the first few steps of proving the specification for the `$yield` function from Figure 3. Recall from §2 we want this specification to have the form $\{I\} \ [\text{invoke } \$\text{addyield}] \ \{I\}$. More precisely, given an invariant I , a fraction q and a frame F , we wish to prove that

$$I \multimap \text{addyield} \xrightarrow{\text{wf}} cl_{\text{yield}} \multimap \text{addrt} \xrightarrow{\text{wtag}}_q ([\] \rightarrow [\]) \multimap \left\{ \begin{array}{l} \text{wp } [\text{invoke } \text{addyield}] ; F \left\{ w \ F', \quad w = \text{immV } [\] * F = F' * I * \right. \\ \left. \text{addyield} \xrightarrow{\text{wf}} cl_{\text{yield}} * \text{addrt} \xrightarrow{\text{wtag}}_q ([\] \rightarrow [\]) \right\} \end{array} \right\}$$

where addyield is $F.\text{inst.funks}[\text{\$yield}]$ (the address of function `$yield` in the store), addrt is $F.\text{inst.tags}[\text{\$t}]$ (the address of tag `$t` in the store) and cl_{yield} is the closure of yield function (containing the function's actual code, the instance of the coroutine module, the function's type and the type of its local variables). The function closure resource $\text{addyield} \xrightarrow{\text{wf}} cl_{\text{yield}}$ is necessary to reason about the act of invoking, and is given back in the post-condition; the tag resource $\text{addrt} \xrightarrow{\text{wtag}}_q ([\] \rightarrow [\])$ is necessary to run the actual code of the function (recall it contains a

¹These technical reasons pertain to the soundness of the proof rule for binding into frame instructions. Since reducing the body of a local is done using a different frame, the authoritative view associated to the state must be updated. This can only be done in the presence of the associated fragmental view, that is in this case, the frame resource. To do this, the strategy implemented by Iris-Wasm was to bake the presence of the frame resource into the definition of the weakest precondition by enforcing that it must be present every time a step in the execution is taken. As we describe in this paragraph, in Iris-WasmFX, we use a different, simpler strategy.

suspend instruction; we showcase in §4.3 the part where this resource is used) and is also given back in the post-condition.

To prove this specification, we start by applying Iris-Wasm rule WP-INVOKE-NATIVE which mirrors the operational semantics of `invoke` (see §3.2) and has two premises: the first requires us to provide a resource corresponding to the function closure for `$yield`, which we hold. The second gives back that resource, and requires us to prove a weakest precondition statement on the expression the `invoke` reduces to:

$$\text{wp} [\text{frame}_0\{F'\} \text{ block } ([\] \rightarrow [\]) \text{code}_{\text{yield}} \text{ end }] ; F \left\{ w F', \begin{array}{l} w = \text{immV } [\] * F = F' * \\ I * \text{addr}_{\text{yield}} \xrightarrow{\text{wf}} \text{cl}_{\text{yield}} \end{array} \right\}$$

where F' is the frame under which the code $\text{code}_{\text{yield}}$ is to be executed, containing the coroutine module's instance and no local variables.

Now we apply rule WP-BIND-FRAME above, which brings us to proving

$$\text{wp} [\text{block } ([\] \rightarrow [\]) \text{code}_{\text{yield}}] ; F' \{w F'', \Phi(w, F'')\}$$

with

$$\Phi(w, F'') = \text{wp} [\text{frame}_0\{F''\} w \text{ end }] ; F \left\{ w F', \begin{array}{l} w = \text{immV } [\] * F = F' * \\ I * \text{addr}_{\text{yield}} \xrightarrow{\text{wf}} \text{cl}_{\text{yield}} \end{array} \right\}$$

i.e. we have focused on the inside of the frame instruction and are now reasoning under frame F' , the one with the closure's environment; once we will reach a value, we can reason about this value placed back into the frame instruction under the top-level F frame. This is where our treatment of the Wasm frame differs from Iris-Wasm, where frame resources $\xrightarrow{\text{Fr}} F$ would be passed around in a sometimes convoluted way, especially when applying the WP-BIND-FRAME rule.

Then we apply Iris-Wasm rule WP-BLOCK, bringing us to reasoning about $[\text{label}_0\{[\]\} \text{code}_{\text{yield}} \text{ end}]$ instead of the block instruction; and finally, applying rule WP-BIND-LABEL from §4.1, our goal becomes

$$\text{wp } \text{code}_{\text{yield}} ; F' \{w F'', \text{wp } \text{label}_0\{[\]\} w \text{ end } ; F'' \{w F'', \Phi(w, F'')\}\}$$

i.e. we are now entirely focused on the code of `$yield`. As displayed in Figure 3, that code is just one line: `[suspend $t]`. We finish the proof in §4.3 and in Appendix E.1.1.

Other Approaches. Adjoint separation logic [Wagner et al. 2025] elegantly deals with this problem with frames by using modalities that make a given frame available or lock it away.

C Syntactic Typing and Type Safety in Detail

In this section, we introduce WasmFX's syntactic type system. When defining WasmFXCert, we give new typing rules for administrative instructions, and state and prove a type safety theorem.

Wasm defines a simple syntactic type system for all its instructions, of the form

$$C \vdash b : ts_1 \rightarrow ts_2$$

where C is a typing context (see Figure 8) that keeps track of the type of all the objects reachable in the current module, including all the labels that can be branched to and the return type of the current function; b is the basic instruction being typechecked; ts_1 is the list of the types of the values expected to be on the stack for the instruction to execute safely; and ts_2 is the list of the types of the values deposited back on the stack after execution. For example, for all contexts C and all numerical types nt , $C \vdash nt.add : [nt; nt] \rightarrow [nt]$.

We show the typing rules for the new instructions of WasmFX in Figure 9.

$$\begin{array}{c}
 \boxed{C \vdash cc : ts} \\
 \hline
 \frac{C.\text{tags}[i] = ts_1 \rightarrow ts_2 \quad C.\text{label}[i'] = ts_1 \mathbin{++} (ts_2 \rightarrow ts).\text{cont} \quad C.\text{tags}[i] = [] \rightarrow ts}{C \vdash \text{on } i \ i' : ts} \quad C \vdash \text{on } i \ \text{switch} : ts \\
 \\
 \boxed{C \vdash b : ft} \\
 \hline
 \frac{\frac{C.\text{types}[i] = ft}{C \vdash \text{cont.new } i : ft.\text{func} \rightarrow ft.\text{cont}} \quad \frac{C.\text{tags}[i] = ft}{C \vdash \text{suspend } i : ft}}{C.\text{types}[i] = ft \quad C.\text{tags}[i'] = [] \rightarrow ts \quad ft = (ts_1 \mathbin{++} (ts_2 \rightarrow ts).\text{cont}) \rightarrow ts} \\
 \hline
 C \vdash \text{switch } i \ i' : (ts_1 \mathbin{++} ft.\text{cont}) \rightarrow ts_2 \\
 \hline
 \frac{C.\text{types}[i] = ft \quad ft = ts_1 \rightarrow ts_2 \quad \forall cc \in ccs, C \vdash cc : ts_2}{C \vdash \text{resume } i \ ccs : (ts_1 \mathbin{++} ft.\text{cont}) \rightarrow ts_2}
 \end{array}$$

Fig. 9. The typing rules for WasmFX

When defining WasmFXCert, we have included the first result on type safety of WasmFX, in the form of a theorem proved in the Rocq proof assistant. To formulate this theorem, we had to provide the first formal definitions of typing rules for the new administrative instructions and for the new parts of the Wasm store. While Phipps-Costin et al. [2023] define a typing rule for prompt, we have identified that their presentation is incomplete: it has a typo, and their strategy of stripping the context is too weak to actually complete the proof of type preservation. Our typing rules are presented in Figure 10.

Two main challenges arise when defining these typing rules. The first is that, unlike in plain Wasm 1.0, not all values automatically typecheck in all typing contexts, since function references and continuation references require the presence of a well-typed object in the typing context. This means that one needs to add typechecking hypotheses to many lemmas, and carefully keep track in proofs of which values are known to typecheck. Moreover, it also means that the store cannot be arbitrarily changed while ensuring all Wasm values keep on being well-typed (in 1.0, values always have a numerical type so they typecheck trivially as long as of the correct integer type). Hence one needs to fine-tune the definition of the $\text{store_extension}(S, S')$ predicate, which describes the way in which the store can evolve during execution and which satisfies that values that typecheck in S still typecheck in S' .

1373	$S, C \vdash dcc : ts$	
1374		
1375	$\frac{S.tags[addr] = ts_1 \rightarrow ts_2 \quad C.label[i] = ts_1 ++ (ts_2 \rightarrow ts).cont}{S, C \vdash on\ addr\ i : ts}$	
1376		
1377	$\frac{S.tags[addr] = [] \rightarrow ts}{S, C \vdash on\ addr\ switch : ts}$	
1378		
1379		
1380	$S, C \vdash es : ft$	
1381		
1382	$\frac{S.conts[addr] = (cont\ ft\ _ \vee dagger\ ft) \quad S.tags[addr] = ts_1 \rightarrow ts_2 \quad S, C \vdash vs : [] \rightarrow ts_1}{S, C \vdash ref.cont\ addr : [] \rightarrow ft.cont \quad S, C \vdash suspend.addr\ vs\ addr : [] \rightarrow ts_2}$	
1383		
1384		
1385	$\frac{S.tags[addr'] = [] \rightarrow ts \quad ft = (ts_1 ++ (ts_2 \rightarrow ts).cont) \rightarrow ts \quad S.conts[addr] = (cont\ ft\ _ \vee dagger\ ft) \quad S, C \vdash vs : [] \rightarrow ts_1}{S, C \vdash switch.addr\ vs\ addr\ ft\ addr' : [] \rightarrow ts_2}$	
1386		
1387	$\frac{S, \emptyset \vdash es : [] \rightarrow ts \quad \forall dcc \in dcs, S, C \vdash dcc : ts}{S, C \vdash prompt_{ts}\{dcs\}\ es\ end : [] \rightarrow ts}$	
1388		
1389		
1390		
1391	$S \vdash cinst\ ok$	
1392		
1393	$\frac{S, \emptyset \vdash vs : [] \rightarrow ts_1 \quad H[vs] = es \quad S, \emptyset \vdash es : [] \rightarrow ts_2}{S \vdagger dagger\ ft\ ok \quad S \vdagger cont\ (ts_1 \rightarrow ts_2)\ H\ ok}$	
1394		
1395		
1396	$S \vdash einst\ ok$	
1397	$\frac{S, \emptyset \vdash e.efields : [] \rightarrow ts \quad S.tags[e.etag] = ts \rightarrow []}{S \vdash e\ ok}$	
1398		
1399		
1400	$\vdash S\ ok$	
1401		
1402	functions	
1403	tables	
1404	memories as in WasmCert	$\forall c \in S.conts, S \vdash c\ ok \quad \forall e \in S.exns, S \vdash e\ ok$
1405	globals	
1406		$\vdash S\ ok$

Fig. 10. Our new typing rules for the administrative instructions and store elements of WasmFX

The second difficulty is the choice of a typing context for the body of a prompt instruction. This choice is ultimately irrelevant since in practice, all continuations start as function calls, and hence after the first few steps of execution, every continuation will be of the form $[frame_n\{F\}\ H\ end]$, and the body of a frame instruction is typed with a specifically built typing context in the Wasm typing rules. However, the presence of this frame instruction immediately under every prompt is not enforced syntactically, and is not true for the first few steps of execution when a new continuation is created using the `cont.new` instruction. Our solution is to mandate that the body of a prompt should typecheck in the empty typing context, as illustrated in Figure 10. This necessitated replacing all tag and type annotations with their explicit values (as is done implicitly in WasmFX during the typing phase), and proving lemmas about typing expressions in the empty typing context, such as the following:

LEMMA C.1. *For all stores S , typing contexts C , administrative instructions es , and function types ft ,*

$$S, \emptyset \vdash es: ft \implies S, C \vdash es: ft$$

Type Safety Theorem. As is standard in Wasm's typing system, we can define a context $C(S, F)$ out of a store S and a frame F by setting $C.locs$ to be the types of the values $F.locs$ and filling all other fields of C by reading the type annotations in the store of the elements at the addresses specified by instance $F.inst$.

We can now formulate the progress and preservation lemmas:

LEMMA C.2 (TYPE PRESERVATION). *If $\vdash S \text{ ok}$ and $S, C(S, F) \vdash es: [] \rightarrow ts$ and $(S, F, es) \hookrightarrow (S', F', es')$, then $\vdash S' \text{ ok}$ and $S', C(S', F') \vdash es': [] \rightarrow ts$*

LEMMA C.3 (TYPE PROGRESS). *If $\vdash S \text{ ok}$ and $S, C(S, F) \vdash es: [] \rightarrow ts$, then either es reduces, or it is a constant value, or it is stuck on a host call, or it is stuck on an unhandled suspend, switch, or exception throw.*

And together, these give the type safety result:

THEOREM C.4 (TYPE SAFETY). *If $\vdash S \text{ ok}$ and $S, C(S, F) \vdash es: [] \rightarrow ts$ and $(S, F, es) \hookrightarrow^* (S', F', es')$, then either es' reduces, or it is a constant value, or it is stuck on a host call, or it is stuck on an unhandled suspend, switch, or exception throw.*

D Bind Rule for Prompt

Since the resume instruction reduces to a prompt instruction, the crux of the reasoning is how to use protocols to bind the body of a prompt instruction. As discussed, the simple EWP-LABEL rule would be unsound due to the presence of effects. Instead, we prove the following rule:

EWP-PROMPT

$$\begin{array}{l}
 (1) \quad (\forall addr, addr \notin dcs \implies \Psi(addr) = \Psi'(addr)) * \\
 (2) \quad CExp(es) * (\forall F, \neg\Phi(\text{trapV}, F)) * \neg\Phi'(\text{trapV}) * \\
 (3) \quad ewp\ es; \emptyset \langle \Psi \rangle \{w\ F, \Phi(w, F)\} * \\
 (4) \quad (\forall w, \Phi(w, \emptyset) \multimap ewp[\text{prompt}_{ts}\{dcs\}\ w\ \text{end}]; \emptyset \langle \Psi' \rangle \{w\ _, \Phi'(w)\}) * \\
 (5) \quad \frac{\forall dcc \in dcs, \langle \Psi \rangle \{ \Phi \} dcc; ts_2 \langle \Psi' \rangle \{ \lambda v\ _, \Phi'(v) \}}{ewp[\text{prompt}_{ts}\{dcs\}\ es\ \text{end}]; F \langle \Psi' \rangle \{w\ F', \Phi'(w)\} * F = F'}
 \end{array}$$

This rule has the structure of a usual bind rule: line (3) focuses on the body es of the prompt instruction, and line (4) injects the resulting value w into the prompt. It allows using a different protocol family for the body es than for the prompt instruction itself, to account for the fact that the prompt specifies the intended behaviour for all the effects handled in its clauses dcs . Line (1) enforces that the protocols families Ψ and Ψ' only differ in the effects handled by the prompt clauses; line (5) ties them together in those remaining places, as we will explain below.

Let us now focus on line (2): aside from enforcing that all postconditions do not hold of the trap value trapV (a different, simpler rule can be invoked for expressions that trap), it requires the body es of the prompt to satisfy predicate $CExp$, which we define as follows:

$$\begin{array}{ll}
 CExp(es) \triangleq & (\exists vs\ addr\ ft, es = vs \mathrel{++} [\text{ref.func}\ addr; \text{ref.call}\ ft]) \vee \quad (a) \\
 & (\exists vs\ addr, es = vs \mathrel{++} [\text{invoke}\ addr]) \vee \quad (b) \\
 & (\exists vs\ n\ F\ es', es = vs \mathrel{++} [\text{frame}_n\{F\}\ es'\ \text{end}]) \vee \quad (c) \\
 & (\exists vs, es = vs \mathrel{++} [\text{trap}]) \vee \quad (d) \\
 & (\exists vs, es = vs) \vee \quad (e) \\
 & (\exists vs\ vs'\ k\ ft, es = vs \mathrel{++} [\text{call_host}\ vs'\ k\ ft]) \quad (f)
 \end{array}$$

i.e. es is of one of 6 forms: (a) values followed by a call to a function reference, (b) values followed by an `invoke` instruction (i.e. the intermediate representation before a call instruction is replaced by the function's actual body), (c) values followed by a `frame` instruction (representing a function body being executed), (d) values followed by a `trap` instruction, (e) es is constant, or (f) values followed by a `call_host` instruction (a stuck instruction representing a call to a function defined by the host language).

As previously mentioned, our continuation resource enforces all continuations to be of a special form, which means that when plugged with a value, the resulting expression is always of form (a) or (c). The other four cases have been added to guarantee closure under reduction: if $CExp(es)$ and es reduces to es' , then $CExp(es')$. The most crucial property verified by expressions of these forms is that they reduce in a frame-agnostic way: if $CExp(es)$ and (S, F, es) reduces to (S', F', es') , then $F = F'$ and for any other frame F'' , (S, F'', es) reduces to (S', F'', es') .

Line (5) of rule EWP-PROMPT specifies the behaviour of the program when an effect handled by one of the clauses of the prompt is performed using the `suspend` instruction. In that case, recall that reduction rule REDUCE-SUSPEND mandates that a continuation is created out of the state of the stack, and the prompt reduces to a `br` instruction with the label specified by the clause from the prompt instruction, with the effect's payload and the newly created continuation on the stack. Hence we wish to require that in that case, we hold an extended weakest precondition statement for this $vs \mathrel{++} [\text{ref.cont}\ kaddr; \text{br}\ ilab]$ expression, where $kaddr$ is the address of the newly created

continuation. Little is known about the form of that continuation, other than that the suspend instruction will have followed a protocol from family Ψ . We define the *clause triple* from line (5) (which closely follows that of Hazel [de Vilhena and Pottier 2021, Fig. 7]), as:

$$\begin{aligned} \langle \Psi \rangle \{ \Phi \} \text{ on } taddr \text{ ilab}; ts \langle \Psi' \rangle \{ \Phi' \} = \\ \exists ts_1 \ ts_2 \ q, taddr \vdash_{\text{wtag}}^q (ts_1 \rightarrow ts_2) * \\ \forall vs \ kaddr \ H, kaddr \vdash_{\text{wcont}} (ts_2 \rightarrow ts, H) \multimap taddr \vdash_{\text{wtag}}^q (ts_1 \rightarrow ts_2) \multimap \\ \uparrow (\Psi \ taddr) (\text{immV } vs) (\lambda w. \triangleright \text{ewp } H[w]; \emptyset \langle \Psi \rangle \{ w \ F, \Phi(w, F) \}) \multimap \\ \text{ewp } vs ++ [\text{ref.cont } kaddr; \text{br } ilab]; \emptyset \langle \Psi' \rangle \{ w \ F, \Phi'(w, F) \} \end{aligned}$$

When establishing this premise, we must prove the extended weakest precondition statement on the br instruction using these resources: a continuation resource corresponding to the new continuation, and knowledge that the effect was performed in a way that follows the protocol $\Psi \ taddr$. Since the continuation H is universally quantified, the only way to reason about resuming the continuation again is by knowing that the protocol is followed.

For example, when $\Psi \ taddr$ is of the form $!x(v)\{P\} \ ?y(w)\{Q\}$, this definition simplifies to

$$\begin{aligned} \exists ts_1 \ ts_2 \ q, taddr \vdash_{\text{wtag}}^q (ts_1 \rightarrow ts_2) * \\ \forall kaddr \ H \ x, kaddr \vdash_{\text{wcont}} (ts_2 \rightarrow ts, H) \multimap taddr \vdash_{\text{wtag}}^q (ts_1 \rightarrow ts_2) \multimap \\ P \multimap (\forall y. Q \multimap \triangleright \text{ewp } H[w]; \emptyset \langle \Psi \rangle \{ w \ F, \Phi(w, F) \}) \multimap \\ \text{ewp } v ++ [\text{ref.cont } kaddr; \text{br } ilab]; \emptyset \langle \Psi' \rangle \{ w \ F, \Phi'(w, F) \} \end{aligned}$$

i.e. we must show that breaking with value v is safe, knowing that P holds, and that we can safely resume the continuation with w as long as we can provide Q .

Frame Agnosticity. The last two lines of the definition of the clause triple refer to an empty frame.

In theory, according to the operational semantics, the frame to be used in this position is the frame as it would be at the suspension site (for the last line) and re-resumption site (for the next-to-last line), but when scrutinising a prompt instruction, the state of the frame at those later times is unknown. Fortunately, the choice of the frame to use at these points is in practice irrelevant since the inside of the prompt instruction reduces in a frame-agnostic way, hence we can choose whatever frame we want. We find it convenient to use the empty frame everywhere (including line (3) in rule EWP-PROMPT), so that all the frames inside the prompt instruction match one another, simplifying reasoning.

To bridge the frame F in the conclusion of EWP-PROMPT with the empty frame in premise (3), we prove the following lemma, which says that if an expression es satisfying $\text{CExp}(es)$ is safe to run under the empty frame, then the same expression is safe to run with any frame, resulting in an unchanged frame, with the same post-condition on the resulting value:

$$\frac{\text{EWP-EMPTY-FRAME} \quad \text{CExp}(es) * \text{ewp } es; \emptyset \langle \Psi \rangle \{ w _, \Phi(w) \}}{\text{ewp } es; F \langle \Psi \rangle \{ w \ F', \Phi(w) \} * F = F'}$$

Resume Rule. The rule for resume is largely the same as for EWP-PROMPT:

EWP-RESUME

- $$\begin{array}{l}
(1) \quad F.\text{inst.types}[i] = ts_1 \rightarrow ts_2 * |vs| = |ts_1| * \text{translate_clauses}(F.\text{inst}, ccs) = dcs * \\
(2) \quad (\forall \text{addr}, \text{addr} \notin dcs \implies \Psi(\text{addr}) = \Psi'(\text{addr})) * \text{addr} \xrightarrow{\text{wcont}} (ts_1 \rightarrow ts_2, H) * \\
(3) \quad (\forall F, \neg\Phi(\text{trapV}, F)) * \neg\Phi'(\text{trapV}) * \triangleright \text{ewp } H[vs]; \emptyset \langle \Psi \rangle \{w \ F, \Phi(w, F)\} * \\
(4) \quad \triangleright (\forall w, \Phi(w, \emptyset) \multimap \text{ewp} [\text{prompt}_{ts_2} \{dcs\} \ w \ \text{end}]; \emptyset \langle \Psi' \rangle \{v _, \Phi'(v)\}) * \\
(5) \quad \triangleright (\forall dcs \in dcs, \langle \Psi \rangle \{ \Phi \} \ dcs; \ ts_2 \langle \Psi' \rangle \{ \lambda v _, \Phi'(v) \}) \\
\hline
\text{ewp } vs \text{ ++ } [\text{ref.cont } \text{addr}; \text{resume } i \ ccs]; F \langle \Psi' \rangle \{v \ F', \Phi'(v)\} * F' = F
\end{array}$$

The additional premises are the continuation resource on line (2), and the premises on line (1) which translate the annotations of the resume instruction and check that the adequate number of values are being taken from the stack. The extra premises are necessary to apply reduction rule REDUCE-RESUME, which transforms the resume into a prompt instruction. Since a step is taken, we can add a later modality \triangleright to lines (4) and (5). The rest of the premises are those of EWP-PROMPT, allowing to bind into the actual code of the continuation being resumed. Because the continuation resource enforces that the context H is safe, the rule does not require establishing $\text{CExp}(H[vs])$.

$\exists \text{addr}_{\text{yield}} \text{addr}_{\text{par}} \text{cl}_{\text{yield}} \text{cl}_{\text{par}},$
 $J_{\text{yield}} * J_{\text{par}} * \text{ (we own the closures for \$par and \$yield, ...)}$
 $\forall P_1 Q_1 P_2 Q_2 I, \exists \Psi,$
 $\text{ (... and we have specifications for both functions)}$
 $\text{ (for function \$yield:)}$
 $\square \left(\begin{array}{l} \forall F, I \multimap J_{\text{yield}} \multimap \\ \text{ewp} [\text{invoke } \text{addr}_{\text{yield}}] ; F \langle \Psi \rangle \left\{ w F', \begin{array}{l} w = \text{immV} [] * F = F' * \\ I * J_{\text{yield}} \end{array} \right\} \end{array} \right) *$
 $\text{ (and for function \$par:)}$
 $\forall \text{addr}_{f_1} \text{addr}_{f_2} \text{cl}_1 \text{cl}_2,$
 $\left(\begin{array}{l} \forall F, P_1 \multimap P_2 \multimap I \multimap J_1 \multimap J_2 \multimap J_{\text{yield}} \multimap J_{\text{par}} \multimap \\ \text{ (precondition of specification for \$par contains a specification for \$f}_1\text{:)} \\ \square \left(\begin{array}{l} \forall F', P_1 \multimap I \multimap J_1 \multimap J_{\text{yield}} \multimap \\ \text{ewp} \left[\begin{array}{l} \text{ref.func } \text{addr}_{f_1}; \\ \text{ref.call} \end{array} \right] ; F' \langle \Psi \rangle \left\{ w F'', \begin{array}{l} w = \text{immV} [] * F' = F'' * \\ Q_1 * I * J_1 * J_{\text{yield}} \end{array} \right\} \end{array} \right) \multimap \\ \text{ (precondition of specification for \$par contains a specification for \$f}_2\text{:)} \\ \square \left(\begin{array}{l} \forall F', P_2 \multimap I \multimap J_2 \multimap J_{\text{yield}} \multimap \\ \text{ewp} \left[\begin{array}{l} \text{ref.func } \text{addr}_{f_2}; \\ \text{ref.call} \end{array} \right] ; F' \langle \Psi \rangle \left\{ w F'', \begin{array}{l} w = \text{immV} [] * F' = F'' * \\ Q_2 * I * J_2 * J_{\text{yield}} \end{array} \right\} \end{array} \right) \multimap \\ \text{ewp} \left[\begin{array}{l} \text{ref.func } \text{addr}_{f_1}; \\ \text{ref.func } \text{addr}_{f_2}; \\ \text{invoke } \text{addr}_{\text{par}} \end{array} \right] ; F \langle \perp \rangle \left\{ w F', \begin{array}{l} w = \text{immV} [] * F = F' * Q_1 * Q_2 * \\ I * J_1 * J_2 * J_{\text{yield}} * J_{\text{par}} \end{array} \right\} \end{array} \right) *$
 Where J_{yield} is a shorthand for $\text{addr}_{\text{yield}} \xrightarrow{\text{wf}} \text{cl}_{\text{yield}}$, J_{par} is a shorthand for $\text{addr}_{\text{par}} \xrightarrow{\text{wf}} \text{cl}_{\text{par}}$,
 J_1 is a shorthand for $\text{addr}_{f_1} \xrightarrow{\text{wf}} \text{cl}_1$, and J_2 is a shorthand for $\text{addr}_{f_2} \xrightarrow{\text{wf}} \text{cl}_2$

Fig. 11. The specification for the coroutines module

E Case Studies in Detail

E.1 Coroutine Library

We now revisit our running coroutine example from §2, focusing on proving the implementation's specification (§E.1.1) and the simple client we introduced in Figure 13 (§E.1.2).

A reader familiar with modern separation logics might notice that one could give a more implicit specification, without this explicit invariant I , using for example Iris invariants [Jung et al. 2018, §2.2]. We present the specification in CSL style for clarity.

E.1.1 Library implementation. The code for the coroutine library module is shown in Figure 3. We presented an informal version of its specification in §2. Now that we have introduced the full machinery of Iris-WasmFX, we show the full specification in Figure 11.

Compared to the snapshot given in §2, there are two main differences:

- (1) We have inlined the definition of the Hoare triple and replaced the weakest precondition statement with our extended weakest precondition
- (2) We have added all relevant function closure resources in the pre- and post-conditions

In essence, the meaning remains the same: in order to $\text{\$yield}$, one must provide the invariant I and that invariant is restituted when the function returns; and in order to run $\text{\$par}$ on two functions $\text{\$f}_1$ and $\text{\$f}_2$, one must have established specifications for both functions.

$$\begin{aligned}
LI = & \exists H_a, \\
& \left(\begin{aligned} & I * J * \$current \xrightarrow{wcont} H_a * \\ & \left(\begin{aligned} & \left(I \multimap J \multimap \triangleright ewp H_a[()]; \emptyset \langle \Psi_0 \rangle \left\{ w F, \begin{array}{l} w = \text{immV } [] * Q_1 * \\ I * \$f_1 \xrightarrow{wf} cl_1 * J \end{array} \right\} \right) * \\ & \left(\$next_is_done = 0 * \exists H_b, \$next \xrightarrow{wcont} H_b * \right. \\ & \quad \left. \left(\begin{aligned} & \left(I \multimap J \multimap \triangleright ewp H_b[()]; \emptyset \langle \Psi_0 \rangle \left\{ w F, \begin{array}{l} w = \text{immV } [] * Q_2 * \\ I * \$f_2 \xrightarrow{wf} cl_2 * J \end{array} \right\} \right) \\ & \vee (\$next_is_done = 1 * Q_2 * \$f_2 \xrightarrow{wf} cl_2) \end{aligned} \right) \end{aligned} \right) * \\ & \left(\begin{aligned} & I * J * \$current \xrightarrow{wcont} H_a * \\ & \left(\begin{aligned} & \left(I \multimap J \multimap \triangleright ewp H_a[()]; \emptyset \langle \Psi_0 \rangle \left\{ w F, \begin{array}{l} w = \text{immV } [] * Q_2 * \\ I * \$f_2 \xrightarrow{wf} cl_2 * J \end{array} \right\} \right) * \\ & \left(\$next_is_done = 0 * \exists H_b, \$next \xrightarrow{wcont} H_b * \right. \\ & \quad \left. \left(\begin{aligned} & \left(I \multimap J \multimap \triangleright ewp H_b[()]; \emptyset \langle \Psi_0 \rangle \left\{ w F, \begin{array}{l} w = \text{immV } [] * Q_1 * \\ I * \$f_1 \xrightarrow{wf} cl_1 * J \end{array} \right\} \right) \\ & \vee (\$next_is_done = 1 * Q_1 * \$f_1 \xrightarrow{wf} cl_1) \end{aligned} \right) \end{aligned} \right) * \end{aligned} \right) \\
& \text{where } J \text{ is } \text{addyield} \xrightarrow{wf} cl_{\text{yield}}.
\end{aligned}
\end{aligned}$$

Fig. 12. Loop invariant for proving the specification of the `$par` function

Note the quantification: the invariant I , as well as the pre- and post-conditions P_1 , Q_1 , P_2 and Q_2 of `$f1` and `$f2`, can be chosen later by the client. This makes the specification more general and useful for more clients. However, the actual addresses and the actual closures of the functions `$yield` and `$par`, as well as the metaprotocol Ψ used by `$yield`, are hidden behind existentials, meaning that a client has no access to the code itself. When the client wishes to reason about invocation of one of these library functions, all they can do is apply this specification. By hiding the metaprotocol Ψ behind an existential quantifier, we can enforce that the only way for the client to use effects in a way that can be specified is by making a call to `$yield`.

Proving the Specification. To prove this specification, it suffices to show that the coroutines module satisfies the existentials with the actual addresses and closures, and to define Ψ as the metaprotocol Ψ_0 described in §4.2. The function resources for `$yield` and `$par` are provided by the host language’s program logic (see §3.5) during the instantiation process.

Then it remains to prove the two specifications; we gave an outline of the proof for `$yield` in Appendix B and in §4.3. Note that the resource for the tag `$t` does not feature in the specification. This is because this resource is not useful for the client and thus is not shown in the specification. When proving the specification, however, the resource is provided by the instantiation process together with the function closure resources.

For `$par`, the proof is slightly longer and requires establishing a loop invariant for the loop, and applying the EWP-RESUME rule. We show the loop invariant LI in Figure 12.

The invariant is a disjunction where the left case corresponds to the case where local variable `$current` is the continuation obtained originally from function `$f1` and `$next` is the continuation obtained originally from function `$f2`; and the right disjunct is vice-versa, `$current` is `$f2` and `$next` is `$f1`. In both cases, we assert that we own the invariant I , the function closure resource J , and a continuation H_a at the address specified by local variable `$current`, and we require that:

- it is sufficient to give I and J to be able to safely run $H_a[()]$, and
- one of the following is true:

- the variable `$next_is_done` is **false**, which means that the continuation in variable `$next` is not done executing: we require ownership of a continuation H_b at the address specified by `$next`, and holding I and J must be enough to safely run $H_b[()];$ or
- the variable `$next_is_done` is **true**, which means that the continuation in variable `$next` has terminated: we must hold the postcondition of the corresponding function.

Let us go through the code in Figure 3 line by line.

Lines 13 to 20 create the continuations; using rule EWP-CONTNEW, we get the continuation resources corresponding to calling `$f1` and `$f2`. To enter the loop on line 21, we must establish the loop invariant, LI (left disjunct, and left disjunct in the inner disjunct). To fulfil the extended weakest precondition premises in LI , we use the specifications for `$f1` and `$f2` since at this moment the continuations in `$current` and `$next` contain a call to `$f1` and `$f2`.

Then we reason about the inside of the loop: first we resume continuation `$current` on line 24. For this, we apply rule EWP-RESUME. Most premises are either trivial or given directly by the loop invariant LI (e.g. the continuation resource on line (2) and the extended weakest precondition on line (3)). The two missing premises are lines (4) and (5).

For line (4), we must establish the termination case. We assume the postcondition of the function that just terminated (i.e. $Q_i * I * \$f_i \vdash^{\text{wf}} c_i$ for the correct $i \in \{1, 2\}$) and prove that we can safely run lines 25 to 31 (right after the resume). On lines 25-26, we check the value of `$next_is_done`. If it is 1, we exit the function; in virtue of our loop invariant LI , we know that the other $Q_j * \$f_j \vdash^{\text{wf}} c_j$ also holds and we can conclude the proof. If `$next_is_done` is 0, then we update its value (lines 27-28), place `$next` in `$current` (lines 29-30), and repeat the loop (line 31). We can do this, since we can reestablish the loop invariant LI (by using the right disjunct in the inner disjunction).

For line (5), we must show that given a continuation resource $\$current \vdash^{\text{wcont}} H$ for some safe context H , given I (from the protocol $\Psi_0(\$t)$), and given that I is enough to be able to safely run $H[()];$, we must prove that it is safe to run lines 32 to 39 (the code to which we branch when the effect is performed). We save the suspended continuation into variable `$current` on line 32, then check the value of `$next_is_done` on lines 33-34. If it is true, then we know `$next` has done executing and we loop without swapping `$next` and `$current`; looping is safe since we can trivially reestablish the loop invariant LI . If `$next_is_done` is false, then we swap `$next` and `$current` (lines 35-28) and loop (line 39), which is safe to do since we can reestablish the loop invariant LI using the left disjunct in the inner disjunction.

A full proof can be found in our Rocq development in file `coroutines_implementation_code.v`.

E.1.2 Client module. We illustrate usage of our coroutine module on a simple client whose code can be seen in Figure 13. For this client, we show the following specification:

$$\begin{aligned} & \text{addrmain} \vdash^{\text{wf}} c_{\text{main}} \text{ —*} \\ & \text{ewp} [\text{invoke addrmain}]; F \langle \perp \rangle \left\{ w, \begin{array}{l} \exists \text{vala valb}, w = \text{immV} [\text{vala}; \text{valb}] * \\ (\text{vala} = 1 \implies \text{valb} = 42) \end{array} \right\} \end{aligned}$$

meaning that as long as we hold the corresponding closure resource, we can run the `$main` function safely, and if it returns, it will return two integers such that if the first is 1, then the second is 42.

The proof of this specification is a classic concurrent separation logic example. We detail it here, as well as in our Rocq development, to illustrate that our logic is strong enough to capture reasoning principles for parallel composition for an implementation based on effects.

The proof relies on applying the specification for `$par`. We show our choice of P_1, Q_1, P_2, Q_2 and I in Figure 14. Our invariant I is a three-case disjunction: either `$x` and `$y` are both still 0, or `$x` has been updated but `$y` not yet, or both variables have been updated. Importantly, it is not possible for `$y` to be 1 if `$x` is not 42. In order to keep track of what variable update has been done already, we

```

1765 1 ;; run with the specfx fork of the reference interpreter with
1766 2 ;; ./wasm coroutin.wat -e '(module instance)' -e '(register "coroutine")' client.wat -e '(
1767   module instance)' -e '(invoke "main")'
1768 3 ;; it should display
1769 4 ;; [0 42] : [i32 i32]
1770 5 (module ;; client
1771 6   (type $func (func))
1772 7   (import "coroutine" "yield" (func $yield)) ;; needs yield and par
1773 8   (import "coroutine" "par" (func $par (param (ref $func)) (param (ref $func))))
1774 9   (global $x (mut i32) (i32.const 0))           ;; shared `data' $x
1775 10  (global $y (mut i32) (i32.const 0))           ;; shared `flag' $y
1776 11  (global $a (mut i32) (i32.const 0))           ;; reader's copy of the value of $y
1777 12  (global $b (mut i32) (i32.const 0))           ;; reader's copy of the value of $x
1778 13  (elem declare funcref (ref.func $f1))         ;; make $f1 available
1779 14  (elem declare funcref (ref.func $f2))         ;; and $f2
1780 15  (func $f1
1781 16    i32.const 42                                ;; write 42 to $x
1782 17    global.set $x
1783 18    call $yield                                  ;; yield
1784 19    i32.const 1                                  ;; set the flag
1785 20    global.set $y)
1786 21  (func $f2
1787 22    global.get $y                                ;; read from $y
1788 23    global.set $a                                ;; save in $a
1789 24    call $yield                                  ;; yield
1790 25    global.get $x                                ;; read from $x
1791 26    global.set $b                                ;; save in $b
1792 27  (func $main (result i32 i32)
1793 28    ref.func $f1                                ;; run $f1 and $f2 concurrently
1794 29    ref.func $f2
1795 30    call $par
1796 31    global.get $a                                ;; look up what the reader saw
1797 32    global.get $b)
1798 33  (export "main" (func $main)))                ;; make $main available outside
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813

```

Fig. 13. A client module using our coroutine library to run message-passing. set and get respectively take their arguments and put their results on the stack.

use *ghost resources* that assist in reasoning but do not directly correspond to a piece of the Wasm state. We use the *one-shot algebra* [Jung et al. 2016], which has two possible states: Pending, and Shot. When allocating a new resource, we get ownership of the resource in the Pending state. We can then later decide to update Pending to Shot, symbolising that a change has taken place. Pending and Shot cannot be simultaneously owned. The resource in the Shot state is duplicable, meaning we can ‘keep a record’ of the fact that the change has taken place. For our example, we allocate two one-shots, which we identify by their *ghost names* γ_x and γ_y . Hence $\{\text{Pending}\}^{\gamma_x}$ means that $\$x$ has not been updated yet, and $\{\text{Shot}\}^{\gamma_x}$ means it has been updated; likewise for γ_y and $\$y$. As displayed

in Figure 14, we place a ghost resource next to each global variable resource in the definition of the invariant I , choosing Shot or Pending depending on the value of the global variable.

When establishing the pre-condition of the specification for $\$par$, we must prove specifications for $\$f_1$ and $\$f_2$. For $\$f_1$, the proof follows this scheme:

- (1) $\{I\}$ Start with I (and vacuous P_1)
`i32.const 42`
`global.set $x`
- (2) $\{I * [\text{Shot}]^{Y_x}\}$ Reestablish I but keep a copy of $[\text{Shot}]^{Y_x}$
`call $yield`
- (3) $\{I * [\text{Shot}]^{Y_x}\}$ Using the specification for $\$yield$
`i32.const 1`
`global.set $y`
- (4) $\{I * [\text{Shot}]^{Y_x}\}$ Reestablish I

We start in (1) with the invariant I (and the pre-condition P_1 , which in this case is vacuous). No matter which disjunct of I is true, it is possible to update $\$x$ to 42 and reestablish the invariant in (2), updating the ghost resource $[\text{Pending}]^{Y_x}$ to $[\text{Shot}]^{Y_x}$ if necessary. Since $[\text{Shot}]^{Y_x}$ is persistent, we can retain a copy of it when reestablishing the invariant I in (2). Next, we apply the specification for $\$yield$, which consumes I and restores it upon return in (3). The resource $[\text{Shot}]^{Y_x}$ in the proof environment is unused and simply carried forward from (2) to (3) by framing. Now, because we own both I and $[\text{Shot}]^{Y_x}$, we know that we cannot be in the first disjunct of I , since it is not possible to simultaneously own $[\text{Pending}]^{Y_x}$ and $[\text{Shot}]^{Y_x}$. Thus, we are in one of the other two disjuncts, which allows us to update $\$y$ and reestablish the invariant in (4), updating the ghost resource $[\text{Pending}]^{Y_y}$ to $[\text{Shot}]^{Y_y}$ as necessary. This concludes the proof for $\$f_1$.

For $\$f_2$, the proof follows this scheme:

- (1) $\{I * P_2\}$ Start with I and P_2
`global.get $y`
`global.set $a` If $\$y$ is 1, keep a copy of $[\text{Shot}]^{Y_y}$
- (2) $\{I * \text{addrb} \xrightarrow{\text{wg}} - * ((\text{addrb} \xrightarrow{\text{wg}} 1 * [\text{Shot}]^{Y_y}) \vee \text{addrb} \xrightarrow{\text{wg}} 0)\}$
`call $yield` Using the specification for $\$yield$
- (3) $\{I * \text{addrb} \xrightarrow{\text{wg}} - * ((\text{addrb} \xrightarrow{\text{wg}} 1 * [\text{Shot}]^{Y_y}) \vee \text{addrb} \xrightarrow{\text{wg}} 0)\}$
`global.get $x`
`global.set $b`
- (4) $\{I * ((\text{addrb} \xrightarrow{\text{wg}} 1 * \text{addrb} \xrightarrow{\text{wg}} 42) \vee (\text{addrb} \xrightarrow{\text{wg}} 0 * \text{addrb} \xrightarrow{\text{wg}} -))\}$

We start in (1) with the invariant I and precondition P_2 . The first step is reading from $\$y$ to update $\$a$. If we are in the third disjunct of I , then $\$a$ gets value 1, and we can keep a copy of $[\text{Shot}]^{Y_y}$ in (2) since that resource is duplicable. If we are in one of the other disjuncts, we know that the new value of $\$a$ is 0. Next, we apply the specification for $\$yield$, which consumes I and gives it back upon return in (3). The other resources in our proof environment are unused by this step and simply carried forward from (2) to (3). Lastly, we read from $\$x$ to update $\$b$. In the case where we own $\text{addrb} \xrightarrow{\text{wg}} 1$, we also own $[\text{Shot}]^{Y_y}$ so we cannot be in the first two disjuncts of I , as one cannot simultaneously own $[\text{Pending}]^{Y_y}$ and $[\text{Shot}]^{Y_y}$. Hence we know that the new value of $\$b$ is 42. In the other case, it does not matter in which disjunct of I we are; in all three cases the read and write are safe to execute, and we need not make a note of the new value of $\$b$. The rest of the proof is simply showing that the proposition in (4) implies the post-condition Q_2 , which is straightforward.

$$\begin{aligned}
P_1 &= \top \\
Q_1 &= [\text{Shot}]^{Y_x} \\
P_2 &= \exists \text{vala valb, } \text{addra} \xrightarrow{\text{wg}} \text{vala} * \text{addrb} \xrightarrow{\text{wg}} \text{valb} \\
Q_2 &= \exists \text{vala valb, } \text{addra} \xrightarrow{\text{wg}} \text{vala} * \text{addrb} \xrightarrow{\text{wg}} \text{valb} * (\text{vala} = 1 \implies \text{valb} = 42) \\
I &= (\text{addrx} \xrightarrow{\text{wg}} 0 * \text{addy} \xrightarrow{\text{wg}} 0 * [\text{Pending}]^{Y_x} * [\text{Pending}]^{Y_y}) \vee \\
&\quad (\text{addrx} \xrightarrow{\text{wg}} 42 * \text{addy} \xrightarrow{\text{wg}} 0 * [\text{Shot}]^{Y_x} * [\text{Pending}]^{Y_y}) \vee \\
&\quad (\text{addrx} \xrightarrow{\text{wg}} 42 * \text{addy} \xrightarrow{\text{wg}} 1 * [\text{Shot}]^{Y_x} * [\text{Shot}]^{Y_y})
\end{aligned}$$

Fig. 14. Resources for the specification of the client of our coroutine library.

Final Theorem. Using the Adequacy Theorem from §4.5 on the specification above, we can prove the following theorem, which makes no mention of Iris:

THEOREM E.1. *If Cor is the module whose code is shown in Figure 3, and Cl is the module whose code is shown in Figure 13, and if the host program*

[instantiate Cor ; instantiate Cl ; invoke addrmain]

terminates on a value w , then there exists two integers vala and valb such that $w = \text{immV } [\text{vala}; \text{valb}]$, and moreover, if vala is 1, then valb is 42.

A full proof can be found in our Rocq development in file `coroutines_client.v`.

E.2 Generator

We define the protocol using a standard trick [Filliâtre and Pereira 2016; Pottier 2017]:

$$!x \text{ xs}(x) \langle \text{Permitted } (x :: \text{xs}) * I \text{ xs} \rangle ?() \langle I (x :: \text{xs}) \rangle$$

where $\text{Permitted } \text{xs}$ states that xs is a reversed prefix of the naturals, and I is a client-chosen predicate that can be used to track which naturals have already been generated. The idea is that when the generator suspends, it sends natural number x to the client. The client has already received numbers xs , enforced with $I \text{ xs}$, so x is the smallest natural that have not already been sent, enforced with $\text{Permitted } (x :: \text{xs})$. Further, if the client wants to resume the generator again, they must ‘admit’ that they have now received x , i.e. $I (x :: \text{xs})$. Using this protocol, we were able to give and prove a precise specification for the `$sum_until` function. Details of specifications and proofs can be found in the Rocq mechanisation.