

Effect Handlers All the Way Down

BRIAN CAMPBELL, University of Edinburgh, United Kingdom

SAM LINDLEY, University of Edinburgh, United Kingdom

WILMER RICCIOTTI, University of Edinburgh, United Kingdom

IAN STARK, University of Edinburgh, United Kingdom

Effect handlers are a popular programming language idiom providing a unified interface to a range of features: exceptions, state, I/O, concurrency, dynamic binding, and more. Effect handlers are typically introduced in high-level languages and implemented through existing lower-level constructions: by contrast, in this work we introduce ASmFX, an abstract assembly language where a handful of instructions directly express the processor-level control-flow requirements for effectful programming. We demonstrate the expressiveness of ASmFX with a compiler from a functional core calculus that generalises all of deep, shallow, and parameterised handlers with multi-shot resumptions. A key contribution is our proof of compiler correctness, obtained by enriching the source program with annotations that track the behaviour of compiled code fragments.

1 Introduction

Effect handlers are a powerful and elegant programming language abstraction providing facilities to define and control a wide range of program behaviours classed as “side effects”. Operationally, they are similar to resumable exception handlers, but they can also be used to model state, I/O, asynchronous programming, concurrency, dynamic binding, and probabilistic programming [6, 8, 9]. Effect handlers are particularly helpful in achieving clean separation of concerns and modularity and have recently been implemented in many research and mainstream languages. Implementations of effect handlers in high-level languages often convert them to existing features like closures, continuations, or prompts [15, 20, 28, 30]; while low-level approaches, on the other hand, often manipulate an existing stack [2, 12, 25]

Our approach in this work is different: instead of translating effect handlers away into other constructs, we provide a foundation for studying their direct implementation by propagating them all the way down to an abstract instruction set architecture with a primitive notion of effect handlers. In particular, we want this language to allow reasoning about effect handlers and concrete low-level implementations of them in the same calculus. We do not know of existing studies of a direct instruction-set representation of effect handlers.

As a starting point for our investigation, we define SFX, a core calculus with effect handlers to act as our high-level effectful source language (Section 3): its design goal is to provide a sufficiently rich effect-handling infrastructure, where resumptions may be single-shot, affine, or multi-shot, where the traditional setting of *deep* handlers can be extended by *parameterised* handlers, or replaceable handlers (such as the *sheep* handlers of WAsmFX [25]). The source language is thus capable of expressing a large range of effects, including but not limited to state, exceptions, nondeterminism, and dynamic scoping. Following prior work [16, 17], we express the semantics of SFX via a CEK-like abstract machine (Section 4).

We then introduce ASmFX (Section 5), an abstract register-based instruction set architecture with primitives to define, install, and invoke effect handlers. The goal of this ISA is to support all forms of effect handler expressible in SFX (and possibly more). Concretely, ASmFX adds to a traditional register-based architecture seven additional instructions, four of which allow the extended control-flow behaviour of effectful programs; the remaining three implement additional load operations

needed for creating special data structures used by ASmFX for representing continuations, handlers, and external code pointers as first-class values. As a reasoning tool, ASmFX does not prescribe particular memory layouts or restrict the amount of data that can be stored in registers. The semantics of the ISA is expressed by means of an abstract machine, which uses, along with main memory, an independent stack-like data structure describing the current effect handler state.

ASmFX provides an interesting computational model in its own right and can be used to study effectful programs written directly in a suitable ISA or as the intermediate language of a compiler for high-level languages. In Section 6 we describe a compiler for SFX targeting ASmFX. The compiler yields both ASmFX assembly code and an augmented version of the original source code annotated to denote which ASmFX code addresses implement each program expression. The augmented (ASFX) code is crucial for stating and proving our main correctness theorems.

We conclude our investigation of the expressive power of ASmFX with a proof of correctness of the SFX compiler, described in Section 7. Due to fundamental differences between declarative and procedural languages, proving directly that compiled ASmFX source code simulates corresponding SFX source code directly is particularly challenging. Instead, we split our proof into two parts, using ASFX as an intermediate step. ASFX provides a generalisation of both SFX and ASmFX by tracking intermediate ASmFX machine states directly in (an extension of) SFX: specifically, ASFX extends the SFX notion of continuation with *epilogues*, which denote fragments of machine code that have no direct representation in the source, but must be executed nonetheless. To our knowledge, this proof technique, and in particular our use of epilogues, is novel. The proof constitutes a key technical contribution of our work.

2 Overview

In order for us to understand the requirements of an assembly-level implementation of effect handlers, we will present a few examples in high-level pseudo-code, highlighting a few different uses of handlers and what we would expect their compiled version to look like.

Before we begin, let us establish terminology for the various layers of an effectful program: we distinguish *external code*, which is any part of the program which may install handlers but does not necessarily run in the context of a handler, and *client code*, which is called by the external code after installing a handler, and will return to the external code that invoked it after terminating normally; *handler code* includes globally defined handlers in the form **handler** h {**return** $x \rightarrow M \mid \#t(y) \rightarrow (r).N \mid \dots$ }, consisting of a *return clause* whose purpose is to post-process the result of a client computation before returning to the external code, and zero or more *operation clauses* expressing the semantics of effectful operations such as $\#t(y)$, where $\#t$ is an operation tag, and y its formal argument, which may resume the client code that invoked the effect by means of a resumption r , or return control to the external code.

An effectful assembly language should provide instructions to transition between these code layers, installing and uninstalling effect handlers as needed.

2.1 Runners

One of the simplest applications of effect handlers is the implementation of *runners* [1]: the functionality of runners corresponds to handlers in which a resumption may only be invoked once by means of a tail call. To illustrate this functionality, we consider an example in which a certain program needs to track a limited resource budget by means of a handler. The handler stores its budget within a mutable variable *budget*.

```
var budget := 100
handler track_budget {
```

```

99      return  $x \rightarrow$  return  $\langle x, budget \rangle$ 
100  | #charge( $n$ )  $\rightarrow$  ( $r$ ).
101      if ( $budget \geq n$ ) then  $budget := budget - n$ ; resume  $r(\langle \rangle)$ 
102      else return  $\langle 0, -1 \rangle$ 
103  }
104  let  $\langle x, balance \rangle \leftarrow$ 
105      handle {
106          let  $\_ \leftarrow$  do #charge(10) in
107          let  $a \leftarrow p()$  in let  $b \leftarrow q()$  in return  $a + b$ 
108      } with track_budget
109  in print_int(balance)

```

The program uses the **handle** keyword to execute client code within the context of the handler *track_budget*, providing access to an effect *#charge(n)*. The client code charges the budget for an amount of 10, then runs two subroutines *p()* and *q()* which may charge the budget as well; finally it sums the results of the two subroutines and returns. The handler *track_budget* provides a return clause that pairs the result x of the client code with the remaining budget; the operation clause for *#charge* performs different actions depending on the available budget: if there is enough budget to grant the request, it subtracts the requested value from the variable *budget* and then resumes the client code; otherwise, it aborts the client code and returns $\langle 0, -1 \rangle$, where 0 is a dummy result, and the negative balance -1 signals to the external code that the client failed to complete.

At an assembly level, to implement the code above, we can assume that the client code is at a memory location denoted by label *budget_client*, and that the handler is implemented by clauses identified by labels *track_budget.ret* and *track_budget.charge*. The client would use instructions *do* and *return* respectively to perform an effect and to return to the external code. Registers $\$a0, \$a1, \dots$ are used to exchange arguments/results between function and operation calls.

```

125 budget_client:
126     $a0 := 10
127     do #charge
128     ; [...] code to call p()
129     $t0 := $a0      ; save result of p() to scratch register
130     ; [...] code to call q()
131     $a0 := $t0 + $a0 ; add result of q() to previously saved result
132     return

```

To execute this code under the handler *track_budget*, our assembly language provides an instruction *resume* to run a continuation under a handler, and instructions *loadk* and *loadh* to load continuations and handlers. To tell the machine where to jump once the handler terminates correctly, we will use instruction *loadl*, initialised to the address of the instruction after *resume*. The external code would therefore be able to run *budget_client* by means of the following code:

```

138     loadk budget_client
139     loadh (return=track_budget.ret, #charge=track_budget.charge)
140     loadl budget_client_done
141     resume
142 budget_client_done:

```

The *loadk* instruction creates a continuation invoking the code at label *budget_client*; similarly, *loadh* creates a handler from its return and operation clauses; and *loadl* a *leave record* pointing to code label *budget_client_done*. The continuation, the handler, and the leave record may be held in distinguished registers $\$kr$, $\$hr$, and $\$lr$.

The `track_budget.ret` clause needs to leave the client result in `$a0` untouched, and copy the current value of `budget` into `$a1`, then yield control to the external code by means of a suitable instruction `exit` (operationally, this will jump to the code address specified by register `$1r`).

```
track_budget.ret:
    $a1 := mem[budget]
    exit
```

When running the `track_budget.charge` clause, the machine automatically loads the current handler, leave record, and the continuation to the client into the corresponding registers: this allows the client to be resumed by a simple resume instruction, without any additional `loadh`, `loadl`, or `loadk` instructions. To abort the client and yield control to the external code, we may be able to use the same `exit` instruction used in the return clause:

```
track_budget.charge:
    $t0 := mem[budget]
    if ($t0 < $a0) branch abort
    $t0 := $t0 - $a0
    mem[budget] := $t0
    resume
abort:
    $a0 := 0
    $a1 := -1
    exit
```

2.2 Parameterised handlers

The budget example we gave uses a mutable global variable to store the budget balance; while this variable should only be used by the handler, there is no way to prevent unauthorised code from altering it. A common idiom with effect handlers is to represent the private state of a handler as a *parameter*, which is initialised when the handler is installed and can be updated before invoking a client resumption. By using parameterised handlers, the budget example can be rewritten as follows:

```
handler track_budget(budget) {
    return x → return ⟨x, budget⟩
    | #charge(n) → (r).
    if (budget ≥ n) then resume r(⟨⟩) with track_budget(budget - n)
    else return ⟨0, -1⟩
}

let ⟨x, balance⟩ ←
    handle {
        let _ ← do #charge(10) in
        let a ← p() in let b ← q() in return a + b
    } with track_budget(100)
in print_int(balance)
```

We would like to support parameterised handlers at the assembly level, using registers to pass the parameters. The `loadh` instruction can accept an additional state operand declaring a list of registers containing the handler state that must be saved. For instance, the external code would load the `track_budget` handler with an initial budget of 100 as follows:

```
$s0 := 100
loadh (return=track_budget.ret, #charge=track_budget.charge), state:($s0)
```

This ensures that whenever the handler clauses are run, the register `$s0` contains the same value it contained when the `loadh` instruction was executed. Similarly, the `track_budget.charge` code, before resuming the client, would update the state and refresh the handler:

```
track_budget.charge:
; $s0 contains the current budget
if ($s0 < $a0) branch abort
$s0 := $s0 - $a0
loadh (return=track_budget.ret, #charge=track_budget.charge), state: ($s0)
resume
; ...
```

Refreshing the handler is crucial to inform the machine that the updated content of `$s0` should be used the next time the handler code is run.

2.3 Producer/consumer

In some cases, when implementing an effect, we may want to resume a client under a different handler from the one that was originally invoked. Consider the following example consisting of two functions *prod* and *cons*: the first uses the `#send` operation to send two integer values over a shared channel, while the second invokes a `#recv` operation to get two integer values from the same channel and return their sum.

```
// Producer sends a couple of values and then stops
fun prod(_) {
  do #send(1); do #send(2); return ⟨⟩
}

// Consumer receives a couple of values and then returns their total
fun cons(_) {
  let x ← do #recv(⟨⟩) in let y ← do #recv(⟨⟩) in return x + y
}
```

The two operations are implemented in two different handlers. The producer handler *ph* parameterised over a consumer resumption *cr* handles `#send` by resuming *cr* under the consumer handler *ch*, which is passed as a parameter the producer resumption *pr*:

```
handler ph(cr) {
  return ⟨⟩ → handle prod(⟨⟩) with ph(cr)
  | #send(x) → (pr).let r ← (cr with ch(pr)) in resume r(x)
}
```

Symmetrically, the consumer handler *ch*, parameterised over a producer resumption *pr*, handles `#recv` by resuming *pr* under the producer handler *ph*, which is passed as a parameter the consumer resumption *cr*:

```
handler ch(pr) {
  return result → return result
  | #recv(⟨⟩) → (cr).let r ← (pr with ph(cr)) in resume r(⟨⟩)
}
```

The main program is started by handling the consumer *cons* under an appropriate handler: ideally, this would be *ch*, but the producer has not started yet, so we do not have a producer resumption to pass as a parameter yet; instead, we use an initial consumer handler *ih* which handles the `#recv` operation by starting *prod* rather than calling a producer resumption.

```
// Program starts the consumer under an initial handler
fun main(_) {
```

```

246   handle cons(⟨⟩) with ih(⟨⟩)
247 }
248 // An initial version of the consumer handler, which starts the producer rather than resuming it
249 handler ih(⟨⟩) {
250     return x → return x
251     | #recv(⟨⟩) → (cr).handle prod(⟨⟩) with ph(cr)
252 }

```

When control is transferred to an operation clause, the resumption received by that clause would normally continue execution under the handler defining that operation, according to *deep handler* semantics. However, one can view a resumption as a pair consisting of a head continuation and of the handler delimiting it, and allow the handler to be replaced. The **resume** keyword works by calling the continuation after reinstalling the delimiting handler. The **with** keyword allows us to replace the delimiting handler with a handler of our choice, recovering the semantics of *sheep handlers* [25]. At the machine level, a resumption can be represented as a continuation value and a handler value (which need to be loaded into the registers \$kr and \$hr to be executed), and the handler replacement functionality of **with** can be expressed by means of the `loadh` instruction. For example, we could implement the `#send` clause of handler *ph* by means of the following assembly:

```

263 ph.send:
264     ; handler parameter cr (consumer resumption) in ($s0, $s1)
265     ; value to be sent in $a0, producer resumption pr in ($kr, $hr)
266     $t0 := $kr
267     $t1 := $hr
268     $kr := $s0 ; continuation of cr
269     $s0 := $t0 ; continuation of pr
270     $s1 := $t1 ; delimiting handler of pr
271     ; load $hr with ch(pr)
272     loadh (return=ch.ret, #recv=ch.recv), state:($s0, $s1)
273     ; resume (cr with ch(pr))
274     resume

```

2.4 Multi-shot resumptions

Our final example illustrates more complex handlers by evaluating propositional formulas under a non-deterministic choice of variables. We use a Boolean valued effect *#flip* to generate values for the propositional variables, then combine them with standard Boolean operators to evaluate the formula of interest. The tautology $(A \wedge B) \vee \neg A \vee \neg B$ can then be expressed by the following code:

```

280 fun prop(⟨⟩) {
281     let a ← do #flip(⟨⟩) in
282     let b ← do #flip(⟨⟩) in
283     return (a && b) || not a || not b
284 }

```

The value of the propositional variables depends on the handler implementing the *#flip* effect. We can have trivial handlers always returning **true** or **false**, or randomly returning either value (given a source of randomness). Most interestingly, we are not limited to one truth value, but we may see what happens when the propositional variables are assigned all possible combinations of truth values. The following *tautology* handler implements *#flip* by resuming the (Boolean-valued) client twice, once with **true** and once with **false**, and then taking the conjunction of the results of the two resumption calls. The final result will be **true** if and only if the client evaluates to **true** for all the truth assignments of the propositional variables, i.e. if the client represents a tautology.

```

293 handler tautology(⟨⟩) {
294

```

```

295     return  $x \rightarrow$  return  $x$ 
296 | #flip( $\langle \rangle$ )  $\rightarrow$  ( $r$ ).
297     let  $x \leftarrow$  resume  $r(\text{false})$  in
298     let  $y \leftarrow$  resume  $r(\text{true})$  in
299     return  $x \ \&\& \ y$ 
300 }
301 handle prop( $\langle \rangle$ ) with tautology( $\langle \rangle$ )

```

This example differs from the previous ones in two ways: first, the resumption is invoked more than once; and second, the resumption call is not the last action in the operation clause — it is not a tail call. This means that when the resumption (really, the client computation) completes its execution, it must not return to the external code, but to the operation clause that invoked it; in other words, when performing a non-tail resumption, the handler code that invoked the resumption becomes the external code with respect to the client code it resumed.

An assembly-level machine needs to be made aware of this fact to correctly implement the control flow. In AsmFX we have allocated a distinguished register $\$lr$ to contain the address of the code that needs to be executed after a client completes its work. Note that the handler code invoking the resumption will want to preserve some registers, and in particular the original value of $\$lr$, which still contains the address the handler must jump to upon termination: these registers can be specified as an additional save parameter to `loadl`, and will be restored to their original value when an `exit` instruction is performed.

We can then implement the `#flip` operation as follows:

```

316 tautology.flip:
317     $a0 := 0      ; false
318     loadl l1, save:($lr)
319     resume       ; will jump to l1 when the resumption ends
320 l1:
321     ; $lr restored to original value
322     $t0 := $a0    ; save result
323     $a0 := 1      ; true
324     loadl l2, save:($lr,$t0)
325     resume       ; will jump to l2 when the resumption ends
326 l2:
327     ; $lr and $t0 restored to the original value
328     $a0 := $t0 && $a0
329     exit         ; return  $r(\text{false}) \ \&\& \ r(\text{true})$ 

```

3 A source calculus for effect handlers

The first step towards our goal of studying the compilation of effect handlers is to introduce a core declarative language expressive enough to easily implement the idioms we are interested in. For our purposes, it makes sense to restrict ourselves to programs using only first order global functions, since techniques to convert nested functions to global definitions are well established (see for instance [18, 23]) and we do not intend to discuss them in this work; we do however support simple and mutual recursion. For similar reasons, we will use globally defined handlers. Based on these requirements, we define the source calculus SFX, whose syntax is given by the following grammar (where we have highlighted non-standard expressions).

Value types	$A, B ::= \tau \mid \langle \bar{A} \rangle \mid \boxed{A \xrightarrow{C} D}$	Effect types	$E ::= \{\#t : A \rightarrow \bar{B}\}$
Comp. types	$C, D ::= A!E$	Top-level types	$G ::= A \Rightarrow C \mid (A)C \Rightarrow D$

Open values	U, V, W	$::=$	$x \mid c \mid \langle \bar{V} \rangle$
Computations	M, N	$::=$	$\mathbf{return} \ V \mid \oplus(V) \mid f(V) \mid \mathbf{if}(V, M, N)$ $\mid \mathbf{let} \ x \leftarrow M \mathbf{in} \ N \mid \mathbf{unpack} \ \bar{x} \leftarrow V \mathbf{in} \ M$ $\mid \mathbf{do} \ \#t(V) \mid \mathbf{handle} \ M \mathbf{with} \ h(V)$ $\mid \mathbf{newbroom}(f) \mathbf{with} \ h(V) \mid U \mathbf{with} \ h(V) \mid \mathbf{resume} \ U(V)$
Literal values	u, v, w	$::=$	$c \mid \langle \bar{v} \rangle \mid (\kappa, h(v))$
Environments	γ	$::=$	$\bar{x} \mapsto \bar{v}$
Continuations	κ, κ'	$::=$	$0 \mid (x).M[\gamma] \cdot \kappa \mid h(v) \cdot \kappa$
Functions	F	$::=$	$\mathbf{fun} \ f(x) \{M\}$
Handlers	H	$::=$	$\mathbf{handler} \ h(x) \{\mathbf{return} \ y \rightarrow M \mid \#t(z) \rightarrow (r).N\}$
Specifications	Σ	$::=$	$\bar{F}; \bar{H}$
Programs	P	$::=$	$\Sigma; M$

As syntactic sugar, we use a vector notation such as \bar{x}_n to indicate the sequence $x_0 \dots x_{n-1}$; the same notation, without subscript, indicates a sequence whose length is unspecified because it is irrelevant or can be inferred from the context.

SFX is a fine-grain call-by-value source language with a basic effect type system; terms fall into two categories: *values*, representing what the program computes, and *computations* describing how values are computed. Value types A comprise primitive types τ , n -ary tuples of value types $\langle \bar{A} \rangle$, and the novel *broom* type $A \xrightarrow{C} D$ generalising functions to resumptions, where a type A argument produces a computation of type C under a handler giving a result of type D . Computation types $C, D = A!E$ combine a return type A with an effect type E specifying the effects that computations of this type may invoke. Effect types assign signatures $A \rightarrow B$ to a finite number of operation tags $\#t$, where A is its input type and B is its output type. If E is empty then by abuse of notation we write computation type $A!\{\}$ simply as A . Handlers are *parameterised* [19, 26], so a handler type is of the form $(A)C \Rightarrow D$ where A is the type of the parameter, C is the type of the computation being handled, and D is the type of the resulting computation. With resumptions and handlers limited to a single argument parameter we use tuples to used to express n -ary versions.

Value terms U, V, W comprise variables x , constants c , and n -tuples: these are open, since they may contain variables that reference values stored in an environment. We distinguish them from *literal* values u, v, w , which are closed, and beside constants and tuples can also be resumptions consisting of a continuation paired with a distinguished handler in the form $(\kappa, h(v))$ — we call these *broom* resumptions, or just *brooms*, because like the tool used for sweeping they consist of a head κ and a handle $h(v)$. Programs can only refer explicitly to value terms: in particular, broom values cannot appear in a program, except as variables of continuation type. At runtime, value terms will evaluate to literal values.

Besides returning a value, primitive operations, conditionals, sequencing (**let**), tuple decomposition (**unpack**), function application ($f(V)$, where f is the name of a globally defined function), we provide computation terms **do** $\#t(V)$ to perform an effect $\#t$ with input V . Computations can also handle effects with the syntax **handle** M **with** $h(V)$, where M is the client computation, h the name of a globally defined handler, and V its actual parameter, representing its initial state.

Our reason for the “broom” terminology is to manage all of deep, shallow, and sheep handlers: a broom resumption has a head (the continuation) and a handle (its effect handler); and in some cases either component may be replaced to refresh the broom. Resuming a broom invokes the continuation and also installs the handler: syntax **resume** $U(V)$ invokes broom U with argument V . Brooms most often represent a suspended computation that has requested performance of an effect, being passed to the corresponding operation clause of the closest matching handler installed. It is

also possible to convert a function f into a broom by means of the syntax **newbroom**(f) **with** $h(V)$: this broom will take an argument as input, install the handler $h(V)$, and then apply f to the argument received. Fitting a broom U with a new handle $h(V)$ is written as U **with** $h(V)$, allowing us to implement sheep handlers.

Top-level handlers **handler** $h(x) \{ \text{return } y \rightarrow M \mid \#t(z) \rightarrow (r).N \}$ accept a parameter x and provide operation clauses for each effect $\#t$ they handle: the effect receives an input argument z , and additionally a broom r that allows the handling code to resume execution within the client, by *resuming* r with a value which will appear, to the client, as the result of evaluating the effect $\#t(z)$; the handler also provides a return clause receiving in the variable y the value computed by the client. Given such a handler, we will use the following metalinguistic operations to refer to its bound variables: $\text{dom}(h)$ for the parameter variable x , $\text{dom}(h_{\text{return}})$ and $\text{dom}(h_{\#t})$ for the formal arguments of the return and operation clauses (respectively, y and z), and $\text{res}(h_{\#t})$ for the continuation argument r of the operation clause. The syntax to define top-level functions is self-explanatory.

Continuations κ are sequences of continuation frames: the base case $\mathbf{0}$ expresses program termination. Pure continuation frames $(x).M[\gamma]$ consist of a closure $M[\gamma]$ abstracted over a variable x : given a literal value, say v , returned by a computation and stored in variable x , the pure frame continues as M , where free occurrences of x have been replaced by v , and any other free variable is defined by γ ; pure computations are used in the intermediate steps of sequencing. Handler continuation frames $h(v)$ consist of a handler identified by its name h , applied to a literal value v representing its state; when a computation is running in the context of such a handler frame, the effects defined by h are available, and if that computation returns a value u , u is passed to the **return** clause of h .

Programs consist of any number of mutually defined handlers and functions plus a main computation that may use them.

3.1 Typing

Typing judgements for the source calculus include four main forms for values, computations, continuations, and handlers/functions. We also provide two minor judgements for environments and for the program. Judgements that depend on the specification Σ are annotated with Σ on the turnstile (\vdash_Σ). However, since Σ does not change in recursive typing assumptions, we will often assume it is fixed for a given program and omit it from the subscript. Figure 1 presents the typing rules.

Value typing. Value typing $\Gamma \vdash_\Sigma V : A$ or $\Gamma \vdash_\Sigma v : A$ assigns a value type A to an open value V or to a literal value v (even though we distinguish the two cases syntactically, they are typed under the same judgement and the same rules) under a typing context Γ associating variables to their value type. The interesting rule T-BROOM says that a broom $(\kappa, h(v))$ of type $A \xrightarrow{C} D$ consists of a continuation head κ taking a value of type A and returning a computation of type C , and a handle which is its delimiting handler h of type $(B)C \Rightarrow D$ instantiated with a value v of type B : $h(v)$ must transform the computation type C returned by κ to another computation type D .

Computation typing. Computation typing uses the judgement $\Gamma \vdash_\Sigma M : C$. The interesting rules are those for invoking and handling effects. An effectful operation **do** $\#t(V)$ (rule T-Do) has type $B!E$ if the effect type E assigns the type $A \rightarrow B$ to $\#t$ and V has type A . Rule T-HANDLE assigns the type D to **handle** M **with** $h(V)$ when M has type C , the handler's parameter V has type A , and that h is defined in the signature as a handler of type $(A)C \Rightarrow D$ linking all three together.

T-VAR	T-VAL	T-TUPLE	T-PRIM
$\frac{x : A \in \Gamma}{\Gamma \vdash_{\Sigma} x : A}$	$\frac{c \in \text{Val}(\tau)}{\Gamma \vdash_{\Sigma} c : \tau}$	$\frac{\Gamma \vdash_{\Sigma} \bar{V} : \bar{A}}{\Gamma \vdash_{\Sigma} \langle \bar{V} \rangle : \langle \bar{A} \rangle}$	$\frac{\Gamma \vdash_{\Sigma} V : A \quad \text{Type}(\oplus) = A \rightarrow B}{\Gamma \vdash_{\Sigma} \oplus(V) : B}$
T-BROOM	T-RET	T-APP	
$\frac{\vdash_{\Sigma} \kappa : A \rightarrow C \quad \Gamma \vdash_{\Sigma} v : B \quad \Sigma \vdash h : (B)C \Rightarrow D}{\Gamma \vdash_{\Sigma} (\kappa, h(v)) : A \xrightarrow{C} D}$	$\frac{\Gamma \vdash_{\Sigma} V : A}{\Gamma \vdash_{\Sigma} \text{return } V : A!E}$	$\frac{\Gamma \vdash V : A \quad \Sigma \vdash f : A \Rightarrow C}{\Gamma \vdash f(V) : C}$	
T-IF	T-LET	T-UNPACK	
$\frac{\Gamma \vdash_{\Sigma} V : \text{Bool} \quad \Gamma \vdash_{\Sigma} M : C \quad \Gamma \vdash_{\Sigma} N : C}{\Gamma \vdash_{\Sigma} \text{if}(V, M, N) : C}$	$\frac{\Gamma \vdash_{\Sigma} M : A!E \quad \Gamma, x : A \vdash_{\Sigma} N : B!E}{\Gamma \vdash_{\Sigma} \text{let } x \leftarrow M \text{ in } N : B!E}$	$\frac{\Gamma \vdash_{\Sigma} V : \langle \bar{A} \rangle \quad \Gamma, x : \bar{A} \vdash_{\Sigma} M : C}{\Gamma \vdash_{\Sigma} \text{unpack } \bar{x} \leftarrow V \text{ in } M : C}$	
T-DO	T-HANDLE		
$\frac{\Gamma \vdash_{\Sigma} V : A \quad E(\#t) = A \rightarrow B}{\Gamma \vdash_{\Sigma} \text{do } \#t(V) : B!E}$	$\frac{\Gamma \vdash_{\Sigma} M : C \quad \Gamma \vdash_{\Sigma} V : A \quad \Sigma \vdash h : (A)C \Rightarrow D}{\Gamma \vdash_{\Sigma} \text{handle } M \text{ with } h(V) : D}$		
T-WITH	T-RESUME		
$\frac{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D'}{\Gamma \vdash_{\Sigma} U \text{ with } h(V) : A \xrightarrow{C} D'}$	$\frac{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : A}{\Gamma \vdash_{\Sigma} \text{resume } U(V) : D}$		
T-NEWBROOM	T-ENV		
$\frac{\Sigma \vdash f : A \Rightarrow C \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D}{\Gamma \vdash_{\Sigma} \text{newbroom}(f) \text{ with } h(V) : A \xrightarrow{C} D}$	$\frac{\text{dom}(\gamma) = \text{dom}(\Gamma) \quad (\vdash_{\Sigma} \gamma(x) : \Gamma(x))_{x \in \text{dom}(\gamma)}}{\vdash_{\Sigma} \gamma : \Gamma}$		
T-KID	T-KLET	T-KHANDLE	
$\frac{}{\vdash 0 : A \rightarrow A!E}$	$\frac{\vdash \gamma : \Gamma \quad \Gamma, x : A \vdash N : B!E \quad \vdash \kappa : B \rightarrow B!E}{\vdash (x).N[\gamma] \cdot \kappa : A \rightarrow B!E}$	$\frac{\Sigma \vdash h : (A)B!E \Rightarrow B'!E' \quad \vdash v : A \quad \vdash \kappa : B' \rightarrow B''!E'}{\vdash h(v) \cdot \kappa : A \rightarrow B''!E'}$	
T-HANDLER			
$\frac{C = B!((\#t_i : B_i \rightarrow B'_i)_i \cup E) \quad D = B'!(E' \cup E) \quad x : A, y : B \vdash_{\Sigma} M : B'!E' \quad (x : A, z_i : B_i, r_i : B'_i \rightarrow D \vdash_{\Sigma} N_i : B'!E'_i) \quad (\text{handler } h(x) \{ \text{return } y \rightarrow M \mid (\#t_i(z_i) \rightarrow (r_i).N_i)_i \}) \in \Sigma}{\Sigma \vdash h : (A)C \Rightarrow D}$			
T-FUN	T-PROGRAM		
$\frac{x : A \vdash_{\Sigma} M : C \quad (\text{fun } f(x) \{M\}) \in \Sigma}{\Sigma \vdash f : A \Rightarrow C}$	$\frac{\vdash_{\Sigma} M : C \quad (\Sigma \vdash \Sigma(f) : A_f \Rightarrow C_f)_{f \in \text{dom}(\Sigma)} \quad (\Sigma \vdash \Sigma(h) : (A_h)C_h \Rightarrow D_h)_{h \in \text{dom}(\Sigma)}}{\vdash \Sigma; M : C}$		

Fig. 1. SFX: typing rules.

To type handle replacement U **with** $h(V)$ on a broom U with type $A \xrightarrow{C} D$, the handler h must have a type $(B)C \Rightarrow D'$ where C matches, and the parameter V must have the expected type B ; then the broom after replacing the handle will have an updated type $A \xrightarrow{C} D'$.

Rule T-RESUME assigns to **resume** $U(V)$ the type D whenever U has a broom type $A \xrightarrow{C} D$ and V has type A . Finally rule T-NEWBROOM lifts a function f from type $A \Rightarrow C$ to a broom type $A \xrightarrow{C} D$ by means of a handler h of type $(B)C \Rightarrow D$ instantiated on a value v of type B . Note that since brooms have a first class type, **newbroom** allows us to convert a function into a value that can be freely passed around as an argument or returned.

Continuation typing. The judgement $\vdash_{\Sigma} \kappa : A \rightarrow C$ states that the continuation κ receives a value of type A and continues as a computation of type C . Continuations include the identity continuation 0 , which takes a value of any type A and returns a (trivial) computation of type $A!E$, and composite continuations obtained by prepending to an existing continuation a pure frame or a handler frame. The typing judgement for pure frames extension T-KLET ensures that the computation N in the frame is well-typed under an environment $\Gamma, x : A$, where x is the abstracted variable, that the environment γ has type Γ , and that the base continuation κ has an input type matching values returned by N and that its output effect matches the effects available to N . Extension by handler frames (rule T-KHANDLE) requires h to be the name of a valid handler of type $(A)B!E \Rightarrow B'!E'$; the actual parameter v must have type A in the empty context, and the base continuation κ must have type $B' \rightarrow B''!E$, where B' matches the type of values returned by h and E' matches the output effect type of h .

Handler/function typing. Handlers and functions must be global items defined in the static specification Σ . The judgement T-HANDLER states that to give h the type $(A)C \Rightarrow D$ we need to be able to express C as $B!\{(\#t_i : B_i \rightarrow B'_i)_i\} \cup E$ and D as $B'!(E' \cup E)$, where E is the type of forwarded effects, i.e. client effects that are *not* handled by h , the $\#t_i$ are the effectful operations handled by h , and E' is the type of additional effects that may be raised by h . The rest of the judgement states that the definition of h in Σ must provide a return clause M and an operation clause N_i for each handled operation $\#t_i$, all of type $B'!E'$ matching the output type of the handler; the typing contexts for M and N_i assign types to the handler parameter x , to the return parameter y (for M) and to the operation argument z_i and the resumption r_i (for N_i). The typing of functions is standard.

4 An abstract machine for SFX

To provide a semantics for SFX, we develop an abstract machine in the style of the CEK machine [10]. The machine thus manipulates triples of control-environment-continuation, where the control is a computation or a closed value.

The machine is an adaptation of previous work on deep and shallow effect handlers [17]; however, our language presents a few differences, which are reflected into the definition:

- We lack type polymorphism and effect polymorphism and do not need to account for them in the machine.
- We include brooms as first-class terms of the language, with their own type. Brooms are structurally similar to the continuation data structure of the abstract machine both in this work and in [17]; however, while in our work SFX brooms and continuations of the abstract machine are structurally the same, their semantics differs subtly: brooms appear as the argument of a **resume**, where they are applied to a value: since the argument is a value, it cannot invoke an effect. On the other hand, abstract machine continuations behave like evaluation contexts: a machine triple $\langle M, \gamma, \kappa \rangle$ must be understood as the term $\kappa[M[\gamma]]$, where all the unhandled effects invoked in $M[\gamma]$ can be handled by κ . The two notions are related, in the sense that resuming a continuation κ with a value V is semantically equivalent to applying the evaluation context represented by κ to the trivial computation **return** V .
- We do not differentiate between deep and shallow handlers; instead we have a single kind of handlers (implemented similarly to deep handlers), and we use the handle replacement operation of brooms to express sheep handlers.

To save a few transition rules that are not relevant to the topic of this paper, we choose to assume that functions are translated to trivial handlers by a pre-processing operation:

$$\begin{aligned} \text{fun } f(x) \{M\} &\rightsquigarrow \text{handler } f(\langle \rangle) \{\text{return } x \rightarrow M\} \\ f(V) &\rightsquigarrow \text{handle return } V \text{ with } f(\langle \rangle) \end{aligned}$$

Closures are only needed in pure continuations. When we have a value V and an environment γ , we substitute eagerly via the operation $\llbracket V \rrbracket \gamma$. Such eager substitutions are trivial and always yield a literal value v .

The full machine is given in Figure 2: it consists of a dozen rules, but those implementing primitive operations, conditionals, and tuple unpacking are trivial. Sequencing **let** $x \leftarrow M$ **in** N is also standard, but it is worth noting that it proceeds by evaluating M under the continuation extended by a pure frame $(x).N[\gamma]$, where the variable x , which will receive the result of the evaluation of M , abstracts over the closure $N[\gamma]$ — this behaviour is expressed by the dual rule evaluating the closed value v when the continuation starts with a pure frame: notice that this rule reinstates the environment stored in the closure for N , which may have been destroyed by the intermediate evaluation steps for M .

Similarly, to evaluate a handled computation **handle** M **with** $h(V)$, we proceed by evaluating M under a continuation extended with a handler frame $h(\llbracket V \rrbracket \gamma)$. The dual rule evaluating the closed value v when the continuation starts with a handler frame proceeds to evaluate the **return** clause of the handler h , completely replacing the environment with only two elements: the actual value w of the handler parameter, and the actual return value v received as input.

The evaluation of an effectful operation **do** $\#t(V)$ depends on the current continuation, which must be of the form $\kappa^{-t} \cdot h^{+t} \cdot \kappa'$, where the superscript $-t$ means that κ does not handle $\#t$, and the superscript $+t$ expresses the fact that the handler h specified in this frame must provide a clause for $\#t$: in other words, h is the first handler in scope which will handle $\#t$. To execute this operation, we proceed to evaluate $h_{\#t}$ (the body of the operation clause for $\#t$ defined by h); the continuation is split: the base continuation κ' is kept as the current continuation, whereas the initial fragment $(\kappa^{-t}, h^{+t}(w))$ is stored in the environment as the value of the resumption variable $\text{res}(h_{\#t})$, along with the values for the handler parameter and for argument of $\#t$. The original environment γ is discarded, but may be reinstated by subsequent resumptions.

The expression **newbroom**(f) **with** $h(V)$ immediately produces the broom corresponding to the function f : this consists of a pure frame $(x).f_{\text{return}}$ (remember that functions are encoded as the return clause of a trivial handler), delimited by the handler $h(\llbracket V \rrbracket \gamma)$. The handle replacement expression U **with** $h(W)$ is evaluated by obtaining the broom $\llbracket U \rrbracket \gamma$, discarding its handle, and replacing it with $h(\llbracket W \rrbracket \gamma)$.

To evaluate a simple resumption **resume** $U(V)$, we apply the environment γ to U to obtain the corresponding literal continuation $(\kappa', h(w))$: then we proceed by evaluating **return** V , where the continuation κ has been extended by $\kappa' \cdot h(w)$.

We evaluate **return** V by converting it to the closed value $\llbracket V \rrbracket \gamma$.

5 An abstract assembly language with effect handlers

Having formalised the high-level effect idioms we intend to support by means of SFX and its abstract machine, it is now time to develop a low-level target language in the style of the typical instruction set architectures supported in hardware, but providing instructions to install effect handlers and perform effects: we call this language AsmFX.

AsmFX comprises a small set of instructions operating on registers, a main memory containing data and code, and an *effect context*, special storage manipulated as a stack and separated from main memory. Though close to typical instruction sets, AsmFX leaves several features abstract:

control	env.	continuation	\hookrightarrow	control	env.	continuation
$\oplus(V)$	γ	κ		$\llbracket \oplus \rrbracket (\llbracket V \rrbracket \gamma)$	γ	κ
if (V, M, N) where: $\llbracket V \rrbracket \gamma = \text{true}$	γ	κ		M	γ	κ
if (V, M, N) where: $\llbracket V \rrbracket \gamma = \text{false}$	γ	κ		N	γ	κ
let $x \leftarrow M$ in N	γ	κ		M	γ	$(x).N[\gamma] \cdot \kappa$
unpack $\overline{x_n} \leftarrow \langle \overline{V_n} \rangle$ in N	γ	κ		N	$(x_n \mapsto \llbracket V_n \rrbracket \gamma)$ $\cdot \gamma$	κ
handle M with $h(V)$	γ	κ		M	γ	$h(\llbracket V \rrbracket \gamma) \cdot \kappa$
do $\#t(V)$	γ	$\kappa^{-t} \cdot h^{+t}(w) \cdot \kappa'$		$h_{\#t}$	$(res(\#t) \mapsto (\kappa^{-t}, h^{+t}(w))$ $\cdot (dom(\#t) \mapsto \llbracket V \rrbracket \gamma)$ $\cdot (dom(h) \mapsto W))$	κ'
newbroom(f) with $h(V)$	γ	κ		$((x).f_{\text{return}}[],$ $h(\llbracket V \rrbracket \gamma))$ where: $x = dom(f_{\text{return}})$	γ	κ
U with $h(W)$ where: $\llbracket U \rrbracket \gamma = (\kappa', _)$	γ	κ		$(\kappa', h(\llbracket W \rrbracket \gamma))$	γ	κ
resume $U(V)$ where: $\llbracket U \rrbracket \gamma = (\kappa', h(w))$	γ	κ		$\llbracket V \rrbracket \gamma$	γ	$\kappa' \cdot h(w)$ $\cdot \kappa$
return V	γ	κ		$\llbracket V \rrbracket \gamma$	γ	κ
v	γ	$(x).N[\gamma'] \cdot \kappa$		N	$(x \mapsto v) \cdot \gamma'$	κ
v	γ	$h(w) \cdot \kappa$		h_{return}	$(dom(h_{\text{return}}) \mapsto v)$ $\cdot (dom(h) \mapsto w)$	κ

Fig. 2. SFX abstract machine

- Every register and every memory location is large enough to store any value, including certain complex values with special architectural meaning that we will call *records*, used to implement control flow.
- There are as many registers as may be needed by any program. Registers are referred by their name, starting with a dollar sign '\$'. While all registers can contain any value, we will use naming conventions to divide them in different categories for different purposes. The register \$pc is the program counter and it contains the memory address of the next instruction that will be executed.
- Only instructions explicitly needed for implementing effect handlers are included. Other standard instructions including unconditional jumps, conditional branches, arithmetic and logical operations, register moves, or load and store from main memory are left unspecified, but are intended to be available as needed, with their natural semantics.

Main memory is an array of words, and memory addresses a, a', \dots are indices in the array specified as natural numbers. Effects are specified by tags $\#tag$; the special tag *return* identifies the return clause of handlers. We now describe the syntax of ASMFx:

Handlers:	H	$::=$	$\text{return} \mapsto a_{\text{return}}, \#tag \mapsto a_{\#tag}$
Records:	R	$::=$	$R_{\mathcal{H}} \mid R_{\mathcal{K}} \mid R_{\mathcal{L}}$
Handler rec.:	$R_{\mathcal{H}}$	$::=$	$\mathcal{H}(H, \$r_h = v_h)$
Continuation rec.:	$R_{\mathcal{K}}$	$::=$	$\mathcal{K}(a_r, \$r_c = v_c, D_r)$
Leave rec.:	$R_{\mathcal{L}}$	$::=$	$\mathcal{L}(a_l, \$r_e = v_e)$
Values:	v	$::=$	$R \mid a \mid I \mid \dots$
Instructions:	I	$::=$	$\text{loadh } H, \text{state}(\overline{\$r_h}) \mid \text{loadk } a_r, \text{save}(\overline{\$r_c}) \mid \text{loadl } a_l, \text{save}(\overline{\$r_e})$ $\mid \text{resume} \mid \text{do } \#tag, \text{save}(\overline{\$r_c}) \mid \text{return} \mid \text{exit}$
Effect contexts:	C, D	$::=$	$\overline{\mathcal{H}(H, \$r_h = v_h, R_{\mathcal{L}})}$

Records. To implement effect handlers at the assembly level, we introduce representations of the key linguistic concepts we need as special values that can be directly stored in registers: in our terminology, these are called *records*. Records are similar to closures, and will be used to represent three kinds of objects:

- Effect handlers, by means of records in the form $\mathcal{H}(H, (\overline{\$r_h = v_h}))$, where H is a finite partial map from operation names to memory addresses of code to handle that operation, including one address providing code for returning from the handler; the second subterm is a finite partial map from registers names to values, used to store the handler state.
- Continuations, by means of records $\mathcal{K}(a_r, (\$r_c = v_c), D_r)$, where a_r is the address of the code that will be executed when resuming the continuation, and the other two parameters express the local environment that needs to be reinstated when control is passed to the continuation: this local environment consists of a finite map from register names to their values, and a delimited effect context expressing additional handlers that will need to be added to the effect context to execute the continuation.
- Leave records, in the form $\mathcal{L}(a_l, (\$r_e = v_e))$, indicate what code should be executed after leaving the scope of a handler; they are similar to continuations, except that they do not provide a delimited effect context because when we leave the scope of a handler we never need to install additional handlers.

Effect context. The effect context is a stack of *handler frames* representing the active handlers at any given moment at runtime. Handler frames contain the same information as a handler record, pairing it with a leave record indicating what code should be run after returning from the handler.

Instructions. AsmFX introduces three instructions `loadh`, `loadk`, `loadl` to load handler, continuation, and leave records into registers. The three instructions use, respectively, registers `$hr`, `$kr`, and `$lr` as their targets: these registers have an architectural meaning in that their content can affect the control flow; it is however possible to copy the content of these registers to any other register, and vice-versa, using general-purpose instructions.

`loadh $H, state : (\$r')$` loads a handler record into `$hr`. The handler specification H contains pairs of tag names and code addresses implementing the handler operation specified by that tag; the pseudo-tag *return* is always present. The content of registers $\overline{\$r'}$, specifying the current state of the handler, is saved within the record.

`loadk $a_r, save : (\$r')$` loads a continuation record into register `$kr` by specifying the address a_r of its code and the registers $\overline{\$r'}$ whose value must be saved now and restored if the continuation is run and terminates regularly. The continuation is created with a trivial (empty) delimited context.

`loadl $a_l, save : (\$r')$` loads a leave record into register `$lr`, where a_l is the address of the code that will be run when the leave record is triggered, and the save registers $\overline{\$r'}$ contain the values that need to be restored to execute that code.

The remaining four instructions deal with control flow. The instruction `resume` invokes the continuation stored in the special register `$kr`, wrapping it in the handler specified by the special register `$hr`. Note that the two registers together essentially perform the same function as brooms in SFX, with `$kr` acting as the head, and `$hr` as the handle.¹ In order to execute this instruction, the register `$lr` must also contain the record referencing the code that will be executed if and when the resumed continuation terminates. To resume the continuation, the machine pushes into the context a new handler frame obtained by combining the contents of `$hr` and `$lr`, followed by the

¹To be more precise, AsmFX, unlike SFX, enables a slightly more liberal policy by allowing the code to reference and replace heads and handles independently. By a cultural reference to the British sitcom *Only Fools and Horses*, we describe the pair of `$kr` and `$hr` as *Trigger's broom*.

delimited context from $\$kr$, then restores the contents of the registers saved in $\$kr$ and jumps to the code address of the continuation.

$\text{do } \#tag, \text{save} : (\overline{\$r})$ performs an effectful operation by jumping to the corresponding handler code; after the jump, the registers $\$kr$, $\$hr$, and $\$lr$ are loaded with the continuation record for the suspended computation that triggered the effect (including saved values for the registers $\overline{\$r}$, the handler record that is handling the effect, and the leave record associated with that handler; the handler state registers are also restored from the handler record. When (and if) the operation returns a value, it will use the continuation in $\$kr$ to jump back to the next instruction after do , restoring the contents of the save registers $\overline{\$r}$.

return operates similarly to do but it always matches the handler at the top of the context and when jumping to its return clause it does not need to provide a resumption. This means it loads $\$lr$ and the state registers from the handler frame, and then jumps to the address specified by the return clause $H(\#return)$.

exit completes the exit procedure from a handler by triggering the leave record in $\$lr$ and restoring the registers saved in it.

instruction	arguments	context	\hookrightarrow	new context	reg. updates
loadk	a_r save : $(\overline{\$r_c})$	C		C	$\$kr = \mathcal{K}(a_r, (\overline{\$r_c} = \Theta(\overline{\$r_c})), [])$
loadh	H state : $(\overline{\$r_h})$	C		C	$\$hr = \mathcal{H}(H, (\overline{\$r_h} = \Theta(\overline{\$r_h})))$
loadl	a_l save : $(\overline{\$r_e})$	C		C	$\$lr = \mathcal{L}(a_l, (\overline{\$r_e} = \Theta(\overline{\$r_e})))$
resume		C where: $\$hr = \mathcal{H}(H, (\overline{\$r_h} = v_h))$ $\$kr = \mathcal{K}(a_r, (\overline{\$r_c} = v_c), D_r)$ $\$lr = R_L$		D_r $\cdot \mathcal{H}(H, (\overline{\$r_h} = v_h), R_L)$ $\cdot C$	$\$pc = a_r$ $\$r_c = v_c$
do	$\#tag$ save : $(\overline{\$r_c})$	D_r^{-tag} $\cdot \mathcal{H}(H^{+tag}, (\overline{\$r_h} = v_h), R_L)$ $\cdot C$		C	$\$kr = \mathcal{K}(\Theta(\$pc), (\overline{\$r_c} = \Theta(\overline{\$r_c})), D_r^{-tag})$ $\$hr = \mathcal{H}(H^{+tag}, (\overline{\$r_h} = v_h))$ $\$lr = R_L$ $\$pc = H^{+tag}(\#tag)$ $\overline{\$r_h} = v_h$
return		$\mathcal{H}(H, (\overline{\$r_h} = v_h), R_L)$ $\cdot C$		C	$\$pc = H(\#return)$ $\overline{\$r_h} = v_h$ $\$lr = R_L$
exit		C where: $\$lr = \mathcal{L}(a_l, (\overline{\$r_e} = v_e))$		C	$\$pc = a_l$ $\$r_e = v_e$

Domain of contexts:

$\text{dom}([]) = \emptyset$ $\text{dom}(\mathcal{H}(H, \overline{\$r_h} = v_h, R_L) \cdot C) = \text{dom}(H) \cup \text{dom}(C)$ $\text{dom}(\#t_1 \mapsto a_1, \dots, \#t_n \mapsto a_n) = \{\#t_1, \dots, \#t_n\}$.
 D^{-t} indicates any delimited context such that $\#t \notin \text{dom}(D^{-t})$. H^{+t} denotes any handler definition where $\#t \in \text{dom}(H^{+t})$.

Fig. 3. ASMFx machine transitions.

The semantics of ASMFx is expressed by the abstract machine in Figure 3. A machine state is specified by a tuple $\langle \Xi, \Theta, C \rangle$, where Ξ , the main memory, maps addresses to values, which may be data or instructions; Θ , the register file, maps register names to their current value; and C , the effect context, is a stack of handler frames specifying what effect implementations are available to the code currently running. When it is ready to execute an instruction, the machine fetches it from memory Ξ at the address stored in the $\$pc$ register (it is assumed that the memory at that address contains a valid instruction, otherwise the machine will be stuck); at this point the machine immediately increments $\$pc$ to ensure it points to the next instruction, then it matches the current instruction (including its arguments) and the effect context C against the first three columns of the transition table: execution involves updating the context C and the registers file Θ according to the

content of the last two columns: among other things, this allows an instruction to perform a jump by updating the register $\$pc$.

The instructions we describe do not use the main memory Ξ except as storage for the code being executed, but a concrete instance of AsmFX can read memory values into registers and write arbitrary values into memory by defining its own load/store instructions.

6 From SFX to AsmFX

We have defined an assembly language with an intuitive notion of effect handlers and the instructions we deemed essential to write effectful programs. Though these instructions, after careful analysis of their machine semantics, may appear to reasonably implement the semantics of effect handlers when considered in isolation, to provide a solid foundation for this claim, we will show that SFX programs can be compiled to AsmFX. For this abstract discussion, we can conveniently express the compiler by means of recursive procedures taking SFX syntactic forms as input, and returning an AsmFX memory image as output. For reasons that will be clear when we will prove the correctness of the compiler, we assume that, along with producing the compiled program, the compiler will also reflect the result of the compilation in the source code in the form of *annotations* stating, for each subexpression, the memory location ℓ of the AsmFX code that evaluates it: this means source values, computations, or programs V, M, P will be augmented into annotated versions $\hat{V}, \hat{M}, \hat{P}$; we call the language of augmented SFX expressions ASFX. We will refrain from giving the tedious details of this backwards annotation (entirely analogous to the debugging symbols produced by most compilers), and we will assume that it is correct (in the sense that each annotation corresponds to the address of the code produced when compiling the corresponding subexpression), and that by erasing annotations we get back the original source. We use $\mu(\cdot)$ and $\nu(\cdot)$ to denote the starting and ending AsmFX location for a given augmented expression, and $|\cdot|$ to indicate the erasure of an augmented expression into the corresponding plain SFX expression.

Fact 6.1. *Let Ξ_P be the AsmFX memory image containing the compiled version of the SFX program P , and \hat{P} be the augmented version of P produced by the compiler: then for every computation \hat{M} or value \hat{V} appearing in \hat{P} , we have that the portion of Ξ_P delimited by $[\mu(\hat{M}), \nu(\hat{M})]$ (resp. $[\mu(\hat{V}), \nu(\hat{V})]$) contains the compiled code for $|\hat{M}|$ (resp. $|\hat{V}|$).*

Sometimes we will also write explicitly the annotations of expressions using the syntax $\hat{M}@\ell$ to indicate that $\mu(\hat{M}) = \ell$. A formal grammar of ASFX is included in the appendix.

Due to limitations on the type of data that can be stored in a single AsmFX register, we will use the notation $\|V\|$ or $\|A\|$ to denote the number of registers needed to store a certain SFX value V , or values of a certain type A .

By convention, the compiler will partition the register set into *variable* registers $\$x_i$ (where x is an SFX variable), *argument* registers $\$a_i$ (holding function inputs or the value of an expression), *state* registers $\$s_i$ (used for handler state), and separate special registers $\$hr, \$kr, \$lr, \pc with architectural meaning. Registers in the first three classes are available in unlimited number.

Compiling computations. To denote the compilation of SFX computations M , we use the syntax $[M]^{\mathbb{I}\mathbb{L}} \Rightarrow \mathbb{L}'$. The operation takes two auxiliary parameters: a set of input registers \mathbb{I} whose value is used in the computation and must be preserved at the end of the compiled code, corresponding to free variables in the source computation M (in particular, each free variable will correspond to zero, one, or more registers depending on its type); and the used label set \mathbb{L} , which is needed to avoid collisions when generating new AsmFX code labels. The result of running the compiled computation is always placed in an initial segment of the argument registers $\overline{\$a}$.

$[M]^{\mathbb{L}} \Rightarrow \mathbb{L}'$ produces a main output, i.e. a list of AsmFX instructions annotated with labels representing the result of compiling E , and a secondary output (the updated list of used labels \mathbb{L}' , which is always a superset of \mathbb{L}). The definition of $[M]^{\mathbb{L}} \Rightarrow \mathbb{L}'$ is given, for any expression M , by using syntax such as the following:

$$\text{let } S = \text{exp} \dots \\ \dots [M]^{\mathbb{L}_M, \mathbb{L}_M} \Rightarrow \mathbb{L}'_M \dots$$

The **let** syntax tells the compiler to evaluate a certain meta-expression exp once and introduces a metavariable S to indicate the result of said evaluation: we use this syntax to introduce fresh code labels, or to express sets of registers or labels in a short form. The recursive compilation of M inlines the compiled code for M and introduces a metavariable \mathbb{L}'_M for the extended set of used labels. The ellipses may contain AsmFX instructions and labels. It is always assumed that the secondary output of the compilation is the one assigned at the end of the compilation process for a given expression to the metavariable \mathbb{L}' . We also assume an operation $\varphi(n, \mathbb{L})$, producing n labels fresh with respect to the used label set \mathbb{L} . The concrete choice of labels is irrelevant and consistent renaming is assumed to denote the same code (we use labels for the sake of readability, but each label must be considered notation for an absolute memory index).

Space constraints prevent us from presenting the details of the compiler here: instead, we will comment the code implementing the compilation of **do** and **handle**.

$$[\text{do } \#t(V)]^{\mathbb{L}} \Rightarrow \mathbb{L}' := \quad [V] \\ \text{do } \#t, \text{save} : \mathbb{I} \\ \text{let } \mathbb{L}' = \mathbb{L}$$

To compile **do** $\#t(V)$, we first compile V , whose value will be then found in $\$a$, then we use the **do** instruction (the *save* clause states that the value of the registers \mathbb{I} used in the client must be saved now and restored if the effectful operation returns to this code).

$$[\text{handle } M \text{ with } h(V)]^{\mathbb{L}} \Rightarrow \mathbb{L}' := \quad \begin{array}{l} \text{let } l_{cli}, l_{exit} = \varphi(2, \mathbb{L}) \\ \text{let } \mathbb{L}_M = \mathbb{L} \cup \{l_{cli}, l_{exit}\} \\ \text{let } \mathbb{I}_h = \$s_{||V||} \\ [V] \\ \mathbb{I}_h \leftarrow \$a_{||V||} \\ \text{loadh } \$hr, [h], \text{state} : \mathbb{I}_h \\ \text{loadk } \$kr, l_{cli}, \text{save} : () \end{array} \quad \begin{array}{l} \text{loadl } \$lr, l_{exit}, \text{save} : \mathbb{I} \\ \text{resume} \\ l_{cli} : \\ [M]^{\mathbb{L}_M} \Rightarrow \mathbb{L}'_M \\ \text{return} \\ l_{exit} : \\ \text{let } \mathbb{L}' = \mathbb{L}'_M \end{array}$$

To execute a computation under a new handler by means of **handle** M **with** $h(V)$, we compile V , copy its evaluation in the $\$s$ registers, and use it together with $[h]$ to create a handler and place it in $\$hr$. We also create a continuation pointing to the code for M (l_{cli}) and place it in $\$kr$, and a leave record for the code to be executed after this computation (l_{exit}), with saved registers corresponding to the input set \mathbb{I} . After these three registers have been loaded, we invoke the client by resuming the continuation we created, using the instruction **resume**. Notice that the client code at location l_{cli} is terminated by a **return** instruction, so that the result of M will be delivered to the return clause of the newly installed handler: we call this a **return** epilogue. Tracking epilogues is essential to prove the correctness of the compiler.

Compiling handlers. Handlers consist of a computation expression for each return and operation clause: it is easy to compile each of them separately using the same procedure we use for all other computations. Additionally, for each clause, the compiler takes care of moving the actual parameter of the handler from the state registers \bar{s} to variable registers \bar{x} corresponding to the formal parameter x used by the SFX handler definition. The compiler will also insert code to move the arguments

of each return clause from the argument registers \bar{a} to variable registers \bar{y} corresponding to the formal argument y of that clause in the SFX definition, and to move the broom resumption from the special registers $\$kr, \hr to the appropriate variable registers $\$r_1, \r_2 .

The handler result (returned when exiting the handler, either normally through the return clause or abnormally through an operation clause) is stored in an initial segment of the argument registers $\$a$. Each clause is terminated by an `exit` instruction which triggers the leave register; this instruction also constitutes an epilogue that needs to be accounted for in the soundness proof.

For each clause **return**, $\#t_i$ of a handler h , we will generate a unique label $h_{\text{return}}, h\#t_i$ pointing to its code: this allows us to compile a handler specification $[h]$ as

$$\text{return} \mapsto h_{\text{return}}, \#t_1 \mapsto h\#t_1, \dots$$

7 Soundness of compilation

As a final step in our effort to assess the expressiveness of AsmFX, we need to prove the correctness of our compiler. The most basic correctness property would state that for any SFX program P and its AsmFX counterpart $[P]$, if P terminates returning a value v , then $[P]$ also terminates and a machine representation of v can be found in an initial segment of the $\$a$ registers; however, we are also interested in proving the correctness of partial executions, by showing that the transition system induced by the AsmFX machine simulates the one induced by the SFX abstract machine.

The proof is not entirely straightforward. The main difficulty lies in the fact that executing an SFX program consumes part of the program itself, blurring the correspondence between source and compiled object code beyond recognisability: the code resulting from the recompilation of a partially executed program will not match, even partially, the code of the original program, because parts of the code are completely lost and thus ignored by the compiler. To match partially executed source code to partially executed AsmFX code we need a more sophisticated instrument.

Our solution is to prove soundness by reasoning not on the plain source code, but on its ASFX version. We refine the SFX abstract machine to operate directly on ASFX programs, which now carry along location information and other information about execution history. The execution of an ASFX program differs from that of the corresponding SFX program especially in the form of its continuations: the identity continuation and all frames are annotated with a start location; additionally, pure frames $((x).\hat{M})@l$ contain a naked computation instead of a computation closure; we also introduce *epilogue frames* **jump**, **exit**, **leave** representing AsmFX instructions that must be executed in order to link the various portions of compiled code together. The closures missing from pure frames resurface as an additional environment \hat{y} component of ASFX brooms.

Much like we did for SFX, we formalise the semantics of ASFX by means of a CEK-like abstract machine. The main difference between the two abstract machines is that the ASFX machine is aware of location annotations and propagates them during its execution, ensuring that all relevant information is preserved. The definition of continuations also differs in that it tracks certain intermediate execution steps that are implicit in SFX but explicit in AsmFX.

For example, the rule for **do** $\#t(\hat{V})$ is now refined to the following form:

$$\begin{aligned} \langle \text{do } \#t(\hat{V}), \hat{y}, \hat{\kappa}^{-t} \cdot h^{+t}(\hat{w}, R_{\mathcal{L}}) @ l \cdot \hat{y}' \rangle &\hookrightarrow \\ \langle h_{\#t}, (\text{LEAVE} \mapsto R_{\mathcal{L}}) \cdot (\text{res}(\#t) \mapsto (\hat{\kappa}^{-t}, h^{+t}(\hat{w}) @ l, \hat{y})) \cdot \\ (\text{dom}(\#t) \mapsto \llbracket \hat{V} \rrbracket \hat{y}) \cdot (\text{dom}(h) \mapsto \hat{w}), \text{exit}@v(h_{\#t}) \cdot \hat{\kappa}' \rangle \end{aligned}$$

In this rule, the AsmFX-like leave record $R_{\mathcal{L}}$ is manipulated explicitly and bound to the pseudovisible LEAVE (meaningful to the abstract machine, but not directly accessible to the code). Furthermore, an **exit** epilogue is added to the output continuation, expressing the fact that if the evaluation of the operation clause $h_{\#t}$ completes without calling a resumption, we need to execute

Values and environments	
$(\kappa, h(w)) \leq (\hat{\kappa}, h(\hat{w})@l, \hat{\gamma}_\kappa)$	$:= \kappa \leq (\hat{\gamma}_\kappa, \hat{\kappa}) \wedge w \leq \hat{w}$
$v \leq \hat{v}$	$:= v = \hat{v} \quad (\text{if } v \neq k)$
$\gamma \leq \hat{\gamma}$	$:= \forall x \in \text{dom}(\gamma). \gamma(x) \leq \hat{\gamma}(x)$
Continuations	
$\hat{\gamma}(\text{LEAVE}) = (_, \hat{\gamma}_\kappa)$	
$\mathbf{0} \leq (\hat{\gamma}, \mathbf{0}@l)$	$\kappa \leq (\hat{\gamma}, \hat{\kappa})$
$\kappa \leq (\hat{\gamma}, \mathbf{jump}@l \cdot \hat{\kappa})$	$\kappa \leq (\hat{\gamma}_\kappa, \hat{\kappa})$
$\kappa \leq (\hat{\gamma}, \mathbf{exit}@l \cdot \hat{\kappa})$	$\kappa \leq ((\text{LEAVE} \mapsto \mathcal{L}(l+2, [])) \cdot \hat{\gamma}, \hat{\kappa})$
$\kappa \leq (\hat{\gamma}, \mathbf{leave}@l \cdot \hat{\kappa})$	
$\forall v. \kappa \leq ((x \mapsto v) \cdot \hat{\gamma}, \hat{\kappa}) \quad \gamma_N \leq \hat{\gamma} \quad N = \hat{N} $	$\kappa \leq (\hat{\gamma}_\kappa, \hat{\kappa})$
$((x).N[\gamma_N]) \cdot \kappa \leq (\hat{\gamma}, ((x).\hat{N})@l \cdot \hat{\kappa})$	$h(w) \cdot \kappa \leq (\hat{\gamma}, h(w, \mathcal{L}(l_h, \hat{\gamma}_\kappa))@l \cdot \hat{\kappa})$
Configurations	
$\langle M, \gamma, \kappa \rangle \leq \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle := \hat{M} = M \wedge \gamma \leq \hat{\gamma} \wedge \kappa \leq (\hat{\gamma}, \hat{\kappa})$	$\langle v, \gamma, \kappa \rangle \leq \langle \hat{v}, \hat{\gamma}, \hat{\kappa} \rangle := v \leq \hat{v} \wedge \kappa \leq (\hat{\gamma}, \hat{\kappa})$

Fig. 4. Scope-preserving simulation

some ASmFX code (an `exit` instruction) that will dispose of the handler h and return to the external code. Due to space constraints, the full definition of the abstract machine and an explanation of its rationale is deferred to the appendix.

7.1 Well-scoping

We prove that single step transitions in SFX are simulated by multistep transitions in ASFX. Since ASFX epilogues essentially correspond to no operation at all in SFX, the main challenge is to show that in spite of a different approach to handling environments, ASFX provides a correct implementation of scopes. The simulation relation, summarised in Figure 4, essentially expresses a well-scoping invariant.

When an SFX configuration s is simulated by an ASFX configuration \hat{s} , we write $s \leq \hat{s}$. The relation is defined by means of auxiliary definitions for the simulation of environments and continuations. An interesting fact about the simulation of continuation is that SFX continuations are simulated by pairs of ASFX continuations and environments: this is explained by the fact that SFX pure frames are closed with respect to an environment, while ASFX frames are not.

The environment simulation $\gamma \leq \hat{\gamma}$ expresses the fact that whenever x is in the domain of γ , it is also in the domain of $\hat{\gamma}$ and the values for x in the two environments “agree”. When $\gamma(x)$ is not a broom, $\hat{\gamma}(x)$ agrees with it if by erasing annotations we get the same value; if $\gamma(x)$ is a broom $(\kappa, h(w))$, then $\hat{\gamma}(x)$ must be a broom $\hat{\kappa}, h(w)@l, \hat{\gamma}_\kappa$ such that κ is simulated by $(\hat{\gamma}_\kappa, \hat{\kappa})$.

The continuation simulation $\kappa \leq \hat{\kappa}$ is defined by means of an inductive judgement, with a base rule for the identity continuation and recursive rules for each type of frame. Identity continuations are simulated by identity continuations regardless of the environment. In the interesting case of `exit` frames, the active augmented environment is swapped with the one contained in the `LEAVE` pseudovalue (which must be defined), reflecting the semantics of the ASmFX instruction `exit`, which will be executed if the `exit` frame is activated. This kind of rule ensures that while the environment of the SFX machine and that of the ASFX machine may temporarily differ in non-trivial ways (specifically when the control is a value), the ASFX machine will eventually synchronise its environment.

The other interesting case is that of pure frames. If both continuations start with a pure frame, in order for the simulation to hold, the computations N and \hat{N} in each pure frame must match (meaning that by erasing \hat{N} we get N). Furthermore, the SFX pure frame is closed by an environment γ_N : that environment must be simulated by the ASFX environment $\hat{\gamma}$, reflecting the fact that in ASFX we do not need an environment in the pure frame, because the active environment is already the correct one. If these conditions hold, the simulation is valid if there is a simulation between the two remaining parts of the two continuations; however on the ASFX side we need to extend the environment with a value for the variable x expected by the pure continuation, and the simulation must hold independently of the value chosen.

This definition allows us to prove the theorem:

THEOREM 7.1. *For all SFX configurations s, s' such that $s \hookrightarrow s'$, for all ASFX configurations \hat{s} such that $s \preccurlyeq \hat{s}$, there exists an ASFX configuration \hat{s}' such that $s' \preccurlyeq \hat{s}'$ and $\hat{s} \hookrightarrow \hat{s}'$.*

The proof, which is detailed in the appendix, is by case analysis on the transition $s \hookrightarrow s'$ and on the relation $s \preccurlyeq \hat{s}$.

7.2 Control flow simulation

To complete the proof of correctness of the compiler, we now need to establish a simulation relation between ASFX and AsmFX. We consider AsmFX configurations $a = \langle \Xi, \Theta, C \rangle$, whose elements are memory stores Ξ (containing the program), register files Θ (finite maps from registers to values), and effect contexts C . Since a property of our compiler is that the memory store for a given source program is established at compilation time and stays read-only during execution, we will allow ourselves to omit it from AsmFX configurations and write $\langle \Theta, C \rangle$ for $\langle \Xi, \Theta, C \rangle$.

A source configuration \hat{s} for a program \hat{P} should correspond to an AsmFX configuration a on $[\hat{P}]$: actually, for every \hat{s} there are multiple such a , because of the plethora of registers available in the architecture: some of the registers will not have a counterpart in the source configuration and will thus be irrelevant. Therefore, the translation to AsmFX of a source configuration \hat{s} , which we denote by $[\hat{s}]$, will be the set of AsmFX configurations that are a model for \hat{s} up to relevant registers. If we write $a \models \hat{s}$ to mean that a models \hat{s} , then we can define $[\hat{s}] := \{a : a \models \hat{s}\}$. We define what it means to be a model separately for the register file and effect context components: the register file must be a model for the whole configuration; on the other hand, the effect context only needs to model the continuation.

$$\langle \Theta, C \rangle \models \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle := (\Theta \models \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle) \wedge (C \models \hat{\kappa})$$

We can define what it means for an effect context to model a continuation by means of a direct translation that only needs to consider the handler frames:

$$\begin{aligned} C \models \hat{\kappa} &:= C = [\hat{\kappa}] \\ [\hat{\kappa}] &:= \begin{cases} [h(\hat{w}, \mathcal{R}_{\mathcal{L}})] @ \ell \cdot [\hat{\kappa}'] & \text{if } \hat{\kappa} = h(\hat{w}, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{\kappa}' \\ [\hat{\kappa}'] & \text{if } \hat{\kappa} = \hat{\zeta} \cdot \hat{\kappa}' \text{ and } \hat{\zeta} \neq h(_) \end{cases} \\ [h(\hat{w}, \mathcal{R}_{\mathcal{L}}) @ \ell] &:= \mathcal{H}([h], (\$s = [\hat{w}]), [\mathcal{R}_{\mathcal{L}}]) \\ [\mathcal{L}(\ell, \hat{\gamma})] &:= \mathcal{L}(\ell, [\hat{\gamma}]) \\ [(x \mapsto \hat{v})] &:= (\overline{\$x} = [\hat{v}]) \end{aligned}$$

A register file models a source configuration if every variable defined in the source environment is backed by (any number of) registers containing the AsmFX representation of that variable's value; if the program counter matches the address of the beginning of the computation being evaluated or, when the computation has ended and we are returning a value, when it matches the address of

Values	$\hat{v}\checkmark := \forall \hat{k} \in \text{subval}(\hat{v}) : \hat{k}\checkmark$
Environments	$\hat{y}\checkmark := \forall x \in \text{dom}(\hat{y}) : \hat{y}(x)\checkmark$
Brooms	$\hat{k}\checkmark := \forall (\hat{k}, h(\hat{w})@ \ell, \hat{y}) = \hat{k}, \forall \hat{k}'\checkmark, \forall \hat{y}' \supset \hat{k}' : \\ \hat{y} \supset (\hat{k} \cdot h(\hat{w}, \mathcal{L}(\mu(\hat{k}'), \hat{y}'))@ \ell \cdot \hat{k}') \wedge (\hat{k} \cdot h(\hat{w}, \mathcal{L}(\mu(\hat{k}'), \hat{y}'))@ \ell \cdot \hat{k}')\checkmark$
Coherence	$\hat{y} \supset \hat{k} := \hat{y}\checkmark \wedge \hat{y}(\text{LEAVE}) = (\ell, \hat{y}') \implies \ell = \text{laddr}(\hat{k}) \wedge \hat{y}' \supset \text{lcont}(\hat{k})$
Continuations	
$\frac{}{\mathbf{0}@ \ell \checkmark} \quad \frac{\hat{k}\checkmark \quad \Xi(\ell) = \mathbf{j} \mu(\hat{k})}{\mathbf{jump}@ \ell \cdot \hat{k}\checkmark} \quad \frac{\hat{k}\checkmark \quad \Xi(\ell) = \mathbf{exit}}{\mathbf{exit}@ \ell \cdot \hat{k}\checkmark}$	
$\frac{\hat{k}\checkmark \quad \Xi(\ell) = \mathbf{loadl}(\ell + 2) \quad \hat{k} = \mathbf{jump}@(\ell + 1) \cdot ((x).\hat{N})@ \mu(\hat{N}) \cdot \mathbf{exit}@ \nu(\hat{N}) \cdot h(\hat{w})@(\ell + 2) \cdot \hat{k}_0}{\mathbf{leave}@ \ell \cdot \hat{k}\checkmark}$	
$\frac{\hat{k}\checkmark \quad \Xi(\ell) = \overline{\$x} \leftarrow \overline{\$a}; [\hat{N}] \quad \nu(\hat{N}) = \mu(\hat{k})}{((x).\hat{N})@ \ell \cdot \hat{k}\checkmark} \quad \frac{\hat{k}\checkmark \quad \Xi(\ell) = \mathbf{return} \quad \hat{y} \supset \hat{k}}{h(\hat{w}, \mathcal{L}(\mu(\hat{k}), \hat{y}))@ \ell \cdot \hat{k}\checkmark}$	
Leave operators	
$\text{lcont}(\hat{k}) := \begin{cases} \text{undefined} & \text{if } \hat{k} = \mathbf{0}@ \ell \text{ or } \hat{k} = h(_)@ \ell \cdot \hat{k}' \\ \hat{k}' & \text{if } \hat{k} = \mathbf{exit}@ \ell \cdot \hat{k}' \\ \text{lcont}(\hat{k}') & \text{if } \hat{k} = \hat{\zeta} \cdot \hat{k}' \text{ and } \hat{\zeta} \neq h(_)@ \ell \end{cases}$	
$\text{laddr}(\hat{k}) := \mu(\text{lcont}(\hat{k}))$	
Configurations	
$\langle \hat{M}, \hat{y}, \hat{k} \rangle \checkmark := \Xi(\mu(\hat{M})) = [\hat{M}] \wedge \nu(\hat{M}) = \mu(\hat{k}) \wedge \hat{y} \supset \hat{k} \wedge \hat{k}\checkmark$	
$\langle \hat{v}, \hat{y}, \hat{k} \rangle \checkmark := \hat{v}\checkmark \wedge \hat{y} \supset \hat{k} \wedge \hat{k}\checkmark$	

Fig. 5. Augmented source validity.

the beginning of the continuation. Note that the pseudovalue LEAVE is always mapped to the register \$1r, and every variable x is mapped to as many $\overline{\$x}$ registers as it is necessary.

$$\begin{aligned} \Theta \models \langle \hat{M}, \hat{y}, \hat{k} \rangle &:= \Theta(\$pc) = \mu(\hat{M}) \wedge \forall x \in \text{dom}(\hat{y}). \overline{\Theta(\$x)} = [\hat{y}(x)] \\ \Theta \models \langle \hat{v}, \hat{y}, \hat{k} \rangle &:= \Theta(\$pc) = \mu(\hat{k}) \wedge \forall x \in \text{dom}(\hat{y}). \overline{\Theta(\$x)} = [\hat{y}(x)] \end{aligned}$$

The soundness property we seek states that the modelling relation is a simulation: whenever we have a source transition $\hat{s} \hookrightarrow \hat{s}'$ and an AsmFX configuration $a \models \hat{s}$, then there exists a transition chain $a \xrightarrow{*} a'$ where $a' \models \hat{s}'$. However, this property cannot be proved for an arbitrary \hat{s} : augmented source configurations contain metadata that encodes part of the execution history; while the compiler produces correct metadata, and the abstract machine ensures that the metadata is propagated correctly, arbitrary configurations could be unsound and not allow the simulation we need for the soundness theorem. We thus need to restrict ourselves to *valid* configurations (denoted by $\hat{s}\checkmark$) and prove, as part of our theorem, that validity is preserved by transitions.

The validity predicate is defined in Figure 5: its role is to check that all the ASFX annotations are consistent with the compiled AsmFX program. A configuration is valid if its environment \hat{y} and continuation \hat{k} are valid and additionally the environment and the continuation are coherent (notation: $\hat{y} \supset \hat{k}$). Furthermore, we require that in a valid configuration whose control is a computation

\hat{M} , the ASmFX code at location $\mu(\hat{M})$ corresponds to the compiled code for \hat{M} , and that if instead the control is a value v , the value itself is valid.

A value is valid if any broom \hat{k} syntactically contained in it is valid; an environment is valid if all the values it assigns to its domain are valid. A broom is valid if, when we resume in any valid environment γ' and continuation κ' such that the two are coherent, we get a valid continuation and, furthermore, the environment in the broom is coherent with the extended continuation.

What it means for a continuation to be valid is expressed by an inductive predicate, with a base rule for the identity continuation and a recursive rule for each frame type. These rules ensure that the annotation for each frame corresponds to suitable ASmFX code: for instance, the annotation for **jump** frames must reference a jump to the rest of the continuation. The theorem we prove states that transitions starting at valid ASFX configurations preserve validity and can be simulated:

THEOREM 7.2. *Let P be a source program, and $\Xi = \lfloor P \rfloor$. For all source configurations \hat{s}, \hat{s}' such that $\hat{s} \checkmark$ (valid with respect to Ξ), if $\hat{s} \hookrightarrow \hat{s}'$, then $\hat{s}' \checkmark$ and for all ASmFX configurations a such that $a \models \hat{s}$ there exists a' such that $a \xrightarrow{*} a'$ and $a' \models \hat{s}'$.*

The proof, detailed in the appendix, is by case analysis on the transition $\hat{s} \hookrightarrow \hat{s}'$.

7.3 Soundness

It is straightforward to show that the annotations produced by compilation are correct and therefore the initial configuration of the ASFX machine is related by simulation to the corresponding initial configurations of the SFX machine and the ASmFX machine.

LEMMA 7.3. *If \hat{M} is the result of annotating M according to the compiler, $|\hat{M}| = M$ and, consequently, $\langle M, [], 0 \rangle \leq \langle \hat{M}, [], 0 @ \mu(\hat{M}) \rangle$.*

LEMMA 7.4. *Let $\hat{P} = \hat{M}; \hat{\Sigma}$. If $\lfloor \hat{P} \rfloor = \Xi$, then $\Xi(\mu(\hat{M})) = \lfloor \hat{M} \rfloor$. Consequently $\langle \hat{M}, [], 0 @ \mu(\hat{M}) \rangle \checkmark$ (with respect to Ξ).*

LEMMA 7.5. *If $\Theta = [\$pc \mapsto \mu(\hat{M})]$, then $\Xi \models \langle \hat{M}, [], 0 @ \mu(\hat{M}) \rangle$. Consequently, $\langle \Xi, [] \rangle \vdash \langle \hat{M}, [], 0 @ \mu(\hat{M}) \rangle$.*

The soundness theorem then arises from Theorems 7.1 and 7.2 as a corollary:

COROLLARY 7.6. *Let $P = M; \Sigma$ be an SFX program and $s = \langle M, [], 0 \rangle$. Let $\hat{P} = \hat{M}; \hat{\Sigma}$ be P annotated by the compiler, $\Theta = [\$pc \mapsto \mu(\hat{M})]$, and $a = \langle \Theta, [] \rangle$. Then if $s \xrightarrow{*} s'$, there exist a', \hat{s}' such that $s' \leq \hat{s}'$ and $a' \models \hat{s}'$.*

8 Discussion

We showed that ASmFX can correctly implement a large class of effect handlers. In devising our compiler we focused on the general case. However, there are many opportunities to provide alternate more efficient strategies in specific cases (exceptions, runners, tail-resumptive handlers, single-shot resumptions, etc.), whether identified by special syntax, static analysis, or dynamically at runtime. For instance, specialised compilation of tail-resumptive handlers could invoke a broom without setting up a new leave record (similarly to how tail-calls to functions are optimised to reuse the caller's stack frame). A similar optimisation is performed by existing effect handler implementations including Effekt, Koka, and libmprompt.

Another obvious optimisation concerns the implementation of functions. For the purposes of this paper it was convenient to treat functions as syntactic sugar for trivial handlers consisting of just a return clause: this trick is standard in the literature, however it may obfuscate the means by which functions are implemented. When we implement a function call as a **handle** expression, the compiler emits the code to set up a new handler frame, which includes a leave record: while the

handler functionality is trivial, the leave record plays the role of the call stack, thanks to its ability to store the caller's return address and environment (including the previous leave record). Indeed, it is easy to give a direct translation of functions to AsmFX that does not make use of handlers: all that is needed to implement functions in AsmFX is the functionality of the leave register \$lr, with the instructions `loadl` and `exit` used to push and pop stack frames.

A desirable feature of the AsmFX implementation of handlers is that it naturally provides a certain level of memory protection: the state of a handler is stored in the effect context and cannot be directly accessed by other code fragments (including client code or other handlers); our compiler does not actively wipe registers when transitioning in and out of handler code, so (read only) information leak is still possible, but it is possible in principle to write a refined compiler that would ensure no such leak happens. Additionally, we believe AsmFX could be useful as an intermediate language, particularly for further translation to *hardware capability architectures* such as CHERI [29]. We conjecture that fine-grained compartmentalisation provided in such architectures could be used to provide an implementation of effect contexts and handler, continuation, and leave records ensuring that information can only be accessed by the code that owns it (and particularly separation between code that uses effects and code that implements them). Inspired by the *Cerise* program logic [11] and its assembler for an idealised capability architecture, we plan to eventually provide an effectful *CeriseFX* to explore the relationship between handlers and machine-level capabilities.

WasmFX [25] extends WebAssembly (Wasm) [14], a low-level portable bytecode-driven stack machine, with support for effect handlers. Though the design of AsmFX is inspired by WasmFX, the two systems have quite different characters and purposes: in particular, WasmFX is rightly matched to the existing Wasm stack model, where AsmFX aims to be agnostic about the underlying implementation of its primitives. OCaml 5 [28] incorporates a high-performance implementation of effect handlers based on low-level stack manipulation operations. The intended semantics of OCaml 5 effect handlers is given by a CEK-like abstract machine much like that of SFX. As with WasmFX, the OCaml 5 implementation is closely tied with underlying features of the host language, in this case delimited continuations and multicore fibers. C libraries such as `libseff` [2] and `libmpeff` [21] provide low-level effect handler implementations in C based respectively on coroutines and on multiprompt delimited control. It would be interesting to try to more formally relate systems such as WasmFX, OCaml 5, `libseff` and `libmprompt` with AsmFX.

Muhcu et al. [24] present an implementation technique for multi-shot resumptions in the context of *named* handlers [5, 31] and stack-allocated mutable state. Their work partially overlaps with ours as they too describe a compilation technique (building on prior work [22]); while we directly compile a high-level language with standard (unnamed) effect handlers to the effectful AsmFX, their work starts with an intermediate language with multi-prompt delimited control implementing named handlers (which they call “lexically-scoped handlers”) and compiles it to an abstract assembly without primitive effect handlers, but with explicit memory allocation. A direct comparison between their work and AsmFX is challenging: not only are the two works based on different semantics for effect handlers, but they also have different purposes: AsmFX is concerned primarily with expressiveness and soundness, whereas [24] is geared towards implementation efficiency. Extending AsmFX to support named handlers would however be interesting: to achieve that, we could provide variants of the `resume` and `do` instructions respectively returning or accepting a handler name as an additional argument; the name may concretely be an index into the effect context (allowing efficient addressing of a handler frame), but the code should have no access to its representation.

Our proof of correctness of the SFX compiler relies on information on execution history that needs to be logically preserved step by step by our execution models. The ASFX language we devised to enrich source programs with annotations providing evidence for such information is reminiscent of other work in the field of *provenance* [4], studying metadata on the origin, history,

and derivation of information, and techniques to propagate it through computation [7, 13]. This is not accidental: it is easy to imagine that a compiler such as the one we have described could be written in a provenance-tracking language (for instance [27]): the compiler would then produce object code annotated with provenance information about the source code that yielded it; by reflecting that information back into SFX, we would get ASFX. Another related area is *justification logic* [3], a modal proof system where propositions can carry evidence of their validity. We have shown that these related concepts of origin, history, evidence, justification are not only important for their epistemological value, but provide an insight on logical constraints on the relationship between the input and output of a computation (in our case, source code and object code), which can be useful to validate the correctness of an algorithm: we regard this as one of the key contributions of this paper.

Acknowledgments

XXX: remember to acknowledge Daniel contributing to the awesome title!

Data-Availability Statement

We are making no distinct submission for Artifact Evaluation: the key artifact is AsmFX itself and the written proofs in the paper.

References

- [1] Danel Ahman and Andrej Bauer. 2020. Runners in Action. In *Programming Languages and Systems*, Peter Müller (Ed.). Springer International Publishing, Cham, 29–55. doi:10.1007/978-3-030-44914-8_2
- [2] Mario Alvarez-Picallo, Teodoro Freund, Dan R. Ghica, and Sam Lindley. 2024. Effect Handlers for C via Coroutines. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 358 (Oct. 2024), 28 pages. doi:10.1145/3689798
- [3] Sergei Artemov. 2008. Justification Logic. In *Proceedings of the 11th European Conference on Logics in Artificial Intelligence* (Dresden, Germany) (*JELIA '08*). Springer-Verlag, Berlin, Heidelberg, 1–4. doi:10.1007/978-3-540-87803-2_1
- [4] David A. Bearman and Richard H. Lytle. 1985. The Power of the Principle of Provenance. *Archivaria* 21 (Jan. 1985), 14–27. <https://archivaria.ca/index.php/archivaria/article/view/11231>
- [5] Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Binders by day, labels by night: effect instances via lexically scoped handlers. *Proc. ACM Program. Lang.* 4, POPL, Article 48 (Dec. 2019), 29 pages. doi:10.1145/3371116
- [6] Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. 2019. Pyro: Deep Universal Probabilistic Programming. *J. Mach. Learn. Res.* 20 (2019), 28:1–6.
- [7] James Cheney, Laura Chiticariu, and Wang-Chiew Tan. 2009. Provenance in Databases: Why, How, and Where. *Foundations and Trends® in Databases* 1, 4 (2009), 379–474. doi:10.1561/1900000006
- [8] Stephen Dolan, Spiros Eliopoulos, Daniel Hillerström, Anil Madhavapeddy, KC Sivaramakrishnan, and Leo White. 2017. Concurrent systems programming with effect handlers. In *TFP*.
- [9] Stephen Dolan, Leo White, KC Sivaramakrishnan, Jeremy Yallop, and Anil Madhavapeddy. 2015. Effective Concurrency through Algebraic Effects. OCaml. http://kcsrk.info/papers/effects_ocaml15.pdf
- [10] Matthias Felleisen and Daniel P. Friedman. 1986. Control operators, the SECD-machine, and the λ -calculus. In *Formal Description of Programming Concepts III*. North-Holland, Amsterdam, 193–217.
- [11] Aina Georges, Armaël Guéneau, Thomas Van Strydonck, Amin Timany, Alix Trieu, Dominique Devriese, and Lars Birkedal. 2023. Cerise: Program Verification on a Capability Machine in the Presence of Untrusted Code. *J. ACM* 71 (09 2023). doi:10.1145/3623510
- [12] Dan Ghica, Sam Lindley, Marcos Maroñas Bravo, and Maciej Piróg. 2022. High-level effect handlers in C++. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 183 (Oct. 2022), 29 pages. doi:10.1145/3563445
- [13] Todd J. Green, Grigoris Karvounarakis, and Val Tannen. 2007. Provenance semirings. In *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (Beijing, China) (*PODS '07*). Association for Computing Machinery, New York, NY, USA, 31–40. doi:10.1145/1265530.1265535
- [14] Andreas Haas, Andreas Rossberg, Derek Schuff, Ben Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon Zakai, and J.F. Bastien. 2017. Bringing the Web up to Speed with WebAssembly. *SIGPLAN Notices* 52, 6 (June 2017), 185–200.

- [15] Daniel Hillerström. 2015. *Handlers for Algebraic Effects in Links*. Master's thesis. University of Edinburgh, Scotland. http://project-archive.inf.ed.ac.uk/msc/20150206/msc_proj.pdf
- [16] Daniel Hillerström and Sam Lindley. 2016. Liberating effects with rows and handlers. In *Proceedings of the 1st International Workshop on Type-Driven Development, TyDe@ICFP 2016, Nara, Japan, September 18, 2016*, James Chapman and Wouter Swierstra (Eds.). ACM, 15–27. doi:10.1145/2976022.2976033
- [17] Daniel Hillerström, Sam Lindley, and Robert Atkey. 2020. Effect handlers via generalised continuations. *J. Funct. Program.* 30 (2020), e5. doi:10.1017/S0956796820000040
- [18] Thomas Johnsson. 1985. Lambda lifting: Transforming programs to recursive equations. In *Functional Programming Languages and Computer Architecture*, Jean-Pierre Jouannaud (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 190–203.
- [19] Ohad Kammar, Sam Lindley, and Nicolas Oury. 2013. Handlers in action. In *ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25 - 27, 2013*, Greg Morrisett and Tarmo Uustalu (Eds.). ACM, 145–158. doi:10.1145/2500365.2500590
- [20] Daan Leijen. 2014. Koka: Programming with Row Polymorphic Effect Types. *Electronic Proceedings in Theoretical Computer Science* 153 (June 2014), 100–126. doi:10.4204/eptcs.153.8
- [21] Daan Leijen and KC Sivaramakrishnan. 2023. libmprompt and libmpeff. <https://github.com/koka-lang/libmprompt>.
- [22] Cong Ma, Zhaoyi Ge, Edward Lee, and Yizhou Zhang. 2024. Lexical Effect Handlers, Directly. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 330 (Oct. 2024), 29 pages. doi:10.1145/3689770
- [23] Yasuhiko Minamide, Greg Morrisett, and Robert Harper. 1996. Typed closure conversion. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (St. Petersburg Beach, Florida, USA) (POPL '96). Association for Computing Machinery, New York, NY, USA, 271–283. doi:10.1145/237721.237791
- [24] Serkan Muehcu, Philipp Schuster, Michel Steuwer, and Jonathan Immanuel Brachthäuser. 2025. Multiple Resumptions and Local Mutable State, Directly. *Proc. ACM Program. Lang.* 9, ICFP, Article 260 (Aug. 2025), 30 pages. doi:10.1145/3747529
- [25] 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
- [26] Gordon D. Plotkin and Matija Pretnar. 2013. Handling Algebraic Effects. *Log. Methods Comput. Sci.* 9, 4 (2013), 36 pages. doi:10.2168/LMCS-9(4:23)2013
- [27] Wilmer Ricciotti. 2017. A core calculus for provenance inspection. In *Proceedings of the 19th International Symposium on Principles and Practice of Declarative Programming* (Namur, Belgium) (PPDP '17). Association for Computing Machinery, New York, NY, USA, 187–198. doi:10.1145/3131851.3131871
- [28] K. C. Sivaramakrishnan, Stephen Dolan, Leo White, Tom Kelly, Sadiq Jaffer, and Anil Madhavapeddy. 2021. Retrofitting effect handlers onto OCaml. In *PLDI '21: 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, Virtual Event, Canada, June 20-25, 2021*, Stephen N. Freund and Eran Yahav (Eds.). ACM, 206–221. doi:10.1145/3453483.3454039
- [29] Robert N. M. Watson, Peter G. Neumann, Jonathan Woodruff, Michael Roe, Hesham Almatary, Jonathan Anderson, John Baldwin, Graeme Barnes, David Chisnall, Jessica Clarke, Brooks Davis, Lee Eisen, Nathaniel Wesley Filardo, Franz A. Fuchs, Richard Grisenthwaite, Alexandre Joannou, Ben Laurie, A. Theodore Markettos, Simon W. Moore, Steven J. Murdoch, Kyndylan Nienhuis, Robert Norton, Alexander Richardson, Peter Rugg, Peter Sewell, Stacey Son, and Hongyan Xia. 2023. *Capability Hardware Enhanced RISC Instructions: CHERI Instruction-Set Architecture (Version 9)*. Technical Report UCAM-CL-TR-987. University of Cambridge, Computer Laboratory. doi:10.48456/tr-987
- [30] Ningning Xie, Jonathan Immanuel Brachthäuser, Daniel Hillerström, Philipp Schuster, and Daan Leijen. 2020. Effect handlers, evidently. *Proc. ACM Program. Lang.* 4, ICFP, Article 99 (Aug. 2020), 29 pages. doi:10.1145/3408981
- [31] Yizhou Zhang and Andrew C. Myers. 2019. Abstraction-safe effect handlers via tunneling. *Proc. ACM Program. Lang.* 3, POPL, Article 5 (Jan. 2019), 29 pages. doi:10.1145/3290318

A An AsmFX compiler for SFX

We give here a detailed definition of the SFX compiler that we only sketched in Section 6. The compiler consists of three procedures, for values, computations, and handlers.

Compiling values. To express the compiler, we start by giving a procedure to compile SFX open values: the operation `[V]` (Figure 6) produces code which will load an initial segment of the \$a registers of suitable length with the AsmFX representation of that value. We use blue symbols

to differentiate meta-operations such as the compiler and its definition from the object syntax it manipulates.

Compilation of values uses an auxiliary definition $[V]^n$ where the subvalues are loaded starting at register $\$a_n$. The representation of constants $[c]$ is assumed as primitive. In the case of free variables x , we assume that the registers $\$x_0, \dots, \$x_{\|x\|-1}$ contain the AsmFX representation of the value of x , and we move their content to the argument registers.

Compiling computations. The table in Figure 7 describes the compilation rules for each possible syntactic form of computations. Before discussing some relevant cases, we note the following preliminary facts:

- We assume that every SFX primitive operation \oplus is correctly implemented by an AsmFX instruction $[\oplus]$, taking as many register arguments as needed to contain its input, and similarly storing its output in an a suitable number of output registers.
- Booleans are represented as integers in the obvious way.
- Instructions j *target_code* and bz *target_code*, $\$r$ representing an unconditional jump to *target_code* and a branch to *target_code* if $\$r$ is zero are available.
- The operation $\phi(n, \mathbb{L})$ produces n labels fresh with respect to the used label set \mathbb{L} . The concrete choice of labels is irrelevant and consistent renaming is assumed to denote the same code (we use labels for the sake of readability, but each label must be considered notation for an absolute memory index).

We discuss the compilation rules in detail. To compile a **return**, we simply compile the returned value. A primitive operation $\oplus(\underline{V})$ is compiled by compiling V first, then executing the instruction $[\oplus]$ implementing \oplus using the $\$a$ registers both as input and as output. Conditionals $\text{if}(V, M, N)$ require us to compile V first and then use a bz “branch if zero” instruction to discriminate between true and false values: in one case we continue with the code for M , in the other case we branch to l_{else} and the code for N ; at the end, both branches meet again at label l_{end} . A **let** $x \leftarrow M$ **in** N is compiled by compiling M first: after executing the compiled code for M , we will copy the result of that computation from the $\$a$ register where it has been stored into the $\$x$ registers, where N expects to find it; finally we compile N . The compilation rule for **unpack** simply needs compile its value argument and then move the result of its evaluation to the registers for the \bar{x}^k variables (the rule makes use of an extended vector notation $\overline{\$x_{m,n}}$ which stands for n registers starting from $\$x_m$, i.e. $\$x_m, \$x_{m+1}, \dots, \$x_{m+n-1}$).

Now we consider computations for the effectful sublanguage. The U **with** $h(V)$ computation is compiled by compiling V first and storing the result of its evaluation in an initial segment of the $\$s$ registers; we use these registers together with the handler specification $[h]$ to load the handler $h(V)$ into register $\$hr$; then we compile U : its value will be stored in $\$a_0, \a_1 ; finally, we replace the content of $\$a_1$ with the content of $\$hr$.

$$\begin{aligned}
 [V] &:= [V]^0 \\
 [x]^n &:= \$a_n \dots \$a_{n+\|x\|-1} \leftarrow \$x_0 \dots \$x_{\|x\|-1} \\
 [c]^n &:= \$a_n \dots \$a_{n+\|c\|-1} \leftarrow [c] \\
 [\langle \bar{V}_p \rangle]^n &:= [V_0]^n; [V_1]^{n+\|V_0\|}; \dots; [V_{p-1}]^{n+\sum_{i=0, \dots, p-2} \|V_i\|}
 \end{aligned}$$

Fig. 6. Compilation SFX values to AsmFX.

1275		$\llbracket \oplus(V) \rrbracket^{\mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}' :=$		
1276		$\frac{\llbracket V \rrbracket}{\$a_{\ out(\oplus)\ } \leftarrow \llbracket \oplus \rrbracket \$a_{\ V\ }}$		
1277	$\llbracket \text{return } V \rrbracket^{\mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}' :=$	$\llbracket V \rrbracket$		
1278	$\llbracket \text{let } \mathbb{L}' = \mathbb{L} \rrbracket$			
1279				
1280	$\llbracket \text{let } x \leftarrow M \rrbracket^{\mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}' :=$			
1281	$\llbracket \text{in } N \rrbracket^{\mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}' :=$			
1282	$\llbracket \text{let } \mathbb{I}_M = \overline{\$x_{\ M\ }} \rrbracket$			
1283	$\llbracket M \rrbracket^{\mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}'_M$			
1284	$\mathbb{I}_M \leftarrow \$a_{\ M\ }$			
1285	$\llbracket N \rrbracket^{\mathbb{I}_M \mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}'$			
1286				
1287	$\llbracket \text{unpack } \overline{x^k} \leftarrow V \rrbracket^{\mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}' :=$			
1288	$\llbracket \text{in } M \rrbracket^{\mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}' :=$			
1289	$\llbracket \text{let } \mathbb{I}_M = \mathbb{I} \cup \{ \$x^k_{\ x_n\ } \} \rrbracket$			
1290	$\llbracket V \rrbracket$			
1291	$\overline{\$x^0_{\ x^0\ }} \leftarrow \overline{\$a_{\ x^0\ }}$			
1292	\vdots			
1293	$\overline{\$x^i_{\ x^i\ }} \leftarrow \overline{\$a_{\ \Sigma \ x^{i-1}\ , \Sigma \ x^i\ }}$			
1294	\vdots			
1295	$\overline{\$x^{n-1}_{\ x^{n-1}\ }} \leftarrow \overline{\$a_{\ \Sigma \ x^{n-2}\ , \Sigma \ x^{n-1}\ }}$			
1296	$\llbracket M \rrbracket^{\mathbb{I}_M \mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}'$			
1297				
1298	$\llbracket U \text{ with } h(V) \rrbracket^{\mathbb{L}\mathbb{L}} \Rightarrow \mathbb{L}' :=$			
1299	$\llbracket V \rrbracket$			
1300	$\overline{\$s_{\ V\ }} \leftarrow \overline{\$a_{\ V\ }}$			
1301	$\text{loadh } \llbracket h \rrbracket, \text{state} : (\overline{\$s_{\ V\ }})$			
1302	$\llbracket U \rrbracket$			
1303	$\$a_1 \leftarrow \hr			
1304	$\text{let } \mathbb{L}' = \mathbb{L}$			
1305				
1306				
1307				
1308				
1309				
1310				
1311				
1312				
1313				
1314				
1315				
1316				
1317				
1318				
1319				
1320				
1321				
1322				
1323				

Fig. 7. Compilation of SFX computations to AsmFX.

To create a continuation from a function by means of **newbroom**(f) **with** hV , we need to create a suitable broom and place it in register $\$a_0, \a_1 . We first compile V , place its content in the $\$s$ registers, and use it together with $\llbracket h \rrbracket$ to create a handler which will be loaded into $\$hr$ first and then $\$a_1$. Then we construct a continuation pointing to a label l_{res} which will contain code to invoke f : the continuation is initially stored in $\$kr$, then moved into $\$a_0$. Finally we jump to l_{end} , the end of the code. At location l_{res} , we to call f we actually need to invoke its return clause: this is done by jumping unconditionally to the f_{return} address, after installing in $\$lr$ a leave record indicating what code should be executed when f terminates — f is executed within the handler $h(V)$, therefore at the end of its execution it should invoke the return clause for $h(V)$: that corresponds to the instruction **return**, at the label l_{exit} .

To invoke an effect by means of **do** $\#t(V)$, we first compile V : this places the result of its evaluation into the $\$a$, where the operation handler expects to find it. To invoke the operation handler, we

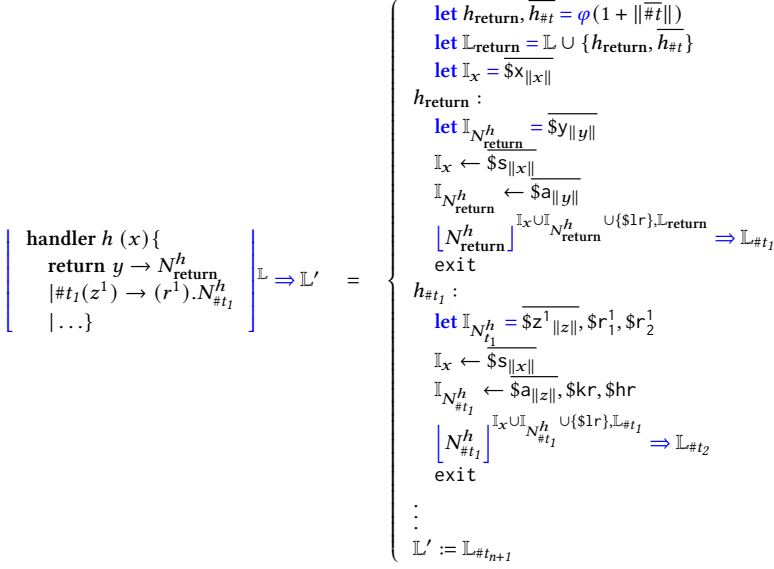


Fig. 8. Compilation of SFX handlers to AsmFX.

just execute the instruction $\text{do } \#t, \text{save} : \mathbb{I}$ (saving all the input registers will allow the resumption for the $\#t$ operation to restore the original client environment).

To execute a computation under a new handler by means of **handle** M with $h(V)$, we compile V , copy its evaluation in the $\$s$ registers, and use it together with $[h]$ to create a handler and place it in $\$hr$. We also create a continuation pointing to the code for M (l_{cli}) and place it in $\$kr$, and a leave record for the code to be executed after this computation (l_{exit}), with saved registers corresponding to the input set \mathbb{I} . After these three registers have been loaded, we invoke the client by resuming the continuation we created, using the instruction **resume**. Notice that the client code at location l_{cli} is terminated by a return instruction, so that the result of M will be delivered to the return clause of the newly installed handler: we call this a **return** epilogue. Tracking epilogues is essential to prove the correctness of the compiler.

A **resume** $U(V)$ works similarly: registers $\$kr$ and $\$hr$ are filled with the result of the evaluation of U ; the result of the evaluation of V is put in registers $\$a$; we create a leave record for the code l_{exit} to be run after executing this computation (again, the input registers \mathbb{I} must be added to the save area of $\$1r$). Finally, we invoke the resume instruction.

Compiling handlers. The procedure to compile handlers is detailed in Figure 8. The return and operation clauses of each handler are compiled separately using the same algorithm defined for computations. In particular, for each handler h with clauses for operations $\#t_i$ we generate unique global labels $h_{\text{return}}, \overline{h_{\#t_i}}$ for the code implementing the operations. Each parameterised handler with x as its formal parameter will receive its state in an initial segment of the state registers $\$s$ of suitable size, and will allocate variable registers $\$x$ corresponding to the SFX parameter variable x to copy the state into (each clause performs this copy separately).

The return clause h_{return} and all operation clauses $\overline{h_{\#t_i}}$ receive their arguments in an initial segment of the argument registers $\$a$ of suitable size: those are also copied into variable registers $\$y, \overline{\z^1} corresponding to the formal arguments y (for the return clause) and z_i (for the $\#t_i$ operation

clause). Operation clauses will also copy the content of registers $\$kr, \hr (which contain the broom resumption for the operation) to variable registers $\$r_1^i, r_2^i$ corresponding to the SFX broom variable r^i .

Then the SFX computation for each clause $N_{\text{return}}^h, N_{\#t_i}^h$ is compiled by providing the suitable set of input registers, which must contain the leave register $\$lr$, as this should not be accidentally overwritten.

The handler result (returned when exiting the handler, either normally through the return clause or abnormally through an operation clause) is stored in an initial segment of the argument registers $\$a$. Each clause is terminated by an `exit` instruction which triggers the leave register; this instruction also constitutes an epilogue that needs to be accounted for in the soundness proof.

B Soundness of compilation

As a final step in our effort to assess the expressiveness of AsmFX, we need to prove the correctness of our compiler. The most basic correctness property would state that for any SFX program P and its AsmFX counterpart $\llbracket P \rrbracket$, if P terminates returning a value v , then $\llbracket P \rrbracket$ also terminates and a machine representation of v can be found in an initial segment of the $\$a$ registers.

In fact, we are also interested in proving the correctness of partial executions: to show this, we need to prove that the transition system induced by the AsmFX machine correctly simulates the transition system induced by the SFX abstract machine.

The proof is not entirely straightforward. The main difficulty lies in the fact that executing an SFX program consumes part of the program itself, blurring the correspondence between source and compiled object code beyond recognisability: the code resulting from the recompilation of a partially executed program will not match, even partially, the code of the original program, because parts of the code are completely lost and thus ignored by the compiler. To match partially executed source code to partially executed AsmFX code we need a more sophisticated instrument than the compiler itself.

Our solution is to annotate key parts of the source terms with the memory locations where they were originally placed by the compiler. If we have enough annotations, we may reconstruct a location-preserving compiled code which will partially match the original AsmFX code for the full program.

B.1 Augmenting SFX

We introduce the augmented source language ASFX: a conservative extension of SFX embedding location information and other information about execution history. Compared to SFX, ASFX sports an enriched syntax and a specific abstract machine clearly derived from its SFX counterpart, but presents some key differences that makes it mimic AsmFX code more closely. For our goals, we do not need to provide a type system for ASFX.

The syntax of ASFX is shown in Figure 9. Augmented values annotate SFX values with a single location referencing the code responsible for loading the value in the appropriate registers (an instance of $\llbracket V \rrbracket$). Augmented computations \hat{M} are isomorphic to a plain M , but every subcomputation (including \hat{M} itself, save for the case of `return`) is annotated with a start location. For every \hat{V} and \hat{M} we define their start locations $\mu(V)$ and $\mu(M)$ as its top-level annotation (in particular, $\mu(\text{return } V @ \ell) = \mu(V @ \ell) = \ell$). We also define the end location of a value or computation as the compiled length of that value or computation plus its start location:

$$v(\hat{V}) = \mu(\hat{V}) + \text{length}(\llbracket V \rrbracket)$$

$$v(\hat{M}) = \mu(\hat{M}) + \text{length}(\llbracket M \rrbracket)$$

Open values:	$\hat{U}, \hat{V}, \hat{W}$	$::=$	$V@l$
Computations:	\hat{M}, \hat{N}	$::=$	$\text{return } \hat{V} \mid \oplus(\hat{V})@l \mid \text{if}(\hat{V}, \hat{M}, \hat{N})@l$ $\mid (\text{let } x \leftarrow \hat{M} \text{ in } \hat{N})@l \mid (\text{unpack } \bar{x} \leftarrow \hat{V} \text{ in } \hat{M})@l$ $\mid (\text{do } \#t(\hat{V}))@l \mid (\text{handle } \hat{M} \text{ with } h(\hat{V}))@l$ $\mid (\text{resume } \hat{U}(\hat{V}))@l \mid (\hat{U} \text{ with } h(\hat{W}))@l$ $\mid (\text{newbroom}@_{\ell_{\text{exit}}}(f) \text{ with } h(\hat{W}))@l$
Literal values:	$\hat{u}, \hat{v}, \hat{w}$	$::=$	$c \mid \langle \hat{v} \rangle \mid \hat{k}$
Environments:	$\hat{\gamma}$	$::=$	$x \mapsto \hat{v}$
Literal cont.:	\hat{k}	$::=$	$(\hat{k}, h(\hat{v}), \hat{\gamma})$
Continuations:	$\hat{\kappa}, \hat{\kappa}'$	$::=$	$0@l \mid \hat{\zeta}@l \cdot \hat{\kappa}$
Frames:	$\hat{\zeta}, \hat{\zeta}'$	$::=$	$\varepsilon \mid (x). \hat{M} \mid h(\hat{v}, \mathcal{L}(\ell, \hat{\gamma}))$
Epilogues:	ε	$::=$	$\text{jump} \mid \text{exit} \mid \text{leave}$
Handlers:	\hat{H}	$::=$	$\text{handler } h(x) \{ \text{return } y \rightarrow \hat{M} \mid \#t(z) \rightarrow (r). \hat{N} \}$
Specifications:	$\hat{\Sigma}$	$::=$	\hat{H}
Programs:	P	$::=$	$\hat{\Sigma}; \hat{M}$

Fig. 9. ASFX: syntax.

Compared to SFX, ASFX continuations have some important differences: the identity continuation and all frames are annotated with a start location; additionally, pure frames $((x).\hat{M})@l$ contain a naked computation instead of a computation closure; we also introduce *epilogues* ε , a type of frame that has no counterpart in SFX but allows us to more closely track the execution of the compiled AsmFX code. The start location $\mu(\hat{\kappa})$ of a continuation κ is defined as the start location of its first frame, or in the case of the identity continuation as $\mu(0@l) = l$. We do not define the end location of a frame or of a continuation.

The closures missing from pure frames resurface as an additional environment $\hat{\gamma}$ component of ASFX brooms: this structure must not be understood as a closure, in the sense that the environment does not apply to the whole broom: it is set as the initial environment when the broom is resumed, and it may be extended or entirely removed before the execution of the continuation is completed.

The other syntactic categories of ASFX are analogous to their SFX counterparts, but use the augmented versions of each nested expression.

B.2 An abstract machine for the augmented source language

Much like we did for SFX, we formalise the semantics of ASFX by means of an abstract machine. The main difference between the two abstract machines is that the ASFX machine is aware of location annotations and propagates them during its execution, ensuring that all the relevant information is preserved. As we previously mentioned, the definition of continuations is also different.

The full definition is given in Figure 10. Note that we omit the top-level location of the augmented code because it is irrelevant to the machine, however the subexpressions contain locations that are used in the transition rules.

The environment can now reference a *pseudovisible* LEAVE that has no counterpart in SFX, but is analogous to the leave records of AsmFX. Note the invariant that the evaluation of a certain computation always terminates with an environment that is an extension of the original one: for this reason, the evaluation of primitive operations terminates with the original $\hat{\gamma}$ instead of the empty environment $[]$ — this was irrelevant in SFX, but it is crucial in ASFX to ensure the

control	env.	continuation	\hookrightarrow	control	env.	continuation
$\oplus(\hat{V})$	\hat{y}	\hat{k}		$\llbracket \oplus \rrbracket (\llbracket \hat{V} \rrbracket \hat{y})$	\hat{y}	\hat{k}
if $(\hat{V}, \hat{M}, \hat{N})$ where: $\llbracket \hat{V} \rrbracket \hat{y} = \text{true}$	\hat{y}	\hat{k}		\hat{M}	\hat{y}	$\text{jump}@v(\hat{M}) \cdot \hat{k}$
if $(\hat{V}, \hat{M}, \hat{N})$ where: $\llbracket \hat{V} \rrbracket \hat{y} = \text{false}$	\hat{y}	\hat{k}		\hat{N}	\hat{y}	\hat{k}
let $\bar{x} \leftarrow \hat{M}$ in \hat{N}	\hat{y}	\hat{k}		\hat{M}	\hat{y}	$((x).\hat{N})@v(\hat{M}) \cdot \hat{k}$
unpack $\bar{x}_n \leftarrow \langle \hat{V}_n \rangle$ in \hat{N}	\hat{y}	\hat{k}		\hat{N}	$(\bar{x}_n \mapsto \llbracket \hat{V}_n \rrbracket \hat{y}) \cdot \hat{y}$	\hat{k}
handle \hat{M} with $h(\hat{V})$	\hat{y}	\hat{k}		\hat{M}	$\hat{y} \setminus \text{LEAVE}$	$h(v, \mathcal{R}_{\mathcal{L}})@v(\hat{M}) \cdot \hat{k}$ where: $v = \llbracket \hat{V} \rrbracket \hat{y}$ $\mathcal{R}_{\mathcal{L}} = \mathcal{L}(\mu(\hat{k}), \hat{y})$
do $\#t(\hat{V})$	\hat{y}	$\hat{k}^{-t} \cdot h^{+t}(\hat{w}, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{k}'$		$h_{\#t}$	$(\text{LEAVE} \mapsto \mathcal{R}_{\mathcal{L}}) \cdot (\text{res}(\#t) \mapsto (\hat{k}^{-t}, h^{+t}(\hat{w})@l, \hat{y})) \cdot (\text{dom}(\#t) \mapsto \llbracket \hat{V} \rrbracket \hat{y}) \cdot (\text{dom}(h) \mapsto \hat{w})$	$\text{exit}@v(h_{\#t}) \cdot \hat{k}'$
resume $\hat{U}(\hat{V})$ where: $\llbracket \hat{U} \rrbracket \hat{y} = (\hat{k}', h(\hat{w})@l, \hat{y}')$	\hat{y}	\hat{k}		$\llbracket \hat{V} \rrbracket \hat{y}$	\hat{y}'	$\hat{k}' \cdot h(\hat{w}, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{k}$ where: $\mathcal{R}_{\mathcal{L}} = \mathcal{L}(\mu(\hat{k}), \hat{y})$
\hat{U} with $h(\hat{W})$ where: $\llbracket \hat{U} \rrbracket \hat{y} = (\hat{k}', _@l, \hat{y}')$	\hat{y}	\hat{k}		$(\hat{k}', h(\llbracket \hat{W} \rrbracket \hat{y})@l, \hat{y}')$	\hat{y}	\hat{k}
newbroom $@l(f)$ with $h(\hat{V})$ where: $\text{dom}(f) = x$ $\ell_f = \mu(f_{\text{return}})$ $\ell'_f = v(f_{\text{return}})$ $\ell' = \ell + 1$ $\ell'' = \ell + 2$	\hat{y}	\hat{k}		$(\text{leave}@l \cdot \text{jump}@l' \cdot ((x).f_{\text{return}})@l_f \cdot \text{exit}@l'_f, h(\llbracket \hat{V} \rrbracket \hat{y})@l'', [])$	\hat{y}	\hat{k}
return \hat{V}	\hat{y}	\hat{k}		$\llbracket \hat{V} \rrbracket \hat{y}$	\hat{y}	\hat{k}
\hat{v}	\hat{y}	$\text{jump}@l \cdot \hat{k}$		\hat{v}	\hat{y}	\hat{k}
\hat{v}	\hat{y}	$\text{exit}@l \cdot \hat{k}$		\hat{v}	$\hat{y}(\text{LEAVE}).\text{env}$	\hat{k} where: $\mu(\hat{k}) = \hat{y}(\text{LEAVE}).\text{loc}$
\hat{v}	\hat{y}	$\text{leave}@l \cdot \hat{k}$		\hat{v}	$(\text{LEAVE} \mapsto \mathcal{L}(\ell+2, [])) \cdot \hat{y}$	\hat{k}
\hat{v}	\hat{y}	$((x).\hat{N})@l \cdot \hat{k}$		\hat{N}	$(x \mapsto \hat{v}) \cdot \hat{y}$	\hat{k}
\hat{v}	\hat{y}	$h(\hat{w}, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{k}$		h_{return}	$(\text{LEAVE} \mapsto \mathcal{R}_{\mathcal{L}}) \cdot (\text{dom}(h_{\text{return}}) \mapsto \hat{v}) \cdot (\text{dom}(h) \mapsto \hat{w})$	$\text{exit}@v(h_{\text{return}}) \cdot \hat{k}$

Fig. 10. ASFX: abstract machine

correctness of the evaluation of **let** bindings despite the fact that we removed closures from pure frames.

Beside the change in pure frames, ASFX introduces epilogue frames, used in the target of the rules for **if** (when the guard evaluates to **true**), for **do**, for **resume**, and in the value case when the continuation starts with a handler frame. Epilogues do not correspond to any code in SFX, but do correspond to AsmFX code that needs to be executed after evaluating a certain computation, and before moving to the next one. In the case of the guard of an **if** being evaluated to true, after

evaluating the *then* branch \hat{M} and before moving to the next computation, we need to jump over the code for the *else* branch: this is why we add to the continuation a **jump** epilogue, which will be located at $v(\hat{M})$ (immediately after the code for \hat{M} : this fine grained analysis will help greatly in matching SFX code, where control flow is dictated in part by the syntactic structure, and AsmFX, which is unstructured and whose control flow depends exclusively on explicit instructions.

The rule for **handle** is similar to the one in SFX, but it clears the content of LEAVE because the AsmFX implementation trashes the original content of register $\$1r$ (if any). The new handler frame receives an annotation $v(\hat{M})$, reflecting the fact that for all computation states, the end location of the current computation must match the start location of the current continuation.

Execution of a **do** differs from the corresponding rule in SFX in that the target environment has an updated resumption value, which saves the source environment $\hat{\gamma}$, ready to be restored when the resumption is triggered (SFX does not need to do this because all pure frames have an environment, and all other frames do not care about the environment). We also provide a value for LEAVE, the leave record of the handler being invoked, which is meaningful to the **exit** epilogue in the target continuation: **exit** corresponds to the **exit** instruction at the end of the AsmFX code for a handler operation clause or return clause: when we reach such an instruction, or equivalently when the epilogue **exit** is activated, control is transferred to the code referenced by that leave record.

The computation **resume** $\hat{U}(\hat{V})$ is evaluated similarly to SFX, but the environment is replaced in the target with the one found in the broom U , whereas the original environment is saved in the leave record linked to the broom handle that is being added to the target continuation. Broom handle replacement \hat{U} with \hat{V} is similar to its SFX counterpart, but it needs to take account for the additional environment field of augmented brooms.

The rule for **newbroom**@ $\ell(f)$ with $h(\hat{V})$ is more complex than its SFX counterpart: while it does create a broom whose head is the body of f and whose handle is the specified handle h , it also needs to decorate the head with a number epilogues for administrative instructions that are executed either in the AsmFX code for **newbroom** (**leave**, **jump**) or at the end of the function f (**exit**).

The transition for **return** simply substitutes the environment into the return value to obtain a closed value. Then we have four different transitions for states examining a closed value, one for each possible frame at the start of the continuation:

- If the continuation starts with a **jump** epilogue, the epilogue is simply removed from the continuation, reflecting the fact that we have jumped to the next frame.
- In the case of an **exit** epilogue, we similarly remove the epilogue from the continuation, but we also replace the environment according to the active leave record, reflecting the execution of an **exit** instruction.
- A **leave** epilogue corresponds to an AsmFX **loadl** instruction loading a distinguished address (two positions after the instruction itself) into the current leave record.
- When the continuation starts with a pure frame, we are passing the value under consideration as an argument to the computation \hat{N} enclosed in the pure frame: the corresponding transition rule works similarly as in SFX, however since augmented pure frames do not come with an environment, we must assume that the source environment is compatible with \hat{N} .
- Finally if we are returning a value to a handler frame, we are really invoking its return clause: just like in SFX, we will continue execution with the return branch of that handler. Compared with SFX, we need to load the leave record of the handler into LEAVE, and also extend the target continuation with an **exit**, to denote the fact that when the evaluation of

the return clause terminates, one has to move to code external to the handler itself. Again, this corresponds to an `exit` instruction placed at the end of the return clause.

B.3 Well-scoping

We can prove that single step transitions in SFX are simulated by multistep transitions in ASFX. Since ASFX epilogues essentially correspond to no operation at all in SFX, the main challenge of this proof is to show that in spite of a different approach to handling environments, the ASFX provides a correct implementation of scopes. The simulation relation, summarised in Figure 4, essentially expresses a well-scoping invariant.

When an SFX configuration s is simulated by an ASFX configuration \hat{s} , we write $s \leq \hat{s}$. The relation is defined by means of auxiliary definitions for the simulation of values, environments and continuations. There are two different cases based on whether the control element of the two configurations is a computation or a closed value: in the first case, $\langle M, \gamma, \kappa \rangle \leq \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle$ means that if we erase all annotations from \hat{M} , we obtain M ; additionally, the augmented environment $\hat{\gamma}$ simulates the simple environment γ , and the pair of augmented environment and continuation $(\hat{\gamma}, \hat{\kappa})$ simulates the simple continuation κ . If instead we are in the value case, we have that $\langle v, \gamma, \kappa \rangle \leq \langle \hat{v}, \hat{\gamma}, \hat{\kappa} \rangle$ if and only if the augmented value \hat{v} simulates the simple value v , and the pair of augmented environment and continuation $(\hat{\gamma}, \hat{\kappa})$ simulates γ (however γ need not be simulated by $\hat{\gamma}$). It may seem odd that the continuation simulation requires, on the augmented side, an environment: this is explained by the fact that SFX pure frames are closed with respect to an environment, while ASFX frames are not.

The value simulation $v \leq \hat{v}$, for non-broom values, simply states that erasure of the augmented value results in the simple value; however, broom simulation $(\kappa, h(w)) \leq (\hat{\kappa}, h(\hat{w}))@l, \hat{\gamma}_\kappa$ is recursively defined: it requires that κ be simulated by $(\hat{\gamma}_\kappa, \hat{\kappa})$ and that w be simulated by \hat{w} .

The environment simulation $\gamma \leq \hat{\gamma}$ is defined pointwise for all variables in the domain of γ ; $\hat{\gamma}$ is allowed to be defined on more variables.

The continuation simulation $\kappa \leq (\hat{\gamma}, \hat{\kappa})$ is defined by means of an inductive judgement, with a base rule for the identity continuation and recursive rules for each type of frame. A **jump** epilogue frame preserves the simulation; an **exit** frame swaps the active augmented environment with the one contained in the **LEAVE** pseudovisible (which must be defined). A **leave** epilogue requires the next continuation $\hat{\kappa}$ to simulate the original plain continuation κ in an environment with an updated leave record. In the case of handler frames, both continuations must start with matching handler frames: if that is the case, the simulation holds if there is a simulation between the remaining parts of the two continuations; however on the ASFX side of the simulation, we will need to swap the environment with the one contained in the leave record of the handler frame, reflecting the fact that when we exit a handler, the environment is restored from that leave record.

Finally, the most interesting case is that of pure frames. If both continuations start with a pure frame, in order for the simulation to hold, the computations N and \hat{N} in each pure frame must match (meaning that by erasing \hat{N} we get N). Furthermore, the SFX pure frame is closed by an environment γ_N : that environment must be simulated by the ASFX environment $\hat{\gamma}$, reflecting the fact that in ASFX we do not need an environment in the pure frame, because the active environment is already the correct one. If these conditions hold, the simulation is valid if there is a simulation between the two remaining parts of the two continuations; however on the ASFX side we need to extend the environment with a value for the variable x expected by the pure continuation, and the simulation must hold independently of the value chosen.

We now move to proving the simulation theorem.

B.3.1 Proof of Theorem 7.1. For all SFX configurations s, s' such that $s \hookrightarrow s'$, for all ASFX configurations \hat{s} such that $s \leq \hat{s}$, there exists an ASFX configuration \hat{s}' such that $s' \leq \hat{s}'$ and $\hat{s} \hookrightarrow \hat{s}'$.

In the proof, we will require the following auxiliary results.

LEMMA B.1. Suppose $\kappa \leq (\hat{\gamma}, \hat{\kappa})$ and $\hat{\gamma} \subseteq \hat{\gamma}'$: then we have $\kappa \leq (\hat{\gamma}', \hat{\kappa})$.

PROOF. By induction on the well-scopedness judgement. \square

LEMMA B.2. If $\gamma \leq \hat{\gamma}$, then $\gamma \leq \hat{\gamma} \setminus \text{LEAVE}$.

PROOF. Trivial by unfolding the definitions. \square

LEMMA B.3. Suppose $\kappa_1 \leq (\hat{\gamma}_1, \hat{\kappa}_1)$, $\kappa_2 \leq (\hat{\gamma}_2, \hat{\kappa}_2)$, and $|\hat{w}| = w$. Then we have $\kappa_1 \cdot h(w) \cdot \kappa_2 \leq (\hat{\gamma}_1, \kappa_1 \cdot h(w, \mathcal{L}(\mu(\hat{\kappa}_2), \hat{\gamma}_2)) @ \ell \cdot \hat{\kappa}_2)$.

PROOF. By induction on $\kappa_1 \leq (\hat{\gamma}_1, \hat{\kappa}_1)$. In the base case $\kappa_1 = \mathbf{0}$ and $\hat{\kappa}_1 = \mathbf{0} @ \ell$, we use the hypothesis $\kappa_2 \leq (\hat{\gamma}_2, \hat{\kappa}_2)$ to prove the thesis in the form $h(w) \cdot \kappa_2 \leq (\hat{\gamma}_1, h(w, \mathcal{L}(\mu(\hat{\kappa}_2), \hat{\gamma}_2)) @ \ell \cdot \hat{\kappa}_2)$. \square

LEMMA B.4. If $\kappa \leq (\hat{\gamma}, \varepsilon \cdot \hat{\kappa})$, then there exists $\hat{\gamma}^*$ such that $\kappa \leq (\hat{\gamma}^*, \hat{\kappa})$ and for all \hat{v} we have a transition $\langle \hat{v}, \hat{\gamma}, \varepsilon \cdot \hat{\kappa} \rangle \hookrightarrow \langle \hat{v}, \hat{\gamma}^*, \hat{\kappa} \rangle$.

PROOF. By case analysis on $\kappa \leq (\hat{\gamma}, \varepsilon \cdot \hat{\kappa})$. \square

Now we prove the main result.

PROOF OF THEOREM 7.1. We proceed by case analysis on $s \hookrightarrow s'$ and $s \leq \hat{s}$, where $s = \langle M, \gamma, \kappa \rangle$ and $s' = \langle M', \gamma', \kappa' \rangle$. For each case, the SFX transition forces s and s' to have a certain shape, and similarly the source simulation forces \hat{s} to have a shape that adapts to s . We will (deterministically) follow the ASFX machine on \hat{s} for one or more steps, obtaining a target configuration \hat{s}' : then we will need to show that $s' \leq \hat{s}'$.

$M = \oplus(V)$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle \llbracket \oplus \rrbracket (\llbracket V \rrbracket \gamma), \gamma, M \rangle$.

We prove $\hat{s} = \langle \oplus(\hat{V}) @ \ell, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}| = V$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. We choose $\hat{s}' = \langle \llbracket \oplus \rrbracket (\llbracket \hat{V} \rrbracket \hat{\gamma}), \hat{\gamma}, \hat{\kappa} \rangle$ and we prove $\llbracket \oplus \rrbracket (\llbracket \hat{V} \rrbracket \hat{\gamma}) = \llbracket \oplus \rrbracket (\llbracket V \rrbracket \gamma)$, to obtain $s' \leq \hat{s}'$.

$M = \text{if}(V, M_1, M_2)$. This runs differently depending on whether $\llbracket V \rrbracket \gamma = \text{true}$ or **false**. First, in both cases, we prove $\hat{s} = \langle \text{if}(\hat{V}, \hat{M}_1, \hat{M}_2) @ \ell, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}| = V$, $|\hat{M}_1| = M_1$, $|\hat{M}_2| = M_2$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$.

Then, if $\llbracket V \rrbracket \gamma = \text{true}$, we have $s' = \langle M', \gamma', \kappa' \rangle = \langle M_1, \gamma, \kappa \rangle$. We choose $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = \langle \hat{M}_1, \hat{\gamma}, \text{jump} @ \nu(\hat{M}_1) \cdot \hat{\kappa} \rangle$ and we prove that $\kappa \leq (\hat{\gamma}', \hat{\kappa}')$ by applying the appropriate rule in Figure 4, to obtain $s' \leq \hat{s}'$.

If instead $\llbracket V \rrbracket \gamma = \text{false}$, we have $s' = \langle M', \gamma', \kappa' \rangle = \langle M_2, \gamma, \kappa \rangle$. We choose $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = \langle \hat{M}_2, \hat{\gamma}, \hat{\kappa} \rangle$ and we obtain $s' \leq \hat{s}'$ trivially.

$M = \text{let } x \leftarrow M_1 \text{ in } M_2$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_1, \gamma, (x).M_2[\gamma] \cdot \kappa \rangle$.

We prove $\hat{s} = \langle (\text{let } x \leftarrow \hat{M}_1 \text{ in } \hat{M}_2) @ \ell, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{M}_1| = M_1$, $|\hat{M}_2| = M_2$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. We choose $\hat{s}' = \langle \hat{M}_1, \hat{\gamma}, ((x).\hat{M}_2) @ \nu(\hat{M}_1) \cdot \hat{\kappa} \rangle$ and we prove $|\hat{M}_2| = M_2$ and $(x).M_2[\gamma] \cdot \kappa \leq (\hat{\gamma}, ((x).\hat{M}_2) @ \nu(\hat{M}_1) \cdot \hat{\kappa})$ (the latter uses Lemma B.1). This immediately implies $s' \leq \hat{s}'$.

$M = \text{unpack } \overline{x_n} \leftarrow \overline{\langle V_n \rangle} \text{ in } M_0$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_0, (\overline{x_n} \mapsto \llbracket V_n \rrbracket \gamma) \cdot \gamma, \kappa \rangle$.

We prove $\hat{s} = \langle \text{unpack } \overline{x_n} \leftarrow \overline{\langle \hat{V}_n \rangle} \text{ in } \hat{M}_0 @ \ell, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}_n| = \overline{V_n}$, $|\hat{M}_0| = M_0$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. We choose $\hat{s}' = \langle \hat{M}_0, (\overline{x_n} \mapsto \llbracket V_n \rrbracket \hat{\gamma}) \cdot \hat{\gamma}, \hat{\kappa} \rangle$ and we prove $|\hat{M}_0| = M_0$ and $\kappa \leq (\hat{\gamma}', \hat{\kappa})$ (the latter uses Lemma B.1, knowing that $\hat{\gamma} \subseteq \hat{\gamma}'$). This immediately implies $s' \leq \hat{s}'$.

$M = \text{handle } M_0 \text{ with } h(V)$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_0, \gamma, h(\llbracket V \rrbracket \gamma) \cdot \kappa \rangle$.

We prove $\hat{s} = \langle \text{handle } \hat{M}_0 \text{ with } h(\hat{V}) @ \ell, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}| = V$, $|\hat{M}_0| = M_0$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. We choose $\hat{s}' = \langle \hat{M}_0, \hat{\gamma} \setminus \text{LEAVE}, h(\llbracket \hat{V} \rrbracket \hat{\gamma}, \mathcal{L}(\mu(\hat{\kappa}), \hat{\gamma})) @ \nu(\hat{M}_0) \cdot \hat{\kappa} \rangle$ and we prove $\gamma \leq \hat{\gamma} \setminus \text{LEAVE}$ (by Lemma B.2), $\kappa \leq (\hat{\gamma} \setminus \text{LEAVE}, h(\llbracket \hat{V} \rrbracket \hat{\gamma}, \mathcal{L}(\mu(\hat{\kappa}), \hat{\gamma})) @ \nu(\hat{M}_0) \cdot \hat{\kappa})$ (by the appropriate rule in Figure 4). This immediately implies $s' \leq \hat{s}'$.

$M = \text{do } \#t(V)$. In this case, we have:

$$\kappa = \kappa_1^{-t} \cdot h^{+t}(w) \cdot \kappa_2$$

$$s' = \langle M', \gamma', \kappa' \rangle = \langle M_{\#t}^h, (\text{res}(h_{\#t}) \mapsto (\kappa_1^{-t}, h(w))) \cdot (\text{dom}(h) \mapsto w) \cdot (\text{dom}(h_{\#t}) \mapsto \llbracket V \rrbracket \gamma), \kappa_2 \rangle$$

We prove $\hat{s} = \langle (\text{do } \#t(\hat{V})) @ \ell, \hat{\gamma}, \hat{\kappa}_1^{-t} \cdot h^{+t}(\hat{w}, \mathcal{L}(\ell_2, \hat{\gamma}_2)) @ \ell_h \cdot \hat{\kappa}' \rangle$ where $|\hat{V}| = V$, $\kappa_1^{-t} \leq (\hat{\gamma}, \hat{\kappa}_1^{-t})$, $\kappa_2 \leq (\hat{\gamma}_2, \hat{\kappa}_2)$. We choose $\hat{\gamma}' = (\text{LEAVE} \mapsto \mathcal{L}(\ell_2, \hat{\gamma}_2)) \cdot (\text{res}(h_{\#t}) \mapsto (\hat{\kappa}_1^{-t}, h(\hat{w}) @ \ell_h, \hat{\gamma})) \cdot (\text{dom}(h) \mapsto \hat{w}) \cdot (\text{dom}(h_{\#t}) \mapsto \llbracket \hat{V} \rrbracket \hat{\gamma})$, $\hat{s}' = \langle \hat{M}_{\#t}^h, \hat{\gamma}', \text{exit} @ \nu(\hat{M}_{\#t}^h) \cdot \hat{\kappa}_2 \rangle$ and we prove $|\hat{M}_{\#t}^h| = M_{\#t}^h$ (by the definition of the compilation procedure), $\gamma' \leq \hat{\gamma}'$, and $\gamma' \leq (\hat{\gamma}', \text{exit} @ \nu(\hat{M}_{\#t}^h) \cdot \hat{\kappa}_2)$ (by applying the appropriate rule in Figure 4). This immediately implies $s' \leq \hat{s}'$.

$M = \text{resume } U(V) \text{ and } \llbracket U \rrbracket \gamma = (\kappa_U, h(w))$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle \llbracket V \rrbracket \gamma, \gamma, \kappa_U \cdot h(w) \cdot \kappa \rangle$.

We prove $\hat{s} = \langle \text{resume } \hat{U}(\hat{V}) @ \ell, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{U}| = U$, $|\hat{V}| = V$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. Note that in order for $\llbracket U \rrbracket \gamma$ to evaluate to a continuation literal, it must be that $U = x_U$, thus $\llbracket U \rrbracket \gamma = \gamma(x_U)$. Then $\hat{U} = x_U$ and $\llbracket \hat{U} \rrbracket \hat{\gamma} = \hat{\gamma}(x_U)$ and, by applying the hypotheses, we prove $\hat{\gamma}(x_U) = (\hat{\kappa}_U, h(\hat{w}) @ \ell_U, \hat{\gamma}_U)$ such that $\kappa_U \leq (\hat{\gamma}_U, \hat{\kappa}_U)$ and $w \leq \hat{w}$.

Choose $\hat{s}' = \langle \llbracket \hat{V} \rrbracket \hat{\gamma}, \hat{\gamma}_U, \hat{\kappa}_U \cdot h(\hat{w}, \mathcal{L}(\mu(\hat{\kappa}), \hat{\gamma})) @ \ell_U \cdot \hat{\kappa} \rangle$. By Lemma B.3 we prove $\kappa_U \cdot h(w) \cdot \kappa \leq (\hat{\gamma}_U, \hat{\kappa}_U \cdot h(\hat{w}, \mathcal{L}(\mu(\hat{\kappa}), \hat{\gamma})) @ \ell_U \cdot \hat{\kappa})$.

Furthermore, we note $\llbracket V \rrbracket \gamma \leq \llbracket \hat{V} \rrbracket \hat{\gamma}$. This is all we need to prove $s' \leq \hat{s}'$.

$M = U \text{ with } h(W) \text{ and } \llbracket U \rrbracket \gamma = (\kappa_U, _)$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle (\kappa_U, h(\llbracket W \rrbracket \gamma)), \gamma, \kappa \rangle$.

We prove $\hat{s} = \langle (\hat{U} \text{ with } h(\hat{W})) @ \ell, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{U}| = U$, $|\hat{W}| = W$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. Note that in order for $\llbracket U \rrbracket \gamma$ to evaluate to a continuation literal, it must be that $U = x_U$, thus $\llbracket U \rrbracket \gamma = \gamma(x_U)$. Then $\hat{U} = x_U$ and $\llbracket \hat{U} \rrbracket \hat{\gamma} = \hat{\gamma}(x_U)$ and, by applying the hypotheses, we prove $\hat{\gamma}(x_U) = (\hat{\kappa}_U, h(\hat{w}) @ \ell_U, \hat{\gamma}_U)$ such that $\kappa_U \leq (\hat{\gamma}_U, \hat{\kappa}_U)$ and $w \leq \hat{w}$.

Choose $\hat{s}' = \langle (\hat{\kappa}_U, h(\hat{w}) @ \ell_U, \hat{\gamma}_U), \hat{\gamma}, \hat{\kappa} \rangle$. Clearly, $\hat{s} \hookrightarrow \hat{s}'$; then prove $s' \leq \hat{s}'$ by unfolding the definition.

$M = \text{newbroom}(f) \text{ with } h(V) \text{ where } \text{dom}(f_{\text{return}}) = x$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle ((x).f_{\text{return}}[]), h(\llbracket V \rrbracket \gamma), \gamma, \kappa \rangle$.

We prove $\hat{s} = \langle \text{newbroom} @ \ell(f) \text{ with } h(\hat{V}), \hat{\gamma}, \hat{\kappa} \rangle$, where $|\hat{V}| = V$; then we choose $\hat{s}' = \langle (\hat{\kappa}_f, h(\llbracket \hat{V} \rrbracket \hat{\gamma}) @ (\ell + 2), [], \hat{\gamma}', \hat{\kappa}') \rangle$ where

$$\hat{\kappa}_f = \text{leave} @ \ell \cdot \text{jump} @ (\ell + 1) \cdot ((x).f_{\text{return}}) @ \mu(f_{\text{return}}) \cdot \text{exit} @ \nu(f_{\text{return}})$$

We prove that $((x).f_{\text{return}}[]) \cdot h(\llbracket V \rrbracket \gamma) \leq [], \hat{\kappa}_f \cdot h(\llbracket \hat{V} \rrbracket \hat{\gamma}) @ (\ell + 2)$ by repeated applications of the rules defining \leq .

$M = \text{return } V$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle \llbracket V \rrbracket \gamma, \gamma, M \rangle$.

We prove $\hat{s} = \langle \text{return } \hat{V}, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}| = V$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. We choose $\hat{s}' = \langle \llbracket \hat{V} \rrbracket \hat{\gamma}, \hat{\gamma}, \hat{\kappa} \rangle$ and we obtain $s' \leq \hat{s}'$.

$M = v \text{ and } \kappa = (x).M_0[\gamma_0] \cdot \kappa_0$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_0, (x \mapsto v) \cdot \gamma_0, \kappa_0 \rangle$.

We prove $\hat{s} = \langle \hat{v}, \hat{\gamma}, \hat{\kappa}_{\text{epi}} \cdot ((x).\hat{M}_0) @ \ell \cdot \hat{\kappa}_0 \rangle$ such that $v \leq \hat{v}$, $(x).M_0[\gamma_0] \cdot \kappa_0 \leq (\hat{\gamma}, \hat{\kappa}_{\text{epi}} \cdot ((x).\hat{M}_0) @ \ell \cdot \hat{\kappa}_0)$ where $\hat{\kappa}_{\text{epi}}$ consists of epilogue frames only.

By multiple applications of Lemma B.4, we prove that the well-scopedness condition on the continuations implies that there exists an environment $\hat{\gamma}_0$ such that $(x).M_0[\gamma_0] \cdot \kappa_0 \leq (\hat{\gamma}_0, ((x).\hat{M}_0) @ \ell \cdot \hat{\kappa}_0)$

and $\hat{s} \xrightarrow{*} \langle \hat{v}, \hat{\gamma}_0, ((x).\hat{M}_0)@l \cdot \hat{\kappa}_0 \rangle$: by taking one ASFX step more, we find ourselves in configuration $\langle \hat{M}_0, (x \mapsto \hat{v}) \cdot \hat{\gamma}_0, \hat{\kappa}_0 \rangle$, which we take to be our \hat{s}' . By case analysis on the well-scopedness condition for $((x).\hat{M}_0)@l \cdot \hat{\kappa}_0$, we obtain $\gamma_0 \leq \hat{\gamma}_0$, $|\hat{M}_0| = M_0$, and $\kappa_0 \leq ((x \mapsto \hat{v}) \cdot \gamma_0, \hat{\kappa}_0)$. This is enough to prove that $s' \leq \hat{s}'$.

$M = v$ and $\kappa = h(w) \cdot \kappa_0$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_{\text{return}}^h, \gamma', \kappa_0 \rangle$, where $\gamma' = (\text{dom}(h_{\text{return}}) \mapsto v) \cdot (\text{dom}(h) \mapsto w)$.

We prove $\hat{s} = \langle \hat{v}, \hat{\gamma}, \hat{\kappa}_{\text{epi}} \cdot h(w, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{\kappa}_0 \rangle$ such that $\mathcal{R}_{\mathcal{L}} = \mathcal{L}(\mu(\hat{\kappa}_0), \hat{\gamma}_0)$, $h(w) \cdot \kappa_0 \leq (\hat{\gamma}, \hat{\kappa}_{\text{epi}} \cdot h(w, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{\kappa}_0)$ where $\hat{\kappa}_{\text{epi}}$ consists of epilogue frames only, and $v \leq \hat{v}$.

By multiple applications of Lemma B.4, we prove that the well-scopedness condition on the continuations implies that there exists an environment $\hat{\gamma}^*$ such that $h(w) \cdot \hat{\kappa}_0 \leq (\hat{\gamma}^*, h(w, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{\kappa}_0)$ and $\hat{s} \xrightarrow{*} \langle \hat{v}, \hat{\gamma}^*, ((x).\hat{M}_0)@l \cdot \hat{\kappa}_0 \rangle$: by taking one ASFX step more, we find ourselves in configuration $\langle \hat{M}_{\text{return}}^h, \hat{\gamma}', \hat{\kappa}_0 \rangle$, where $\hat{\gamma}' = (\text{LEAVE} \mapsto \mathcal{R}_{\mathcal{L}}) \cdot (\text{dom}(h_{\text{return}}) \mapsto \hat{v}) \cdot (\text{dom}(h) \mapsto \hat{w})$: we take this last ASFX configuration to be \hat{s}' . By case analysis on the well-scopedness condition for $h(w, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{\kappa}_0$, we obtain $\kappa_0 \leq (\hat{\gamma}_0, \hat{\kappa}_0)$; we also prove $\gamma' \leq \hat{\gamma}'$ by unfolding the definitions, and $|\hat{M}_{\text{return}}^h| = M_{\text{return}}^h$ by the definition of compilation. This is enough to prove that $s' \leq \hat{s}'$. \square

B.4 Control flow simulation

To complete the proof of correctness of the compiler, we now need to establish a simulation relation between ASFX and AsmFX. We consider AsmFX configurations $a = \langle \Xi, \Theta, C \rangle$, whose elements are memory stores Ξ (containing the program), register files Θ (finite maps from registers to values), and effect contexts C . Since a property of our compiler is that the memory store for a given source program is established at compilation time and stays read-only during execution, we will allow ourselves to omit it from AsmFX configurations and write $\langle \Theta, C \rangle$ for $\langle \Xi, \Theta, C \rangle$.

A source configuration \hat{s} for a program \hat{P} should correspond to an AsmFX configuration a on $|P|$: actually, for every \hat{s} there are multiple such a , because of the plethora of registers available in the architecture: some of the registers will not have a counterpart in the source configuration and will thus be irrelevant. Therefore, the translation to AsmFX of a source configuration \hat{s} , which we denote by $\llbracket \hat{s} \rrbracket$, will be the set of AsmFX configurations that are a model for \hat{s} up to relevant registers. If we write $a \models \hat{s}$ to mean that a models \hat{s} , then we can define $\llbracket \hat{s} \rrbracket := \{a : a \models \hat{s}\}$. We define what it means to be a model separately for the register file and effect context components: the register file must be a model for the whole configuration; on the other hand, the effect context only needs to model the continuation.

$$\langle \Theta, C \rangle \models \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle := (\Theta \models \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle) \wedge (C \models \hat{\kappa})$$

We can define what it means for an effect context to model a continuation by means of a direct translation that only needs to consider the handler frames:

$$C \models \hat{\kappa} \quad := \quad C = \llbracket \hat{\kappa} \rrbracket$$

$$\llbracket \hat{\kappa} \rrbracket \quad := \quad \begin{cases} \llbracket h(w, \mathcal{R}_{\mathcal{L}}) \rrbracket @l \cdot \llbracket \hat{\kappa}' \rrbracket & \text{if } \hat{\kappa} = h(w, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{\kappa}' \\ \llbracket \hat{\kappa}' \rrbracket & \text{if } \hat{\kappa} = \hat{\xi} \cdot \hat{\kappa}' \text{ and } \hat{\xi} \neq h(_) \end{cases}$$

$$\llbracket h(w, \mathcal{R}_{\mathcal{L}})@l \rrbracket \quad := \quad \mathcal{H}(\llbracket h \rrbracket, (\overline{\$s} = \llbracket w \rrbracket), \llbracket \mathcal{R}_{\mathcal{L}} \rrbracket)$$

$$\llbracket \mathcal{L}(l, \hat{\gamma}) \rrbracket \quad := \quad \mathcal{L}(l, \llbracket \hat{\gamma} \rrbracket)$$

$$\llbracket \overline{(x \mapsto v)} \rrbracket \quad := \quad \overline{(\$x = \llbracket v \rrbracket)}$$

A register file models a source configuration if every variable defined in the source environment is backed by (any number of) registers containing the ASmFX representation of that variable's value; if the program counter matches the address of the beginning of the computation being evaluated or, when the computation has ended and we are returning a value, when it matches the address of the beginning of the continuation. Note that the pseudovalue LEAVE is always mapped to the register \$1r, and every variable x is mapped to as many $\$x$ registers as it is necessary.

$$\begin{aligned} \Theta \models \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle &:= \Theta(\$pc) = \mu(\hat{M}) \wedge \forall x \in \text{dom}(\hat{\gamma}). \overline{\Theta(\$x)} = \lfloor \hat{\gamma}(x) \rfloor \\ \Theta \models \langle v, \hat{\gamma}, \hat{\kappa} \rangle &:= \Theta(\$pc) = \mu(\hat{\kappa}) \wedge \forall x \in \text{dom}(\hat{\gamma}). \overline{\Theta(\$x)} = \lfloor \hat{\gamma}(x) \rfloor \end{aligned}$$

The soundness property we seek states that the modelling relation is a simulation: whenever we have a source transition $\hat{s} \hookrightarrow \hat{s}'$ and an ASmFX configuration $a \models \hat{s}$, then there exists a transition chain $a \xrightarrow{*} a'$ where $a' \models \hat{s}'$. However, this property cannot be proved for an arbitrary \hat{s} : augmented source configurations contain metadata that encodes part of the execution history; while the compiler produces correct metadata, and the abstract machine ensures that the metadata is propagated correctly, arbitrary configurations could be unsound and not allow the simulation we need for the soundness theorem. We thus need to restrict ourselves to *valid* configurations (denoted by $\hat{s}\checkmark$) and prove, as part of our theorem, that validity is preserved by transitions.

The validity predicate is defined in Figure 5: its role is to check that all the ASmFX annotations are consistent with the compiled ASmFX program. A configuration is valid if its environment $\hat{\gamma}$ and continuation $\hat{\kappa}$ are valid and additionally the environment and the continuation are coherent (notation: $\hat{\gamma} \subset \hat{\kappa}$). Furthermore, we require that in a valid configuration whose control is a computation \hat{M} , the ASmFX code at location $\mu(\hat{M})$ corresponds to the compiled code for \hat{M} , and that if instead the control is a value v , the value itself is valid.

A value is valid if any broom \hat{k} syntactically contained in it is valid; an environment is valid if all the values it assigns to its domain are valid. A broom is valid if, when we resume in any valid environment γ' and continuation κ' such that the two are coherent, we get a valid continuation and, furthermore, the environment in the broom is coherent with the extended continuation.

What it means for a continuation to be valid is expressed by an inductive predicate, with a base rule for the identity continuation and a recursive rule for each type of frame. These rules ensure that the annotation for each frame corresponds to suitable ASmFX code: the annotation for **jump** frames must refer to a jump linking to the start of the rest of the continuation; in the case of **leave**, the annotation points to code loading a new leave record; in the case of **exit**, the annotation must refer to an `exit` instruction; the annotation of pure frames corresponds to an intermediate instruction of the compilation rule for sequencing; the annotation for handler frames refers to the return instruction that will invoke that handler's return clause.

B.4.1 Proof of Theorem 7.2. Let P be a source program, and $\Xi = \lfloor P \rfloor$. For all source configurations \hat{s}, \hat{s}' such that $\hat{s}\checkmark$ (valid with respect to Ξ), if $\hat{s} \hookrightarrow \hat{s}'$, then $\hat{s}'\checkmark$ and for all ASmFX configurations a such that $a \models \hat{s}$ there exists a' such that $a \xrightarrow{*} a'$ and $a' \models \hat{s}'$.

To prove the theorem, we will need a few basic results on validity.

LEMMA B.5. If $\hat{\gamma}\checkmark$, then $\llbracket \hat{V} \rrbracket \hat{\gamma}\checkmark$.

PROOF. Routine induction on \hat{V} . □

LEMMA B.6. Let $\hat{\kappa} = \hat{\kappa}_1 \cdot h(\hat{w}, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{\kappa}_0$. Suppose $\hat{\kappa}\checkmark$ and $\hat{\gamma} \subset \hat{\kappa}$. Then we have:

- $\hat{w}\checkmark$
- $\hat{\kappa}_0\checkmark$
- $(\hat{\kappa}_1, h(\hat{w})@l, \hat{\gamma})\checkmark$

- $\mathcal{R}_{\mathcal{L}} = \mathcal{L}(\mu(\hat{\kappa}_0), \hat{\gamma}_0)$ such that $\hat{\gamma}_0 \supset \hat{\kappa}_0$

PROOF. By induction on $\hat{\kappa}_1$ with a case analysis on $\hat{\kappa} \checkmark$. □

LEMMA B.7. If $(\hat{\kappa}, h(\hat{w})@l, \hat{\gamma}) \checkmark$ then $(\hat{\kappa}, h'(\hat{w}')@l, \hat{\gamma}) \checkmark$.

PROOF. The definition of $\hat{\kappa} \checkmark$ is based on coherence, which uses operations *laddr* and *lcont*: the values of these operations are the same for both continuations (the proof is by induction on $\hat{\kappa}$). □

LEMMA B.8. Let $\hat{\kappa} = (\hat{\kappa}, h(\hat{w})@l, \hat{\gamma})$. If $\hat{\kappa} \checkmark$, $\hat{\kappa}_0 \checkmark$ and $\hat{\gamma}_0 \supset \hat{\kappa}_0$, then $(\hat{\kappa} \cdot h(w, \mathcal{L}(\mu(\hat{\kappa}_0), \hat{\gamma}_0))@l \cdot \hat{\kappa}_0) \checkmark$.

PROOF. By the definition of $\hat{\kappa} \checkmark$. □

PROOF OF THEOREM 7.2. The proof proceeds by cases on the possible reductions $\langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle \hookrightarrow \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle$. First we show that the transition preserves validity, and then we are able to show that all ASMFx models of the source configuration deterministically transition (in multiple steps) to a model of the target configuration.

Case $\hat{M} = \oplus(\hat{V})$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = \langle \llbracket \oplus \rrbracket (\llbracket \hat{V} \rrbracket \hat{\gamma}), \hat{\gamma}, \hat{\kappa} \rangle$. Since the environment and the continuation do not change, we need to check that the result of the evaluation is valid: this is trivial because we have assumed that $\llbracket \oplus \rrbracket$ only outputs valid values.

If $a \models \hat{s}$, we know that the ASMFx machine is about to execute $\llbracket \oplus(\hat{V}) \rrbracket$. By direct inspection of the code, we see that execution until $v(\hat{M})$ leads us to a state a' where the $\$a$ registers have been loaded with the value $\llbracket \oplus \rrbracket (\llbracket \hat{V} \rrbracket \hat{\gamma})$ and everything else is unchanged. Then $a' \models \hat{s}'$ and $a \xrightarrow{*} a'$.

Case $\hat{M} = \text{if}(\hat{V}, \hat{M}_1, \hat{M}_2)$. In this case we perform different operations depending on whether $\llbracket \hat{V} \rrbracket \hat{\gamma}$ is **true** or **false**.

If it is **true**, then $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = \langle \hat{M}_b \hat{\gamma}, \text{jump}@v(\hat{M}_1) \cdot \hat{\kappa} \rangle$. To prove $\hat{\kappa}' \checkmark$, we show that $\hat{\kappa} \checkmark$ and $\Xi(v(\hat{M}_1)) = j \mu(\hat{\kappa})$ (by the validity of \hat{s} and a direct inspection of $\llbracket \hat{M} \rrbracket$).

It is then trivial to prove that $\Xi(\mu(\hat{M}_1)) = \llbracket \hat{M}_1 \rrbracket$, $v(\hat{M}_1) = \mu(\text{jump}@v(\hat{M}_1) \cdot \hat{\kappa})$, and $\hat{\gamma} \supset \hat{\kappa}'$, therefore $\hat{s}' \checkmark$.

If $a \models \hat{s}$, we can execute the ASMFx code deterministically up to $\llbracket \hat{M}_1 \rrbracket$ (since the if condition evaluated to **true**, the branch-if-zero instruction *bz* will not branch) and take the resulting configuration as a' : we can easily show that $a' \models \hat{s}'$, and $a \xrightarrow{*} a'$ holds trivially.

If instead $\llbracket \hat{V} \rrbracket \hat{\gamma}$ evaluates to **false**, then $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = \langle \hat{M}_b \hat{\gamma}, \hat{\kappa} \rangle$. The continuation and environment do not change therefore, to prove that $\hat{s}' \checkmark$, we only need to show that $\Xi(\mu(\hat{M}_2)) = \llbracket \hat{M}_2 \rrbracket$ and $v(\hat{M}_2) = \mu(\hat{\kappa})$: these are immediate consequences of $\hat{s} \checkmark$ after unfolding the definition of \hat{M} .

If $a \models \hat{s}$, we can execute the ASMFx code deterministically up to $\llbracket \hat{M}_2 \rrbracket$ (since the if condition evaluated to **false**, the branch-if-zero instruction will branch to *l_{else}*) and take the resulting configuration as a' : we can easily show that $a' \models \hat{s}'$, and $a \xrightarrow{*} a'$ holds trivially.

Case $\hat{M} = \text{let } x \leftarrow \hat{M}_1 \text{ in } \hat{M}_2$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = \langle \hat{M}_1, \hat{\gamma}, ((x).\hat{M}_2)@v(\hat{M}_1) \cdot \hat{\kappa} \rangle$. To prove $\hat{\kappa}' \checkmark$, we show that $\hat{\kappa} \checkmark$, $\Xi(v(\hat{M}_1)) = \overline{\$x} \leftarrow \overline{\$a}; \llbracket \hat{M}_2 \rrbracket$ and $v(\hat{M}_2) = \mu(\hat{\kappa})$ (all by the validity of \hat{s} ; the second and last property also require a direct inspection of $\llbracket \hat{M} \rrbracket$).

It is then trivial to prove that $\Xi(\mu(\hat{M}_1)) = \llbracket \hat{M}_1 \rrbracket$, $v(\hat{M}_1) = \mu(((x).\hat{M}_2)@v(\hat{M}_1)) \cdot \hat{\kappa}$, and $\hat{\gamma} \supset \hat{\kappa}'$, therefore $\hat{s}' \checkmark$.

If $a \models \hat{s}$, we can choose $a' = a$: we can easily show that $a \models \hat{s}'$, and $a \xrightarrow{*} a$ holds trivially.

Case $M = \text{unpack } \overline{x_n} \leftarrow \overline{\langle \hat{V}_n \rangle}$ in \hat{M}_0 . In this case $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \langle \hat{M}_0, (x_n \mapsto \overline{\llbracket \hat{V}_n \rrbracket} \hat{y}) \cdot \hat{y}, \hat{k} \rangle$. The continuation has not changed, and it is easy to show that, since $\hat{y} \supset \hat{k}$, then $\hat{y}' \supset \hat{k}$ too (because the value of LEAVE , if any, has not changed with the transition).

It is then trivial to prove that $\Xi(\mu(\hat{M}_0)) = \llbracket \hat{M}_0 \rrbracket$ and $v(\hat{M}_0) = \mu(\hat{k})$ (by direct inspection of $\llbracket \hat{M} \rrbracket$); therefore $\hat{s}' \checkmark$.

If $a \models \hat{s}$, we execute deterministically the code for $\llbracket \hat{M} \rrbracket$ up to and excluding $\llbracket \hat{M}_0 \rrbracket$, obtaining a configuration a' where the registers corresponding to variables $\overline{x_n}$ have been loaded with the subvalues from the evaluation of \hat{V} . We can then easily show that $a' \models \hat{s}'$, and $a \xrightarrow{*} a'$ holds trivially.

Case $\hat{M} = \text{handle } \hat{M}_0 \text{ with } h(\hat{V})$. In this case

$$\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \langle \hat{M}_0, \hat{y} \setminus \text{LEAVE}, h(\llbracket \hat{V} \rrbracket \hat{y}, \mathcal{L}(\mu(\hat{k}), \hat{y})) @ v(\hat{M}_0) \cdot \hat{y} \rangle$$

To prove $\hat{k}' \checkmark$, we show that $\hat{k} \checkmark$, $\Xi(v(\hat{M}_0)) = \text{return}$, $\hat{y} \supset \hat{k}$ (all by the validity of \hat{s} , the second property also requires a direct inspection of $\llbracket \hat{M} \rrbracket$).

It is then trivial to prove that $\Xi(\mu(\hat{M}_0)) = \llbracket \hat{M}_0 \rrbracket$, $v(\hat{M}_0) = \mu(h(\llbracket \hat{V} \rrbracket \hat{y}, \mathcal{L}(\mu(\hat{k}), \hat{y})) @ v(\hat{M}_0) \cdot \hat{k})$, and $\hat{y} \setminus \text{LEAVE} \supset \hat{k}'$, therefore $\hat{s}' \checkmark$.

If $a \models \hat{s}$, we perform as many steps as we need to reach the recursive compilation $\llbracket \hat{M}_0 \rrbracket$. It is easy to show that the resulting configuration a' is a model of \hat{s}' .

Case $\hat{M} = \text{do } \#t(\hat{V})$ and $\hat{k} = \hat{k}^{-t} \cdot h^{+t}(\hat{w}, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{k}_0$. In this case

$$\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \left\langle \begin{array}{l} \hat{M}_{h_{\#t}} \\ (\text{LEAVE} \mapsto \mathcal{R}_{\mathcal{L}}) \cdot (\text{res}(\#t) \mapsto (\hat{k}^{-t}, h^{+t}(\hat{w}) @ \ell, \hat{y})) \\ \cdot (\text{dom}(h_{\#t}) \mapsto \llbracket \hat{V} \rrbracket \hat{y}) \cdot (\text{dom}(h) \mapsto \hat{w}) \\ \text{exit} @ v(\hat{M}_{h_{\#t}}) \cdot \hat{k}_0 \end{array} \right\rangle$$

To prove $\hat{k}' \checkmark$, we show that $\hat{k} \checkmark$ and $\Theta(v(\hat{M}_{h_{\#t}})) = \text{exit}$ (by the validity of \hat{s} and, for the second property, by direct inspection of $\llbracket h_{\#t} \rrbracket$).

It is then trivial to prove that $\Xi(\mu(\hat{M}_{h_{\#t}})) = \llbracket \hat{M}_{h_{\#t}} \rrbracket$, $v(\hat{M}_{h_{\#t}}) = \mu(\text{exit} @ v(\hat{M}_{h_{\#t}}) \cdot \hat{k})$. To prove $\hat{y}' \supset \hat{k}'$ we show:

- $\mathcal{R}_{\mathcal{L}} = (\mu(\hat{k}_0), \hat{y}_0)$ such that $\hat{y}_0 \supset \hat{k}_0$ (by $\hat{s} \checkmark$ and Lemma B.6)
- $(\hat{k}^{-t}, h^{+t}(\hat{w}) @ \ell, \hat{y}) \checkmark$ (by $\hat{s} \checkmark$ and Lemma B.6)
- $\llbracket \hat{V} \rrbracket \hat{y} \checkmark$ (by $\hat{s} \checkmark$ and Lemma B.5)
- $\hat{w} \checkmark$ (by $\hat{s} \checkmark$ and Lemma B.6).

Then it follows that $\hat{s}' \checkmark$.

If $a \models \hat{s}$, we perform as many steps as we need to reach the recursive compilation $\llbracket \hat{M}_{\#t} \rrbracket$. It is easy to show that the resulting configuration a' is a model of \hat{s}' .

Case $\hat{M} = \text{resume } \hat{U}(\hat{V})$. Assume $\hat{k} := \llbracket \hat{U} \rrbracket \hat{y} = (\hat{k}_U, h_U(\hat{w}) @ \ell, \hat{y}_U)$. Then:

$$\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = (\llbracket \hat{V} \rrbracket \hat{y}, \hat{y}_U, \hat{k}_U \cdot h_U(\hat{w}, \mathcal{L}(\mu(\hat{k}), \hat{y}))) @ \ell \cdot \hat{k}$$

To prove $\hat{s}' \checkmark$, we first show that $\llbracket \hat{V} \rrbracket \hat{y} \checkmark$ (by Lemma B.5). Then we note $\hat{y} \supset \hat{k}$, $\hat{k} \checkmark$ by the validity hypothesis on the source configuration. Furthermore, in order for $\llbracket \hat{U} \rrbracket \hat{y}$ to evaluate to the broom \hat{k} , we must have $\hat{U} = x_U$ for some variable x_U : therefore $\hat{k} = \hat{y}(x_U)$, from which we prove $\hat{k} \checkmark$. Using these results, we apply Lemma B.8 to prove $\hat{y}_U \supset \hat{k}'$ and $\hat{k}' \checkmark$.

If $a \models \hat{s}$, we execute deterministically the entire compiled code for $\text{resume } \hat{U}(\hat{V})$: it is easy to show that the resulting ASAFX state a' is such that $a' \models \hat{s}'$. In particular, we can verify that just

before executing the resume instruction, the register file holds the following values:

$$\begin{aligned} \$lr &\mapsto \mathcal{L}(\mu(\hat{\kappa}), \lfloor \hat{y} \rfloor) \\ \$kr &\mapsto \mathcal{K}(\ell_k, \lfloor \hat{y}_U \rfloor, \lfloor \hat{\kappa}_U \rfloor) \\ \$hr &\mapsto \mathcal{H}(\lfloor h \rfloor, (\$s = \lfloor \hat{w} \rfloor)) \end{aligned}$$

where $\ell_k = \mu(\hat{\kappa}_U \cdot h(\hat{w})@ \ell)$. By executing the next instruction, we therefore move into a configuration a' where the register $\$pc$ holds the address ℓ_k , the register file has been updated to match the environment \hat{y}_U , and the effect context has been extended with $\lfloor \hat{\kappa}_U \rfloor \cdot \mathcal{H}(\lfloor h \rfloor, (\$s = \lfloor \hat{w} \rfloor), \mathcal{L}(\mu(\hat{y}_0), \lfloor \hat{y} \rfloor))$. We can thus easily show that $a' \models \hat{s}'$.

Case $\hat{M} = \hat{U}$ with $h(\hat{W})$. Assume $\hat{k} := \llbracket \hat{U} \rrbracket \hat{y} = (\hat{\kappa}_U, _@ \ell, \hat{y}_U)$. Then:

$$\hat{s}' = \langle \hat{M}', \hat{y}', \hat{\kappa}' \rangle = \langle (\hat{\kappa}_U, h(\llbracket \hat{W} \rrbracket \hat{y}), \hat{y}_U), \hat{y}, \hat{\kappa} \rangle$$

To prove $\hat{s}' \checkmark$, we unfold the definition and use Lemma B.8.

If $a \models \hat{s}$, we execute deterministically the entire code for $\llbracket \hat{M} \rrbracket$. It is easy to show that the resulting configuration a' is a model of \hat{s}' .

Case $\hat{M} = \text{newbroom}@ \ell(f)$ with $h(\hat{V})$. We have $\hat{s}' = \langle (\hat{\kappa}_f, h(\llbracket \hat{V} \rrbracket \hat{y})@(\ell + 2), \lfloor \rfloor), \hat{y}', \hat{\kappa}' \rangle$ where

$$\hat{\kappa}_f = \text{leave}@ \ell \cdot \text{jump}@(\ell + 1) \cdot ((x).f_{\text{return}})@ \mu(f_{\text{return}}) \cdot \text{exit}@v(f_{\text{return}})$$

To prove $\hat{s}' \checkmark$, we unfold the definition and in particular we verify that $(\hat{\kappa}_f \cdot h(\llbracket \hat{V} \rrbracket \hat{y})@(\ell + 2)) \checkmark$ by repeated applications of the rules defining augmented source validity (this requires us to verify that the locations annotating the continuation contain the expected instructions).

If $a \models \hat{s}$, we execute deterministically the entire code for $\llbracket \hat{M} \rrbracket$. It is easy to show that the resulting configuration a' is a model of \hat{s}' .

Case $\hat{M} = \text{return } \hat{V}$. In this case $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{\kappa}' \rangle = (\llbracket \hat{V} \rrbracket \hat{y}, \hat{y}, \hat{\kappa})$. Knowing by hypothesis that $\hat{s} \checkmark$, to prove $\hat{s}' \checkmark$ we only need to show, by Lemma B.5, that $\llbracket \hat{V} \rrbracket \hat{y} \checkmark$.

If $a \models \hat{s}$, we perform the code $\llbracket \text{return } \hat{V} \rrbracket = \llbracket \hat{V} \rrbracket$, which loads an initial segment $\overline{\$a}$ of the argument registers with the evaluation of \hat{V} . We can then show that the resulting state a' is a model of \hat{s}' .

Case $\hat{M} = \hat{v}$ and $\hat{\kappa} = \text{jump}@ \ell \cdot \hat{\kappa}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{\kappa}' \rangle = (\hat{v}, \hat{y}, \hat{\kappa}_0)$. To prove $\hat{s}' \checkmark$ knowing $\hat{s} \checkmark$, it is enough to show that $\hat{\kappa}_0 \checkmark$ by inversion on $\hat{\kappa} \checkmark$.

If $a \models \hat{s}$, we know the next instruction is $j \mu(\hat{\kappa}_0)$. By executing that instruction, we move into a configuration a' that is the same as a except for the register $\$pc$ holding the address $\mu(\hat{\kappa}_0)$: then we can easily prove $a' \models \hat{s}'$.

Case $\hat{M} = \hat{v}$ and $\hat{\kappa} = \text{exit}@ \ell \cdot \hat{\kappa}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{\kappa}' \rangle = (\hat{v}, \hat{y}(\text{LEAVE}).env, \hat{\kappa}_0)$. Knowing by hypothesis that $\hat{s} \checkmark$, we show that $\hat{\kappa}_0 \checkmark$ by inversion on $\hat{\kappa} \checkmark$.

To show that $\hat{y}(\text{LEAVE}).env \supset \hat{\kappa}_0$, we note $\hat{y} \supset \hat{\kappa}$ which, by its definition, implies the thesis.

Then, knowing $\hat{v} \checkmark$, we have immediately $\hat{s}' \checkmark$.

If $a \models \hat{s}$, we know the next instruction is an exit and the current value of register $\$lr$ is $\mathcal{L}(\mu(\hat{y}_0), \lfloor \hat{y}(\text{LEAVE}).env \rfloor)$. Then, by executing that instruction, we move into a configuration a' where the register $\$pc$ holds the address $\mu(\hat{y}_0)$, and the register file has been updated to match the environment $\hat{y}(\text{LEAVE}).env$. We can thus easily show that $a' \models \hat{s}'$.

Case $\hat{M} = \hat{v}$ and $\hat{k} = \text{leave}@l \cdot \hat{k}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{k}' \rangle = (\hat{v}, (\text{LEAVE} \mapsto \mathcal{L}(\ell+2, [])) \cdot \hat{\gamma}, \hat{k}_0)$. Knowing by hypothesis that $\hat{s} \checkmark$, we show that $\hat{k}_0 \checkmark$ by inversion on $\hat{k} \checkmark$.

To show that $(\text{LEAVE} \mapsto \mathcal{L}(\ell+2, [])) \cdot \hat{\gamma} \supset \hat{k}_0$, we note, by inversion on validity, that $\hat{k}_0 = \text{jump}@(\ell+1) \cdot ((x). \hat{N}) @ \mu(\hat{N}) \cdot \text{exit}@v(\hat{N}) \cdot h(\hat{w}) @ (\ell+2) \cdot \hat{k}_1$. We can thus see that $\text{lcont}(\hat{k}_0) = \ell+2$, which is all we need.

Then, knowing $\hat{v} \checkmark$, we have immediately $\hat{s}' \checkmark$.

If $a \vdash \hat{s}$, we know the next instruction is `loadl $\ell+2$` . Then, by executing that instruction, we move into a configuration a' where the register `$lr` holds the record $\mathcal{L}(\ell+2, [])$ and all the other registers match the environment $\hat{\gamma}$. We can thus easily show that $a' \models \hat{s}'$.

Case $\hat{M} = \hat{v}$ and $\hat{k} = ((x). \hat{M}_0) @ l \cdot \hat{k}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{k}' \rangle = (\hat{M}_0, (x \mapsto \hat{v}) \cdot \hat{\gamma}, \hat{k}_0)$.

To prove $\hat{s}' \checkmark$, knowing $\hat{s} \checkmark$, we trivially show $\hat{k}_0 \checkmark$ (by inversion on $\hat{k} \checkmark$) and then proceed to show that $\hat{\gamma} \supset \hat{k}$ implies $((x \mapsto \hat{v}) \cdot \hat{\gamma}) \supset \hat{k}_0$ (LEAVE is defined in one environment if and only if it is defined in the other, and if it is its value is the same; furthermore, the conditions it has to satisfy in one continuation are exactly the same as those needed in the other). Therefore, $\hat{s}' \checkmark$ holds.

If $a \models \hat{s}$, we know that the register `$pc` points to code that will deterministically move data from registers `$a` (holding the value $\llbracket \hat{v} \rrbracket$) into registers `$x`, then proceed with the code for $\llbracket \hat{M}_0 \rrbracket$. We can then easily show that by executing the first few instructions until we reach code address $\mu(\hat{M}_0)$, we obtain a configuration a' where the registers `$x` have been loaded with $\llbracket \hat{v} \rrbracket$. Then, if we note that $\llbracket \hat{k} \rrbracket = \llbracket \hat{k}' \rrbracket$ (since pure frames are ignored when compiling a source continuation), we prove $a' \models \hat{s}'$.

Case $\hat{M} = \hat{v}$ and $\hat{k} = h(\hat{w}, \mathcal{R}_\mathcal{L}) @ l \cdot \hat{k}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{k}' \rangle = (h_{\text{return}}, (\text{LEAVE} \mapsto \mathcal{R}_\mathcal{L}) \cdot (\text{dom}(h_{\text{return}}) \mapsto \hat{v}) \cdot (\text{dom}(h) \mapsto \hat{w}), \hat{k}_0)$.

Knowing $\hat{s} \checkmark$, to prove $\hat{s}' \checkmark$ we start by proving that $(\text{exit}@v(h_{\text{return}}) \cdot \hat{k}_0) \checkmark$ which requires us to show that $\Xi(v(h_{\text{return}})) = \text{exit}$ (trivial by inspection of $\llbracket h_{\text{return}} \rrbracket$).

To prove $\hat{\gamma}' \supset \hat{k}'$, it is sufficient to show that $\mathcal{R}_\mathcal{L} = \mathcal{L}(\mu(\hat{k}_0), \hat{\gamma}_0)$ such that $\hat{\gamma}_0 \supset \hat{k}_0$. This is proved by inversion on $\hat{k} \checkmark$.

Noting that $\Xi(\mu(h_{\text{return}})) = \llbracket h_{\text{return}} \rrbracket$ (by the definition of Ξ), we have $\hat{s}' \checkmark$.

If $a \models \hat{s}$, we know that the next instruction is a `return`, that the initial segment of argument registers `$a` holds $\llbracket \hat{v} \rrbracket$, and that the effect context is $\mathcal{H}(\llbracket h \rrbracket, (\overline{\$s} = \llbracket \hat{w} \rrbracket), \llbracket \mathcal{R}_\mathcal{L} \rrbracket) \cdot \llbracket \hat{k}_0 \rrbracket$. Then by executing `return` we load the encoding of \hat{w} into registers `$s`, the encoding of $\mathcal{R}_\mathcal{L}$ into register `$lr`, and we jump to the address μh_{return} . We execute the first few instructions of $\llbracket h_{\text{return}} \rrbracket$ until we reach \hat{M}_{return} — these instructions move data between registers and ensure that the register file corresponds to $\hat{\gamma}'$. It is thus easy to show that the configuration a' we have reached is such that $a' \models \hat{s}'$. \square