

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 provide a uniform abstraction for defining computational effects such as backtracking, concurrency, dynamic binding, exceptions, I/O, and state. They are typically treated as a high-level feature, implemented by compilation to existing low-level mechanisms. In contrast to this we introduce ASmFX, an abstract assembly language with direct support for effect handlers, with the aim of characterising their relationship to machine architecture. We demonstrate the expressiveness of ASmFX by compiling a functional core calculus supporting deep, shallow, and parameterised handlers with multi-shot resumptions into ASmFX. We establish soundness of the compiler by factoring it through an intermediate calculus that annotates computations in the source with *epilogues* – fragments of target code that must run when the computation completes, but which have no direct source-level counterpart.

1 Introduction

Effect handlers are a powerful and elegant programming language abstraction for defining and controlling 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, 9, 10]. Effect handlers are particularly helpful in achieving clean separation of concerns and modularity and have lately been incorporated into many research and mainstream languages. Implementations of effect handlers in high-level languages typically convert them to existing control features like closures, continuations, or prompts [16, 20, 28, 30]; while more low-level approaches often manipulate an existing stack [2, 13, 25].

Our approach in this work is different: instead of translating effect handlers away into pre-existing constructs, we preserve them as a primary control feature from source language all the way down to assembly, giving an instruction set architecture for an abstract register machine where effect handlers are primitive. With this we aim to characterise effect handlers at the lowest level: what they require of machine architecture and what they provide for code that runs on it.

We propose this assembly language as a foundation to explore the correspondence between source-level properties of effect handlers and machine-level features, notably those for memory protection and control-flow integrity. Contemporary processor architectures have multiple such mechanisms: shadow stacks, branch tracking, pointer authentication, protection keys, and permission overlays all try to efficiently manage concepts of non-interference and controlled isolation [8, 22]. Our vision is that characterising the architectural requirements of effect handlers will enable compilation that maps source-level separation of concerns into these kind of machine-enforced protections.

The paper structure follows the passage from high-level language feature to assembly-level architecture, culminating in a compilation soundness result relating the two. We begin by setting our context for this in Section 2 with representative examples of effect handler use: their source-level presentation, corresponding assembly-level implementation, and what this requires of the abstract machine architecture.

In Section 3 we define SFX, a core calculus with effect handlers, to act as our high-level effectful source language. Its design goal is to provide a sufficiently rich effect-handling infrastructure,

Authors' Contact Information: [Brian Campbell](mailto:Brian.Campbell@ed.ac.uk), University of Edinburgh, United Kingdom, Brian.Campbell@ed.ac.uk; [Sam Lindley](mailto:Sam.Lindley@ed.ac.uk), University of Edinburgh, United Kingdom, Sam.Lindley@ed.ac.uk; [Wilmer Ricciotti](mailto:Wilmer.Ricciotti@ed.ac.uk), University of Edinburgh, United Kingdom, research@wilmer-ricciotti.net; [Ian Stark](mailto:Ian.Stark@ed.ac.uk), University of Edinburgh, United Kingdom, Ian.Stark@ed.ac.uk.

50 where resumptions may be single-shot, affine, or multi-shot; and where the traditional setting of
51 *deep* handlers can be extended by *parameterised* handlers, or replaceable handlers (such as the *sheep*
52 handlers of WASMFX [25]). This source language is capable of expressing a large range of effects,
53 including but not limited to state, exceptions, nondeterminism, and dynamic scoping. Following
54 prior work [17, 18], we express the semantics of SFX via a CEK-like abstract machine (Section 3.3).

55 We then introduce AsmFX (Section 4), an abstract register-based instruction set architecture
56 (ISA) with primitives to define, install, and invoke effect handlers. The goal of this abstract ISA is to
57 support all forms of effect handler expressible in SFX (and possibly more). Concretely, AsmFX adds
58 to a traditional register-based architecture seven new instructions, four of which allow the extended
59 control-flow behaviour of effectful programs; the remaining three implement load operations
60 needed for creating special data structures used by AsmFX for representing continuations, handlers,
61 and external code pointers as first-class values. As a reasoning tool, AsmFX does not prescribe
62 particular memory layouts or restrict the amount of data that can be stored in registers. The
63 semantics of AsmFX is expressed by means of an abstract machine, which uses, along with main
64 memory, an independent stack-like data structure describing the current effect handler state.

65 Having given the source and target language, we next present a compiler from SFX to AsmFX
66 (Section 5). AsmFX provides an interesting computational model in its own right and can be used
67 to study effectful programs written directly in a suitable ISA or as the intermediate language
68 of a compiler for high-level languages. The compiler yields both AsmFX assembly code and an
69 augmented version of the original source code annotated to denote which AsmFX code addresses
70 implement each program expression. The augmented (ASFX) code is crucial for stating and proving
71 our main correctness theorems.

72 The culmination of our investigation of the expressive power of AsmFX, and a central contribution
73 of the paper, is the proof of correctness of the SFX to AsmFX compiler (Section 6). The proof
74 — constructing a transition simulation between SFX and AsmFX— is challenging due to code
75 rearrangements between source and target languages. At any given execution step this simulation
76 depends on execution history and the AsmFX memory locations of each compiled program fragment.
77 To manage this we split the proof into two simulations, with an augmented source language ASFX
78 as midpoint. We speculate that this technique can be useful in soundness proofs for other compilers
79 as well. ASFX annotates SFX with details of the relationship between source and compiled code. It
80 also extends the notion of continuation with *epilogues*, fragments of machine code that have no
81 direct representation in the source program but are needed for correct AsmFX control flow.

82 We close in Section 7 with a discussion: how this work can be taken forward, and details of
83 how it relates to other effect handler research. In anticipation of that, though, we address here one
84 immediate potential confusion: ASmFX and WASMFX.

85
86
87 *Comparison with WASMFX.* Though they have similar names, ASmFX and WASMFX [25] are quite
88 different languages designed for quite different purposes. WASMFX builds on WebAssembly (Wasm
89 for short), which is not in fact an assembly language but, like JVM or .NET, a typed platform-
90 independent bytecode for a stack-based virtual machine. To this WASMFX adds a specific variety of
91 single-shot effect handlers, building on the existing structured control flow mechanism. In contrast,
92 ASmFX is an abstract assembly language that provides a template for extending register machines
93 with effect handlers as primitive. It starts from a raw machine language with no control flow and
94 adds low-level features sufficient to implement a range of different kinds of effect handler: *deep*,
95 *shallow*, *sheep*; *single-shot* or *multishot*. In particular it includes no built-in instructions for calling
96 and returning from functions—though, as it turns out, handlers can be used to implement these.

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 our high-level SFX language, highlighting a few different uses of handlers and what we would expect their compiled version to look like.

A handler is defined by a statement **handler** h $\{\text{return } x \rightarrow N \mid \#t(y) \rightarrow (r).N' \mid \dots\}$, where h is its identifier. To run a computation M in the context of handler h we use the syntax **handle** M **with** h . We say that M is the *client code* of that particular handle statement. The code within the definition of h is *handler code*, and we refer to any other code outside the handle statement as *external code*.

The client code M may perform the effectful operations $\text{do } \#t(V)$ defined by h , where V is an input argument. What happens when this operation is performed is expressed by the corresponding operation clause $\#t(y)$ in the definition of h , where y is the operation's formal argument; this operation clause also receives as input a resumption r which can be used to resume the client that invoked $\#t$; alternatively, the operation can **return** a value to the external code, aborting the client. When the client M successfully evaluates to a value, this value is first handed to the **return** clause of h , which can post-process it before returning to the external code.

An effectful assembly language should provide instructions to transition between external, client, and handler code, 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 has a *budget* parameter that holds an integer representing the available resources.

```

125 handler track_budget(budget) {
126     return x → return ⟨x, budget⟩
127     | #charge(n) → (r).
128     if (budget ≥ n) then
129         let r' ← r with track_budget(budget - n) in
130         resume r'(<>)
131     else return ⟨0, -1⟩
132 }

```

```

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

```

The program uses the **handle** keyword to execute client code within the context of the handler *track_budget* (initialised with a budget of 100), 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 resumes the client code while reinstalling itself with the updated budget — this operation is performed by replacing the *budget* parameter for the tracker handler in the resumption r using the expression **with** *track_budget*(*budget* - n), and then resuming the updated resumption r' ; 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.

148 At an assembly level, the code above is not implemented by structural nesting, but by flat snippets
 149 of code that reference each other using labels. For this example, our new assembly instructions
 150 will need to record the handler installed in the main code so that client code can search for the
 151 implementation of the *#charge* effect, and also store the handler's state, *budget*.

```

152 track_budget.ret:
153     ; $a0 already contains the result x
154     $a1 := mem[budget]
155     exit
156
157 track_budget.charge:
158     ; $s0 contains the current budget
159     if ($s0 < $a0) branch abort
160     $s0 := $s0 - $a0
161     loadh (return=track_budget.ret, #charge=track_budget.charge), state:($s0)
162     resume
163 abort:
164     $a0 := 0
165     $a1 := -1
166     exit
167
168 budget_client:
169     $a0 := 10
170     do #charge
171     ; [...] code to call p()
172     $t0 := $a0 ; save result of p() to scratch register
173     ; [...] code to call q()
174     $a0 := $t0 + $a0 ; add result of q() to previously saved result
175     return
176
177 main:
178     loadk budget_client
179     $s0 := 100
180     loadh (return=track_budget.ret, #charge=track_budget.charge), state:($s0)
181     loadl client_done
182     resume
183 client_done:
184     ; [...] code to print the balance returned in $a1
  
```

180 Execution starts from the `main` label near the bottom. It runs the `budget_client` under the tracker
 181 handler, whose two clauses are specified by snippets `track_budget.ret` and `track_budget.charge`.
 182 To achieve this, it loads into special registers `$kr` and `$hr` a continuation for the client using an
 183 instruction `loadk budget_client`, and the tracker handler with the instruction `loadh`: its operands
 184 specify the code snippets for the return and charge clauses, and the *handler state* — a list of registers
 185 whose value is saved now and restored when the handler is invoked. Here the state consists of the
 186 register `$s0`, holding a budget initialised to 100. Finally, before invoking the client, we need to save
 187 the external context that will need to be restored once the client terminates: this information is held
 188 in a *leave record* contained in a special register `$lr`, which is loaded by the instruction `loadl`. A
 189 leave record contains the address of the external code that will be executed when the client is done,
 190 and optionally the values of some registers: in this case, we only load the address `client_done`.

191 After setting up continuation, handler, and leave record, we use the instruction `resume`, which
 192 will record the new handler and run the continuation. In particular, in our case, it will run the client
 193 under the budget tracker handler.

194 The `budget_client` uses instructions `do` and `return` respectively to perform an effect and to
 195 return to the external code. Registers `$a0`, `$a1`, ... are used to exchange arguments/results between
 196

197 function and operation calls. Since we are running under the tracker handler, `do #charge` jumps
 198 to `track_budget.charge`; when running an operation clause, the machine automatically loads
 199 the current handler, leave record, and the continuation to the client into the corresponding special
 200 registers: this allows the client to be resumed by a simple `resume` instruction, without any additional
 201 `loadh`, `loadl`, or `loadk` instructions. In this example, before resuming the client, we need to update
 202 the budget: this is achieved by updating the state register `$s0` and reloading the handler by means
 203 of `loadh`. To abort the computation without resuming the client, when the budget is not enough,
 204 we use the instruction `exit`, which restores the external context specified by the leave record.

205 The `track_budget.ret` clause leaves the client result in `$a0` untouched, and copies the current
 206 value of budget into `$a1`, then yields control to the external code again by means of `exit`.
 207

208 2.2 Handler Replacement

209 In some cases, when implementing an effect, we may want to resume a client under a different
 210 handler from the one that was originally invoked. Consider the following example consisting of
 211 two functions `prod` and `cons`: the first uses the `#send` operation to send two integer values over
 212 a shared channel, while the second invokes a `#recv` operation to get two integer values from the
 213 same channel and return their sum.
 214

```

215
216 handler ph(cr) {                                     // Producer sends two values then stops
217   return ⟨⟩ → handle prod(⟨⟩) with ph(cr)          fun prod(_) {
218   | #send(x) → (pr).                               do #send(1); do #send(2); return ⟨⟩
219     let r ← (cr with ch(pr)) in resume r(x)        }
220 }
221
222 handler ch(pr) {                                     // Consumer gets two values then returns their sum
223   return result → return result                   fun cons(_) {
224   | #recv(x) → (cr).                               let x ← do #recv(⟨⟩) in
225     let r ← (pr with ph(cr)) in resume r(⟨⟩)      let y ← do #recv(⟨⟩) in return x + y
226 }
227
228 // An initial version of the consumer handler      // Program starts the consumer under an initial handler
229 handler ih(_) {                                     fun main(_) {
230   return x → return x                             handle cons(⟨⟩) with ih(⟨⟩)
231   | #recv(x) → (cr).handle prod(⟨⟩) with ph(cr) }
232 }
233
```

232 The two operations are implemented in two different handlers. The producer handler `ph` param-
 233 eterised over a consumer resumption `cr` handles `#send` by resuming `cr` under the consumer handler
 234 `ch`, which is passed as a parameter the producer resumption `pr`. Symmetrically, the consumer
 235 handler `ch`, parameterised over a producer resumption `pr`, handles `#recv` by resuming `pr` under
 236 the producer handler `ph`, which is passed as a parameter the consumer resumption `cr`. In this way
 237 the handlers interleave the execution of the two functions, using handler replacement to switch
 238 between the two complementary interfaces.

239 The main program is started by handling the consumer `cons` under an appropriate handler: ideally,
 240 this would be `ch`, but the producer has not started yet, so we do not have a producer resumption to
 241 pass as a parameter yet; instead, we use an initial consumer handler `ih` which handles the `#recv`
 242 operation by starting `prod` under handler `ph` rather than calling a producer resumption. Note that
 243 the return clause of `ph` will restart `prod` if it terminates, guaranteeing that a `#recv` effect will always
 244 be matched by a `#send`.
 245

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:

```

246 ph.send:
247   ; handler parameter cr (consumer resumption) in ($s0, $s1)
248   ; value to be sent in $a0, producer resumption pr in ($kr, $hr)
249   $t0 := $kr
250   $t1 := $hr
251   $kr := $s0 ; continuation of cr
252   $s0 := $t0 ; continuation of pr
253   $s1 := $t1 ; delimiting handler of pr
254   ; replace handler by loading $hr with ch(pr)
255   loadh (return=ch.ret,#recv=ch.recv),state:($s0,$s1)
256   ; resume at $kr under the handler in $hr, i.e. (cr with ch(pr))
257   resume

```

Note that the continuations in the assembly are not merely an address to restart execution from, but must also contain all the state necessary to restart them from the handler where they are captured, including the live registers and any intermediate handlers that have been installed.

2.3 Multishot 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:

```

276
277 handler tautology(⟨⟩) {
278   return x → return x
279   | #flip(⟨⟩) → (r).
280   let x ← resume r(false) in
281   let y ← resume r(true) in
282   return x && y
283 }
284
285 fun prop(⟨⟩) {
286   let a ← do #flip(⟨⟩) in
287   let b ← do #flip(⟨⟩) in
288   return (a && b) || not a || not b
289 }
290
291 handle prop(⟨⟩) with tautology(⟨⟩)

```

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.

This example differs from the previous ones in two ways: first, the resumption is invoked more than once and we must be able to restore the resumption's state both times; 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 directly to the external code, but to the remainder of the operation clause that invoked it.

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:

```

306 tautology.flip:
307     $a0 := 0    ; false
308     loadl l1, save:($lr)
309     resume     ; will jump to l1 when the resumption ends
310 l1:
311     ; $lr restored to original value
312     $t0 := $a0 ; save result
313     $a0 := 1    ; true
314     loadl l2, save:($lr,$t0)
315     resume     ; will jump to l2 when the resumption ends
316 l2:
317     ; $lr and $t0 restored to the original value
318     $a0 := $t0 && $a0
319     exit     ; return r(false) && r(true)

```

3 A Source Calculus for Effect Handlers

In this section we introduce SFX, a first-order fine-grain call-by-value functional language extended with support for a rather general form of effect handlers. (The restriction to first-order fine-grain call-by-value is not a limitation as we can use standard techniques such as let-insertion and closure conversion to elaborate a higher-order call-by-value surface language into SFX.)

3.1 Syntax

The syntax of SFX is given by the following grammar, in which the novel features are highlighted.

Value types	$A, B ::= \tau \mid \langle \bar{A} \rangle \mid A \xrightarrow{C} D$	Effect types	$E ::= \overline{\{\#t : A \rightarrow B\}}$
Comp. types	$C, D ::= \frac{A!E}{A}$	Top-level types	$G ::= A \Rightarrow C \mid (A)C \Rightarrow D$
Typing contexts	$\Gamma ::= \bar{x} : A$		
Values	$U, V, W ::= x \mid c \mid \langle \bar{V} \rangle$		
Computations	$M, N ::= \mathbf{return} V \mid \oplus(V) \mid f(V) \mid \mathbf{if}(V, M, N) \mid \mathbf{let} x \leftarrow M \mathbf{in} N \mid \mathbf{unpack} \bar{x} \leftarrow V \mathbf{in} M \mid \mathbf{do} \#t(V) \mid \mathbf{handle} M \mathbf{with} h(V) \mid \mathbf{resume} U(V) \mid \mathbf{newbroom}(f) \mathbf{with} h(V) \mid U \mathbf{with} h(V)$		
Runtime values	$u, v, w ::= c \mid \langle \bar{v} \rangle \mid (\kappa, h(v))$		
Environments	$\gamma ::= \bar{x} \mapsto \bar{v}$		
Continuations	$\kappa, \kappa' ::= \mathbf{0} \mid (x).M[\gamma] \cdot \kappa \mid h(v) \cdot \kappa$		
Functions	$F ::= \mathbf{fun} f(x) \{M\}$		
Handlers	$H ::= \mathbf{handler} h(x) \{\mathbf{return} y \rightarrow M \mid \overline{\#t(z) \rightarrow (r).N}\}$		

Specifications	$\Sigma ::= \overline{F}; \overline{H}$
Programs	$P ::= \Sigma; M$

We use vector notation such as \overline{x}_n to denote the sequence $x_0 \dots x_{n-1}$. We often omit the n subscript when the length 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 (\overline{A}) , and *resumption* types $A \xrightarrow{C} D$ consisting of a pair of a continuation, which takes a value of type A performs a computation of type C , and a handler, which transforms a computation of type C into one of type D . We discuss the motivation for the design of resumptions in SFX later in this section. A computation type $A!E$ combines 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 the input type and B is the output type. If E is empty then 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. Functions of type $A \Rightarrow C$ take a value as input and return a computation. With resumptions, handlers, and functions limited to a single argument parameter we use tuples to express n -ary versions. Typing contexts Γ assign value types to local variables.

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. Programs can only refer explicitly to value terms: in particular, resumption values cannot appear in a program, except as variables of a resumption type. At runtime, value terms evaluate to runtime values.

Computation terms M, N include standard forms for returning a value (**return**), primitive operations (\oplus), application of top-level functions, conditionals, sequencing (**let**), tuple decomposition (**unpack**). The form **do** $\#t(V)$ performs an effect $\#t$ with payload V , **handle** M **with** $h(V)$ handles computation M with globally defined handler h and argument V , and **resume** $U(V)$ resumes resumption U with argument V .

We call our novel kind of resumptions *broom resumptions*, or just *brooms*, because like the tool used for sweeping they consist of a head κ and a handle $h(v)$. Our motivation for designing broom resumptions is that they allow us to easily support all of deep, shallow, and sheep handlers. A broom resumption has a head (the continuation) and a handle (its effect handler), and either component may be replaced. Resuming a broom invokes the continuation and installs the handler: **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 using the **newbroom**(f) **with** $h(V)$ construct: 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. We can simulate shallow handlers by replacing the handle with a trivial handler that has no operations and the identity return clause.

Unlike value terms, runtime values are closed. A runtime value can be a constant c , a tuple (\overline{v}) or a broom resumption consisting of a continuation paired with a distinguished handler of the form $(\kappa, h(v))$. Environments γ are substitutions mapping variables to runtime values. Continuations κ are sequences of continuation frames: the base case $\mathbf{0}$ expresses termination. Pure continuation frames $(x).M[\gamma]$ consist of a closure $M[\gamma]$ abstracted over a variable x : given a runtime 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 runtime 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 .

Top-level functions are standard. Top-level handlers **handler** $h(x) \{\text{return } y \rightarrow M \mid \#t(z) \rightarrow (r).N\}$ take a parameter x and include a return clause along with with a collection of operation clauses. A return clause **return** $y \rightarrow M$ binds the final return value of the handled computation to y in M . An operation clause $\#t(z) \rightarrow (r).N$ handles the operation t binding the operation argument to z and the resumption to r in N . 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 resumption argument r of the operation clause. Programs consist of any number of mutually defined functions and handlers plus a main computation that may use them.

3.2 Typing

Typing judgements for the source calculus include four main forms for values $\Gamma \vdash_{\Sigma} V : A$, computations $\Gamma \vdash_{\Sigma} M : C$, continuations $\vdash_{\Sigma} \kappa : A \rightarrow C$, and handlers/functions $\Sigma \vdash h : (A)C \Rightarrow D$. We also provide two minor judgements for environments and for the program. The specification Σ is extended with a top-level type G for each handler and function, and the judgements that depend on it are annotated with Σ on the turnstile (\vdash_{Σ}). Figure 1 presents the typing rules.

This paper deals with control flow and the results we present do not depend on the type system, which is almost entirely standard for an effectful functional language. For this reason we will not discuss the details of typing, except for brooms, which are a novel concept that the type system can help us understand.

In most other proposals, a language supports either a single form of effect handler (e.g. sheep handlers in WASMFX [25]) or provides different syntactic constructs for different kinds of handlers (such as deep and shallow handlers in [18]): in this latter case, the two handlers differ in the type and semantics of the resumption that is generated and provided to the code performing an effect: a deep resumption consists of a continuation delimited by the active handler *including the handler itself*; a shallow resumption provides the same continuation, but omits the delimiting handler; a sheep resumption is the same as a shallow resumption, but the language forces the user to provide a new enclosing handler when invoking the resumption.

It is clear that, in spite of the idiosyncratic syntax of those systems, we do not really have different kinds of handlers, but rather different kinds of resumptions. In our source language, we want to provide the simplest syntax to support both deep and sheep handlers; as noted, what we really need to achieve this goal is to design a notion of resumption that subsumes both deep and shallow resumptions. The obvious way to do this would be to encode deep resumptions by pairing a shallow resumption of type $A \rightarrow C$ with an enclosing handler of type $C \Rightarrow D$. However, to construct such a pair, we would need handlers to be first class program expressions (a change that we regard as both profound and unnecessary, since users do not normally expect to manipulate handlers as program values). A simpler solution is to deal with resumptions not as composite values of type $\langle A \rightarrow C, C \Rightarrow D \rangle$, but as abstract values of a special type which we call *brooms*. Concretely, a broom $(\kappa, h(v))$ of type $A \xrightarrow{C} D$ still consists of a shallow resumption 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

<p>442 $\frac{\text{T-VAR}}{x : A \in \Gamma}$</p> <p>443 $\frac{\text{T-VAL}}{c \in \text{Val}(\tau)}$</p> <p>444 $\frac{\text{T-TUPLE}}{\Gamma \vdash_{\Sigma} \bar{V} : \langle \bar{A} \rangle}$</p> <p>445 $\frac{\text{T-PRIM}}{\Gamma \vdash_{\Sigma} V : A \quad \text{Type}(\oplus) = A \rightarrow B}$</p>	<p>446 $\frac{\text{T-BROOM}}{\Gamma \vdash_{\Sigma} \kappa : A \rightarrow C \quad \Gamma \vdash_{\Sigma} v : B \quad \Sigma \vdash h : (B)C \Rightarrow D}$</p> <p>447 $\frac{\text{T-RET}}{\Gamma \vdash_{\Sigma} V : A}$</p> <p>448 $\frac{\text{T-APP}}{\Gamma \vdash V : A \quad \Sigma \vdash f : A \Rightarrow C}$</p>	<p>449 $\frac{\text{T-IF}}{\Gamma \vdash_{\Sigma} V : \text{Bool} \quad \Gamma \vdash_{\Sigma} M : C \quad \Gamma \vdash_{\Sigma} N : C}$</p> <p>450 $\frac{\text{T-LET}}{\Gamma \vdash_{\Sigma} M : A!E \quad \Gamma, x : A \vdash_{\Sigma} N : B!E}$</p> <p>451 $\frac{\text{T-UNPACK}}{\Gamma \vdash_{\Sigma} V : \langle \bar{A} \rangle \quad \Gamma, x : \bar{A} \vdash_{\Sigma} M : C}$</p> <p>452 $\frac{\text{T-DO}}{\Gamma \vdash_{\Sigma} V : A \quad E(\#t) = A \rightarrow B}$</p> <p>453 $\frac{\text{T-HANDLE}}{\Gamma \vdash_{\Sigma} M : C \quad \Gamma \vdash_{\Sigma} V : A \quad \Sigma \vdash h : (A)C \Rightarrow D}$</p>	<p>454 $\frac{\text{T-WITH}}{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D'}$</p> <p>455 $\frac{\text{T-RESUME}}{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : A}$</p> <p>456 $\frac{\text{T-NEWBROOM}}{\Sigma \vdash f : A \Rightarrow C \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D}$</p> <p>457 $\frac{\text{T-ENV}}{\text{dom}(\gamma) = \text{dom}(\Gamma) \quad (\vdash_{\Sigma} \gamma(x) : \Gamma(x))_{x \in \text{dom}(\gamma)}}$</p>
<p>458 $\frac{\text{T-KID}}{\vdash_{\Sigma} \mathbf{0} : A \rightarrow A!E}$</p> <p>459 $\frac{\text{T-KLET}}{\vdash_{\Sigma} \gamma : \Gamma \quad \Gamma, x : A \vdash_{\Sigma} N : B!E}$</p> <p>460 $\frac{\text{T-KHANDLE}}{\Sigma \vdash h : (A)B!E \Rightarrow B'!E' \quad \vdash_{\Sigma} v : A}$</p>	<p>461 $\frac{\text{T-RESUME}}{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : A}$</p> <p>462 $\frac{\text{T-NEWBROOM}}{\Sigma \vdash f : A \Rightarrow C \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D}$</p> <p>463 $\frac{\text{T-ENV}}{\text{dom}(\gamma) = \text{dom}(\Gamma) \quad (\vdash_{\Sigma} \gamma(x) : \Gamma(x))_{x \in \text{dom}(\gamma)}}$</p>	<p>464 $\frac{\text{T-LET}}{\Gamma \vdash_{\Sigma} M : A!E \quad \Gamma, x : A \vdash_{\Sigma} N : B!E}$</p> <p>465 $\frac{\text{T-UNPACK}}{\Gamma \vdash_{\Sigma} V : \langle \bar{A} \rangle \quad \Gamma, x : \bar{A} \vdash_{\Sigma} M : C}$</p> <p>466 $\frac{\text{T-DO}}{\Gamma \vdash_{\Sigma} V : A \quad E(\#t) = A \rightarrow B}$</p> <p>467 $\frac{\text{T-HANDLE}}{\Gamma \vdash_{\Sigma} M : C \quad \Gamma \vdash_{\Sigma} V : A \quad \Sigma \vdash h : (A)C \Rightarrow D}$</p>	<p>468 $\frac{\text{T-WITH}}{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D'}$</p> <p>469 $\frac{\text{T-RESUME}}{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : A}$</p> <p>470 $\frac{\text{T-NEWBROOM}}{\Sigma \vdash f : A \Rightarrow C \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D}$</p> <p>471 $\frac{\text{T-ENV}}{\text{dom}(\gamma) = \text{dom}(\Gamma) \quad (\vdash_{\Sigma} \gamma(x) : \Gamma(x))_{x \in \text{dom}(\gamma)}}$</p>
<p>472 $\frac{\text{T-HANDLER}}{\Sigma \vdash h : (A)C \Rightarrow D}$</p> <p>473 $\frac{\text{T-FUN}}{x : A \vdash_{\Sigma} M : C \quad (\text{fun } f(x) \{M\}) \in \Sigma}$</p> <p>474 $\frac{\text{T-PROGRAM}}{(\Sigma \vdash \Sigma(f) : A_f \Rightarrow C_f)_{f \in \text{dom}(\Sigma)} \quad \vdash_{\Sigma} M : C}$</p>	<p>475 $\frac{\text{T-RESUME}}{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : A}$</p> <p>476 $\frac{\text{T-NEWBROOM}}{\Sigma \vdash f : A \Rightarrow C \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D}$</p> <p>477 $\frac{\text{T-ENV}}{\text{dom}(\gamma) = \text{dom}(\Gamma) \quad (\vdash_{\Sigma} \gamma(x) : \Gamma(x))_{x \in \text{dom}(\gamma)}}$</p>	<p>478 $\frac{\text{T-LET}}{\Gamma \vdash_{\Sigma} M : A!E \quad \Gamma, x : A \vdash_{\Sigma} N : B!E}$</p> <p>479 $\frac{\text{T-UNPACK}}{\Gamma \vdash_{\Sigma} V : \langle \bar{A} \rangle \quad \Gamma, x : \bar{A} \vdash_{\Sigma} M : C}$</p> <p>480 $\frac{\text{T-DO}}{\Gamma \vdash_{\Sigma} V : A \quad E(\#t) = A \rightarrow B}$</p> <p>481 $\frac{\text{T-HANDLE}}{\Gamma \vdash_{\Sigma} M : C \quad \Gamma \vdash_{\Sigma} V : A \quad \Sigma \vdash h : (A)C \Rightarrow D}$</p>	<p>482 $\frac{\text{T-WITH}}{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D'}$</p> <p>483 $\frac{\text{T-RESUME}}{\Gamma \vdash_{\Sigma} U : A \xrightarrow{C} D \quad \Gamma \vdash_{\Sigma} V : A}$</p> <p>484 $\frac{\text{T-NEWBROOM}}{\Sigma \vdash f : A \Rightarrow C \quad \Gamma \vdash_{\Sigma} V : B \quad \Sigma \vdash h : (B)C \Rightarrow D}$</p> <p>485 $\frac{\text{T-ENV}}{\text{dom}(\gamma) = \text{dom}(\Gamma) \quad (\vdash_{\Sigma} \gamma(x) : \Gamma(x))_{x \in \text{dom}(\gamma)}}$</p>

Fig. 1. SFX: Typing Rules

(B)C \Rightarrow D instantiated with a value v of type B ; however, since the handler cannot be accessed independently, we save ourselves a more radical restructuring of the type system.

Operationally, a broom of type $A \xrightarrow{C} D$ can be used as a deep resumption $A \rightarrow D$ or, when combined with a new handler of type $(B)C \Rightarrow D$ by means of the **with** operator, as a sheep resumption of type $A \rightarrow D'$:

488 `resume $r(V)$` // deep resumption
 489 `let $r' \leftarrow (r \text{ with } h'(W))$ in resume $r'(V)$` // sheep resumption

3.3 An Abstract Machine for SFX

To provide a semantics for SFX, we develop an abstract machine in the style of the CEK machine [11]. 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 [18]; however, our language presents a few differences, which are reflected into the definition:

- We omit type polymorphism and effect polymorphism and do not need to account for them in the machine.
- 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:

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

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 runtime value v .

The machine is given by the following transition rules:

$$\begin{aligned} \langle \oplus(V), \gamma, \kappa \rangle &\hookrightarrow \langle \llbracket \oplus \rrbracket (\llbracket V \rrbracket \gamma), \gamma, \kappa \rangle \\ \langle \mathbf{if}(V, M, N), \gamma, \kappa \rangle &\hookrightarrow \langle M, \gamma, \kappa \rangle && (\llbracket V \rrbracket \gamma = \mathbf{true}) \\ \langle \mathbf{if}(V, M, N), \gamma, \kappa \rangle &\hookrightarrow \langle N, \gamma, \kappa \rangle && (\llbracket V \rrbracket \gamma = \mathbf{false}) \\ \langle \mathbf{let} \ x \leftarrow M \ \mathbf{in} \ N, \gamma, \kappa \rangle &\hookrightarrow \langle M, \gamma, (x).N[\gamma] \cdot \kappa \rangle \\ \langle \mathbf{unpack} \ \overline{x_n} \leftarrow \langle \overline{V_n} \rangle \ \mathbf{in} \ N, \gamma, \kappa \rangle &\hookrightarrow \langle N, (\overline{x_n} \mapsto \llbracket \overline{V_n} \rrbracket \gamma) \cdot \gamma, \kappa \rangle \\ \langle \mathbf{handle} \ M \ \mathbf{with} \ h(V), \gamma, \kappa \rangle &\hookrightarrow \langle M, \gamma, h(\llbracket V \rrbracket \gamma) \cdot \kappa \rangle \\ \langle \mathbf{do} \ \#t(V), \gamma, \kappa^{-t} \cdot h^{+t}(w) \cdot \kappa' \rangle &\hookrightarrow \langle h_{\#t}, \mathbf{DoEnv}_{h_{\#t}}((\kappa^{-t}, h^{+t}(w)), \llbracket V \rrbracket \gamma, w), \kappa' \rangle \\ \langle \mathbf{newbroom}(f) \ \mathbf{with} \ h(V), \gamma, \kappa \rangle &\hookrightarrow \langle ((x).f.\mathbf{return} [], h(\llbracket V \rrbracket \gamma)), \gamma, \kappa \rangle && (x = \mathbf{dom}(f.\mathbf{return})) \\ \langle U \ \mathbf{with} \ h(W), \gamma, \kappa \rangle &\hookrightarrow \langle (\kappa', h(\llbracket W \rrbracket \gamma)), \gamma, \kappa \rangle && (\llbracket U \rrbracket \gamma = (\kappa', _)) \\ \langle \mathbf{resume} \ U(V), \gamma, \kappa \rangle &\hookrightarrow \langle \llbracket V \rrbracket \gamma, \gamma, \kappa' \cdot h(w) \cdot \kappa \rangle && (\llbracket U \rrbracket \gamma = (\kappa', h(w))) \\ \langle \mathbf{return} \ V, \gamma, \kappa \rangle &\hookrightarrow \langle \llbracket V \rrbracket \gamma, \gamma, \kappa \rangle \\ \langle v, \gamma, (x).N[\gamma'] \cdot \kappa \rangle &\hookrightarrow \langle N, (x \mapsto v) \cdot \gamma', \kappa \rangle \\ \langle v, \gamma, h(w) \cdot \kappa \rangle &\hookrightarrow \langle h_{\mathbf{return}}, \mathbf{RetEnv}_h(v, w), \kappa \rangle \\ \mathbf{DoEnv}_{h_{\#t}}(u, v, w) &= (\mathbf{res}(h_{\#t}) \mapsto u) \cdot (\mathbf{dom}(h_{\#t}) \mapsto v) \cdot (\mathbf{dom}(h) \mapsto w) \\ \mathbf{RetEnv}_h(u, v) &= (\mathbf{dom}(h_{\mathbf{return}}) \mapsto u) \cdot (\mathbf{dom}(h) \mapsto v) \end{aligned}$$

The transitions implementing primitive operations, conditionals, and tuple unpacking are trivial. Sequencing $\mathbf{let} \ x \leftarrow M \ \mathbf{in} \ N$ 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 $\mathbf{handle} \ M \ \mathbf{with} \ h(V)$, 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 \mathbf{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 $\text{do } \#t(V)$ 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 $\text{newbroom}(f)$ **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 resumption $\text{resume } U(V)$, we apply the environment γ to U to obtain the corresponding broom $(\kappa', h(w))$, then we proceed by applying the environment γ to the argument V to obtain a closed value that we can pass to the extended continuation $\kappa' \cdot h(w) \cdot \kappa$.

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

4 An Abstract Assembly Language with Effect Handlers

In this section, we introduce ASMF_X, an abstract assembly language in the style of typical instruction set architectures supported in hardware, but providing instructions to install effect handlers and perform effects. ASMF_X comprises a small set of instructions operating on registers, a main memory containing data and code, and an *effect context*, which is special storage manipulated as a stack and separated from main memory. Though close to typical instruction sets, ASMF_X deliberately leaves several features abstract:

- 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 distinguish different categories of register for different purposes. The register $\$pc$ is the program counter and it contains the memory address of the next instruction to be run.
- Only instructions explicitly needed for implementing effect handlers are included. Other standard instructions including unconditional jumps, conditional branches, arithmetic and logical operations, register moves, and load and store from main memory are left unspecified, but are intended to be available as needed, with their usual semantics.

Main memory is an array of words, and memory addresses ℓ, ℓ', \dots are indices in the array specified as natural numbers. Effects are specified by tags $\#tag$; the special tag *return* identifies the

589 handler return clauses. We now describe the syntax of AsmFX:

590	Handlers	H	$::=$	$return \mapsto \ell_{return}, \#tag \mapsto \ell_{\#tag}$
591	Records	R	$::=$	$R_{\mathcal{H}} \mid R_{\mathcal{K}} \mid R_{\mathcal{L}}$
592	Handler records	$R_{\mathcal{H}}$	$::=$	$\mathcal{H}(H, \overline{\$r_h = v_h})$
593	Continuation records	$R_{\mathcal{K}}$	$::=$	$\mathcal{K}(\ell_r, \overline{\$r_c = v_c}, D_r)$
594	Leave records	$R_{\mathcal{L}}$	$::=$	$\mathcal{L}(\ell_l, \overline{\$r_e = v_e})$
595	Values	v	$::=$	$R \mid \ell \mid I \mid \dots$
596	Instructions	I	$::=$	$loadh\ H, state : (\overline{\$r_h}) \mid loadk\ \ell_r, save : (\overline{\$r_c}) \mid loadl\ \ell_l, save : (\overline{\$r_e})$
597				$\mid resume \mid do\ \#tag, save : (\overline{\$r_c}) \mid return \mid exit$
598				
599	Effect contexts	C, D	$::=$	$\overline{\mathcal{H}(H, \overline{\$r_h = v_h}, R_{\mathcal{L}})}$

600 *Records.* To implement effect handlers at the assembly level, we introduce representations of the
 601 key linguistic concepts we need as special values that can be directly stored in registers: in our
 602 terminology, these are called *records*. Records are similar to closures, and will be used to represent
 603 three kinds of objects:
 604

- 605 • Effect handlers, by means of records of the form $\mathcal{H}(H, (\overline{\$r_h = v_h}))$, where H is a finite partial
 606 map from operation names to memory addresses for the code to handle the corresponding
 607 operation, $\overline{\$r_h = v_h}$ is a finite partial map from register names to values used to store the
 608 handler state. H always includes a distinguished *#return* entry which maps to the address
 609 of the code to be run when the handler returns.
- 610 • Continuations, by means of records $\mathcal{K}(\ell_r, (\overline{\$r_c = v_c}), D_r)$, where ℓ_r is the address of the code
 611 that will be run when resuming the continuation, and the other two parameters express the
 612 local environment that needs to be reinstated when control is passed to the continuation:
 613 this local environment consists of a finite map from register names to their values, and a
 614 delimited effect context containing additional handlers that must be added to the effect
 615 context before invoking the continuation.
- 616 • Leave records, in the form $\mathcal{L}(\ell_l, (\overline{\$r_e = v_e}))$, indicate what code should be run after leaving
 617 the scope of a handler; they are similar to continuations, except that they do not provide a
 618 delimited effect context because when we leave the scope of a handler we never need to
 619 install additional handlers.

620
 621 *Effect Context.* The effect context is a stack of *handler frames* representing the active handlers. A
 622 handler frame is a handler record augmented with a leave record that denotes the code to be run
 623 after returning from the handler.

624
 625 *Instructions.* AsmFX introduces three instructions *loadh*, *loadk*, *loadl* to load handler, continu-
 626 ation, and leave records into registers. The three instructions use, respectively, registers $\$hr$, $\$kr$,
 627 and $\$lr$ as their targets: these registers have an architectural meaning in that their content can
 628 affect the control flow; it is however possible to copy the content of these registers to any other
 629 register, and vice-versa, using general-purpose instructions.

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

634 *loadk* $\ell_r, save : (\overline{\$r_c})$ loads a continuation record into register $\$kr$ by specifying the address ℓ_r of
 635 its code and the registers $\overline{\$r_c}$ whose value must be saved now and restored if the continuation is
 636 run and terminates normally. The continuation is created with a trivial (empty) delimited context.

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

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

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

657 `return` operates similarly to `do` but it always matches the handler at the top of the context and
 658 when jumping to its return clause it does not need to provide a resumption. This means it loads
 659 $\$lr$ and the state registers from the handler frame, and then jumps to the address specified by the
 660 return clause $H(\#return)$.

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

663 The semantics of ASMF_X is expressed by the abstract machine in Figure 2. A machine state is
 664 specified by a tuple $\langle \Xi, \Theta, C \rangle$, where Ξ , the main memory, maps addresses to values, which may
 665 be data or instructions; Θ , the register file, maps register names to their current value; and C , the
 666 effect context, is a stack of handler frames specifying what effect implementations are available
 667 to the code currently running. When it is ready to execute an instruction, the machine fetches it
 668 from memory Ξ at the address stored in the $\$pc$ register (it is assumed that the memory at that
 669 address contains a valid instruction, otherwise the machine will be stuck); at this point the machine
 670 immediately increments $\$pc$ to ensure it points to the next instruction, then it matches the current
 671 instruction (including its arguments) and the effect context C against the first three columns of the
 672 transition table: execution involves updating the context C and the registers file Θ according to the
 673 content of the last two columns: among other things, this allows an instruction to perform a jump
 674 by updating the register $\$pc$.

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

679 5 Compiling SFX to ASMF_X

680 In this section we specify a compiler from SFX to ASMF_X demonstrating that ASMF_X is a suitable
 681 target for implementing high-level effect handlers.

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

instruction	arguments	initial context	updated context	register updates
687 loadk	ℓ_r $save : (\overline{\$r_c})$	C	C	$\$kr = \mathcal{K}(\ell_r, (\overline{\$r_c} = \Theta(\overline{\$r_c})), [])$
689 loadh	H $state : (\overline{\$r_h})$	C	C	$\$hr = \mathcal{H}(H, (\overline{\$r_h} = \Theta(\overline{\$r_h})))$
691 loadl	ℓ_l $save : (\overline{\$r_e})$	C	C	$\$lr = \mathcal{L}(\ell_l, (\overline{\$r_e} = \Theta(\overline{\$r_e})))$
693 resume		C with: $\$hr = \mathcal{H}(H, (\overline{\$r_h} = v_h))$ $\$kr = \mathcal{K}(\ell_r, (\overline{\$r_c} = v_c), D_r)$ $\$lr = R_{\mathcal{L}}$	D_r $\cdot \mathcal{H}(H, (\overline{\$r_h} = v_h), R_{\mathcal{L}})$ $\cdot C$	$\$pc = \ell_r$ $\$r_c = v_c$
696 do	$\#tag$ $save : (\overline{\$r_c})$	D_r^{-tag} $\cdot \mathcal{H}(H^{+tag}, (\overline{\$r_h} = v_h), R_{\mathcal{L}})$ $\cdot C$	C	$\$kr = \mathcal{K}(\Theta(\$pc),$ $\quad (\overline{\$r_c} = \Theta(\overline{\$r_c})),$ $\quad D_r^{-tag})$ $\$hr = \mathcal{H}(H^{+tag}, (\overline{\$r_h} = v_h))$ $\$lr = R_{\mathcal{L}}$ $\$pc = H^{+tag}(\#tag)$ $\$r_h = v_h$
702 return		$\mathcal{H}(H, (\overline{\$r_h} = v_h), R_{\mathcal{L}})$ $\cdot C$	C	$\$pc = H(\#return)$ $\$r_h = v_h$ $\$lr = R_{\mathcal{L}}$
705 exit		C with: $\$lr = \mathcal{L}(\ell_l, (\overline{\$r_e} = v_e))$	C	$\$pc = \ell_l$ $\$r_e = v_e$

Domain of contexts:

$$dom([]) = \emptyset \quad dom(\mathcal{H}(H, \overline{\$r_h} = v_h, R_{\mathcal{L}}) \cdot C) = dom(H) \cup dom(C) \quad dom(\#t_1 \mapsto \ell_1, \dots, \#t_n \mapsto \ell_n) = \{\#t_1, \dots, \#t_n\}.$$

D^{-t} indicates any delimited context such that $\#t \notin dom(D^{-t})$. H^{+t} denotes any handler definition where $\#t \in dom(H^{+t})$.

Fig. 2. ASMFx Machine Transitions

We 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 become clear when we 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 back into the source code in the form of *annotations*. These annotations record, for each subexpression, the memory location ℓ of the ASMFx code that evaluates it. We let $\hat{V}, \hat{M}, \hat{P}$ range over annotated source values, computations, and programs respectively. We call the language of augmented SFX expressions ASFX. We omit the tedious details of the backwards annotation process (which is entirely analogous to the output of debugging symbols by many compilers), and 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 obtain the original source. We write $\mu(\cdot)$ and $\nu(\cdot)$ for the starting and ending ASMFx location of a given augmented expression, and $|\cdot|$ for the erasure of an augmented expression into the corresponding plain SFX expression.

PROPOSITION 5.1. *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 explicitly write annotations on expressions using the syntax $\hat{M}@l$ to indicate that $\mu(\hat{M}) = l$. A formal grammar of ASFX is included in the appendix. We 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 partitions the register set into *variable* registers $\$x_i$ (where x

<pre> 736 737 738 [return V] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ 739 [V] 740 let $\mathbb{L}' = \mathbb{L}$ 741 742 [let x ← M 743 in N] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ 744 let $\mathbb{I}_M = \overline{\\$x_{\ M\ }}$ 745 [M] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}'_M$ 746 $\mathbb{I}_M \leftarrow \overline{\\$a_{\ M\ }}$ 747 [N] $\mathbb{I} \cup \mathbb{I}_M, \mathbb{L}'_M \Rightarrow \mathbb{L}'$ 748 749 [unpack $\overline{x^k} \leftarrow V$ 750 in M] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ 751 let $\mathbb{I}_M = \mathbb{I} \cup \{\overline{\\$x^k_{\ x_n\ }}\}$ 752 [V] 753 $\overline{\\$x^0_{\ x^0\ }} \leftarrow \overline{\\$a_{\ x^0\ }}$ 754 : 755 $\overline{\\$x^i_{\ x^i\ }} \leftarrow \overline{\\$a_{\ \Sigma \ x^{i-1}\ , \Sigma \ x^{n-1}\ }}$ 756 : 757 $\overline{\\$x^{n-1}_{\ x^{n-1}\ }} \leftarrow \overline{\\$a_{\ \Sigma \ x^{n-2}\ , \Sigma \ x^{n-1}\ }}$ 758 [M] $\mathbb{I}_M, \mathbb{L} \Rightarrow \mathbb{L}'$ 759 760 [U with h(V)] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ 761 [V] 762 $\overline{\\$s_{\ V\ }} \leftarrow \overline{\\$a_{\ V\ }}$ 763 loadh [h], state : ($\overline{\\$s_{\ V\ }}$) 764 [U] 765 $\\$a_1 \leftarrow \\hr 766 let $\mathbb{L}' = \mathbb{L}$ </pre>	<pre> [⊕(V)] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ [V] $\overline{\\$a_{\ out(\oplus)\ }} \leftarrow \overline{[\oplus] \overline{\\$a_{\ V\ }}}$ let $\mathbb{L}' = \mathbb{L}$ 741 [do #t(V)] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ [V] do #t, save :\mathbb{I} let $\mathbb{L}' = \mathbb{L}$ 742 [resume U(V)] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ let $l_{exit} = \varphi(1, \mathbb{L})$ [U] $\\$kr \leftarrow \\a_0 $\\$hr \leftarrow \\a_1 [V] loadl $\\$lr, l_{exit}$, save :$\mathbb{I}$ resume $l_{exit} :$ let $\mathbb{L}' = \mathbb{L}$ 743 [newbroom f] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ with h(V) let $l_{res}, l_{exit}, l_{end} = \varphi(3, \mathbb{L})$ let $\mathbb{L}' = \mathbb{L} \cup \{l_{res}, l_{exit}, l_{end}\}$ [V] $\overline{\\$s_{\ V\ }} \leftarrow \overline{\\$a_{\ V\ }}$ loadh [h], state : ($\overline{\\$s_{\ V\ }}$) loadk l_{res} $\\$a_0, \\$a_1 \leftarrow \\$kr, \\hr j l_{end} $l_{res} :$ loadl l_{exit} j freturn $l_{exit} :$ return l_{end} </pre>	<pre> [if(V, M, N)] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ let $l_{end}, l_{else} = \varphi(2, \mathbb{L})$ let $\mathbb{L}_M = \mathbb{L} \cup \{l_{end}, l_{else}\}$ [V] bz $l_{else}, \\$a_0$ [M] $\mathbb{I}, \mathbb{L}_M \Rightarrow \mathbb{L}'_M$ j l_{end} $l_{else} :$ [N] $\mathbb{I}, \mathbb{L}'_M \Rightarrow \mathbb{L}'$ $l_{end} :$ 744 [handle M] $\mathbb{I}, \mathbb{L} \Rightarrow \mathbb{L}' :=$ with h(V) let $l_{cli}, l_{exit} = \varphi(2, \mathbb{L})$ let $\mathbb{L}_M = \mathbb{L} \cup \{l_{cli}, l_{exit}\}$ let $l_h = \overline{\\$s_{\ V\ }}$ [V] $\overline{l}_h \leftarrow \overline{\\$a_{\ V\ }}$ loadh $\\$hr, [h]$, state :$\overline{l}_h$ loadk $\\$kr, l_{cli}$, save :() loadl $\\$lr, l_{exit}$, save :$\mathbb{I}$ resume $l_{cli} :$ [M] $\mathbb{I}, \mathbb{L}_M \Rightarrow \mathbb{L}'_M$ return $l_{exit} :$ let $\mathbb{L}' = \mathbb{L}'_M$ </pre>
--	--	---

Fig. 3. Compilation of SFX Computations to ASMFx

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. We assume an unlimited number of registers in each of the first three classes.

Compiling Computations. To denote the compilation of SFX computations M , we use the syntax $[M]_{\mathbb{I}, \mathbb{L}}^{\mathbb{I}, \mathbb{L}} \Rightarrow \mathbb{L}'$. 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 $\overline{\$a}$.

We define the compiler by way of rules of the form $[M]_{\mathbb{I}, \mathbb{L}}^{\mathbb{I}, \mathbb{L}} \Rightarrow \mathbb{L}' := \overline{I}$, where the primary output \overline{I} is a list of ASMFx instructions annotated with labels representing the result of compiling M , and the secondary output \mathbb{L}' is an updated set of used labels which must extend \mathbb{L} . The instructions \overline{I}

are defined using syntax such as the following:

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

where **let** is meta-level let-binding, which we use as a concise notation for binding fresh code labels, sets of registers, and sets of labels. 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. We 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 morally each label simply represents a distinct absolute memory index.)

The compiler is given in Figure 3. It is defined by structural recursion on computations, making use of an auxiliary operation on values V to generate code that loads an initial segment of the $\$a$ registers of suitable length with the AsmFX representation of V .

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

For each primitive operation \oplus , we assume a corresponding AsmFX instruction $[\oplus]$ with the appropriate number of register inputs and outputs.

We discuss two interesting cases of the compiler, corresponding to performing an effect and installing a handler. To compile $\text{do } \#t(V)$, 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).

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`. The client code at location l_{cli} is terminated by a return instruction, so the result of running M will be delivered to the return clause of the newly installed handler.

Remark. The return instruction appended to the compiled code for the client of a **handle** computation presents a significant challenge for the proof of soundness of the compiler, because the compiled **handle** code is partitioned into two fragments: one that runs before the client code, and another that runs afterwards. For a full proof of soundness, we cannot just process subcomputations in order, but must remember all outer computations that have not run to completion yet. In Section 6 we introduce a novel technique to carry out this kind of proof.

The return instruction is a form of *epilogue*. Epilogues are assembly fragments that must be run after running the code generated for a certain computation, but before moving on to the next computation. The purpose of epilogues is to make the implicit control flow of SFX explicit, as demanded by AsmFX. Epilogues are related to continuation frames: the return epilogue created when compiling a **handle** corresponds to an operation that must be performed when an SFX handler frame is activated. Other epilogues used in the compilation of **if**, **newbroom**, and in handler clauses

do not directly map to SFX continuation frames, so require a sophisticated analysis which seems difficult to express solely in terms of the original SFX program or the generated AsmFX code.

Compiling Handlers. The compilation of handlers is defined by the following rule:

$$\begin{array}{l}
 \llbracket \text{handler } h(x) \{ \text{return } y \rightarrow M \mid (\#t_i(z^i) \rightarrow (r^i).N_i)_{i=1,\dots,n} \} \rrbracket^L \Rightarrow L' := \\
 \begin{array}{l}
 \text{let } h_{\text{return}}, \overline{h_{\#t}} = \varphi(1 + \|\overline{\#t}\|) \\
 \text{let } \mathbb{L}_{\text{return}} = \mathbb{L} \cup \{h_{\text{return}}, \overline{h_{\#t}}\} \\
 \text{let } \mathbb{L}_x = \$x_{\|x\|} \\
 h_{\text{return}} : \\
 \text{let } \mathbb{L}_M = \$y_{\|y\|} \\
 \mathbb{L}_x \leftarrow \$S_{\|x\|} \\
 \mathbb{L}_M \leftarrow \$a_{\|y\|} \\
 \llbracket M \rrbracket^{\mathbb{L}_x \cup \mathbb{L}_M \cup \{\$1r\}, \mathbb{L}_{\text{return}}} \Rightarrow \mathbb{L}_{\#t_1} \\
 \text{exit}
 \end{array}
 \quad
 \begin{array}{l}
 h_{\#t_1} : \\
 \text{let } \mathbb{L}_{N_1} = \$z^1_{\|z^1\|}, \$r_1, \$r_2 \\
 \mathbb{L}_x \leftarrow \$S_{\|x\|} \\
 \mathbb{L}_{N_1} \leftarrow \$a_{\|z^1\|}, \$kr, \$hr \\
 \llbracket N_1 \rrbracket^{\mathbb{L}_x \cup \mathbb{L}_{N_1} \cup \{\$1r\}, \mathbb{L}_{\#t_1}} \Rightarrow \mathbb{L}_{\#t_2} \\
 \text{exit} \\
 h_{\#t_2} : \\
 \vdots \\
 \text{let } L' = \mathbb{L}_{\#t_{n+1}}
 \end{array}
 \end{array}$$

For each return clause of a handler h , the compiler generates a unique label h_{return} . Similarly, for each operation clause handling operation $\#t$ it generates a unique label $h_{\#t}$. A handler h that handles operations $\#t_1, \dots, \#t_n$ thus generates the handler specification $\llbracket h \rrbracket = \text{return} \mapsto h_{\text{return}}, \#t_1 \mapsto h_{\#t_1}, \dots, \#t_n \mapsto h_{\#t_n}$. The generated label for each clause is attached to the code generated for the body of the clause. For each clause, as well as generating the code for the body of the clause, the compiler takes care of moving the actual parameter of the handler from the state registers $\$s$ to variable registers $\$x$ corresponding to the formal parameter x used by the SFX handler definition. The compiler also inserts code to move arguments of each clause from the argument registers $\$a$ to variable registers $\$y$ corresponding to the formal argument y of the clause in the SFX definition, and in the case of operation clauses 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 must be accounted for in the soundness proof.

6 Soundness of Compilation

In this section we prove that our compiler soundly compiles SFX into AsmFX. 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. However, we are also interested in the correctness of partial executions, and so we prove a stronger property 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. Consider for example the SFX execution (where we assume handler h has an identity return clause (`return x` \rightarrow `return x`) and we use the shorthands $u := \llbracket U \rrbracket \gamma, v := \llbracket V \rrbracket \gamma$, and $(M; N) := \text{let } _ \leftarrow M \text{ in } N$:

$$\langle \langle \text{handle return } U \text{ with } h(V); M \rangle, \gamma, \kappa \rangle \quad (1)$$

$$\hookrightarrow \langle \langle \text{handle return } U \text{ with } h(V) \rangle, \gamma, ((_).M[\gamma]) \cdot \kappa \rangle \quad (2)$$

$$\hookrightarrow \langle \langle \text{return } U \rangle, \gamma, h(v) \cdot ((_).M[\gamma]) \cdot \kappa \rangle \quad (3)$$

$$\hookrightarrow \langle u, \gamma, h(v) \cdot ((_).M[\gamma]) \cdot \kappa \rangle \quad (4)$$

$$\hookrightarrow \langle \text{return } x \rangle, (y \mapsto v) \cdot (x \mapsto u), ((_).M[\gamma]) \cdot \kappa \rangle \quad (5)$$

$$\hookrightarrow \langle u, (y \mapsto v) \cdot (x \mapsto u), ((_).M[\gamma]) \cdot \kappa \rangle \quad (6)$$

$$\hookrightarrow \langle M, \gamma, \kappa \rangle \tag{7}$$

For each SFX configuration $\langle M, \gamma, \kappa \rangle$, we can compile the code M ; then, for each transition $\langle M, \gamma, \kappa \rangle \hookrightarrow \langle M', \gamma', \kappa' \rangle$, we plan to show that the AsmFX machine can move (in any number of steps) from the compiled version of M to the compiled version of M' (where some conditions over environments and continuations need to be satisfied). By direct inspection of the compiled code for configuration (1), we see that after a few AsmFX instructions we reach the compiled code for (2), as desired. The reasoning to reach configuration (3) is slightly more complex because when we compile a **handle** statement in AsmFX we create and resume a continuation: however we can still see that we will deterministically end up at the compiled code for **return** U ; additionally, we pleasantly notice that the new AsmFX effect context includes a frame for the compiled version of handler $h(v)$, establishing a correspondence with the SFX continuation.

The first complication arises in the transition to configuration (5): that corresponds to invoking the **return** clause of h . In SFX, this happens because the first continuation frame is the $h(v)$ handler, but in AsmFX we need to execute a **return** instruction. That instruction was indeed generated when compiling the **handle** statement and is the first instruction after evaluating the client code. But how do we know, looking at the compiled code for (4), that the next instruction is **return**? There must be an invariant relating SFX and AsmFX configurations in such a way that if we are about to activate frame $h(v)$ in SFX, the program counter in AsmFX points to a **return** instruction: this means that we need to be able to relate SFX expressions to the AsmFX code locations corresponding to them. Additionally, this invariant involving code locations, which is used at configuration (5), must have been established when the $h(v)$ frame was originally installed, at (2).

At (6), we find another issue. To reach configuration (7) we must activate the pure frame $((_) . M[\gamma])$. However, this frame has no explicit counterpart in AsmFX. What happens is that the compiled code for the **return** clause of h is terminated by an **exit** instruction, which activates the leave record in register $\$1r$, and if that record was set up correctly, AsmFX will jump to the code location corresponding to M and, at the same time, restore the registers corresponding to the environment γ . The compiled code ultimately works thanks to the careful implementation of steps (1), (2), (4), (6) which enforce the correspondence between program expressions and code locations, and between continuation frames and explicit control flow instructions such as **exit**. Moreover, a machine representation of the environment γ will be moved around through registers, handler frames, and finally the active leave record.

6.1 Outline of the Proof

One of the key insights from the example above is that, in order for us to prove the soundness of compilation, we need to track compiled code location information through the SFX execution. To this end, we will use an abstract machine operating not on plain SFX computations, but with their ASFX counterpart (introduced in Section 5) which contains code locations. Besides tracking locations, the ASFX machine will also log other information about execution history, and manipulate environments in a slightly different way (closer to the compiled AsmFX code).

ASFX configurations differ from those in SFX also because of the form of their continuations. The identity continuation and all frames are annotated with a start location and 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 component $\hat{\gamma}$ of ASFX brooms.

Due to space constraints, the full definition of the abstract machine and a detailed explanation of its rules are deferred to the appendix. Instead, we illustrate how our SFX execution example changes when translated to the ASFX machine.

$$\langle\langle(\text{handle return } \hat{U} \text{ with } h(\hat{V}))@_{\ell_1}; \hat{M}\rangle, \hat{\gamma}, \hat{\kappa}\rangle \quad (1)$$

$$\hookrightarrow \langle\langle\text{handle return } \hat{U} \text{ with } h(\hat{V})@_{\ell_1}\rangle, \hat{\gamma}, ((_) . \hat{M})@_{\ell_2} \cdot \hat{\kappa}\rangle \quad (2)$$

$$\hookrightarrow \langle\langle\text{return } \hat{U}\rangle, \hat{\gamma}, h(\hat{v}, \mathcal{R}_{\mathcal{L}})@_{\ell_3} \cdot ((_) . \hat{M})@_{\ell_2} \cdot \hat{\kappa}\rangle \quad (3)$$

$$\hookrightarrow \langle\hat{u}, \hat{\gamma}, h(\hat{v}, \mathcal{R}_{\mathcal{L}})@_{\ell_3} \cdot ((_) . \hat{M})@_{\ell_2} \cdot \hat{\kappa}\rangle \quad (4)$$

$$\hookrightarrow \langle\text{return } x, (\text{LEAVE} \mapsto \mathcal{R}_{\mathcal{L}}) \cdot (y \mapsto \hat{v}) \cdot (x \mapsto \hat{u}), \text{exit}@_{\ell_4} \cdot ((_) . \hat{M})@_{\ell_2} \cdot \hat{\kappa}\rangle \quad (5)$$

$$\hookrightarrow \langle\hat{u}, (\text{LEAVE} \mapsto \mathcal{R}_{\mathcal{L}}) \cdot (y \mapsto \hat{v}) \cdot (x \mapsto \hat{u}), \text{exit}@_{\ell_4} \cdot ((_) . \hat{M})@_{\ell_2} \cdot \hat{\kappa}\rangle \quad (6)$$

$$\hookrightarrow \langle\hat{u}, \hat{\gamma}, ((_) . \hat{M})@_{\ell_2} \cdot \hat{\kappa}\rangle \quad (7)$$

$$\hookrightarrow \langle\hat{M}, \hat{\gamma}, \hat{\kappa}\rangle \quad (8)$$

In the configurations above, ℓ_1 and ℓ_2 are locations for the beginning and the end of the `handle` statement; ℓ_3 points to the end of the client code, which corresponds to a `return` instruction; ℓ_4 is the location at the end of the return clause of h , which contains an `exit` instruction. An ASFX leave record $\mathcal{R}_{\mathcal{L}} := (\ell_2, \hat{\gamma})$ is generated to augment the handler frame in (3) and propagated afterwards.

Compared to the SFX example, when we trigger the handler frame at (4), we know that the current instruction is indeed a `return` because the frame is annotated with ℓ_3 . At the same time, we move $\mathcal{R}_{\mathcal{L}}$ to a pseudo-variable `LEAVE` (reflecting AsmFX register `$1r` into ASFX) for later use. At configuration (6), we know that the AsmFX program counter points to an `exit` instruction, because the continuation starts with an *exit epilogue frame*, which was not present in SFX. When we move to (7) (corresponding to the execution of `exit`), we restore the environment $\hat{\gamma}$ from the `LEAVE` pseudo-variable: in this way, even though the ASFX pure frame $((_) . \hat{M})$ does not contain an environment, we ensure that at the time we need to run \hat{M} the active environment will be $\hat{\gamma}$. This means that ASFX uses environments quite differently from SFX, and more similarly to AsmFX.

Instead of providing a direct soundness proof for the compiler, we found it convenient to split it into two parts using ASFX as the midpoint. First, we show that there is a simulation from SFX to ASFX expressed by an inductively defined relation $s \leq \hat{s}$ matching SFX configurations s to ASFX configurations \hat{s} . Then we show that there is a simulation from the *valid* subset of ASFX to AsmFX, which relates ASFX configurations \hat{s} and AsmFX configurations a and is denoted by $a \vDash \hat{s}$. The valid subset of ASFX essentially consists of those configurations that are correctly annotated (e.g. for any subterm \hat{M} , the AsmFX memory at location $\mu(\hat{M})$ contains the compiled code for \hat{M}). Valid configurations are expressed by an inductive predicate denoted by $\hat{s} \checkmark$.

The meaning of the proofs can be explained as follows. SFX and ASFX have a similar control flow, but manipulate environments differently, so in order for the latter to simulate the former we need to show that every variable will be evaluated in similar environments (a property that we call *well-scoping*). On the other hand, to prove that the ASFX control flow is correctly simulated by AsmFX, we need some deep reasoning involving the computation history that led to a certain ASFX configuration. This historical information, collected in the validity predicate, is enough to show that ASFX and AsmFX have corresponding transitions.

Formally, we prove the following theorems:

THEOREM 6.1. *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 $\hat{s} \xrightarrow{*} \hat{s}'$.*

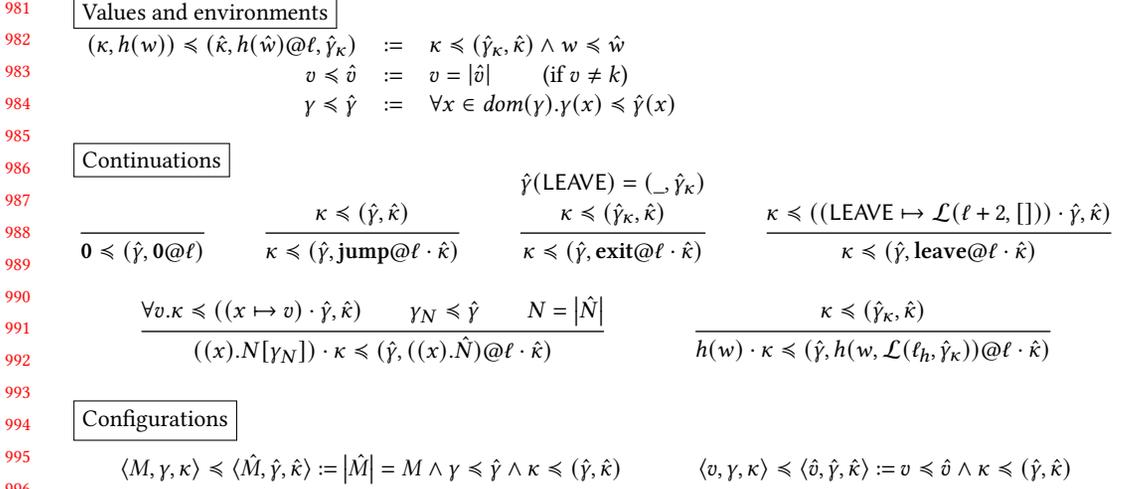


Fig. 4. Scope-Preserving Simulation

PROOF SKETCH. By case analysis on the transition $s \hookrightarrow s'$ and on the relation $s \leq \hat{s}$. The details can be found in the appendix. \square

THEOREM 6.2. *Let P be an SFX program, and $\Xi = \lfloor P \rfloor$. For all ASFX 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 ASMX configurations a such that $a \vDash \hat{s}$ there exists a' such that $a \xrightarrow{*} a'$ and $a' \vDash \hat{s}'$.*

PROOF SKETCH. By case analysis on the transition $\hat{s} \hookrightarrow \hat{s}'$ and on the predicate $\hat{s} \checkmark$. The full proof is in the appendix. \square

We now discuss the definitions of $s \leq \hat{s}$, $\hat{s} \checkmark$, and $a \vDash \hat{s}$ which are crucial to the soundness proof.

6.2 Well-Scoping

As we noted, ASFX differs from SFX because of the different treatment of environments and because of the presence of epilogue frames. However, ASFX epilogues essentially correspond to no operation at all in SFX, therefore the main challenge is to show that even though ASFX switches environments in a quite different way to SFX, its approach is still entirely coherent. The simulation relation, summarised in Figure 4, expresses a well-scoping invariant of ASFX environments with respect to SFX environments under any execution.

The relation is defined by means of auxiliary definitions for the simulation of environments and continuations. An interesting property of the simulation of continuations 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, whereas 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 obtain the same value; if $\gamma(x)$ is a broom $(\kappa, h(w))$, then $\hat{\gamma}(x)$ must be a broom $\hat{\kappa}, h(w)@_{\ell}, \hat{\gamma}_{\kappa}$ such that κ is simulated by $(\hat{\gamma}_{\kappa}, \hat{\kappa})$.

The continuation simulation $\kappa \leq \hat{\kappa}$ is defined by 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 replaced by the one contained in the LEAVE pseudovisible (which must be defined), reflecting the semantics of the ASMFx instruction `exit`, which is executed if the `exit` frame is activated. This kind of rule ensures that though 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 erasing \hat{N} yields 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 must 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.

6.3 Control Flow Simulation

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 often 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 $\llbracket \hat{P} \rrbracket$. In fact, for every \hat{s} there are multiple such a , because of the plethora of registers available in the architecture (some registers have no counterpart in the source configuration and are thus 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 model \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. In contrast, the effect context need only 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 need only consider the handler frames:

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

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 the beginning of the continuation. The pseudovisible LEAVE is always mapped to the register

1079	Values	$\hat{v}\checkmark := \forall \hat{k} \in \text{subval}(\hat{v}) : \hat{k}\checkmark$
1080		
1081	Environments	$\hat{y}\checkmark := \forall x \in \text{dom}(\hat{y}) : \hat{y}(x)\checkmark$
1082		
1083	Brooms	$\hat{k}\checkmark := \forall (\hat{x}, h(\hat{w})@l, \hat{y}) = \hat{k}, \forall \hat{k}'\checkmark, \forall \hat{y}' \supset \hat{k}' :$
1084		$\hat{y} \supset (\hat{x} \cdot h(\hat{w}, \mathcal{L}(\mu(\hat{k}'), \hat{y}'))@l \cdot \hat{k}') \wedge (\hat{x} \cdot h(\hat{w}, \mathcal{L}(\mu(\hat{k}'), \hat{y}'))@l \cdot \hat{k}')\checkmark$
1085	Coherence	$\hat{y} \supset \hat{k} := \hat{y}\checkmark \wedge \hat{y}(\text{LEAVE}) = (l, \hat{y}') \implies l = \text{laddr}(\hat{k}) \wedge \hat{y}' \supset \text{lcont}(\hat{k})$
1086		
1087	Continuations	
1088		$\frac{\hat{k}\checkmark \quad \Xi(l) = \text{j } \mu(\hat{k})}{\text{jump}@l \cdot \hat{k}\checkmark} \qquad \frac{\hat{k}\checkmark \quad \Xi(l) = \text{exit}}{\text{exit}@l \cdot \hat{k}\checkmark}$
1089		
1090		
1091		$\frac{\hat{k}\checkmark \quad \Xi(l) = \text{loadl } (l + 2) \quad \hat{k} = \text{jump}@l + 1 \cdot ((x).\hat{N})@l \cdot \mu(\hat{N}) \cdot \text{exit}@l \cdot h(\hat{w})@l + 2 \cdot \hat{k}_0}{\text{leave}@l \cdot \hat{k}\checkmark}$
1092		
1093		
1094		$\frac{\hat{k}\checkmark \quad \Xi(l) = \overline{\$x} \leftarrow \overline{\$a}; [\hat{N}] \quad v(\hat{N}) = \mu(\hat{k})}{((x).\hat{N})@l \cdot \hat{k}\checkmark} \qquad \frac{\hat{k}\checkmark \quad \Xi(l) = \text{return} \quad \hat{y} \supset \hat{k}}{h(\hat{w}, \mathcal{L}(\mu(\hat{k}), \hat{y}))@l \cdot \hat{k}\checkmark}$
1095		
1096	Leave operators	
1097		
1098		$\text{lcont}(\hat{k}) := \begin{cases} \text{undefined} & \text{if } \hat{k} = \mathbf{0}@l \text{ or } \hat{k} = h(_)@l \cdot \hat{k}' \\ \hat{k}' & \text{if } \hat{k} = \text{exit}@l \cdot \hat{k}' \\ \text{lcont}(\hat{k}') & \text{if } \hat{k} = \hat{\zeta} \cdot \hat{k}' \text{ and } \hat{\zeta} \neq h(_)@l \end{cases}$
1099		
1100		$\text{laddr}(\hat{k}) := \mu(\text{lcont}(\hat{k}))$
1101		
1102	Configurations	
1103		$\langle \hat{M}, \hat{y}, \hat{k} \rangle\checkmark := \Xi(\mu(\hat{M})) = [\hat{M}] \wedge v(\hat{M}) = \mu(\hat{k}) \wedge \hat{y} \supset \hat{k} \wedge \hat{k}\checkmark$
1104		$\langle \hat{v}, \hat{y}, \hat{k} \rangle\checkmark := \hat{v}\checkmark \wedge \hat{y} \supset \hat{k} \wedge \hat{k}\checkmark$
1105		

Fig. 5. Augmented Source Validity

$\$1r$, and every variable x is mapped to as many $\$x$ registers as necessary.

$$\Theta \vDash \langle \hat{M}, \hat{y}, \hat{k} \rangle := \Theta(\$pc) = \mu(\hat{M}) \wedge \forall x \in \text{dom}(\hat{y}). \overline{\Theta(\$x)} = \lfloor \hat{y}(x) \rfloor$$

$$\Theta \vDash \langle \hat{v}, \hat{y}, \hat{k} \rangle := \Theta(\$pc) = \mu(\hat{k}) \wedge \Theta(\$a) = \lfloor \hat{v} \rfloor \wedge \forall x \in \text{dom}(\hat{y}). \overline{\Theta(\$x)} = \lfloor \hat{y}(x) \rfloor$$

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 $a \vDash \hat{s}$, then there exists a transition chain $a \xrightarrow{*} a'$ where $a' \vDash \hat{s}'$. However, this property cannot be proved for an arbitrary \hat{s} as augmented source configurations contain metadata that encode part of the execution history. Though the compiler produces correct metadata, and the abstract machine ensures that the metadata are 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{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 location $v(\hat{M})$ matches location $\mu(\hat{k})$; if instead the control is a value v , the value itself must be 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 obtain 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. The annotation for **jump** frames must reference a jump to the rest of the continuation, whereas **exit** corresponds to an `exit` instruction, and **leave**, generated in brooms created by **newbroom**, points to a `loadl` instruction and additionally requires the rest of the continuation to have a certain shape that mimics the compilation rule for **newbroom**. The validity rule involving pure frames checks that their location matches the relevant part of the compilation rule for **let** and ensures that the computation \hat{N} is correctly linked to the remaining part of the continuation. Finally, handler frames are annotated with a location pointing to a return instruction (triggered to pass a value to the return clause) and the environment stored in their leave record must be coherent with the rest of the continuation.

6.4 Soundness

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

LEMMA 6.3. *If \hat{M} is the result of annotating M according to the compiler, $|\hat{M}| = M$ and, consequently, $\langle M, [], \mathbf{0} \rangle \leq \langle \hat{M}, [], \mathbf{0}@v(\hat{M}) \rangle$.*

LEMMA 6.4. *Let $\hat{P} = \hat{M}; \hat{\Sigma}$. If $[\hat{P}] = \Xi$, then $\Xi(\mu(\hat{M})) = [\hat{M}]$. Consequently $\langle \hat{M}, [], \mathbf{0}@v(\hat{M}) \rangle \checkmark$ (with respect to Ξ).*

LEMMA 6.5. *If $\Theta = [\text{\$pc} \mapsto \mu(\hat{M})]$, then $\Theta \vDash \langle \hat{M}, [], \mathbf{0}@v(\hat{M}) \rangle$. Consequently, we have that $\langle \Theta, [] \rangle \vDash \langle \hat{M}, [], \mathbf{0}@v(\hat{M}) \rangle$.*

The soundness theorem then arises from Theorems 6.1 and 6.2 as a corollary:

COROLLARY 6.6. *Let $P = M; \Sigma$ be an SFX program and $s = \langle M, [], \mathbf{0} \rangle$. Let $\hat{P} = \hat{M}; \hat{\Sigma}$ be P annotated by the compiler, $\Theta = [\text{\$pc} \mapsto \mu(\hat{M})]$, and $a = \langle \Theta, [] \rangle$. Then if $s \xrightarrow{*} s'$, there exist a', s' such that $a \xrightarrow{*} a', s' \leq \hat{s}'$, and $a' \vDash \hat{s}'$.*

7 Discussion

We have demonstrated how to correctly compile a general form of effect handlers into ASmFX. Our correctness proof relies on a novel proof technique based on epilogues and annotated source terms. In devising our compiler we focused on the general case. However, there are ample 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

1177 compiler emits the code to set up a new handler frame, which includes a leave record: while the
1178 handler functionality is trivial, the leave record plays the role of the call stack, thanks to its ability
1179 to store the caller’s return address and environment (including the previous leave record). Indeed,
1180 it is easy to give a direct translation of functions to AsmFX that does not make use of handlers: all
1181 that is needed to implement functions in AsmFX is the functionality of the leave register `$lr`, with
1182 the instructions `loadl` and `exit` used to push and pop stack frames.

1183 A desirable feature of the AsmFX implementation of handlers is that it naturally provides a
1184 certain level of memory protection: the state of a handler is stored in the effect context and cannot
1185 be directly accessed by other code fragments (including client code or other handlers); our compiler
1186 does not actively wipe registers when transitioning in and out of handler code, so (read only)
1187 information leak is still possible, but it is possible in principle to write a refined compiler that would
1188 ensure no such leak happens. Additionally, we believe AsmFX could be useful as an intermediate
1189 language, particularly for further translation to *hardware capability architectures* such as CHERI [29].
1190 We conjecture that fine-grained compartmentalisation provided in such architectures could be
1191 used to provide an implementation of effect contexts and handler, continuation, and leave records
1192 ensuring that information can only be accessed by the code that owns it (and particularly separation
1193 between code that uses effects and code that implements them). Inspired by the *Cerise* program
1194 logic [12] and its assembler for an idealised capability architecture, we plan to eventually provide
1195 an effectful *CeriseFX* to explore the relationship between handlers and machine-level capabilities.

1196 WASMFX [25] extends WebAssembly (Wasm) [15], a portable bytecode-driven stack machine,
1197 with support for effect handlers. AsmFX and WASMFX have quite different purposes. WASMFX
1198 builds on the Wasm stack model, incorporating a range of existing features such as functions
1199 and exceptions, whereas AsmFX strives to be agnostic to the underlying implementation of its
1200 primitives. OCaml 5 [28] incorporates a high-performance implementation of effect handlers based
1201 on low-level stack manipulation operations. The intended semantics of OCaml 5 effect handlers
1202 is given by a CEK-like abstract machine much like that of SFX. As with WASMFX, the OCaml 5
1203 implementation is closely tied with underlying features of the host language. C libraries such
1204 as `libseff` [2] and `libmpeff` [21] provide low-level effect handler implementations in C based
1205 respectively on coroutines and on multiprompt delimited control. It would be interesting to try to
1206 more formally relate systems such as WASMFX, OCaml 5, `libseff` and `libmprompt` with AsmFX.

1207 Muhcu et al.[24] implement multi-shot resumptions in the context of *named* handlers [5, 31] and
1208 stack-allocated mutable state. Their work is related to ours in that they too describe a compilation
1209 technique (building on prior work [23]); while we directly compile a high-level language with
1210 standard (unnamed) effect handlers to the effectful AsmFX, their work starts with an intermediate
1211 language with multi-prompt delimited control implementing named handlers (which they call
1212 “lexically-scoped handlers”) and compiles it to an abstract assembly without primitive effect handlers,
1213 but with explicit memory allocation. Unlike AsmFX, they do not support re-entrant resumptions (a
1214 resumption is re-entrant if it resumes itself within its own body). A direct comparison between their
1215 work and AsmFX is challenging: not only are the two works based on different semantics for effect
1216 handlers, but they also have different purposes: AsmFX is concerned primarily with expressiveness
1217 and soundness, whereas [24] is geared towards implementation efficiency. Extending AsmFX to
1218 support named handlers would however be interesting: to achieve that, we could provide variants of
1219 the `resume` and `do` instructions respectively returning or accepting a handler name as an additional
1220 argument; the name may concretely be an index into the effect context (allowing efficient addressing
1221 of a handler frame), but the code should have no access to its representation.

1222 Our proof of correctness of the SFX compiler relies on information on execution history that
1223 needs to be logically preserved step by step by our execution models. The ASFX language we
1224 devised to enrich source programs with annotations providing evidence for such information is
1225

1226 reminiscent of other work in the field of *provenance* [4], studying metadata on the origin, history,
 1227 and derivation of information, and techniques to propagate it through computation [7, 14]. This is
 1228 not accidental: it is easy to imagine that a compiler such as the one we have described could be
 1229 written in a provenance-tracking language (for instance [27]): the compiler would then produce
 1230 object code annotated with provenance information about the source code that yielded it; by
 1231 reflecting that information back into SFX, we would get ASFX. Another related area is *justification*
 1232 *logic* [3], a modal proof system where propositions can carry evidence of their validity. We have
 1233 shown that these related concepts of origin, history, evidence, justification are not only important
 1234 for their epistemological value, but provide an insight on logical constraints on the relationship
 1235 between the input and output of a computation (in our case, source code and object code), which
 1236 can be useful to validate the correctness of an algorithm.

1237

1238

References

- 1239 [1] Danel Ahman and Andrej Bauer. 2020. Runners in Action. In *Programming Languages and Systems*, Peter Müller (Ed.).
 1240 Springer International Publishing, Cham, 29–55. doi:10.1007/978-3-030-44914-8_2
- 1241 [2] Mario Alvarez-Picallo, Teodor Freund, Dan R. Ghica, and Sam Lindley. 2024. Effect Handlers for C via Coroutines.
 1242 *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 358 (Oct. 2024), 28 pages. doi:10.1145/3689798
- 1243 [3] Sergei Artemov. 2008. Justification Logic. In *Proceedings of the 11th European Conference on Logics in Artificial Intelligence*
 1244 (Dresden, Germany) (*JELIA '08*). Springer-Verlag, Berlin, Heidelberg, 1–4. doi:10.1007/978-3-540-87803-2_1
- 1245 [4] David A. Bearman and Richard H. Lytle. 1985. The Power of the Principle of Provenance. *Archivaria* 21 (Jan. 1985),
 1246 14–27. <https://archivaria.ca/index.php/archivaria/article/view/11231>
- 1247 [5] Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Binders by day, labels by night: effect
 1248 instances via lexically scoped handlers. *Proc. ACM Program. Lang.* 4, POPL, Article 48 (Dec. 2019), 29 pages. doi:10.
 1249 1145/3371116
- 1250 [6] Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit
 1251 Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. 2019. Pyro: Deep Universal Probabilistic Programming. *J.*
 1252 *Mach. Learn. Res.* 20 (2019), 28:1–6.
- 1253 [7] James Cheney, Laura Chiticariu, and Wang-Chiew Tan. 2009. Provenance in Databases: Why, How, and Where.
 1254 *Foundations and Trends® in Databases* 1, 4 (2009), 379–474. doi:10.1561/1900000006
- 1255 [8] Intel Corporation. 2019. *Control-Flow Enforcement Technology Specification*. Document Number 334525-003, Revision 3.
- 1256 [9] Stephen Dolan, Spiros Eliopoulos, Daniel Hillerström, Anil Madhavapeddy, KC Sivaramakrishnan, and Leo White.
 1257 2017. Concurrent systems programming with effect handlers. In *TFP*.
- 1258 [10] Stephen Dolan, Leo White, KC Sivaramakrishnan, Jeremy Yallop, and Anil Madhavapeddy. 2015. Effective Concurrency
 1259 through Algebraic Effects. OCaml. http://kcsrk.info/papers/effects_ocaml15.pdf
- 1260 [11] Matthias Felleisen and Daniel P. Friedman. 1986. Control operators, the SECD-machine, and the λ -calculus. In *Formal*
 1261 *Description of Programming Concepts III*. North-Holland, Amsterdam, 193–217.
- 1262 [12] Aina Georges, Armaël Guéneau, Thomas Van Strydonck, Amin Timany, Alix Trieu, Dominique Devriese, and Lars
 1263 Birkedal. 2023. Cerise: Program Verification on a Capability Machine in the Presence of Untrusted Code. *J. ACM* 71
 1264 (09 2023). doi:10.1145/3623510
- 1265 [13] Dan Ghica, Sam Lindley, Marcos Maroñas Bravo, and Maciej Piróg. 2022. High-level effect handlers in C++. *Proc.*
 1266 *ACM Program. Lang.* 6, OOPSLA2, Article 183 (Oct. 2022), 29 pages. doi:10.1145/3563445
- 1267 [14] Todd J. Green, Grigoris Karvounarakis, and Val Tannen. 2007. Provenance semirings. In *Proceedings of the Twenty-Sixth*
 1268 *ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (Beijing, China) (*PODS '07*). Association
 1269 for Computing Machinery, New York, NY, USA, 31–40. doi:10.1145/1265530.1265535
- 1270 [15] Andreas Haas, Andreas Rossberg, Derek Schuff, Ben Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon Zakai,
 1271 and J.F. Bastien. 2017. Bringing the Web up to Speed with WebAssembly. *SIGPLAN Notices* 52, 6 (June 2017), 185–200.
- 1272 [16] Daniel Hillerström. 2015. *Handlers for Algebraic Effects in Links*. Master’s thesis. University of Edinburgh, Scotland.
 1273 http://project-archive.inf.ed.ac.uk/msc/20150206/msc_proj.pdf
- 1274 [17] Daniel Hillerström and Sam Lindley. 2016. Liberating effects with rows and handlers. In *Proceedings of the 1st*
 1275 *International Workshop on Type-Driven Development, TyDe@ICFP 2016, Nara, Japan, September 18, 2016*, James Chapman
 1276 and Wouter Swierstra (Eds.). ACM, 15–27. doi:10.1145/2976022.2976033
- 1277 [18] Daniel Hillerström, Sam Lindley, and Robert Atkey. 2020. Effect handlers via generalised continuations. *J. Funct.*
 1278 *Program.* 30 (2020), e5. doi:10.1017/S0956796820000040
- 1279 [19] Ohad Kammar, Sam Lindley, and Nicolas Oury. 2013. Handlers in action. In *ACM SIGPLAN International Conference on*
 1280 *Functional Programming, ICFP'13, Boston, MA, USA - September 25 - 27, 2013*, Greg Morrisett and Tarmo Uustalu (Eds.).

- 1275 ACM, 145–158. doi:10.1145/2500365.2500590
- 1276 [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
- 1277 [21] Daan Leijen and KC Sivaramakrishnan. 2023. libmprompt and libmpeff. <https://github.com/koka-lang/libmprompt>.
- 1278 [22] Arm Ltd. 2025. *Arm Architecture Reference Manual for A-profile Architecture*. ARM DDI 0487, Version M.a.a. <https://developer.arm.com/documentation/ddi0487/maa>
- 1279 [23] 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
- 1280 [24] Serkan Muhcu, 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
- 1281 [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
- 1282 [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
- 1283 [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
- 1284 [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
- 1285 [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
- 1286 [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
- 1287 [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

1303 A An AsmFX Compiler for SFX

1304 We give here a detailed definition of the SFX compiler that we only sketched in Section 5. The
1305 compiler consists of three procedures, for values, computations, and handlers.

1306 *Compiling Computations.* The table in Figure 3 describes the compilation rules for each possible
1307 syntactic form of computations. Before discussing some relevant cases, we note the following
1308 preliminary facts:

- 1309 • We assume that every SFX primitive operation \oplus is correctly implemented by an AsmFX
1310 instruction $[\oplus]$, taking as many register arguments as needed to contain its input, and
1311 similarly storing its output in an a suitable number of output registers.
- 1312 • Booleans are represented as integers in the obvious way.
- 1313 • Instructions j *target_code* and bz *target_code*, $\$r$ representing an unconditional jump to
1314 *target_code* and a branch to *target_code* if $\$r$ is zero are available.
- 1315 • The operation $\varphi(n, \mathbb{L})$ produces n labels fresh with respect to the used label set \mathbb{L} . The
1316 concrete choice of labels is irrelevant and consistent renaming is assumed to denote the
1317 same code (we use labels for the sake of readability, but each label must be considered
1318 notation for an absolute memory index).

1319 We discuss the compilation rules in detail. To compile a **return**, we simply compile the returned
1320 value. A primitive operation $\oplus(V)$ is compiled by compiling V first, then executing the instruction

1324 $[\oplus]$ implementing \oplus using the $\overline{\$a}$ registers both as input and as output. Conditionals $\text{if}(V, M, N)$
 1325 require us to compile V first and then use a `bz` “branch if zero” instruction to discriminate between
 1326 true and false values: in one case we continue with the code for M , in the other case we branch to
 1327 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
 1328 compiled by compiling M first: after executing the compiled code for M , we will copy the result
 1329 of that computation from the $\overline{\$a}$ register where it has been stored into the $\overline{\$x}$ registers, where N
 1330 expects to find it; finally we compile N . The compilation rule for `unpack` simply needs compile its
 1331 value argument and then move the result of its evaluation to the registers for the x^k variables (the
 1332 rule makes use of an extended vector notation $\overline{\$x_{m,n}}$ which stands for n registers starting from $\$x_m$,
 1333 i.e. $\$x_m, \$x_{m+1}, \dots, \$x_{m+n-1}$.

1334 Now we consider computations for the effectful sublanguage. The U `with` $h(V)$ computation is
 1335 compiled by compiling V first and storing the result of its evaluation in an initial segment of the $\overline{\$}$
 1336 registers; we use these registers together with the handler specification $[h]$ to load the handler
 1337 $h(V)$ into register $\$hr$; then we compile U : its value will be stored in $\$a_0, \a_1 ; finally, we replace
 1338 the content of $\$a_1$ with the content of $\$hr$.

1339 To create a continuation from a function by means of `newbroom`(f) `with` hV , we need to create
 1340 a suitable broom and place it in register $\$a_0, \a_1 . We first compile V , place its content in the $\overline{\$}$
 1341 registers, and use it together with $[h]$ to create a handler which will be loaded into $\$hr$ first and
 1342 then $\$a_1$. Then we construct a continuation pointing to a label l_{res} which will contain code to invoke
 1343 f : the continuation is initially stored in $\$kr$, then moved into $\$a_0$. Finally we jump to l_{end} , the end
 1344 of the code. At location l_{res} , we to call f we actually need to invoke its return clause: this is done by
 1345 jumping unconditionally to the f_{return} address, after installing in $\$lr$ a leave record indicating what
 1346 code should be executed when f terminates — f is executed within the handler $h(V)$, therefore
 1347 at the end of its execution it should invoke the return clause for $h(V)$: that corresponds to the
 1348 instruction `return`, at the label l_{exit} .

1349 To invoke an effect by means of `do #t`(V), we first compile V : this places the result of its evaluation
 1350 into the $\overline{\$a}$, where the operation handler expects to find it. To invoke the operation handler, we
 1351 just execute the instruction `do #t, save :I` (saving all the input registers will allow the resumption
 1352 for the $\#t$ operation to restore the original client environment).

1353 To execute a computation under a new handler by means of `handle` M `with` $h(V)$, we compile V ,
 1354 copy its evaluation in the $\overline{\$s}$ registers, and use it together with $[h]$ to create a handler and place it
 1355 in $\$hr$. We also create a continuation pointing to the code for M (l_{cli}) and place it in $\$kr$, and a leave
 1356 record for the code to be executed after this computation (l_{exit}), with saved registers corresponding
 1357 to the input set I . After these three registers have been loaded, we invoke the client by resuming
 1358 the continuation we created, using the instruction `resume`. Notice that the client code at location
 1359 l_{cli} is terminated by a `return` instruction, so that the result of M will be delivered to the return
 1360 clause of the newly installed handler: we call this a `return` epilogue. Tracking epilogues is essential
 1361 to prove the correctness of the compiler.

1362 A `resume` $U(V)$ works similarly: registers $\$kr$ and $\$hr$ are filled with the result of the evaluation
 1363 of U ; the result of the evaluation of V is put in registers $\overline{\$a}$; we create a leave record for the code
 1364 l_{exit} to be run after executing this computation (again, the input registers I must be added to the
 1365 save area of $\$lr$). Finally, we invoke the `resume` instruction.

1366
 1367 *Compiling Handlers.* The procedure to compile handlers is detailed in Figure 6. The return and
 1368 operation clauses of each handler are compiled separately using the same algorithm defined for
 1369 computations. In particular, for each handler h with clauses for operations $\#t_i$ we generate unique
 1370 global labels $h_{return}, h_{\#i}$ for the code implementing the operations. Each parameterised handler
 1371
 1372

1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403
 1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421

$$\left[\text{handler } h(x) \left\{ \begin{array}{l} \text{return } y \rightarrow N_{\text{return}}^h \\ \#t_i(z^i) \rightarrow (r^i).N_{\#t_i}^h \\ \dots \end{array} \right\} \right]_{\mathbb{L}} \Rightarrow \mathbb{L}' = \left\{ \begin{array}{l} \text{let } h_{\text{return}}, \overline{h_{\#t}} = \varphi(1 + \|\overline{\#t}\|) \\ \text{let } \mathbb{L}_{\text{return}} = \mathbb{L} \cup \{h_{\text{return}}, \overline{h_{\#t}}\} \\ \text{let } \mathbb{I}_x = \overline{\$x_{\|x\|}} \\ h_{\text{return}} : \\ \quad \text{let } \mathbb{I}_{N_{\text{return}}^h} = \overline{\$y_{\|y\|}} \\ \quad \mathbb{I}_x \leftarrow \overline{\$s_{\|x\|}} \\ \quad \mathbb{I}_{N_{\text{return}}^h} \leftarrow \overline{\$a_{\|y\|}} \\ \quad [N_{\text{return}}^h]_{\mathbb{I}_x \cup \mathbb{I}_{N_{\text{return}}^h} \cup \{\$1r, \mathbb{L}_{\text{return}}\}} \Rightarrow \mathbb{L}_{\#t_1} \\ \quad \text{exit} \\ h_{\#t_1} : \\ \quad \text{let } \mathbb{I}_{N_{\#t_1}^h} = \overline{\$z^i_{\|z\|}, \$r_1^i, \$r_2^i} \\ \quad \mathbb{I}_x \leftarrow \overline{\$s_{\|x\|}} \\ \quad \mathbb{I}_{N_{\#t_1}^h} \leftarrow \overline{\$a_{\|z\|}, \$kr, \$hr} \\ \quad [N_{\#t_1}^h]_{\mathbb{I}_x \cup \mathbb{I}_{N_{\#t_1}^h} \cup \{\$1r, \mathbb{L}_{\#t_1}\}} \Rightarrow \mathbb{L}_{\#t_2} \\ \quad \text{exit} \\ \vdots \\ \mathbb{L}' := \mathbb{L}_{\#t_{n+1}} \end{array} \right.$$

Fig. 6. Compilation of SFX Handlers to AsmFX

with x as its formal parameter will receive its state in an initial segment of the state registers $\overline{\$s}$ of suitable size, and will allocate variable registers $\overline{\$x}$ 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 $h_{\#t_i}$ receive their arguments in an initial segment of the argument registers $\overline{\$a}$ of suitable size: those are also copied into variable registers $\overline{\$y}, \overline{\z^i} 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 $\$1r$, 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 $\overline{\$a}$. 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 $\overline{\$a}$ registers.

In fact, in realistic scenarios, we are also interested in proving the correctness of partial executions to guarantee that even non-terminating programs interact correctly with the environment. This

requires us 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.

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

Fig. 7. ASFX: Syntax

The syntax of ASFX is shown in Figure 7. Augmented values annotate SFX values with a single location referencing the code responsible for loading the value in the appropriate registers (an instance of $[V]$). 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, $\mu(\mathbf{return} V@l) = \mu(V@l) = l$). We also define the end location of a value or computation as the

compiled length of that value or computation plus its start location:

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

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

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 $\mu(\mathbf{0}@l) = l$. 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 8. 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 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.

	control	env.	continuation	\hookrightarrow	control	env.	continuation
1520	$\oplus(\hat{V})$	\hat{y}	\hat{k}		$\llbracket \oplus \rrbracket (\llbracket \hat{V} \rrbracket \hat{y})$	\hat{y}	\hat{k}
1521	if $(\hat{V}, \hat{M}, \hat{N})$	\hat{y}	\hat{k}		\hat{M}	\hat{y}	jump @ $v(\hat{M})$
1522	where:						$\cdot \hat{k}$
1523	$\llbracket \hat{V} \rrbracket \hat{y} = \text{true}$						
1524	if $(\hat{V}, \hat{M}, \hat{N})$	\hat{y}	\hat{k}		\hat{N}	\hat{y}	\hat{k}
1525	where:						
1526	$\llbracket \hat{V} \rrbracket \hat{y} = \text{false}$						
1527	let $x \leftarrow \hat{M}$	\hat{y}	\hat{k}		\hat{M}	\hat{y}	$((x).\hat{N})@v(\hat{M})$
1528	in \hat{N}						$\cdot \hat{k}$
1529	unpack						
1530	$\overline{x_n} \leftarrow \langle \hat{V}_n \rangle$	\hat{y}	\hat{k}		\hat{N}	$(x_n \mapsto \llbracket \hat{V}_n \rrbracket \hat{y})$	\hat{k}
1531	in \hat{N}					$\cdot \hat{y}$	
1532	handle \hat{M}	\hat{y}	\hat{k}		\hat{M}	$\hat{y} \setminus \text{LEAVE}$	$h(v, \mathcal{R}_L)@v(\hat{M}) \cdot \hat{k}$
1533	with $h(\hat{V})$						where: $v = \llbracket \hat{V} \rrbracket \hat{y}$ $\mathcal{R}_L = \mathcal{L}(\mu(\hat{k}), \hat{y})$
1534						$(\text{LEAVE} \mapsto \mathcal{R}_L)$	
1535						$\cdot (\text{res}(\#t) \mapsto$	
1536						$(\hat{k}^{-t}, h^{+t}(\hat{w})@l, \hat{y}))$	exit @ $v(h_{\#t}) \cdot \hat{k}'$
1537	do $\#t(\hat{V})$	\hat{y}	$\hat{k}^{-t} \cdot h^{+t}(\hat{w}, \mathcal{R}_L)@l \cdot \hat{k}'$		$h_{\#t}$	$\cdot (\text{dom}(\#t) \mapsto \llbracket \hat{V} \rrbracket \hat{y})$	
1538						$\cdot (\text{dom}(h) \mapsto \hat{w})$	
1539	resume $\hat{U}(\hat{V})$	\hat{y}	\hat{k}		$\llbracket \hat{V} \rrbracket \hat{y}$	\hat{y}'	\hat{k}'
1540	where:						$\cdot h(\hat{w}, \mathcal{R}_L)@l$
1541	$\llbracket \hat{U} \rrbracket \hat{y} =$						$\cdot \hat{k}$
1542	$(\hat{k}', h(\hat{w})@l, \hat{y}')$						where: $\mathcal{R}_L = \mathcal{L}(\mu(\hat{k}), \hat{y})$
1543	\hat{U} with $h(\hat{W})$	\hat{y}	\hat{k}		$(\hat{k}',$	\hat{y}	\hat{k}
1544	where:				$h(\llbracket \hat{W} \rrbracket \hat{y})@l,$		
1545	$\llbracket \hat{U} \rrbracket \hat{y} =$				$\hat{y}')$		
1546	$(\hat{k}', _ @l, \hat{y}')$						
1547	newbroom @ $l(f)$	\hat{y}	\hat{k}		$(\text{leave}@l$	\hat{y}	\hat{k}
1548	with $h(\hat{V})$				jump @ l'		
1549	where:				$\cdot ((x).\text{return})@l_f$		
1550	$\text{dom}(f) = x$				exit @ $l'_f,$		
1551	$l_f = \mu(f_{\text{return}})$				$h(\llbracket \hat{V} \rrbracket \hat{y})@l'',$		
1552	$l'_f = v(f_{\text{return}})$				$[\]$		
1553	$l'_f = l + 1$						
1554	$l'' = l + 2$						
1555	return \hat{V}	\hat{y}	\hat{k}		$\llbracket \hat{V} \rrbracket \hat{y}$	\hat{y}	\hat{k}
1556	\hat{v}	\hat{y}	jump @ $l \cdot \hat{k}$		\hat{v}	\hat{y}	\hat{k}
1557							
1558	\hat{v}	\hat{y}	exit @ $l \cdot \hat{k}$		\hat{v}	$\hat{y}(\text{LEAVE}).\text{env}$	\hat{k}
1559							where: $\mu(\hat{k})$
1560							$= \hat{y}(\text{LEAVE}).\text{loc}$
1561	\hat{v}	\hat{y}	leave @ $l \cdot \hat{k}$		\hat{v}	$(\text{LEAVE} \mapsto \mathcal{L}(l+2, [\]))$	\hat{k}
1562						$\cdot \hat{y}$	
1563	\hat{v}	\hat{y}	$((x).\hat{N})@l \cdot \hat{k}$		\hat{N}	$(x \mapsto \hat{v}) \cdot \hat{y}$	\hat{k}
1564							
1565	\hat{v}	\hat{y}	$h(\hat{w}, \mathcal{R}_L)@l \cdot \hat{k}$		h_{return}	$(\text{LEAVE} \mapsto \mathcal{R}_L)$	exit @ $v(h_{\text{return}}) \cdot \hat{k}$
1566						$\cdot (\text{dom}(h_{\text{return}}) \mapsto \hat{v})$	
1567						$\cdot (\text{dom}(h) \mapsto \hat{w})$	

Fig. 8. ASFX: Abstract Machine

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{y} , 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 ASAFX 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@ ℓ (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 ASAFX 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 ASAFX `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 ASAFX. Since ASAFX 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 ASAFX 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 ASAFX configuration \hat{s} , we write $s \preccurlyeq \hat{s}$. 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, $\langle M, \gamma, \kappa \rangle \preccurlyeq \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle$ 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 $\langle v, \gamma, \kappa \rangle \preccurlyeq \langle \hat{v}, \hat{\gamma}, \hat{\kappa} \rangle$ if and

1618 only if the augmented value \hat{v} simulates the simple value v , and the pair of augmented environment
 1619 and continuation $(\hat{\gamma}, \hat{\kappa})$ simulates γ (however γ need not be simulated by $\hat{\gamma}$). It may seem odd that
 1620 the continuation simulation requires, on the augmented side, an environment: this is explained by
 1621 the fact that SFX pure frames are closed with respect to an environment, while ASFX frames are
 1622 not.

1623 The value simulation $v \leq \hat{v}$, for non-broom values, simply states that erasure of the augmented
 1624 value results in the simple value; however, broom simulation $(\kappa, h(w)) \leq (\hat{\kappa}, h(\hat{w}))@l, \hat{\gamma}_\kappa$ is recur-
 1625 sively defined: it requires that κ be simulated by $(\hat{\gamma}_\kappa, \hat{\kappa})$ and that w be simulated by \hat{w} .

1626 The environment simulation $\gamma \leq \hat{\gamma}$ is defined pointwise for all variables in the domain of γ ; $\hat{\gamma}$ is
 1627 allowed to be defined on more variables.

1628 The continuation simulation $\kappa \leq (\hat{\gamma}, \hat{\kappa})$ is defined by means of an inductive judgement, with a
 1629 base rule for the identity continuation and recursive rules for each type of frame. A **jump** epilogue
 1630 frame preserves the simulation; an **exit** frame swaps the active augmented environment with the
 1631 one contained in the LEAVE pseudovalue (which must be defined). A **leave** epilogue requires the
 1632 next continuation $\hat{\kappa}$ to simulate the original plain continuation κ in an environment with an updated
 1633 leave record. In the case of handler frames, both continuations must start with matching handler
 1634 frames: if that is the case, the simulation holds if there is a simulation between the remaining parts
 1635 of the two continuations; however on the ASFX side of the simulation, we will need to swap the
 1636 environment with the one contained in the leave record of the handler frame, reflecting the fact
 1637 that when we exit a handler, the environment is restored from that leave record.

1638 Finally, the most interesting case is that of pure frames. If both continuations start with a pure
 1639 frame, in order for the simulation to hold, the computations N and \hat{N} in each pure frame must
 1640 match (meaning that by erasing \hat{N} we get N). Furthermore, the SFX pure frame is closed by an
 1641 environment γ_N : that environment must be simulated by the ASFX environment $\hat{\gamma}$, reflecting the
 1642 fact that in ASFX we do not need an environment in the pure frame, because the active environment
 1643 is already the correct one. If these conditions hold, the simulation is valid if there is a simulation
 1644 between the two remaining parts of the two continuations; however on the ASFX side we need to
 1645 extend the environment with a value for the variable x expected by the pure continuation, and the
 1646 simulation must hold independently of the value chosen.

1647 We now move to proving the simulation theorem.

1648 *B.3.1 Proof of Theorem 6.1. For all SFX configurations s, s' such that $s \hookrightarrow s'$, for all ASFX configura-*
 1649 *tions \hat{s} such that $s \leq \hat{s}$, there exists an ASFX configuration \hat{s}' such that $s' \leq \hat{s}'$ and $\hat{s} \hookrightarrow \hat{s}'$.*

1650 In the proof, we will require the following auxiliary results.

1651 LEMMA B.1. *Suppose $\kappa \leq (\hat{\gamma}, \hat{\kappa})$ and $\hat{\gamma} \subseteq \hat{\gamma}'$: then we have $\kappa \leq (\hat{\gamma}', \hat{\kappa})$.*

1652 PROOF. By induction on the well-scopedness judgement. □

1653 LEMMA B.2. *If $\gamma \leq \hat{\gamma}$, then $\gamma \leq \hat{\gamma} \setminus \text{LEAVE}$.*

1654 PROOF. Trivial by unfolding the definitions. □

1655 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 $\kappa_1 \cdot h(w) \cdot \kappa_2 \leq$
 1660 $(\hat{\gamma}_1, \kappa_1 \cdot h(w, \mathcal{L}(\mu(\hat{\kappa}_2), \hat{\gamma}_2)))@l \cdot \hat{\kappa}_2$.*

1661 PROOF. By induction on $\kappa_1 \leq (\hat{\gamma}_1, \hat{\kappa}_1)$. In the base case $\kappa_1 = 0$ and $\hat{\kappa}_1 = 0@l$, we use the hypothesis
 1662 $\kappa_2 \leq (\hat{\gamma}_2, \hat{\kappa}_2)$ to prove the thesis in the form $h(w) \cdot \kappa_2 \leq (\hat{\gamma}_1, h(w, \mathcal{L}(\mu(\hat{\kappa}_2), \hat{\gamma}_2)))@l \cdot \hat{\kappa}_2$. □

1663 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
 1665 transition $\langle \hat{v}, \hat{\gamma}, \varepsilon \cdot \hat{\kappa} \rangle \hookrightarrow \langle \hat{v}, \hat{\gamma}^*, \hat{\kappa} \rangle$.*

1666

1667 **PROOF.** By case analysis on $\kappa \leq (\hat{\gamma}, \varepsilon \cdot \hat{\kappa})$. □

1668

1669 Now we prove the main result.

1670

1671 **PROOF OF THEOREM 6.1.** We proceed by case analysis on $s \hookrightarrow s'$ and $s \leq \hat{s}$, where $s = \langle M, \gamma, \kappa \rangle$
 1672 and $s' = \langle M', \gamma', \kappa' \rangle$. For each case, the SFX transition forces s and s' to have a certain shape, and
 1673 similarly the source simulation forces \hat{s} to have a shape that adapts to s . We will (deterministically)
 1674 follow the ASFX machine on \hat{s} for one or more steps, obtaining a target configuration \hat{s}' : then we
 1675 will need to show that $s' \leq \hat{s}'$.

1676

1677 $M = \oplus(V)$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle \llbracket \oplus \rrbracket (\llbracket V \rrbracket \gamma), \gamma, M \rangle$.

1678

1679 We prove $\hat{s} = \langle \oplus(\hat{V})@l, \hat{\gamma}, \hat{\kappa} \rangle$ 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 $s' \leq \hat{s}'$.

1680

1681 $M = \text{if}(V, M_1, M_2)$. This runs differently depending on whether $\llbracket V \rrbracket \gamma = \text{true}$ or **false**. First, in
 1682 both cases, we prove $\hat{s} = \langle \text{if}(\hat{V}, \hat{M}_1, \hat{M}_2)@l, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}| = V, |\hat{M}_1| = M_1, |\hat{M}_2| = M_2, \gamma \leq \hat{\gamma}$, and
 1683 $\kappa \leq (\hat{\gamma}, \hat{\kappa})$.

1684

1685 Then, if $\llbracket V \rrbracket \gamma = \text{true}$, we have $s' = \langle M', \gamma', \kappa' \rangle = \langle M_1, \gamma, \kappa \rangle$. We choose $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle =$
 1686 $\langle \hat{M}_1, \hat{\gamma}, \text{jump}@v(\hat{M}_1) \cdot \hat{\kappa} \rangle$ and we prove that $\kappa \leq (\gamma', \kappa')$ by applying the appropriate rule in Figure 4,
 1687 to obtain $s' \leq \hat{s}'$.

1688

1689 If instead $\llbracket V \rrbracket \gamma = \text{false}$, we have $s' = \langle M', \gamma', \kappa' \rangle = \langle M_2, \gamma, \kappa \rangle$. We choose $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle =$
 1690 $\langle \hat{M}_2, \hat{\gamma}, \hat{\kappa} \rangle$ and we obtain $s' \leq \hat{s}'$ trivially.

1691

1692 $M = \text{let } x \leftarrow M_1 \text{ in } M_2$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_1, \gamma, (x).M_2[\gamma] \cdot \kappa \rangle$.

1693

1694 We prove $\hat{s} = \langle (\text{let } x \leftarrow \hat{M}_1 \text{ in } \hat{M}_2)@l, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{M}_1| = M_1, |\hat{M}_2| = M_2, \gamma \leq \hat{\gamma}$, and $\kappa \leq$
 1695 $(\hat{\gamma}, \hat{\kappa})$. We choose $\hat{s}' = \langle \hat{M}_1, \hat{\gamma}, ((x).\hat{M}_2)@v(\hat{M}_1) \cdot \hat{\kappa} \rangle$ and we prove $|\hat{M}_2| = M_2$ and $(x).M_2[\gamma] \cdot \kappa \leq$
 1696 $(\hat{\gamma}, ((x).\hat{M}_2)@v(\hat{M}_1) \cdot \hat{\kappa})$ (the latter uses Lemma B.1). This immediately implies $s' \leq \hat{s}'$.

1697

1698 $M = \text{unpack } \overline{x_n} \leftarrow \overline{V_n} \text{ in } M_0$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_0, \overline{(x_n \mapsto \llbracket V_n \rrbracket \gamma)} \cdot \gamma, \kappa \rangle$.

1699

1700 We prove $\hat{s} = \langle \text{unpack } \overline{x_n} \leftarrow \overline{\hat{V}_n} \text{ in } \hat{M}_0@l, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}_n| = \overline{V_n}, |\hat{M}_0| = M_0, \gamma \leq \hat{\gamma}$, and
 1701 $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. 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
 1702 latter uses Lemma B.1, knowing that $\hat{\gamma} \subseteq \hat{\gamma}'$). This immediately implies $s' \leq \hat{s}'$.

1703

1704 $M = \text{handle } M_0 \text{ with } h(V)$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_0, \gamma, h(\llbracket V \rrbracket \gamma) \cdot \kappa \rangle$.

1705

1706 We prove $\hat{s} = \langle \text{handle } \hat{M}_0 \text{ with } h(\hat{V})@l, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}| = V, |\hat{M}_0| = M_0, \gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$.
 1707 We choose $\hat{s}' = \langle \hat{M}_0, \hat{\gamma} \setminus \text{LEAVE}, h(\llbracket \hat{V} \rrbracket \hat{\gamma}, \mathcal{L}(\mu(\hat{\kappa}), \hat{\gamma}))@v(\hat{M}_0) \cdot \hat{\kappa} \rangle$ and we prove $\gamma \leq \hat{\gamma} \setminus \text{LEAVE}$ (by
 1708 Lemma B.2), $\kappa \leq (\hat{\gamma} \setminus \text{LEAVE}, h(\llbracket \hat{V} \rrbracket \hat{\gamma}, \mathcal{L}(\mu(\hat{\kappa}), \hat{\gamma}))@v(\hat{M}_0) \cdot \hat{\kappa})$ (by the appropriate rule in Figure 4).
 1709 This immediately implies $s' \leq \hat{s}'$.

1710

1711 $M = \text{do } \#t(V)$. In this case, we have:

1712

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

1713

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

1714

1715 We prove $\hat{s} = \langle (\text{do } \#t(\hat{V}))@l, \hat{\gamma}, \hat{\kappa}_1^{-t} \cdot h^{+t}(\hat{w}, \mathcal{L}(\ell_2, \hat{\gamma}_2))@l_h \cdot \hat{\kappa}' \rangle$ where $|\hat{V}| = V, \kappa_1^{-t} \leq (\hat{\gamma}, \hat{\kappa}_1^{-t})$,
 1716 $\kappa_2 \leq (\hat{\gamma}_2, \hat{\kappa}_2)$. We choose $\hat{\gamma}' = (\text{LEAVE} \mapsto \mathcal{L}(\ell_2, \hat{\gamma}_2)) \cdot (\text{res}(h_{\#t}) \mapsto (\hat{\kappa}_1^{-t}, h(\hat{w})@l_h, \hat{\gamma})) \cdot (\text{dom}(h) \mapsto$
 1717 $\hat{w}) \cdot (\text{dom}(h_{\#t}) \mapsto \llbracket \hat{V} \rrbracket \hat{\gamma})$, $\hat{s}' = \langle \hat{M}_{\#t}^h, \hat{\gamma}', \text{exit}@v(\hat{M}_{\#t}^h) \cdot \hat{\kappa}_2 \rangle$ and we prove $|\hat{M}_{\#t}^h| = M_{\#t}^h$ (by the
 1718 definition of the compilation procedure), $\gamma' \leq \hat{\gamma}'$, and $\gamma' \leq (\hat{\gamma}', \text{exit}@v(\hat{M}_{\#t}^h) \cdot \hat{\kappa}_2)$ (by applying the
 1719 appropriate rule in Figure 4). This immediately implies $s' \leq \hat{s}'$.

1720

1721

1722

1716 $M = \mathbf{resume} U(V)$ and $\llbracket U \rrbracket \gamma = (\kappa_U, h(w))$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle \llbracket V \rrbracket \gamma, \gamma, \kappa_U \cdot h(w) \cdot \kappa \rangle$.

1717 We prove $\hat{s} = \langle \mathbf{resume} \hat{U}(\hat{V}) @ \ell, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{U}| = U$, $|\hat{V}| = V$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. Note that in
1718 order for $\llbracket U \rrbracket \gamma$ to evaluate to a continuation value, it must be that $U = x_U$, thus $\llbracket U \rrbracket \gamma = \gamma(x_U)$. Then
1719 $\hat{U} = x_U$ and $\llbracket \hat{U} \rrbracket \hat{\gamma} = \hat{\gamma}(x_U)$ and, by applying the hypotheses, we prove $\hat{\gamma}(x_U) = (\hat{\kappa}_U, h(\hat{w})) @ \ell_U, \hat{\gamma}_U$
1720 such that $\kappa_U \leq (\hat{\gamma}_U, \hat{\kappa}_U)$ and $w \leq \hat{w}$.

1721 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 $\kappa_U \cdot h(w) \cdot \kappa \leq$
1722 $(\hat{\gamma}_U, \hat{\kappa}_U \cdot h(\hat{w}, \mathcal{L}(\mu(\hat{\kappa}), \hat{\gamma})) @ \ell_U \cdot \hat{\kappa})$.

1723 Furthermore, we note $\llbracket V \rrbracket \gamma \leq \llbracket \hat{V} \rrbracket \hat{\gamma}$. This is all we need to prove $s' \leq \hat{s}'$.

1724
1725 $M = U$ with $h(W)$ and $\llbracket U \rrbracket \gamma = (\kappa_U, _)$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle (\kappa_U, h(\llbracket W \rrbracket \gamma)), \gamma, \kappa \rangle$.

1726 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 $\kappa \leq (\hat{\gamma}, \hat{\kappa})$.
1727 Note that in order for $\llbracket U \rrbracket \gamma$ to evaluate to a continuation value, it must be that $U = x_U$, thus
1728 $\llbracket U \rrbracket \gamma = \gamma(x_U)$. Then $\hat{U} = x_U$ and $\llbracket \hat{U} \rrbracket \hat{\gamma} = \hat{\gamma}(x_U)$ and, by applying the hypotheses, we prove
1729 $\hat{\gamma}(x_U) = (\hat{\kappa}_U, h(\hat{w})) @ \ell_U, \hat{\gamma}_U$ such that $\kappa_U \leq (\hat{\gamma}_U, \hat{\kappa}_U)$ and $w \leq \hat{w}$.

1730 Choose $\hat{s}' = \langle (\hat{\kappa}_U, h(\hat{w}), \hat{\gamma}_U), \hat{\gamma}, \hat{\kappa} \rangle$. Clearly, $\hat{s} \hookrightarrow \hat{s}'$; then prove $s' \leq \hat{s}'$ by unfolding the definition.

1731
1732 $M = \mathbf{newbroom}(f)$ with $h(V)$ where $\text{dom}(f_{\text{return}}) = x$. In this case, $s' = \langle M', \gamma', \kappa' \rangle$
1733 $= \langle ((x).f_{\text{return}} \llbracket \cdot \rrbracket), h(\llbracket V \rrbracket \gamma), \gamma, \kappa \rangle$.

1734 We prove $\hat{s} = \langle \mathbf{newbroom} @ \ell(f) \text{ with } h(\hat{V}), \hat{\gamma}, \hat{\kappa} \rangle$, where $|\hat{V}| = V$; then we choose
1735 $\hat{s}' = \langle (\hat{\kappa}_f, h(\llbracket \hat{V} \rrbracket \hat{\gamma})) @ (\ell + 2), \llbracket \cdot \rrbracket, \hat{\gamma}', \hat{\kappa}' \rangle$ where

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

1736
1737 We prove that $((x).f_{\text{return}} \llbracket \cdot \rrbracket) \cdot h(\llbracket V \rrbracket \gamma) \leq \llbracket \cdot \rrbracket, \hat{\kappa}_f \cdot h(\llbracket \hat{V} \rrbracket \hat{\gamma}) @ (\ell + 2)$ by repeated applications of the
1738 rules defining \leq .

1739
1740
1741 $M = \mathbf{return} V$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle \llbracket V \rrbracket \gamma, \gamma, M \rangle$.

1742 We prove $\hat{s} = \langle \mathbf{return} \hat{V}, \hat{\gamma}, \hat{\kappa} \rangle$ where $|\hat{V}| = V$, $\gamma \leq \hat{\gamma}$, and $\kappa \leq (\hat{\gamma}, \hat{\kappa})$. We choose $\hat{s}' = \langle \llbracket \hat{V} \rrbracket \hat{\gamma}, \hat{\gamma}, \hat{\kappa} \rangle$
1743 and we obtain $s' \leq \hat{s}'$.

1744
1745 $M = v$ and $\kappa = (x).M_0[\gamma_0] \cdot \kappa_0$. In this case, $s' = \langle M', \gamma', \kappa' \rangle = \langle M_0, (x \mapsto v) \cdot \gamma_0, \kappa_0 \rangle$.

1746 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 $v \leq \hat{v}$, $(x).M_0[\gamma_0] \cdot \kappa_0 \leq (\hat{\gamma}, \hat{\kappa}_{\text{epi}} \cdot ((x).\hat{M}_0) @ \ell \cdot \hat{\kappa}_0)$
1747 where $\hat{\kappa}_{\text{epi}}$ consists of epilogue frames only.

1748 By multiple applications of Lemma B.4, we prove that the well-scopedness condition on the contin-
1749 uations implies that there exists an environment $\hat{\gamma}_0$ such that $(x).M_0[\gamma_0] \cdot \kappa_0 \leq (\hat{\gamma}_0, ((x).\hat{M}_0) @ \ell \cdot \hat{\kappa}_0)$
1750 and $\hat{s} \xrightarrow{*} \langle \hat{v}, \hat{\gamma}_0, ((x).\hat{M}_0) @ \ell \cdot \hat{\kappa}_0 \rangle$: by taking one ASFX step more, we find ourselves in configuration
1751 $\langle \hat{M}_0, (x \mapsto \hat{v}) \cdot \hat{\gamma}_0, \hat{\kappa}_0 \rangle$, which we take to be our \hat{s}' . By case analysis on the well-scopedness condition
1752 for $((x).\hat{M}_0) @ \ell \cdot \hat{\kappa}_0$, 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
1753 prove that $s' \leq \hat{s}'$.

1754
1755 $M = v$ 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 $\gamma' = (\text{dom}(h_{\text{return}}) \mapsto$
1756 $v) \cdot (\text{dom}(h) \mapsto w)$.

1757 We prove $\hat{s} = \langle \hat{v}, \hat{\gamma}, \hat{\kappa}_{\text{epi}} \cdot h(w, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{\kappa}_0 \rangle$ such that $\mathcal{R}_{\mathcal{L}} = \mathcal{L}(\mu(\hat{\kappa}_0), \hat{\gamma}_0)$,
1758 $h(w) \cdot \kappa_0 \leq (\hat{\gamma}, \hat{\kappa}_{\text{epi}} \cdot h(w, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{\kappa}_0)$ where $\hat{\kappa}_{\text{epi}}$ consists of epilogue frames only, and $v \leq \hat{v}$.

1759 By multiple applications of Lemma B.4, we prove that the well-scopedness condition on the con-
1760 tinuations implies that there exists an environment $\hat{\gamma}^*$ such that $h(w) \cdot \kappa_0 \leq (\hat{\gamma}^*, h(w, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{\kappa}_0)$
1761 and $\hat{s} \xrightarrow{*} \langle \hat{v}, \hat{\gamma}^*, ((x).\hat{M}_0) @ \ell \cdot \hat{\kappa}_0 \rangle$: by taking one ASFX step more, we find ourselves in configuration
1762 $\langle \hat{M}_{\text{return}}^h, \hat{\gamma}^*, \hat{\kappa}_0 \rangle$, 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
1763 ASFX configuration to be \hat{s}' . By case analysis on the well-scopedness condition for $h(w, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{\kappa}_0$,
1764

1765 we obtain $\kappa_0 \leq (\hat{\gamma}_0, \hat{\kappa}_0)$; we also prove $\gamma' \leq \hat{\gamma}'$ by unfolding the definitions, and $|\hat{M}_{\text{return}}^h| = M_{\text{return}}^h$
 1766 by the definition of compilation. This is enough to prove that $s' \leq \hat{s}'$. \square

1767

1768 B.4 Control Flow Simulation

1769 To complete the proof of correctness of the compiler, we now need to establish a simulation relation
 1770 between ASFX and AsmFX. We consider AsmFX configurations $a = \langle \Xi, \Theta, C \rangle$, whose elements are
 1771 memory stores Ξ (containing the program), register files Θ (finite maps from registers to values),
 1772 and effect contexts C . Since a property of our compiler is that the memory store for a given source
 1773 program is established at compilation time and stays read-only during execution, we will allow
 1774 ourselves to omit it from AsmFX configurations and write $\langle \Theta, C \rangle$ for $\langle \Xi, \Theta, C \rangle$.

1775 A source configuration \hat{s} for a program \hat{P} should correspond to an AsmFX configuration a on
 1776 $\llbracket P \rrbracket$: actually, for every \hat{s} there are multiple such a , because of the plethora of registers available
 1777 in the architecture: some of the registers will not have a counterpart in the source configuration
 1778 and will thus be irrelevant. Therefore, the translation to AsmFX of a source configuration \hat{s} , which
 1779 we denote by $\llbracket \hat{s} \rrbracket$, will be the *set* of AsmFX configurations that are a model for \hat{s} up to relevant
 1780 registers. If we write $a \vDash \hat{s}$ to mean that a models \hat{s} , then we can define $\llbracket \hat{s} \rrbracket := \{a : a \vDash \hat{s}\}$. We define
 1781 what it means to be a model separately for the register file and effect context components: the
 1782 register file must be a model for the whole configuration; on the other hand, the effect context only
 1783 needs to model the continuation.

$$1784 \quad \langle \Theta, C \rangle \vDash \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle := (\Theta \vDash \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle) \wedge (C \vDash \hat{\kappa})$$

1786 We can define what it means for an effect context to model a continuation by means of a direct
 1787 translation that only needs to consider the handler frames:

$$1788 \quad C \vDash \hat{\kappa} \quad := \quad C = \llbracket \hat{\kappa} \rrbracket$$

1789

$$1790 \quad \llbracket \hat{\kappa} \rrbracket \quad := \quad \begin{cases} \llbracket h(w, \mathcal{R}_{\mathcal{L}}) @ \ell \rrbracket \cdot \llbracket \hat{\kappa}' \rrbracket & \text{if } \hat{\kappa} = h(w, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{\kappa}' \\ \llbracket \hat{\kappa}' \rrbracket & \text{if } \hat{\kappa} = \hat{\zeta} \cdot \hat{\kappa}' \text{ and } \hat{\zeta} \neq h(_) \end{cases}$$

1792

$$1793 \quad \llbracket h(w, \mathcal{R}_{\mathcal{L}}) @ \ell \rrbracket \quad := \quad \mathcal{H}(\llbracket h \rrbracket, \overline{\$s} = \llbracket w \rrbracket, \llbracket \mathcal{R}_{\mathcal{L}} \rrbracket)$$

1794

$$1795 \quad \llbracket \mathcal{L}(\ell, \hat{\gamma}) \rrbracket \quad := \quad \mathcal{L}(\ell, \llbracket \hat{\gamma} \rrbracket)$$

1796

$$1797 \quad \llbracket \overline{(x \mapsto v)} \rrbracket \quad := \quad \overline{(\$x = \llbracket v \rrbracket)}$$

1798

1799 A register file models a source configuration if every variable defined in the source environment is
 1800 backed by (any number of) registers containing the AsmFX representation of that variable's value;
 1801 if the program counter matches the address of the beginning of the computation being evaluated
 1802 or, when the computation has ended and we are returning a value, when it matches the address of
 1803 the beginning of the continuation. Note that the pseudovvariable LEAVE is always mapped to the
 1804 register $\$1r$, and every variable x is mapped to as many $\$x$ registers as it is necessary.

$$1806 \quad \Theta \vDash \langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle := \Theta(\$pc) = \mu(\hat{M}) \wedge \forall x \in \text{dom}(\hat{\gamma}). \overline{\Theta(\$x)} = \llbracket \hat{\gamma}(x) \rrbracket$$

$$1807 \quad \Theta \vDash \langle v, \hat{\gamma}, \hat{\kappa} \rangle := \Theta(\$pc) = \mu(\hat{\kappa}) \wedge \forall x \in \text{dom}(\hat{\gamma}). \overline{\Theta(\$x)} = \llbracket \hat{\gamma}(x) \rrbracket$$

1808

1809 The soundness property we seek states that the modelling relation is a simulation: whenever we have
 1810 a source transition $\hat{s} \hookrightarrow \hat{s}'$ and an AsmFX configuration $a \vDash \hat{s}$, then there exists a transition chain
 1811 $a \xrightarrow{*} a'$ where $a' \vDash \hat{s}'$. However, this property cannot be proved for an arbitrary \hat{s} : augmented source
 1812 configurations contain metadata that encodes part of the execution history; while the compiler
 1813

1813

1814 produces correct metadata, and the abstract machine ensures that the metadata is propagated
 1815 correctly, arbitrary configurations could be unsound and not allow the simulation we need for the
 1816 soundness theorem. We thus need to restrict ourselves to *valid* configurations (denoted by $\hat{s}\checkmark$) and
 1817 prove, as part of our theorem, that validity is preserved by transitions.

1818 The validity predicate is defined in Figure 5: its role is to check that all the ASFX annotations are
 1819 consistent with the compiled ASMX program. A configuration is valid if its environment $\hat{\gamma}$ and
 1820 continuation $\hat{\kappa}$ are valid and additionally the environment and the continuation are coherent (nota-
 1821 tion: $\hat{\gamma} \circ \hat{\kappa}$). Furthermore, we require that in a valid configuration whose control is a computation
 1822 \hat{M} , the ASMX code at location $\mu(\hat{M})$ corresponds to the compiled code for \hat{M} , and that if instead
 1823 the control is a value v , the value itself is valid.

1824 A value is valid if any broom \hat{k} syntactically contained in it is valid; an environment is valid if
 1825 all the values it assigns to its domain are valid. A broom is valid if, when we resume in any valid
 1826 environment γ' and continuation κ' such that the two are coherent, we get a valid continuation
 1827 and, furthermore, the environment in the broom is coherent with the extended continuation.

1828 What it means for a continuation to be valid is expressed by an inductive predicate, with a
 1829 base rule for the identity continuation and a recursive rule for each type of frame. These rules
 1830 ensure that the annotation for each frame corresponds to suitable ASMX code: the annotation for
 1831 **jump** frames must refer to a jump linking to the start of the rest of the continuation; in the case of
 1832 **leave**, the annotation points to code loading a new leave record; in the case of **exit**, the annotation
 1833 must refer to an `exit` instruction; the annotation of pure frames corresponds to an intermediate
 1834 instruction of the compilation rule for sequencing; the annotation for handler frames refers to the
 1835 return instruction that will invoke that handler's return clause.

1836 *B.4.1 Proof of Theorem 6.2. Let P be a source program, and $\Xi = \lfloor P \rfloor$. For all source configurations
 1837 \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 ASMX configurations a
 1838 such that $a \vDash \hat{s}$ there exists a' such that $a \xrightarrow{*} a'$ and $a' \vDash \hat{s}'$.*

1840 To prove the theorem, we will need a few basic results on validity.

1841 LEMMA B.5. *If $\hat{\gamma}\checkmark$, then $\llbracket \hat{V} \rrbracket \hat{\gamma}\checkmark$.*

1842 PROOF. Routine induction on \hat{V} . □

1844 LEMMA B.6. *Let $\hat{\kappa} = \hat{\kappa}_1 \cdot h(\hat{w}, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{\kappa}_0$. Suppose $\hat{\kappa}\checkmark$ and $\hat{\gamma} \circ \hat{\kappa}$. Then we have:*

- 1845 • $\hat{w}\checkmark$
- 1846 • $\hat{\kappa}_0\checkmark$
- 1847 • $(\hat{\kappa}_1, h(\hat{w}) @ \ell, \hat{\gamma})\checkmark$
- 1848 • $\mathcal{R}_{\mathcal{L}} = \mathcal{L}(\mu(\hat{\kappa}_0), \hat{\gamma}_0)$ such that $\hat{\gamma}_0 \circ \hat{\kappa}_0$

1849 PROOF. By induction on $\hat{\kappa}_1$ with a case analysis on $\hat{\kappa}\checkmark$. □

1851 LEMMA B.7. *If $(\hat{\kappa}, h(\hat{w}) @ \ell, \hat{\gamma})\checkmark$ then $(\hat{\kappa}, h'(\hat{w}') @ \ell, \hat{\gamma})\checkmark$.*

1852 PROOF. The definition of $\hat{k}\checkmark$ is based on coherence, which uses operations *laddr* and *lcont*: the
 1853 values of these operations are the same for both continuations (the proof is by induction on $\hat{\kappa}$). □

1855 LEMMA B.8. *Let $\hat{\kappa} = (\hat{\kappa}, h(\hat{w}) @ \ell, \hat{\gamma})$. If $\hat{k}\checkmark, \hat{\kappa}_0\checkmark$ and $\hat{\gamma}_0 \circ \hat{\kappa}_0$, then $(\hat{\kappa} \cdot h(w, \mathcal{L}(\mu(\hat{\kappa}_0), \hat{\gamma}_0)) @ \ell \cdot \hat{\kappa}_0)\checkmark$.*

1856 PROOF. By the definition of $k\checkmark$. □

1857
 1858 PROOF OF THEOREM 6.2. The proof proceeds by cases on the possible reductions $\langle \hat{M}, \hat{\gamma}, \hat{\kappa} \rangle \hookrightarrow$
 1859 $\langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle$. First we show that the transition preserves validity, and then we are able to show that
 1860 all ASMX models of the source configuration deterministically transition (in multiple steps) to a
 1861 model of the target configuration.

1862

1863 *Case $\hat{M} = \oplus(\hat{V})$.* In this case $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \langle \llbracket \oplus \rrbracket (\llbracket \hat{V} \rrbracket \hat{y}), \hat{y}, \hat{k} \rangle$. Since the environment and
 1864 the continuation do not change, we need to check that the result of the evaluation is valid: this is
 1865 trivial because we have assumed that $\llbracket \oplus \rrbracket$ only outputs valid values.

1866 If $a \vDash \hat{s}$, we know that the ASMFx machine is about to execute $\llbracket \oplus(\hat{V}) \rrbracket$. By direct inspection of
 1867 the code, we see that execution until $v(\hat{M})$ leads us to a state a' where the $\$a$ registers have been
 1868 loaded with the value $\llbracket \oplus \rrbracket (\llbracket \hat{V} \rrbracket \hat{y})$ and everything else is unchanged. Then $a' \vDash \hat{s}'$ and $a \xrightarrow{*} a'$.
 1869

1870 *Case $\hat{M} = \text{if}(\hat{V}, \hat{M}_1, \hat{M}_2)$.* In this case we perform different operations depending on whether
 1871 $\llbracket \hat{V} \rrbracket \hat{y}$ is **true** or **false**.

1872 If it is **true**, then $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \langle \hat{M}_b \hat{y}, \text{jump}@v(\hat{M}_1) \cdot \hat{k} \rangle$. To prove $\hat{k}' \checkmark$, we show that $\hat{k} \checkmark$
 1873 and $\Xi(v(\hat{M}_1)) = \text{j } \mu(\hat{k})$ (by the validity of \hat{s} and a direct inspection of $\llbracket \hat{M} \rrbracket$).

1874 It is then trivial to prove that $\Xi(\mu(\hat{M}_1)) = \llbracket \hat{M}_1 \rrbracket$, $v(\hat{M}_1) = \mu(\text{jump}@v(\hat{M}_1) \cdot \hat{k})$, and $\hat{y} \supset \hat{k}'$,
 1875 therefore $\hat{s}' \checkmark$.
 1876

1877 If $a \vDash \hat{s}$, we can execute the ASMFx code deterministically up to $\llbracket \hat{M}_1 \rrbracket$ (since the if condition eval-
 1878 uated to **true**, the branch-if-zero instruction `bz` will not branch) and take the resulting configuration
 1879 as a' : we can easily show that $a' \vDash \hat{s}'$, and $a \xrightarrow{*} a'$ holds trivially.

1880 If instead $\llbracket \hat{V} \rrbracket \hat{y}$ evaluates to **false**, then $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \langle \hat{M}_f \hat{y}, \hat{k} \rangle$. The continuation and
 1881 environment do not change therefore, to prove that $\hat{s}' \checkmark$, we only need to show that $\Xi(\mu(\hat{M}_2)) =$
 1882 $\llbracket \hat{M}_2 \rrbracket$ and $v(\hat{M}_2) = \mu(\hat{k})$: these are immediate consequences of $\hat{s} \checkmark$ after unfolding the definition of
 1883 \hat{M} .
 1884

1885 If $a \vDash \hat{s}$, we can execute the ASMFx code deterministically up to $\llbracket \hat{M}_2 \rrbracket$ (since the if condition evalu-
 1886 ated to **false**, the branch-if-zero instruction will branch to l_{else}) and take the resulting configuration
 1887 as a' : we can easily show that $a' \vDash \hat{s}'$, and $a \xrightarrow{*} a'$ holds trivially.

1888 *Case $\hat{M} = \text{let } x \leftarrow \hat{M}_1 \text{ in } \hat{M}_2$.* In this case $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \langle \hat{M}_1, \hat{y}, ((x).\hat{M}_2)@v(\hat{M}_1) \cdot \hat{k} \rangle$. To
 1889 prove $\hat{k}' \checkmark$, we show that $\hat{k} \checkmark$, $\Xi(v(\hat{M}_1)) = \overline{\$x} \leftarrow \overline{\$a}; \llbracket \hat{M}_2 \rrbracket$ and $v(\hat{M}_2) = \mu(\hat{k})$ (all by the validity of
 1890 \hat{s} ; the second and last property also require a direct inspection of $\llbracket \hat{M} \rrbracket$).

1891 It is then trivial to prove that $\Xi(\mu(\hat{M}_1)) = \llbracket \hat{M}_1 \rrbracket$, $v(\hat{M}_1) = \mu(((x).\hat{M}_2)@v(\hat{M}_1) \cdot \hat{k})$, and $\hat{y} \supset \hat{k}'$,
 1892 therefore $\hat{s}' \checkmark$.
 1893

1894 If $a \vDash \hat{s}$, we can choose $a' = a$: we can easily show that $a \vDash \hat{s}'$, and $a \xrightarrow{*} a$ holds trivially.

1895 *Case $M = \text{unpack } \overline{x_n} \leftarrow \overline{\langle \hat{V}_n \rangle}$ in \hat{M}_0 .* In this case $\hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \langle \hat{M}_0, (x_n \mapsto \overline{\llbracket \hat{V}_n \rrbracket \hat{y}}) \cdot \hat{y}, \hat{k} \rangle$.
 1896 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
 1897 the value of `LEAVE`, if any, has not changed with the transition).

1898 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 $\llbracket \hat{M} \rrbracket$);
 1899 therefore $\hat{s}' \checkmark$.
 1900

1901 If $a \vDash \hat{s}$, we execute deterministically the code for $\llbracket \hat{M} \rrbracket$ up to and excluding $\llbracket \hat{M}_0 \rrbracket$, obtaining
 1902 a configuration a' where the registers corresponding to variables $\overline{x_n}$ have been loaded with the
 1903 subvalues from the evaluation of \hat{V} . We can then easily show that $a' \vDash \hat{s}'$, and $a \xrightarrow{*} a'$ holds trivially.
 1904

1905 *Case $\hat{M} = \text{handle } \hat{M}_0 \text{ with } h(\hat{V})$.* In this case

$$1906 \quad \hat{s}' = \langle \hat{M}', \hat{y}', \hat{k}' \rangle = \langle \hat{M}_0, \hat{y} \setminus \text{LEAVE}, h(\llbracket \hat{V} \rrbracket \hat{y}, \mathcal{L}(\mu(\hat{k}), \hat{y}))@v(\hat{M}_0) \cdot \hat{y} \rangle$$

1907 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
 1908 property also requires a direct inspection of $\llbracket \hat{M} \rrbracket$).

1909 It is then trivial to prove that $\Xi(\mu(\hat{M}_0)) = \llbracket \hat{M}_0 \rrbracket$, $v(\hat{M}_0) = \mu(h(\llbracket \hat{V} \rrbracket \hat{y}, \mathcal{L}(\mu(\hat{k}), \hat{y}))@v(\hat{M}_0) \cdot \hat{k})$,
 1910 and $\hat{y} \setminus \text{LEAVE} \supset \hat{k}'$, therefore $\hat{s}' \checkmark$.
 1911

If $a \vDash \hat{s}$, 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 $\hat{M} = \mathbf{do} \#t(\hat{V})$ and $\hat{\kappa} = \hat{\kappa}^{-t} \cdot h^{+t}(\hat{w}, \mathcal{R}_{\mathcal{L}})@l \cdot \hat{\kappa}_0$. In this case

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

To prove $\hat{\kappa}' \checkmark$, we show that $\hat{\kappa} \checkmark$ 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 $\Xi(\mu(\hat{M}_{h_{\#t}})) = \llbracket \hat{M}_{h_{\#t}} \rrbracket$, $v(\hat{M}_{h_{\#t}}) = \mu(\text{exit}@v(\hat{M}_{h_{\#t}}) \cdot \hat{\kappa})$. To prove $\hat{\gamma}' \subset \hat{\kappa}'$ we show:

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

Then it follows that $s' \checkmark$.

If $a \vDash \hat{s}$, 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 $\hat{M} = \mathbf{resume} \hat{U}(\hat{V})$. Assume $\hat{k} := \llbracket \hat{U} \rrbracket \hat{\gamma} = (\hat{\kappa}_U, h_U(\hat{w})@l, \hat{\gamma}_U)$. Then:

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

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

If $a \vDash \hat{s}$, we execute deterministically the entire compiled code for $\mathbf{resume} \hat{U}(\hat{V})$: it is easy to show that the resulting ASAFX state a' is such that $a' \vDash \hat{s}'$. In particular, we can verify that just before executing the resume instruction, the register file holds the following values:

$$\begin{array}{ll} \$lr & \mapsto \mathcal{L}(\mu(\hat{\kappa}), \llbracket \hat{\gamma} \rrbracket) \\ \$kr & \mapsto \mathcal{K}(\ell_k, \llbracket \hat{\gamma}_U \rrbracket, \llbracket \hat{\kappa}_U \rrbracket) \\ \$hr & \mapsto \mathcal{H}(\llbracket h \rrbracket, (\overline{\$s} = \llbracket \hat{w} \rrbracket)) \end{array}$$

where $\ell_k = \mu(\hat{\kappa}_U \cdot h(\hat{w})@l)$. 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 $\hat{\gamma}_U$, and the effect context has been extended with $\llbracket \hat{\kappa}_U \rrbracket \cdot \mathcal{H}(\llbracket h \rrbracket, (\overline{\$s} = \llbracket \hat{w} \rrbracket), \mathcal{L}(\mu(\hat{\gamma}_0), \llbracket \hat{\gamma} \rrbracket))$. We can thus easily show that $a' \vDash \hat{s}'$.

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

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

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

If $a \vDash \hat{s}$, 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}' .

1961 Case $\hat{M} = \mathbf{newbroom}@l(f)$ with $h(\hat{V})$. We have $\hat{s}' = \langle (\hat{\kappa}_f, h(\llbracket \hat{V} \rrbracket \hat{\gamma})@(\ell + 2), []) \rangle, \hat{\gamma}', \hat{\kappa}' \rangle$ where

$$1962 \quad \hat{\kappa}_f = \mathbf{leave}@l \cdot \mathbf{jump}@(\ell + 1) \cdot ((x).f_{\text{return}})@(\mu(f_{\text{return}})) \cdot \mathbf{exit}@v(f_{\text{return}})$$

1964 To prove $\hat{s}' \checkmark$, we unfold the definition and in particular we verify that $(\hat{\kappa}_f \cdot h(\llbracket \hat{V} \rrbracket \hat{\gamma})@(\ell + 2)) \checkmark$
 1965 by repeated applications of the rules defining augmented source validity (this requires us to verify
 1966 that the locations annotating the continuation contain the expected instructions).

1967 If $a \vDash \hat{s}$, we execute deterministically the entire code for $\llbracket \hat{M} \rrbracket$. It is easy to show that the resulting
 1968 configuration a' is a model of \hat{s}' .
 1969

1970 Case $\hat{M} = \mathbf{return} \hat{V}$. In this case $\hat{s}' = \langle \llbracket \hat{V} \rrbracket \hat{\gamma}, \hat{\gamma}, \hat{\kappa} \rangle$. Knowing by hypothesis that $\hat{s} \checkmark$,
 1971 to prove $\hat{s}' \checkmark$ we only need to show, by Lemma B.5, that $\llbracket \hat{V} \rrbracket \hat{\gamma} \checkmark$.
 1972

1973 If $a \vDash \hat{s}$, we perform the code $\llbracket \mathbf{return} \hat{V} \rrbracket = \llbracket \hat{V} \rrbracket$, which loads an initial segment $\overline{\$a}$ of the
 1974 argument registers with the evaluation of \hat{V} . We can then show that the resulting state a' is a model
 1975 of \hat{s}' .
 1976

1977 Case $\hat{M} = \hat{v}$ and $\hat{\kappa} = \mathbf{jump}@l \cdot \hat{\kappa}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = (\hat{v}, \hat{\gamma}, \hat{\kappa}_0)$. To prove $\hat{s}' \checkmark$ knowing
 1978 $\hat{s} \checkmark$, it is enough to show that $\hat{\kappa}_0 \checkmark$ by inversion on $\hat{\kappa} \checkmark$.
 1979

1980 If $a \vDash \hat{s}$, we know the next instruction is $\mathbf{j} \mu(\hat{\kappa}_0)$. By executing that instruction, we move into a
 1981 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 $a' \vDash \hat{s}'$.
 1982

1983 Case $\hat{M} = \hat{v}$ and $\hat{\kappa} = \mathbf{exit}@l \cdot \hat{\kappa}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = (\hat{v}, \hat{\gamma}(\mathbf{LEAVE}).env, \hat{\kappa}_0)$. Knowing
 1984 by hypothesis that $\hat{s} \checkmark$, we show that $\hat{\kappa}_0 \checkmark$ by inversion on $\hat{\kappa} \checkmark$.
 1985

1986 To show that $\hat{\gamma}(\mathbf{LEAVE}).env \supset \hat{\kappa}_0$, we note $\hat{\gamma} \supset \hat{\kappa}$ which, by its definition, implies the thesis.

1987 Then, knowing $\hat{v} \checkmark$, we have immediately $\hat{s}' \checkmark$.

1988 If $a \vDash \hat{s}$, we know the next instruction is an \mathbf{exit} and the current value of register $\$lr$ is
 1989 $\mathcal{L}(\mu(\hat{\gamma}_0), \llbracket \hat{\gamma}(\mathbf{LEAVE}).env \rrbracket)$. Then, by executing that instruction, we move into a configuration a'
 1990 where the register $\$pc$ holds the address $\mu(\hat{\gamma}_0)$, and the register file has been updated to match the
 environment $\hat{\gamma}(\mathbf{LEAVE}).env$. We can thus easily show that $a' \vDash \hat{s}'$.
 1991

1992 Case $\hat{M} = \hat{v}$ and $\hat{\kappa} = \mathbf{leave}@l \cdot \hat{\kappa}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = (\hat{v}, (\mathbf{LEAVE} \mapsto \mathcal{L}(\ell + 2, [])) \cdot \hat{\gamma}, \hat{\kappa}_0)$.
 1993 Knowing by hypothesis that $\hat{s} \checkmark$, we show that $\hat{\kappa}_0 \checkmark$ by inversion on $\hat{\kappa} \checkmark$.
 1994

1995 To show that $(\mathbf{LEAVE} \mapsto \mathcal{L}(\ell + 2, [])) \cdot \hat{\gamma} \supset \hat{\kappa}_0$, we note, by inversion on validity, that $\hat{\kappa}_0 =$
 1996 $\mathbf{jump}@(\ell + 1) \cdot ((x).\hat{N})@(\mu(\hat{N})) \cdot \mathbf{exit}@v(\hat{N}) \cdot h(\hat{w})@(\ell + 2) \cdot \hat{\kappa}_1$. We can thus see that $lcont(\hat{\kappa}_0) = \ell + 2$,
 which is all we need.
 1997

1998 Then, knowing $\hat{v} \checkmark$, we have immediately $\hat{s}' \checkmark$.

1999 If $a \vDash \hat{s}$, we know the next instruction is $\mathbf{loadl} \ell + 2$. Then, by executing that instruction, we
 2000 move into a configuration a' where the register $\$lr$ holds the record $\mathcal{L}(\ell + 2, [])$ and all the other
 2001 registers match the environment $\hat{\gamma}$. We can thus easily show that $a' \vDash \hat{s}'$.
 2002

2003 Case $\hat{M} = \hat{v}$ and $\hat{\kappa} = ((x).\hat{M}_0)@l \cdot \hat{\kappa}_0$. In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = (\hat{M}_0, (x \mapsto \hat{v}) \cdot \hat{\gamma}, \hat{\kappa}_0)$.

2004 To prove $\hat{s}' \checkmark$, knowing $\hat{s} \checkmark$, we trivially show $\hat{\kappa}_0 \checkmark$ (by inversion on $\hat{\kappa} \checkmark$) and then proceed to
 2005 show that $\hat{\gamma} \supset \hat{\kappa}$ implies $((x \mapsto \hat{v}) \cdot \hat{\gamma}) \supset \hat{\kappa}_0$ (\mathbf{LEAVE} is defined in one environment if and only if it
 2006 is defined in the other, and if it is its value is the same; furthermore, the conditions it has to satisfy
 2007 in one continuation are exactly the same as those needed in the other). Therefore, $\hat{s}' \checkmark$ holds.
 2008

2009 If $a \vDash \hat{s}$, we know that the register $\$pc$ points to code that will deterministically move data from
 registers $\overline{\$a}$ (holding the value $\llbracket \hat{v} \rrbracket$) into registers $\overline{\$x}$, 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

2010 obtain a configuration a' where the registers $\overline{\$x}$ have been loaded with $\llbracket \hat{v} \rrbracket$. Then, if we note that
 2011 $\llbracket \hat{\kappa} \rrbracket = \llbracket \hat{\kappa}' \rrbracket$ (since pure frames are ignored when compiling a source continuation), we prove $a' \vDash \hat{s}'$.

2012 *Case $\hat{M} = \hat{v}$ and $\hat{\kappa} = h(\hat{w}, \mathcal{R}_{\mathcal{L}}) @ \ell \cdot \hat{\kappa}_0$.* In this case $\hat{s}' = \langle \hat{M}', \hat{\gamma}', \hat{\kappa}' \rangle = (h_{\text{return}}, (\text{LEAVE} \mapsto$
 2013 $\mathcal{R}_{\mathcal{L}}) \cdot (\text{dom}(h_{\text{return}}) \mapsto \hat{v}) \cdot (\text{dom}(h) \mapsto \hat{w}), \hat{\kappa}_0)$.

2014 Knowing $\hat{s} \checkmark$, to prove $\hat{s}' \checkmark$ we start by proving that $(\text{exit}@v(h_{\text{return}}) \cdot \hat{\kappa}_0) \checkmark$ which requires us to
 2015 show that $\Xi(v(h_{\text{return}})) = \text{exit}$ (trivial by inspection of $\llbracket h_{\text{return}} \rrbracket$).

2016 To prove $\hat{\gamma}' \supset \hat{\kappa}'$, it is sufficient to show that $\mathcal{R}_{\mathcal{L}} = \mathcal{L}(\mu(\hat{\kappa}_0), \hat{\gamma}_0)$ such that $\hat{\gamma}_0 \supset \hat{\kappa}_0$. This is proved
 2017 by inversion on $\hat{\kappa} \checkmark$.

2018 Noting that $\Xi(\mu(h_{\text{return}})) = \llbracket h_{\text{return}} \rrbracket$ (by the definition of Ξ), we have $\hat{s}' \checkmark$.

2019 If $a \vDash \hat{s}$, we know that the next instruction is a return, that the initial segment of argument
 2020 registers $\overline{\$a}$ holds $\llbracket \hat{v} \rrbracket$, and that the effect context is $\mathcal{H}(\llbracket h \rrbracket, (\overline{\$s} = \llbracket \hat{w} \rrbracket), \llbracket \mathcal{R}_{\mathcal{L}} \rrbracket) \cdot \llbracket \hat{\kappa}_0 \rrbracket$. Then by
 2021 executing return we load the encoding of \hat{w} into registers $\overline{\$s}$, the encoding of $\mathcal{R}_{\mathcal{L}}$ into register
 2022 $\$lr$, and we jump to the address μh_{return} . We execute the first few instructions of $\llbracket h_{\text{return}} \rrbracket$ until we
 2023 reach \hat{M}_{return} — these instructions move data between registers and ensure that the register file
 2024 corresponds to $\hat{\gamma}'$. It is thus easy to show that the configuration a' we have reached is such that
 2025 $a' \vDash \hat{s}'$. □

2026

2027

2028

2029

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

2042

2043

2044

2045

2046

2047

2048

2049

2050

2051

2052

2053

2054

2055

2056

2057

2058