

# Linking as Evaluation

## Correct Separate Compilation via First-Class Environments

JAM KABEER ALI KHAN, Independent, China

JINHAO TAN, The University of Hong Kong, China

BRUNO C. D. S. OLIVEIRA, The University of Hong Kong, China

Separate compilation is central to modular software: components are compiled independently and linked into an executable. Separate compilation is correct when a linked program behaves like the program formed by merging its sources. This property is less routine to prove than type soundness, even for simple types, because its proof needs a formal account of linking. Verified compilers prove it, but only for static links outside the source language. Some typed calculi make linking an expression, but none, to our knowledge, proves it.

We present  $\lambda_{\text{SCE}}$ , a simply typed calculus in which linking, modules, and environments are all *first-class*. Programmers control module linking through a link operator, which a type-directed elaboration translates into ordinary terms of a core calculus of first-class environments. The core calculus's existing evaluation rules then perform the linking: *linking as evaluation*. Linking independently compiled modules produces exactly the elaboration of their source link, so compiling and linking commute by construction. Separate compilation correctness then needs no proof of its own: it follows directly from semantics and behavior preservation, for static and dynamic links alike. We also prove that a module's exports depend on its provider only through the imported values. We mechanize every result in Lean 4 and provide an OCaml compiler to WebAssembly that uses no dedicated linking instruction.

Additional Key Words and Phrases: Separate Compilation, Linking, Semantics Preservation, First-Class Environments, Compiler Correctness, Lean 4 Mechanization

### 1 Introduction

Type soundness [51, 52] is a widely accepted sanity criterion for statically typed programming languages, supported by well-established proof techniques and mechanization frameworks. Modern software also relies on separate compilation: components are developed, typechecked, and compiled independently, then linked into larger systems [4, 26]. Separate compilation supports modularity, incremental builds, and third-party reuse. The correctness requirement is simple: linking independently compiled components should produce a program with the same behavior as linking their sources before compilation. However, establishing this correspondence is less routine than proving type soundness, even for simple types, because the proof requires a formal account of linking.

Cardelli [7] gives a foundational account of separate compilation and linking. The account represents program fragments as open terms whose free variables are their imports. Fragments are grouped into *linksets*. Linking proceeds by syntactic substitution: it matches an import to a fragment by name and interface, then substitutes that fragment's code for the import. Repeated linking resolves all imports into an executable. SMLSC, a separate compilation extension to Standard ML, follows the same model [48]. Cardelli proves, and SMLSC states, that linking preserves typing. Both link *outside* the language. Neither proves that linked code behaves like the merged sources.

Linking need not remain outside the language. In Units [12], a compound expression connects units' imports and exports, so programs can assemble components at run time. Alice ML expresses linking through functions: components are functions that receive an `import` function for acquiring their imports. Static linking constructs a wrapper component [37]. MixML combines recursive linking with ML-style abstraction [39]. Its semantics and soundness proof go through translation into a core calculus. None of these designs proves that linking separately compiled components

---

Authors' Contact Information: [Jam Kabeer Ali Khan](#), jamkhan@connect.hku.hk, Independent, Hong Kong, China; [Jinhao Tan](#), jinhaot@connect.hku.hk, The University of Hong Kong, Hong Kong, China; [Bruno C. d. S. Oliveira](#), bruno@cs.hku.hk, The University of Hong Kong, Hong Kong, China.

preserves the behavior of the combined source program. Lee and Friedman [25] prove, for an untyped calculus, that compiled linking agrees with source hole filling, a meta-level operation. We know of no typed calculus with linking as an expression that has such a theorem.

Separate compilation correctness has been established for verified compilers. For SepCompCert, Kang et al. show that linking separately compiled C modules refines the behavior of the linked source program [23]. For passes that commute with linking, compiling the linked sources produces the same code as linking the compiled parts. In this easy case, Kang et al. derive separate compilation correctness from whole-program correctness. For CertiCoq, which compiles Coq’s higher-order language Gallina to C, Paraskevopoulou et al. establish separate compilation correctness for the optimizing middle end using compositional logical relations [33]. In both, linking remains outside the source language. Neither theorem covers *dynamic linking*, at run time.

We bring the last two strands together: linking within the language, backed by a proof of separate compilation correctness that covers dynamic links. Our approach is *linking as evaluation*: a source link elaborates into ordinary target terms whose evaluation supplies a module’s imports and combines both modules’ exports into a linked library.

We develop this approach in  $\lambda_{\text{SCE}}$ , a *simply typed* calculus with first-class modules, environments, and linking. First-class environments [49, 50] let programs construct the environment in which a module runs, and use a linked library or module as an environment for computation.  $\lambda_{\text{SCE}}$  elaborates into  $\lambda_E$ , a calculus of first-class environments [50]. A separately compiled module with imports becomes a closed function whose parameter supplies its import environment. Linking constructs this environment from the required exports of another module, applies the function, and combines their exports. Function application, field selection, and environment composition thus perform the linking using the target’s existing evaluation rules. A module may also be *sandboxed*: it then sees only its declared import, never the scope where it is written, and compiles to the same term anywhere. A link passes the client only the selected import, never its own scope: the client runs under the environment it captured, extended with that import.

Programmers control linking: a link is an ordinary expression. The same link can assemble compiled modules at top level, link against a module passed to a function at run time, or draw an import from the current environment. The link’s elaboration uses only its operands’ core terms and interfaces, so the core linker emits that same term from compiled fragments: compiling and linking commute by construction. Commutation is thus SepCompCert’s easy case, holding by design and for dynamic links too.

We first prove that elaboration preserves types: well-typed source links compile to well-typed core terms. Beyond the typing guarantees studied by Cardelli, we define  $\lambda_{\text{SCE}}$ ’s operational semantics and prove that elaboration preserves source semantics. A source program and its elaboration evaluate to values related by elaboration or diverge together. The source link is the combined program. Its elaboration is the core linker’s output, so semantics and behavior preservation yield separate compilation correctness, with no account of linking beyond  $\lambda_{\text{SCE}}$ ’s own semantics. This covers static and dynamic links, at top level or inside functions, including diverging runs. The static case (Corollary 4.13) proves, for  $\lambda_{\text{SCE}}$ , what Cardelli names but does not formalize. We also prove *interface abstraction*: a client’s exports depend on its provider only through the imported values.

Our contributions are:

- **The  $\lambda_{\text{SCE}}$  calculus (§3).** First-class modules and typed links. We prove type safety.
- **Elaboration (§§4.2–4.3).** Type-directed elaboration, preserving types and semantics.
- **Separate compilation correctness (§4.3).** We derive correctness for both static and dynamic linking from semantics and behavior preservation. We also prove interface abstraction, and confinement for sandboxed modules.
- **Extensions (§5.1).** Unions, iso-recursive types, fixpoints, and multi-import linking.

- **Mechanization and implementation (§5.2).**<sup>1</sup> We mechanize all results in Lean 4 [29] and implement SCE-Lang, a prototype compiler written in OCaml. It targets WebAssembly [15] with no linking instruction and supports staged linking in both the core and WebAssembly.

## 2 Overview

This section shows, on small examples, how  $\lambda_{\text{SCE}}$  turns linking into an expression evaluated under first-class environments. Compiling a link then produces the same code as compiling its parts and linking them. We introduce first-class environments (Section 2.1), and review Cardelli’s account of linking and separate compilation (Section 2.2). Next, we present the design of  $\lambda_{\text{SCE}}$  (Section 2.3) and illustrate its use through programming examples (Section 2.4).

The examples use a simplified surface syntax of SCE-Lang, our implementation (Section 5.2).

### 2.1 First-Class Environments

Environments are fundamental to language implementations: they hold the bindings against which code executes. However, environments normally belong to the implementation, not the language, so programs cannot reach or manipulate them. *First-class environments* [13, 36] lift this restriction: environments become ordinary values that programs can build, inspect, extend, and pass to functions. Because the body of a module is a sequence of bindings, a first-class environment can also represent that body.  $\lambda_{\text{SCE}}$  builds on this idea and elaborates into  $\lambda_E$  [50], a call-by-value calculus in which dependent merges [49] build first-class environments.

*Dependent merges.* In  $\lambda_E$ , a record  $\{l = e\}$  supplies a single labeled entry. A merge  $e_1 \circ e_2$  evaluates its operands and joins their values into one environment, so multi-field records are merges of single-field ones. Merges are *dependent*:  $e_1$  is evaluated first, and its value extends the environment in which  $e_2$  runs. In  $\lambda_E$  a typing context is the type of an environment, so the typing rule extends the context  $\Gamma$  with the left operand’s type and pairs the result types:

$$\frac{\Gamma \vdash e_1 : A \quad \Gamma \& A \vdash e_2 : B}{\Gamma \vdash e_1 \circ e_2 : A \& B}$$

For example, the right component of

$$\{x = 2\} \circ \{y = x + 1\}$$

refers to  $x$ . By the rule above, the merge has type  $\{x : \text{Int}\} \& \{y : \text{Int}\}$ . At run time, the value  $\{x = 2\}$  on the left is added to the environment before the right component runs. So  $x$  resolves to 2, and the merge evaluates to  $\{x = 2\} \circ \{y = 3\}$ .

Dependencies between declarations are pervasive in modules, where a definition often refers to earlier definitions of the same module. A dependent merge makes such a sequence of dependent declarations an ordinary expression. The following declarations

```
let x : Int = 2
let y : Int = x + 1
```

are modeled as the merge above, and its type gives the types of the declarations.

*Variables and closures.* Every  $\lambda_E$  expression, such as the merge above, runs under a current, or *ambient*, environment. To resolve a variable,  $\lambda_E$  looks it up in the ambient environment and never substitutes into terms. Terms of  $\lambda_E$  are nameless: a variable is a de Bruijn index [9] giving the position of an entry in the environment. For readability, this section writes names instead of indices.

<sup>1</sup>Implementation and mechanization can be found at <https://github.com/jam-khan/SCE-Lang>




Fig. 1. Double.sce compiled against interface Add.scei.

It also writes  $\varepsilon$  for both the empty environment and its type. A function evaluates to a closure  $\langle v, \lambda x:A. e \rangle$ , which pairs the function with the environment  $v$  of its definition.

*Programmer-controlled environments.* Besides dependent merges, which build first-class environments,  $\lambda_E$  provides constructs that let a program choose the environment a computation runs in. The *query* operator **env** returns the current environment. A *box*  $e_1 \triangleright e_2$  runs  $e_2$  under the local environment produced by  $e_1$ , which replaces the ambient one. Assuming that the ambient environment binds  $x$  to an integer, the expression

$$(\mathbf{env} \text{ , } \{y = 2\}) \triangleright (y + x)$$

is well typed. Here **env** returns the ambient environment, which contains  $x$ . The merge extends it with the local binding  $y = 2$ . The box then runs the body under the extended environment, so both  $y$  and  $x$  are in scope. Dropping the query gives

$$\{y = 2\} \triangleright (y + x)$$

which is ill typed, because the environment that the box supplies binds only  $y$ . A box therefore runs a computation under an environment that the program controls. Its body sees that environment's bindings and no others. The simply typed  $\lambda$ -calculus has no such construct: a  $\lambda$  only extends the environment, so a function body sees every enclosing binding.

With the box construct,  $\lambda_E$  provides a form of sandboxed modules [50]. A module that takes other modules as arguments is boxed under the empty environment, so it acts only through its arguments. The treatment is informal:  $\lambda_E$  has no module syntax, types or rules. Nor do Tan and Oliveira [50] study linking independently developed modules or separate compilation correctness.

## 2.2 Linking Outside The Language

Separate compilation compiles a client module against its provider's interface without the provider's implementation. The provider can be compiled independently and linked with the compiled client to produce an executable. Source files have the extension `.sce`, and interface files, which list a module's exports, have `.scei`. In Figure 1, Double imports Add and calls `Add.op` in double. Double.sce takes the type of `Add.op` from Add.scei, so it compiles before Add.sce is available. The figure shows only this *compilation* step, not how the modules are *linked* or what makes compiling and linking *correct*.

Cardelli [7] formalizes the linking step that the figure omits. Separately compiled modules become *linksets*: compilation produces them, and combination and linking act on them. The underlying calculus is  $F_1$ , a simply typed  $\lambda$ -calculus. A judgment  $E \vdash a : A$  assigns type  $A$  to the term  $a$  under the typing environment  $E$ , a typing context rather than a run-time environment. A linkset gathers

$$\begin{aligned}
& \text{(a) Compilation} \\
& L_{\text{Add}} = \emptyset \mid \text{op} \quad \vdash (\emptyset \vdash a_+ : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}) \\
& \quad \text{unit} \quad \vdash (\emptyset, \text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \vdash 0 : \text{Int}) \\
& L_{\text{Double}} = \emptyset, \text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \mid \text{double} \vdash (\emptyset \vdash \lambda x : \text{Int}. \text{op } x \ x : \text{Int} \rightarrow \text{Int}) \\
& \text{(b) Combination} \\
& L_{\text{Add}} + L_{\text{Double}} = \emptyset \mid \text{op} \quad \vdash (\emptyset \vdash a_+ : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}) \\
& \quad \text{unit} \quad \vdash (\emptyset, \text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \vdash 0 : \text{Int}) \\
& \quad \text{double} \vdash (\emptyset, \text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \vdash \lambda x : \text{Int}. \text{op } x \ x : \text{Int} \rightarrow \text{Int}) \\
& \text{(c) Linking} \\
& \text{link}(L_{\text{Add}} + L_{\text{Double}}) = \emptyset \mid \text{op} \quad \vdash (\emptyset \vdash a_+ : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}) \\
& \quad \text{unit} \quad \vdash (\emptyset \vdash 0 : \text{Int}) \\
& \quad \text{double} \vdash (\emptyset \vdash \lambda x : \text{Int}. a_+ \ x \ x : \text{Int} \rightarrow \text{Int}) \\
& \quad \text{where } a_+ = \lambda x : \text{Int}. \lambda y : \text{Int}. x + y
\end{aligned}$$

Fig. 2. Compiling, combining and linking the linksets of Add and Double, after Cardelli [7].

judgments under distinct names:

$$E_0 \mid x_1 \vdash E_1 \vdash a_1 : A_1 \ \dots \ x_n \vdash E_n \vdash a_n : A_n$$

The typing environment  $E_0$  is the *interface* of the linkset and lists its imports. Each *fragment* exports the term  $a_i$  under the name  $x_i$ . The typing environment  $E_i$  of fragment  $x_i$  gives the types of the fragments that  $a_i$  may refer to. Interfaces are thus typing environments outside the language: no term of  $F_1$  has a typing environment as its type.

Figure 2 shows the three steps on the modules of Figure 1, in adapted notation. Shading marks what each step changes. Cardelli’s modules list imports one variable at a time, not by module, so Double imports `op` directly. Compilation turns a module into a linkset with one fragment per binding, each typed under the earlier bindings. The linkset’s interface lists the module’s imports (a). Combining two linksets, written  $L + L'$ , puts their fragments into one linkset. An import of one linkset that the other exports moves from the interface to the typing environments of the importing fragments (b). Linking, written  $\text{link}(L)$ , then resolves every name in the typing environments by substitution: the moved imports and the earlier bindings. Each such name leaves the typing environment and is replaced by the term exported under that name by a closed fragment, one typed under the empty environment (c).

Cardelli proves that compilation, combination and linking preserve typing. That is, each fragment typechecks, and each import has exactly the type that its provider exports. Cardelli names separate compilation correctness, but  $F_1$  has no evaluation in which to state it. It holds when linked code behaves like the merged sources, one program with both modules. Linking acts on linksets, not on terms of  $F_1$ , so it lies outside the language. A correctness proof thus needs an account of linking.

Cardelli’s account raises three questions:

- Fragments are open terms, and linking rewrites them by substitution. Can a compiled module instead be a closed term that linking runs unchanged?
- Interfaces are typing environments, and linking is an operation on linksets. Can interfaces be ordinary types, and can linking be an expression evaluated during execution?
- Can separate compilation correctness follow from the compiler’s correctness alone, with no account of linking?

Section 2.3 answers them in order, using first-class environments.

|                                                      |                                                                                                                                         |
|------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| (a) Importing file                                   |                                                                                                                                         |
| <code>import M (* sig x : Int; y : Int end *)</code> | $\rightsquigarrow \varepsilon \triangleright (\lambda \{M : \{x : \text{Int}\} \ \& \ \{y : \text{Int}\}\}. \{z = \underline{0}.M.x\})$ |
| <code>let z : Int = M.x</code>                       |                                                                                                                                         |
| (b) Sandboxed functor                                |                                                                                                                                         |
| <code>sandbox functor (I :</code>                    |                                                                                                                                         |
| <code>{ M : sig x : Int; y : Int end }) -&gt;</code> | $\rightsquigarrow \varepsilon \triangleright (\lambda \{M : \{x : \text{Int}\} \ \& \ \{y : \text{Int}\}\}. \{z = \underline{0}.M.x\})$ |
| <code>struct let z : Int = I.M.x end</code>          |                                                                                                                                         |
| (c) Structure                                        |                                                                                                                                         |
| <code>struct let x : Int = 1;</code>                 |                                                                                                                                         |
| <code>let y : Int = x + 1 end</code>                 | $\rightsquigarrow \{x = 1\} \circ \{y = \underline{x} + 1\}$                                                                            |

Fig. 3. Source modules (left) and simplified  $\lambda_E$  elaborations (right, shaded gray). We omit administrative bindings. In (a–b),  $\underline{0} = \text{env}.0$  is the import record, by de Bruijn index. In (c), underlining marks a field lookup.

### 2.3 Linking *Inside* The Language

$\lambda_{\text{SCE}}$  makes linking *first-class*: a program can write a link at top level, place it inside a function, or choose the link’s provider at run time. This design unifies static and dynamic linking in a *typed* setting. It also makes separate compilation correctness a corollary of semantics and behavior preservation.

To illustrate, take Add and Double from Figure 1, now declared within the program. We make Double a *sandboxed* functor requiring an op operation of type  $\text{Int} \rightarrow \text{Int} \rightarrow \text{Int}$ . Any provider exposing a unique op of this type can be linked with Double. First-class linking lets us combine these modules into a library and use that library as the environment for a computation:

```

module Lib = link Add with Double
let main : Int = box Lib in op (double 5) unit    (* 10 *)

```

The link operator supplies Double’s import from Add and combines their exports in Lib. Because structures elaborate to *first-class environments*, Lib serves directly as the environment of main.

*Modules as closed terms.* Lib links two modules declared in the program. A separately compiled file links in the same way, because a file with imports elaborates to the same simplified form as a sandboxed functor, up to parameter renaming. Figure 3(a–b) shows this form on a smaller module: a boxed lambda whose parameter binds the imports. The box hides the ambient environment and the lambda binds every import, so the compiled file is closed. Linking applies it and never rewrites it, unlike the substitution of Figure 2(c). This answers the first question of Section 2.2. In Figure 3(c), a structure elaborates to a first-class environment built by dependent merge: y can refer to x.

*Interfaces as types.* Module interfaces are ordinary types, so functions can accept and return modules. This answers the first half of the second question. We write  $\text{sig } A$  for the type of a structure whose contents have type  $A$ , and  $A \rightarrow_m B$  for a functor with import type  $A$  and result type  $B$ . The type translation  $\llbracket \cdot \rrbracket$  maps the functor and structure types in Figure 3(b–c) as follows:

$$\begin{aligned}
\llbracket \{M : \text{sig } (\{x : \text{Int}\} \ \& \ \{y : \text{Int}\}) \} \rightarrow_m \text{sig } \{z : \text{Int}\} \rrbracket &= \{M : \{x : \text{Int}\} \ \& \ \{y : \text{Int}\}\} \rightarrow \{z : \text{Int}\} \\
\llbracket \text{sig } (\{x : \text{Int}\} \ \& \ \{y : \text{Int}\}) \rrbracket &= \{x : \text{Int}\} \ \& \ \{y : \text{Int}\}
\end{aligned}$$

*Linking as evaluation.* The second question also asks whether linking can be an expression. In  $\lambda_{\text{SCE}}$ , a source link elaborates into ordinary  $\lambda_E$  terms, so the core evaluator links with its existing rules for application, lookup, selection and merge.

$$\begin{array}{c}
\text{295} \quad D = \{\text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}\} \quad P = D \ \& \ \{\text{unit} : \text{Int}\} \quad T = \{\text{double} : \text{Int} \rightarrow \text{Int}\} \\
\text{296} \quad m' = \{\text{op} = \lambda x:\text{Int}. \lambda y:\text{Int}. x + y\} \ \& \ \{\text{unit} = 0\} \quad n' = \varepsilon \triangleright (\lambda D. \{\text{double} = \lambda x:\text{Int}. \text{op } x \ x\}) \\
\text{297} \quad \text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \in \text{sig } P \\
\text{298} \quad \dots \quad \dots \\
\text{299} \quad \frac{}{\varepsilon \vdash \text{Double} : D \rightarrow_m \text{sig } T \rightsquigarrow n'} \text{ELAB-FUN} \quad \frac{}{\varepsilon \vdash \text{Add} : \text{sig } P \rightsquigarrow m'} \text{ELAB-STR} \\
\text{300} \quad \frac{}{\varepsilon \vdash \text{Add} \bowtie \text{Double} : \text{sig } P \ \& \ \text{sig } T \rightsquigarrow \underbrace{(\lambda p:P. \lambda f:(D \rightarrow T). p \ \& \ (f \ \{\text{op} = p.\text{op}\})) \ m' \ n'}_{e_\ell}} \text{ELAB-LINK}
\end{array}$$

Fig. 4. Elaboration of Add linked with Double.

*Link elaboration into  $\lambda_E$ .* Figure 4 elaborates the link of Add (Figure 1) with Double, a sandboxed functor (marked  $\blacksquare$ ) as in Figure 3(b). We write  $e_1 \bowtie e_2$  for link  $e_1$  with  $e_2$ .<sup>2</sup> The judgment  $\varepsilon \vdash e : A \rightsquigarrow e'$  types a source term  $e$  and translates it to a core term  $e'$ . In **ELAB-LINK**, the premise  $\text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \in \text{sig } P$  checks that Add exports  $\text{op}$  exactly once, with the type Double imports.

In  $e_\ell$ , the applications evaluate the provider and client once, in that order, and bind their values to  $p$  and  $f$ . The body selects the  $\text{op}$  field from  $p$  and supplies it to  $f$  as  $\{\text{op} = p.\text{op}\}$ , then merges  $p$  with the client's result. Linking thus becomes ordinary application and environment composition. Because the current environment  $\text{env}$ , which  $\lambda_{\text{SCE}}$  inherits from  $\lambda_E$ , is a first-class value (Section 2.1), it can also be the provider. In a structure that declares  $\text{op}$ , with no other  $\text{op}$  in scope, a later declaration can be  $\text{env} \bowtie \text{Double}$ , which links Double against that  $\text{op}$ .

*Confined linking.* Linking against  $\text{env}$  makes the link site the provider, by explicit choice. First-class environments invite a design that makes this implicit: run the client under the environment where the link is written, extended with its import. The client's meaning would then depend on the link site. A name that the client uses but does not import would resolve to whatever that site binds, and its compiled code would be valid only there. In  $\lambda_{\text{SCE}}$ , a link instead applies the client, as  $e_\ell$  shows. The client runs under the environment its closure captured, extended with the selected import. The link itself adds nothing from its own site. This holds for every client. A functor without the sandbox mark, an *open functor*, still captures the environment where it is written. Its body thus sees the definition site. The box in  $n'$  removes that environment as well. Double's body runs under the empty environment extended with the record that  $f$  receives, so it sees  $\text{op}$  from Add and nothing else. A sandboxed client therefore compiles to the same term wherever it is written. The compiled client runs at any link site, against any provider that exports its import exactly once at the required type, and its exports depend only on that import. Context Independence (Lemma 4.1) and Sandbox and Link Confinement (Theorems 4.16 and 4.17) state these properties.

*Link typing and the core linker.* Elaboration preserves typing (Section 4.2), so a well-typed source link elaborates to a well-typed core term. In the example, the source link has type  $\text{sig } P \ \& \ \text{sig } T$ :  $P$  describes Add's exports ( $\text{op}$ ,  $\text{unit}$ ),  $D$  the import ( $\text{op}$ ), and  $T$  Double's export ( $\text{double}$ ). Type translation gives the core types  $P$  for  $m'$ ,  $D \rightarrow T$  for  $n'$ , and  $P \ \& \ T$  for  $e_\ell$ .

With these core types, the core linker needs no sources. Rule **CL-LINK** checks that  $P$  supplies the client's import and emits  $e_\ell$ , the same term that **ELAB-LINK** produced:

$$\frac{\text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \in P}{m' : P \ \& \ n' : D \rightarrow T \rightsquigarrow_{\text{link}} e_\ell} \text{CL-LINK}$$

<sup>2</sup>The examples assume integer arithmetic primitives, omitted from the formal grammar.

Table 1. Separate compilation with linksets and  $\lambda_{\text{SCE}}$ .

|                      | Linksets [7]          | $\lambda_{\text{SCE}}$                    |
|----------------------|-----------------------|-------------------------------------------|
| Compiled modules     | Open fragments        | Closed terms                              |
| Interfaces           | Typing environments   | Ordinary types                            |
| Linking              | Operation on linksets | Ordinary expression                       |
| Import resolution    | Term substitution     | Argument passing and lookup               |
| Separate compilation | Type preservation     | Type, semantics and behavior preservation |

The rule does not typecheck  $m'$  or  $n'$ . Given their typings, which elaboration guarantees, Link typing (Theorem 4.11) derives  $\varepsilon \vdash e_\ell : P \ \& \ T$ .

*Linking semantics and preservation.* A functor evaluates to a module closure  $\langle v, \text{functor}(A) \ e \rangle$ , which pairs the functor with its captured environment  $v$ . Let  $v_A$  be  $\text{Add}$ 's value,  $c_+$  its addition closure, and  $b$   $\text{Double}$ 's body. With  $\rho = \varepsilon \circ \{\text{op} = c_+\}$  and an empty ambient environment, the following abridged reduction supplies  $\text{op}$  to the client and retains both modules' exports:

$$\begin{aligned}
 \text{Add} \bowtie \text{Double} &\longrightarrow^* v_A \bowtie \langle \varepsilon, \text{functor}(D) \ b \rangle \longrightarrow v_A \circ (\rho \triangleright b) \\
 &\longrightarrow^* v_A \circ \underbrace{\text{struct} \ \{\text{double} = \langle \rho, \lambda x:\text{Int}. \text{op } x \ x \rangle\}}_{v_\ell}
 \end{aligned}$$

The semantics of Section 3.3 gives this reduction both for separately compiled files linked at top level and for modules linked inside a program. Semantics preservation (Theorem 4.8) ensures that  $e_\ell$  evaluates to a value related to  $v_\ell$  by elaboration.

*Separate compilation correctness.* We can now answer the last question: separate compilation correctness is a corollary of semantics and behavior preservation. Table 1 summarizes all three answers.

The upper square of Figure 5 compares the two paths. Rule **ELAB-LINK** elaborates the source link. That link is the merged program, and the source semantics gives its meaning. Rule **CL-LINK** links the independently compiled operands using only their interfaces. Both paths produce the same core term  $e_\ell$ , so elaboration commutes with linking by construction. The lower square is semantics preservation, for terminating runs. With behavior preservation (Theorem 4.9), it yields separate compilation correctness for every well-typed link, at top level or in a program (Corollaries 4.12 and 4.13).




Fig. 5. Correctness diagram.

## 2.4 Programming with First-Class Linking and Environments

Linkset linking (Section 2.2) runs before execution, so it cannot take a provider that a program builds at run time. In  $\lambda_{\text{SCE}}$ , a link is an expression. So a program can load a compiled client at run time and link it against a provider built per call. Figure 6 builds a page renderer whose formatting plugin is such a client. The figure omits the imports of `Sys` and `Loader` and the union declarations. The host supplies logging and falls back to plain text if loading fails.

*Compile the client once.* First, the plugin in panel (a) imports `Cap` for logging and exports `render`. Compilation produces a sandboxed functor. Its body can call `Cap.log` but cannot name the host's

```

393 (a) Compile the plugin separately: fancy.sce to fancy.sceo
394 import Cap : { log : String -> Unit }
395 let render (text : String) : String =
396   let noted : Unit = Cap.log "rendering" in
397   "** " ^ text ^ " **"
398
399 (b) Build a logger, then package it with a message
400 let caps (name : String) =
401   { log = fun (s : String) -> Sys.print ("[" ^ name ^ "]" " ^ s) }
402 let host (cap : { log : String -> Unit }) (text : String) = struct
403   let message : String = text
404   let Cap = cap
405 end
406
407 (c) Load the client, link it, and run under the resulting environment
408 let page (cap : { log : String -> Unit })
409   (path text : String) : String =
410   let saved = host cap text in
411   match Loader.load path with
412   | Plugin p ->
413     let lib = link saved with p in
414     box lib in render message
415   | Failed e -> box saved in message
416 end
417
418 let aliceCap = caps "alice"
419 page aliceCap "fancy.sceo" "hello"
420   (* prints "[alice] rendering"; returns "** hello **" *)
421 page aliceCap "fancy.sceo" "bye"
422   (* prints "[alice] rendering"; returns "** bye **" *)
423 page aliceCap "missing.sceo" "hello" (* returns "hello" *)

```

Fig. 6. A plugin host in three steps. Comments show the output.

Sys, although Sys is in scope where the host links it (Section 2.3). Sandboxed functors elaborate independently of their ambient context (Lemma 4.1).

*Construct the provider.* Next, caps in panel (b) builds a logger that captures the caller's name, and host packages it with a message. Both successful calls reuse aliceCap, although their providers carry different messages. Linking passes only Cap to the plugin, which exports the same render closure in both cases. This illustrates interface abstraction (Theorem 4.14). The linked result still contains message. A no-op logger would silence the plugin without recompiling it.

*Link, then choose the environment.* Finally, page in panel (c) loads the plugin. On success, linking saved with p produces lib. The box evaluates render message under lib, using the plugin's function and the provider's message. The calls return different strings because the host supplies different messages to that function. On failure, the box reads the message from the saved provider.

The declared type of Loader.load gives p the type of panel (a), so the compiler checks the link before loading. At run time, Loader.load returns Plugin p only for a compiled file of that type.

The example illustrates the linking construction covered by separate compilation correctness (Corollary 4.12). Strings, I/O, and artifact loading lie outside the formal calculus.

|               |                                                                                                                                                                                                                                                                                                               |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Types         | $A, B, \Gamma ::= \text{Int} \mid \varepsilon \mid A \rightarrow B \mid A \& B \mid \{l : A\} \mid \text{sig } A \mid A \rightarrow_m B$                                                                                                                                                                      |
| Functor modes | $s ::= \blacksquare \mid \square$                                                                                                                                                                                                                                                                             |
| Expressions   | $e ::= \text{env} \mid i \mid \varepsilon \mid e.n \mid e.l \mid \lambda A. e \mid \langle v, \lambda A. e \rangle \mid e_1 e_2 \mid \{l = e\}$<br>$\mid e_1 \triangleright e_2 \mid e_1 \circ e_2 \mid e_1 \circ\circ e_2 \mid \text{let } e_1 \text{ in } e_2 \mid \text{open } e_1 \text{ in } e_2 \mid m$ |
| Modules       | $m ::= \text{struct } e \mid \text{functor}_s(A) e \mid \langle v, \text{functor}(A) e \rangle \mid e_1 \bullet e_2 \mid e_1 \triangleright e_2$                                                                                                                                                              |
| Values        | $v ::= i \mid \varepsilon \mid \langle v, \lambda A. e \rangle \mid \langle v, \text{functor}(A) e \rangle \mid \text{struct } v \mid v_1 \circ v_2 \mid \{l = v\}$                                                                                                                                           |
| Sugar         | $\underline{l} \triangleq \text{env}.l \quad \underline{n} \triangleq \text{env}.n$                                                                                                                                                                                                                           |

Fig. 7. Syntax of  $\lambda_{\text{SCE}}$ .

### 3 The Source Calculus: $\lambda_{\text{SCE}}$

$\lambda_{\text{SCE}}$  extends  $\lambda_E$  [50] with first-class modules and a linking expression. It keeps first-class environments and dependent merges [49]. The merged program is the source link, which  $\lambda_{\text{SCE}}$ 's semantics evaluates: Section 4.3 states separate compilation against that evaluation. Rule **ELAB-LINK** translates a link to the term that the core linker emits for compiled fragments. By Corollary 3.7, a well-typed program cannot get stuck, even for links in functions or with providers chosen at run time.

#### 3.1 Syntax

Figure 7 shows the syntax of  $\lambda_{\text{SCE}}$ . We use blue for modules and module types, and orange for the other extensions to  $\lambda_E$ . The black constructs are inherited from  $\lambda_E$ . These colors also distinguish the constructs in the typing and reduction rules.

*Types as contexts.* Metavariables  $A$ ,  $B$ , and  $\Gamma$  range over types. Any type can be a typing context. Types include integers ( $\text{Int}$ ), a unit type ( $\varepsilon$ ) for the empty environment, environment types ( $A \& B$ ), arrow types ( $A \rightarrow B$ ), and single-field record types ( $\{l : A\}$ ). An environment type pairs two environment components in order. There is no subtyping. Type equality is structural, so  $\text{Int} \& \varepsilon$  differs from  $\varepsilon \& \text{Int}$ , as  $(\text{Int} \& \text{Int}) \& \varepsilon$  does from  $\text{Int} \& (\text{Int} \& \varepsilon)$ . We associate chains of environment types and merges to the left. A multi-field record type abbreviates a chain of single-field ones.

*Module types.* Any module's type is its interface: an interface type for a structure, and a functor type for a functor. The interface type **sig**  $A$ , written with the keyword **sig**, classifies structures whose body has type  $A$ . The functor type  $A \rightarrow_m B$  classifies functors importing  $A$  and exporting  $B$ .

*Expressions.* Expressions include integers ( $i$ ), a unit expression ( $\varepsilon$ ), applications ( $e_1 e_2$ ), boxes ( $e_1 \triangleright e_2$ ), dependent merges ( $e_1 \circ e_2$ ), and lambda abstractions ( $\lambda A. e$ ). A lambda binder is nameless. Its annotation  $A$  supplies the parameter type when typing the body. The query **env** returns the current environment. Projection ( $e.n$ ) returns the  $n$ -th entry of the value of  $e$ , so a variable with de Bruijn index  $n$  [9] is **env**. $n$ , abbreviated  $\underline{n}$ . Records ( $\{l = e\}$ ) create labeled entries, which selection ( $e.l$ ) retrieves. Multi-field records are merges of single-field ones. The sugar  $\underline{l}$  abbreviates **env**. $l$ . A closure ( $\langle v, \lambda A. e \rangle$ ) saves the environment in which a function was created.

The orange constructs elaborate to core expressions (Section 4). A non-dependent merge ( $e_1 \circ\circ e_2$ ) evaluates both components under the same environment, without making the value of  $e_1$  available to  $e_2$ . The let binding (**let**  $e_1$  **in**  $e_2$ ) adds the value of  $e_1$  to the environment of  $e_2$ . It requires no annotation because typing synthesizes the type of  $e_1$  (Section 3.2). Finally, **open**  $e_1$  **in**  $e_2$  opens a named component:  $e_1$  must have a single-field record type  $\{l : A\}$ , and its field value extends the environment of  $e_2$ . If the field holds a structure,  $e_2$  can then select the structure's exports by label.

*Modules.* The blue constructs form the syntactic category  $m$  of modules, which is part of the expression grammar ( $e ::= \dots \mid m$ ). Their subterms are arbitrary expressions. In particular, the

|     |                                                                                                                                                                                                                                        |                                                                                                                                                                                  |                                                                                  |                                                                                                                             |
|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| 491 | $\boxed{l : A \in B}$                                                                                                                                                                                                                  | (Containment)                                                                                                                                                                    |                                                                                  |                                                                                                                             |
| 492 |                                                                                                                                                                                                                                        |                                                                                                                                                                                  |                                                                                  |                                                                                                                             |
| 493 | CTM-RCD                                                                                                                                                                                                                                | CTM-ANDL                                                                                                                                                                         | CTM-ANDR                                                                         | CTM-SIG                                                                                                                     |
| 494 | $\frac{}{l : A \in \{l : A\}}$                                                                                                                                                                                                         | $\frac{l : A \in B \quad l \notin \text{label}(C)}{l : A \in B \& C}$                                                                                                            | $\frac{l : A \in C \quad l \notin \text{label}(B)}{l : A \in B \& C}$            | $\frac{l : A \in B}{l : A \in \text{sig } B}$                                                                               |
| 495 |                                                                                                                                                                                                                                        |                                                                                                                                                                                  |                                                                                  |                                                                                                                             |
| 496 | $\boxed{\Gamma \vdash e : A \rightsquigarrow e'}$                                                                                                                                                                                      | (Typing and elaboration)                                                                                                                                                         |                                                                                  |                                                                                                                             |
| 497 |                                                                                                                                                                                                                                        |                                                                                                                                                                                  |                                                                                  |                                                                                                                             |
| 498 | ELAB-INT                                                                                                                                                                                                                               | ELAB-TOP                                                                                                                                                                         | ELAB-ENV                                                                         | ELAB-IDX                                                                                                                    |
| 499 | $\frac{}{\Gamma \vdash i : \text{Int} \rightsquigarrow i}$                                                                                                                                                                             | $\frac{}{\Gamma \vdash \varepsilon : \varepsilon \rightsquigarrow \varepsilon}$                                                                                                  | $\frac{}{\Gamma \vdash \text{env} : \Gamma \rightsquigarrow \text{env}}$         | $\frac{\Gamma \vdash e : A \rightsquigarrow e' \quad \text{lookup}(A, n) = B}{\Gamma \vdash e.n : B \rightsquigarrow e'.n}$ |
| 500 |                                                                                                                                                                                                                                        |                                                                                                                                                                                  |                                                                                  |                                                                                                                             |
| 501 | ELAB-LAM                                                                                                                                                                                                                               | ELAB-RCD                                                                                                                                                                         | ELAB-SEL                                                                         |                                                                                                                             |
| 502 | $\frac{}{\Gamma \& A \vdash e : B \rightsquigarrow e'}$                                                                                                                                                                                | $\frac{}{\Gamma \vdash e : A \rightsquigarrow e'}$                                                                                                                               | $\frac{}{\Gamma \vdash e : B \rightsquigarrow e' \quad l : A \in B}$             |                                                                                                                             |
| 503 |                                                                                                                                                                                                                                        |                                                                                                                                                                                  |                                                                                  |                                                                                                                             |
| 504 | $\frac{}{\Gamma \vdash \lambda A. e : A \rightarrow B \rightsquigarrow \lambda[A]. e'}$                                                                                                                                                | $\frac{}{\Gamma \vdash \{l = e\} : \{l : A\} \rightsquigarrow \{l = e'\}}$                                                                                                       | $\frac{}{\Gamma \vdash e.l : A \rightsquigarrow e'.l}$                           |                                                                                                                             |
| 505 | ELAB-APP                                                                                                                                                                                                                               | ELAB-BOX                                                                                                                                                                         | ELAB-MRG                                                                         |                                                                                                                             |
| 506 | $\frac{}{\Gamma \vdash e_1 : A \rightarrow B \rightsquigarrow e'_1}$                                                                                                                                                                   | $\frac{}{\Gamma \vdash e_1 : \Gamma_1 \rightsquigarrow e'_1}$                                                                                                                    | $\frac{}{\Gamma \vdash e_1 : A \rightsquigarrow e'_1}$                           |                                                                                                                             |
| 507 | $\frac{}{\Gamma \vdash e_2 : A \rightsquigarrow e'_2}$                                                                                                                                                                                 | $\frac{}{\Gamma_1 \vdash e_2 : A \rightsquigarrow e'_2}$                                                                                                                         | $\frac{}{\Gamma \& A \vdash e_2 : B \rightsquigarrow e'_2}$                      |                                                                                                                             |
| 508 | $\frac{}{\Gamma \vdash e_1 \ e_2 : B \rightsquigarrow e'_1 \ e'_2}$                                                                                                                                                                    | $\frac{}{\Gamma \vdash e_1 \triangleright e_2 : A \rightsquigarrow e'_1 \triangleright e'_2}$                                                                                    | $\frac{}{\Gamma \vdash e_1 \ , \ e_2 : A \& B \rightsquigarrow e'_1 \ , \ e'_2}$ |                                                                                                                             |
| 509 |                                                                                                                                                                                                                                        |                                                                                                                                                                                  |                                                                                  |                                                                                                                             |
| 510 | ELAB-CLOS                                                                                                                                                                                                                              | ELAB-NMRG                                                                                                                                                                        |                                                                                  |                                                                                                                             |
| 511 | $\frac{\varepsilon \vdash v : \Gamma_1 \rightsquigarrow v' \quad \Gamma_1 \& A \vdash e : B \rightsquigarrow e'}{\Gamma \vdash \langle v, \lambda A. e \rangle : A \rightarrow B \rightsquigarrow \langle v', \lambda[A]. e' \rangle}$ | $\frac{}{\Gamma \vdash e_1 : A \rightsquigarrow e'_1 \quad \Gamma \vdash e_2 : B \rightsquigarrow e'_2}$                                                                         |                                                                                  |                                                                                                                             |
| 512 |                                                                                                                                                                                                                                        | $\frac{}{\Gamma \vdash e_1 \ \text{open } e_2 : A \& B \rightsquigarrow (\text{env} \ , \ ((\underline{0} \triangleright e'_1) \ , \ (\underline{1} \triangleright e'_2))) . 0}$ |                                                                                  |                                                                                                                             |
| 513 | ELAB-LET                                                                                                                                                                                                                               | ELAB-OPEN                                                                                                                                                                        |                                                                                  |                                                                                                                             |
| 514 | $\frac{}{\Gamma \vdash e_1 : A \rightsquigarrow e'_1 \quad \Gamma \& A \vdash e_2 : B \rightsquigarrow e'_2}$                                                                                                                          | $\frac{}{\Gamma \vdash e_1 : \{l : A\} \rightsquigarrow e'_1 \quad \Gamma \& A \vdash e_2 : B \rightsquigarrow e'_2}$                                                            |                                                                                  |                                                                                                                             |
| 515 | $\frac{}{\Gamma \vdash \text{let } e_1 \text{ in } e_2 : B \rightsquigarrow (e'_1 \ , \ e'_2) . 0}$                                                                                                                                    | $\frac{}{\Gamma \vdash \text{open } e_1 \text{ in } e_2 : B \rightsquigarrow (\lambda[A]. e'_2) (e'_1.l)}$                                                                       |                                                                                  |                                                                                                                             |
| 516 |                                                                                                                                                                                                                                        |                                                                                                                                                                                  |                                                                                  |                                                                                                                             |

Fig. 8. Containment, typing, and elaboration of expressions.

operands of module application and linking need not be module forms. A functor may be selected from the environment or computed by another expression.

A structure (**struct**  $e$ ) evaluates its body under the ambient environment and packages the result as the value **struct**  $v$ . A functor (**functor** <sub>$s$</sub> ( $A$ )  $e$ ) abstracts a body over an import of type  $A$ . Like a lambda abstraction, a functor evaluates to a closure, here a module closure ( $\langle v, \text{functor}(A) \ e \rangle$ ). The functor's mode  $s$  determines the captured environment: empty in sandbox mode ( $\blacksquare$ ) and ambient in open mode ( $\square$ ). Module application ( $e_1 \bullet e_2$ ) supplies the functor  $e_1$  with its import  $e_2$ . Modules can be bound, merged, passed to functions, and returned.

Linking ( $e_1 \bowtie e_2$ ) connects a provider  $e_1$  to a client functor  $e_2$ . The client's import is a single-field record  $\{l : A\}$ , as for **open**, and its label  $l$  names an export of the provider. Linking selects that export, supplies it to the client, and merges the provider with the client's result. The provider can be any expression whose type contains the import, such as a function's argument or **env** (ELAB-ENV).

### 3.2 Typing

Figure 8 presents the typing and elaboration rules for expressions. The judgment  $\Gamma \vdash e : A \rightsquigarrow e'$  states that source expression  $e$  has type  $A$  under  $\Gamma$  and elaborates to the  $\lambda_E$  term  $e'$ . Elaborated terms are shaded gray, and  $\llbracket A \rrbracket$  translates a source type to a  $\lambda_E$  type. This subsection explains typing. Section 3.3 explains the term that rule ELAB-LINK generates, and Section 4.2 the type translation and other generated terms.

(Typing and elaboration)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <div style="border: 1px solid black; padding: 5px; margin-bottom: 10px; display: inline-block;"> <math>\Gamma \vdash e : A \rightsquigarrow e'</math> </div> <div style="margin-bottom: 10px;"> <math display="block">\frac{\Gamma \vdash e : B \rightsquigarrow e'}{\Gamma \vdash \text{struct } e : \text{sig } B \rightsquigarrow e'}</math> </div> <div> <math display="block">\frac{\varepsilon \&amp; A \vdash e : B \rightsquigarrow e'}{\Gamma \vdash \text{functor}_{\blacksquare}(A) \ e : A \rightarrow_m B \rightsquigarrow \varepsilon \triangleright (\lambda \llbracket A \rrbracket. e')}</math> </div> | <div style="margin-bottom: 10px;"> <math display="block">\frac{\Gamma \vdash e_1 : A \rightarrow_m B \rightsquigarrow e'_1 \quad \Gamma \vdash e_2 : A \rightsquigarrow e'_2}{\Gamma \vdash e_1 \bullet e_2 : B \rightsquigarrow e'_1 e'_2}</math> </div> <div style="margin-bottom: 10px;"> <math display="block">\frac{\Gamma \&amp; A \vdash e : B \rightsquigarrow e'}{\Gamma \&amp; A \vdash e : B \rightsquigarrow e'}</math> </div> <div> <math display="block">\frac{\varepsilon \vdash v : \Gamma_1 \rightsquigarrow v' \quad \Gamma_1 \&amp; A \vdash e : B \rightsquigarrow e'}{\Gamma \vdash \langle v, \text{functor}(A) \ e \rangle : A \rightarrow_m B \rightsquigarrow \langle v', \lambda \llbracket A \rrbracket. e' \rangle}</math> </div> |
| <div style="margin-bottom: 10px;"> <math display="block">\frac{\Gamma \vdash e_1 : \Gamma_1 \rightsquigarrow e'_1 \quad \Gamma \vdash e_2 : \{l : A\} \rightarrow_m B \rightsquigarrow e'_2 \quad l : A \in \Gamma_1}{\Gamma \vdash e_1 \bowtie e_2 : \Gamma_1 \&amp; B \rightsquigarrow (\lambda \llbracket \Gamma_1 \rrbracket. \lambda(\{l : \llbracket A \rrbracket\} \rightarrow \llbracket B \rrbracket). \underline{1}_9 (\underline{1} \{l = \underline{0}.l\})) e'_1 e'_2}</math> </div>                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Fig. 9. Typing and elaboration of modules.

*Environments and lookup.* The query `env` inspects the context, and has the current context  $\Gamma$  as its type (ELAB-ENV). A projection  $e.n$  has type `lookup(A, n)` when  $e$  has type  $A$  (ELAB-IDX). The partial function `lookup` counts from the right end, so index 0 is the rightmost entry:

$$\text{lookup}(A \& B, 0) = B \quad \text{lookup}(A \& B, n+1) = \text{lookup}(A, n)$$

A box  $e_1 \triangleright e_2$  replaces the ambient context  $\Gamma$  with the type of  $e_1$  when checking  $e_2$  (ELAB-BOX). Lambdas and dependent merges retain  $\Gamma$  and extend it on the right to  $\Gamma \& A$ , where  $A$  is the parameter type or the left operand's type, respectively. By rule ELAB-CLOS, a closure  $\langle v, \lambda A. e \rangle$  uses its captured environment: its body is checked under  $\Gamma_1 \& A$ , where  $v$  has type  $\Gamma_1$  under  $\varepsilon$ .


*Labels and containment.* Projection finds an environment entry by position. Selection and linking instead find a field by label. Both use the containment judgment  $l : A \in B$  of Figure 8: type  $B$  exposes exactly one field named  $l$ , and that field has type  $A$ . The function `label(·)` records the labels exposed by a type:

$$\begin{aligned} \text{label}(\{l : A\}) &= \{l\} & \text{label}(A \& B) &= \text{label}(A) \cup \text{label}(B) \\ \text{label}(\text{sig } A) &= \text{label}(A) & \text{label}(A) &= \emptyset \text{ otherwise} \end{aligned}$$

For an environment type, containment requires  $l$  in exactly one component. A merge may contain duplicate labels, but selecting a duplicated label or linking a client that imports it is ill typed.

*Let and open.* Like selection, `open` reads a field by its label, but it needs no containment. Its operand  $e_1$  must have exactly the type  $\{l : A\}$ . Rule ELAB-OPEN adds only  $A$  to the context. Rule ELAB-LET instead adds the whole synthesized type of  $e_1$ , so for the same  $e_1$  it adds  $\{l : A\}$ . Take  $e_1 = \{M = \text{struct } \{x = 1\}\}$ . After `open`  $e_1$  `in`  $e_2$ ,  $e_2$  selects the structure's export as  $\underline{x}$ . After `let`  $e_1$  `in`  $e_2$ ,  $e_2$  writes  $\underline{M}.x$ . Either way, the first label  $e_2$  selects,  $x$  or  $M$ , must be unique in the context.

*Modules and linking.* Figure 9 gives the module rules. Like a `let` or `open` body, a functor body is checked under an extended context. Its import type  $A$  extends the empty context or the ambient one, depending on the functor's mode. A sandboxed body is checked under  $\varepsilon \& A$ , so it accesses only its import. An open body is checked under  $\Gamma \& A$  and may also use ambient bindings. Both modes give the functor the type  $A \rightarrow_m B$ , where  $B$  is the type of its body. Open mode serves functors over ambient modules. Sandbox mode serves separately compiled clients.

Fig. 10. Small-step semantics of  $\lambda_{\text{SCE}}$ .

Rule **ELAB-LINK** constructs the import from the one provider field of the required name and type. Although  $\lambda_{\text{SCE}}$  has no subtyping, linking thus accepts a provider with additional exports. Module application returns only the client's result, whereas linking returns the provider merged with that result. For example, take `Add` from Figure 1, which exports `op : Int → Int → Int` and `unit : Int`, and a functor `Double'` that imports only `{op : Int → Int → Int}` and exports `double : Int → Int`. Direct application `Double' • Add` is ill typed, because `Add`'s interface differs from the required import record type. The adapted application `Double' • {op = Add.op}` returns only the client's structure, whereas `Add ◀ Double'` needs no adapter. The link also retains `op` and `unit`.

### 3.3 Small-Step Semantics

Figure 10 defines the call-by-value small-step semantics of  $\lambda_{\text{SCE}}$ . The judgment  $v \vdash e \longrightarrow e'$  states that  $e$  takes one step to  $e'$  under the ambient environment  $v$ , which is a value. We write  $v \vdash e \longrightarrow^* e'$  for its reflexive and transitive closure.

*Environments in reduction.* Frames  $F$  specify left-to-right evaluation, and rule **S-FRAME** preserves the ambient environment. Rules **S-ENV**, **S-BOXV**, **S-IDXV** and **S-SELV** return the environment, the

body's value, or the entry found. Boxes and dependent merges change the environment used for their right operands. A box  $e_1 \triangleright e_2$  replaces it with the value of  $e_1$  (**S-BoxR**), whereas a dependent merge extends it with the left operand's value (**S-MrgR**). Rules **S-LETV** and **S-OPENV** reduce **let** and **open** to boxes whose bodies run under the ambient environment extended with the bound value. A non-dependent merge leaves the environment unchanged: frames evaluate both operands under it, and rule **S-NMrgV** then turns the two values into a dependent merge.

A lambda captures the ambient environment in a closure (**S-LAM**). Rule **S-BETA** then runs the body in a box under the captured environment extended with the argument, without substitution. Functors work in the same way (**S-FUN**, **S-MBETA**), except that rule **S-FUN** captures  $\text{amb}_s(v)$ , which is the empty environment for a sandboxed functor. Rules **S-MrgR** and **S-BETA** both extend the environment on the right, so  $\text{lookupV}$  also counts from the right end:

$$\text{lookupV}(v_1 \circ v_2, 0) = v_2$$

$$\text{lookupV}(v_1 \circ v_2, n+1) = \text{lookupV}(v_1, n)$$

For example, take an ambient environment  $v$  and the term  $(\lambda A. \lambda B. e_1 \circ e_2) v_1 v_2$ . Its two beta steps run the body under  $(v \circ v_1) \circ v_2$ . The merge then adds the value  $v_3$  of  $e_1$  for its right operand. So the indices denote the following:

|                     | Environment                           | <u>0</u> | <u>1</u> |
|---------------------|---------------------------------------|----------|----------|
| Left operand $e_1$  | $(v \circ v_1) \circ v_2$             | $v_2$    | $v_1$    |
| Right operand $e_2$ | $((v \circ v_1) \circ v_2) \circ v_3$ | $v_3$    | $v_2$    |

The same index thus denotes different entries in the two operands of a dependent merge. Index 1 is  $v_1$  on the left and  $v_2$  on the right. This difference is what rule **ELAB-LINK** relies on. In its term  $\underline{1} \circ (\underline{1} \{l = \underline{0}.l\})$ , the left 1 denotes the provider. On the right, 1 denotes the client and 0 the provider.

Rule **S-LINKV** is module application followed by a merge: its right operand is the **S-MBETA** reduct of the client applied to  $\{l = v_l\}$ , and the provider is merged on the left. The link evaluates the provider first, then the client. Below, the provider `One` supplies  $x$  to a client `Inc` that adds 1:

$$\text{One} = \text{struct } \{x = 1\}, \quad \text{Inc} = \text{functor} \blacksquare (\{x : \text{Int}\}) (\underline{x} + 1)$$

$$\varepsilon \vdash \text{One} \bowtie \text{Inc} \longrightarrow^* \text{One} \circ ((\varepsilon \circ \{x = 1\}) \triangleright (\underline{x} + 1)) \longrightarrow^* \text{One} \circ 2$$

Under  $\varepsilon \circ \{x = 1\}$ , the box runs the sandboxed client `Inc`, so its body can access only the selected import. The body evaluates to 2, and the link to `One`  $\circ 2$ .

The link above selected  $x$  from `One`. Label selection  $v.l \rightsquigarrow v'$  searches both sides of a merge and looks through structures (**SEL-MrgL**, **SEL-MrgR**, **SEL-STR**). When the type of  $v$  contains  $l$ , only one side has  $l$ , so selection is unique, which makes well-typed evaluation deterministic (Theorem 4.15).

Rule **S-LINKV** uses only selection, a record, a box and merges, all of which  $\lambda_E$  has. Rule **ELAB-LINK** expresses this behavior as a  $\lambda_E$  term (Section 4.2), so  $\lambda_E$  needs no separate linking rule.

### 3.4 Type Safety

We state the typing properties of  $\lambda_{\text{SCE}}$  using the elaboration judgment with its core term hidden:

$$\Gamma \vdash e : A \triangleq \exists e', \Gamma \vdash e : A \rightsquigarrow e'$$

An expression has at most one type under a given context.

**THEOREM 3.1 (INFERENCE UNIQUENESS).** *If  $\Gamma \vdash e : A_1$  and  $\Gamma \vdash e : A_2$ , then  $A_1 = A_2$ .*

A value has the same type under every context.

**LEMMA 3.2 (VALUE WEAKENING).** *If  $\Gamma_1 \vdash v : A$ , then  $\Gamma_2 \vdash v : A$  for every  $\Gamma_2$ .*

Indexed lookup and label selection succeed on a well-typed value and return an entry of the type that lookup or containment gives.

LEMMA 3.3 (LOOKUP PROGRESS). *If  $\varepsilon \vdash v : A$  and  $\text{lookup}(A, n) = B$ , then  $\text{lookupV}(v, n) = v'$  for some value  $v'$  with  $\varepsilon \vdash v' : B$ .*

LEMMA 3.4 (SELECTION PROGRESS). *If  $\varepsilon \vdash v : A$  and  $l : B \in A$ , then  $v.l \rightsquigarrow v'$  for some value  $v'$  with  $\varepsilon \vdash v' : B$ .*

Boxes and dependent merges change the environment during evaluation, so we prove progress and preservation for an arbitrary environment  $v$  of type  $\Gamma$ . Because environments are values, the expression judgment types them too.

THEOREM 3.5 (PROGRESS). *If  $\varepsilon \vdash v : \Gamma$  and  $\Gamma \vdash e : A$ , then  $e$  is a value or  $v \vdash e \longrightarrow e'$  for some  $e'$ .*

THEOREM 3.6 (PRESERVATION). *If  $\varepsilon \vdash v : \Gamma$ ,  $\Gamma \vdash e : A$ , and  $v \vdash e \longrightarrow e'$ , then  $\Gamma \vdash e' : A$ .*

Although variables are de Bruijn indices, the proofs need no term-substitution or index-shifting lemmas: they use environment lookup and value weakening. For linking, Lemma 3.4 supplies the import value at the required type. The typing rules for records, boxes, and merges then establish preservation of the linked result. Instantiating both theorems with  $v = \varepsilon$  and  $\Gamma = \varepsilon$  gives type safety.

COROLLARY 3.7 (TYPE SAFETY). *If  $\varepsilon \vdash e : A$ , then  $e$  is a value or  $\varepsilon \vdash e \longrightarrow e'$  for some  $e'$ , and every such  $e'$  satisfies  $\varepsilon \vdash e' : A$ .*

## 4 Elaboration and Separate Compilation Correctness

Modules in  $\lambda_{\text{SCE}}$  elaborate to terms of the core calculus  $\lambda_E$  [50]. The elaboration of a source link is exactly the term that the core linker emits for the compiled fragments. Compiling and linking therefore commute by construction (the upper square of Figure 5). Separate compilation correctness then follows from semantics preservation (the lower square) and behavior preservation, for static and dynamic links alike. Section 4.1 fixes notation for  $\lambda_E$ , and Section 4.2 presents the elaboration. Section 4.3 proves correctness and interface abstraction.

### 4.1 Target calculus $\lambda_E$

The calculus  $\lambda_E$  has the expression and type forms shown in black in Figure 7. Tan and Oliveira [50] give the typing and semantics of  $\lambda_E$ . We write  $\vdash$  for source judgments and  $\vdash$  for core judgments. Where both calculi appear,  $u$  ranges over core values and  $v$  over source values.

*Typing and Semantics.* The typing judgment  $\Gamma \vdash e : A$  of  $\lambda_E$  uses a type as its context  $\Gamma$ , as in  $\lambda_{\text{SCE}}$ .  $\lambda_E$  is type safe [50], and our mechanization reproves progress and preservation for  $\lambda_E$ .

### 4.2 Type-Directed Elaboration

Elaboration turns a structure into its body, a functor into a function, and a link into a core term built from applications, selection, and merge. The target therefore needs no module constructs. Figures 8 and 9 pair each typing rule with its core term, and this subsection explains those terms.

*Type Translation.* The translation  $\llbracket \cdot \rrbracket$  maps source types to  $\lambda_E$  types. It is homomorphic on the constructors shared with  $\lambda_E$ , removes interface wrappers, and turns functor types into arrows:

$$\llbracket \text{sig } A \rrbracket = \llbracket A \rrbracket \quad \llbracket A \rightarrow_m B \rrbracket = \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket$$

Since contexts are types, the same translation applies to the typing context  $\Gamma$ .

*Core Expressions.* Elaboration preserves the expression forms shared with  $\lambda_E$  and recursively elaborates their subterms. Rules ELAB-LAM and ELAB-CLOS also translate the parameter annotation.

*Derived Expressions.* Let bindings, open expressions, and non-dependent merges elaborate to core environment operations. The examples evaluate under the ambient environment  $\{x = 7\}$ :

| Source expression                                                         | Elaborated term                                                                                          | Value                       |
|---------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|-----------------------------|
| <b>let</b> $\{y = 1\}$ <b>in</b> $\underline{y} + \underline{x}$          | $(\{y = 1\} \circ (\underline{y} + \underline{x})).0$                                                    | 8                           |
| <b>open</b> $\{m = \{y = 1\}\}$ <b>in</b> $\underline{y} + \underline{x}$ | $(\lambda\{y : \text{Int}\}. \underline{y} + \underline{x}) (\{m = \{y = 1\}\}.m)$                       | 8                           |
| $\{x = 1\} \circ \{z = \underline{x} + 1\}$                               | $(\text{env} \circ ((0 \triangleright \{x = 1\}) \circ (1 \triangleright \{z = \underline{x} + 1\}))).0$ | $\{x = 1\} \circ \{z = 8\}$ |

A let binding uses a dependent merge to extend the body's environment with the whole bound value, then projects the body's result (**ELAB-LET**). An open expression instead extends the environment with only the field value of its operand (**ELAB-OPEN**). The operand's synthesized type  $\{l : A\}$  supplies the selection label  $l$  and the lambda annotation  $\llbracket A \rrbracket : m$  and  $\{y : \text{Int}\}$  in the second example. Only the open encoding uses type information.

In a non-dependent merge, the outer **env** saves the ambient environment, which boxes restore for each operand (**ELAB-NMRG**). In the third example, the right operand sees the ambient  $x = 7$ , so  $z = 8$ .

*Modules.* Elaboration removes structure wrappers (**ELAB-STR**) and maps module closures and applications to their core counterparts (**ELAB-MCLOS**, **ELAB-MAPP**). An open functor elaborates to a lambda abstraction that captures the ambient environment (**ELAB-FUN**). A sandboxed functor elaborates to the boxed abstraction  $\varepsilon \triangleright (\lambda \llbracket A \rrbracket. e')$  (**ELAB-FUN**). The box makes the abstraction capture the empty environment. Rule **ELAB-FUN** ignores the ambient context, so a sandboxed functor elaborates the same way wherever it occurs. A sandboxed client therefore compiles once and links in any context:

LEMMA 4.1 (CONTEXT INDEPENDENCE). *If  $\Gamma_1 \vdash \text{functor} \llbracket A \rrbracket (A) e : C \rightsquigarrow e'$ , then*

$$\Gamma_2 \vdash \text{functor} \llbracket A \rrbracket (A) e : C \rightsquigarrow e' \quad \text{for every } \Gamma_2.$$

PROOF. Rule **ELAB-FUN** checks the body under  $\varepsilon \& A$ , independently of the ambient context.  $\square$

*Linking.* A link  $e_1 \bowtie e_2$ , with operand types  $\Gamma_1$  and  $\{l : A\} \rightarrow_m B$ , elaborates by rule **ELAB-LINK** to the following term. The named form identifies the provider  $p$  and client  $f$ :

$$\begin{array}{l|l} \text{Named} & (\lambda(p : \llbracket \Gamma_1 \rrbracket). \lambda(f : \{l : \llbracket A \rrbracket\} \rightarrow \llbracket B \rrbracket). p \circ (f \{l = p.l\})) e'_1 e'_2 \\ \text{Nameless} & (\lambda \llbracket \Gamma_1 \rrbracket. \lambda(\{l : \llbracket A \rrbracket\} \rightarrow \llbracket B \rrbracket). 1 \circ (1 \{l = 0.l\})) e'_1 e'_2 \end{array} \quad (1)$$

We write  $\text{link}_c(\llbracket \Gamma_1 \rrbracket, \{l : \llbracket A \rrbracket\}, \llbracket B \rrbracket, e'_1, e'_2)$  for this term. It generalizes the term  $e_\ell$  that links  $m'$  and  $n'$  in the overview to any import label and interfaces.

Across all rules, elaboration preserves types (Theorem 4.2) and produces a unique core term (Theorem 4.3). With Theorem 3.1, a well-typed expression and its context therefore determine the expression's type and elaboration. The elaboration of a link is thus well defined.

THEOREM 4.2 (TYPE PRESERVATION). *If  $\Gamma \vdash e : A \rightsquigarrow e'$  then  $\llbracket \Gamma \rrbracket \vdash e' : \llbracket A \rrbracket$ .*

THEOREM 4.3 (ELABORATION UNIQUENESS). *If  $\Gamma \vdash e : A_1 \rightsquigarrow e'_1$  and  $\Gamma \vdash e : A_2 \rightsquigarrow e'_2$  then  $e'_1 = e'_2$ .*

### 4.3 Separate Compilation Correctness

A sandboxed client is compiled against only the *interface* of its import, and links with any provider whose type contains it. We derive separate compilation correctness from semantics and behavior preservation, and show that the client's result depends on its provider only through the import.

*Big-Step Semantics.* We use big-step evaluation to state semantics preservation. The judgment  $v \vdash e \Rightarrow v'$  states that  $e$  evaluates to the value  $v'$  under the ambient environment  $v$ . Its rules evaluate subexpressions under the same environments as the small-step rules, and Appendix A gives the rules in full. The rule for a link mirrors rule **S-LINKV**:

$$\text{EVAL-LINK} \quad \frac{v \vdash e_1 \Rightarrow v_1 \quad v \vdash e_2 \Rightarrow \langle v_2, \text{functor}(\{l : A\}) e \rangle \quad v_1.l \rightsquigarrow v_l \quad v_2 \circ \{l = v_l\} \vdash e \Rightarrow v_3}{v \vdash e_1 \bowtie e_2 \Rightarrow v_1 \circ v_3}$$

Big-step evaluation agrees with small-step reduction to a value. The results on terminating runs below use big-step evaluation, and Theorem 4.4 carries them to small-step reduction:

**THEOREM 4.4 (BIG-STEP/SMALL-STEP EQUIVALENCE).**  $v \vdash e \Rightarrow v'$  if and only if  $v \vdash e \longrightarrow^* v'$ .

For  $\lambda_E$ , we write  $u \vdash e \Rightarrow u'$  for big-step evaluation, which the mechanization also proves equivalent to small-step reduction.

*Semantics Preservation.* Elaboration preserves evaluation: a terminating source expression and its elaboration produce values related by elaboration. The proof uses three facts: source values elaborate to core values, and indexed lookup and label selection commute with elaboration. We use the same notation for the corresponding core operations.

**LEMMA 4.5 (VALUE ELABORATION).** If  $\Gamma \vdash v : A \rightsquigarrow e'$ , then  $e'$  is a value.

**LEMMA 4.6 (LOOKUP PRESERVATION).** If  $\varepsilon \vdash v : A \rightsquigarrow u$ ,  $\text{lookup}(A, n) = B$ , and  $\text{lookupV}(v, n) = v'$ , then  $\exists u'$ ,  $\text{lookupV}(u, n) = u'$  and  $\varepsilon \vdash v' : B \rightsquigarrow u'$ .

**LEMMA 4.7 (SELECTION PRESERVATION).** If  $\varepsilon \vdash v : A \rightsquigarrow u$ ,  $l : B \in A$ , and  $v.l \rightsquigarrow v'$ , then  $\exists u'$ ,  $u.l \rightsquigarrow u'$  and  $\varepsilon \vdash v' : B \rightsquigarrow u'$ .

**THEOREM 4.8 (SEMANTICS PRESERVATION).** If  $\Gamma \vdash e : A \rightsquigarrow e'$ ,  $\varepsilon \vdash v : \Gamma \rightsquigarrow u$ , and  $v \vdash e \Rightarrow v'$ , then  $\exists u'$ ,  $u \vdash e' \Rightarrow u'$  and  $\varepsilon \vdash v' : A \rightsquigarrow u'$ .

**PROOF SKETCH.** By induction on the source evaluation, using Lemmas 4.5–4.7 for queries, indexed projections, and field selection. In the linking case, the induction hypotheses and Lemma 4.7 give the core provider, client, and export, and the term (1) runs the client body as rule **EVAL-LINK** does.  $\square$

Theorem 4.8 runs from source to core and speaks only of terminating runs. Theorem 4.9 adds the converse and divergence. Big-step evaluation cannot express a run that never ends, so we write  $v \vdash e \Uparrow$  when  $e$  has an infinite small-step reduction under  $v$ , and likewise  $u \vdash e' \Uparrow$  for  $\lambda_E$ .

**THEOREM 4.9 (BEHAVIOR PRESERVATION).** If  $\Gamma \vdash e : A \rightsquigarrow e'$  and  $\varepsilon \vdash v : \Gamma \rightsquigarrow u$ , then

- (i)  $u \vdash e' \Rightarrow u'$  implies  $v \vdash e \Rightarrow v'$  and  $\varepsilon \vdash v' : A \rightsquigarrow u'$  for some  $v'$ ; and
- (ii)  $v \vdash e \Uparrow$  if and only if  $u \vdash e' \Uparrow$ .

**PROOF SKETCH.** An auxiliary simulation relation matches each source step with zero or more core steps. Zero-step matches strictly decrease a natural-number measure on source terms, so infinite source reductions yield infinite core reductions. Termination and divergence reflection then follow from source type safety, semantics preservation, and core determinism.  $\square$

Fixpoints or iso-recursive types, added in Section 5.1, let terms diverge. Like every theorem in the paper, Theorem 4.9 covers that extended language.

Taking  $\Gamma = \varepsilon$  and  $v = u = \varepsilon$  in Theorems 4.8 and 4.9 gives correctness for whole programs.

**COROLLARY 4.10 (WHOLE-PROGRAM CORRECTNESS).** If  $\varepsilon \vdash e : A \rightsquigarrow e'$ , then  $e$  and  $e'$  have the same behavior under  $\varepsilon$ :

- (a)  $\varepsilon \vdash e \Rightarrow v'$  implies  $\varepsilon \vdash e' \Rightarrow u'$  and  $\varepsilon \vdash v' : A \rightsquigarrow u'$  for some  $u'$ ;
- (b)  $\varepsilon \vdash e' \Rightarrow u'$  implies  $\varepsilon \vdash e \Rightarrow v'$  and  $\varepsilon \vdash v' : A \rightsquigarrow u'$  for some  $v'$ ; and
- (c)  $\varepsilon \vdash e \Uparrow$  if and only if  $\varepsilon \vdash e' \Uparrow$ .

**Core Linking.** Corollary 4.10 covers a program compiled as a whole. The *core linker* instead links separately compiled fragments. It receives two core fragments and their interfaces, not their sources. We model it by the judgment

$$e_1 : I_1 \bowtie e_2 : \{l : I_l\} \rightarrow I_2 \rightsquigarrow_{\text{link}} e_\ell$$

The judgment links a provider  $e_1$  of interface  $I_1$  and a client  $e_2$  into the term  $e_\ell$ . Here the client imports a field  $l$  of type  $I_l$  and exports  $I_2$ . The interfaces  $I_1$ ,  $I_l$  and  $I_2$  play the roles of  $\llbracket \Gamma_1 \rrbracket$ ,  $\llbracket A \rrbracket$  and  $\llbracket B \rrbracket$  in term (1). Rule **CL-LINK** checks containment and emits the term of rule **ELAB-LINK**:

CL-LINK

$$\frac{l : I_l \in I_1}{e_1 : I_1 \bowtie e_2 : \{l : I_l\} \rightarrow I_2 \rightsquigarrow_{\text{link}} \text{link}_c(I_1, \{l : I_l\}, I_2, e_1, e_2)}$$

The judgment uses only the fragments and their interfaces, so it has no context index and does not check the fragments' types. Theorem 4.11 assumes their typing under any core context  $\Gamma_c$ .

**THEOREM 4.11 (LINK TYPING).** *If  $e_1 : I_1 \bowtie e_2 : \{l : I_l\} \rightarrow I_2 \rightsquigarrow_{\text{link}} e_\ell$ ,  $\Gamma_c \vdash e_1 : I_1$ , and  $\Gamma_c \vdash e_2 : \{l : I_l\} \rightarrow I_2$ , then  $\Gamma_c \vdash e_\ell : I_1 \& I_2$ .*

The result keeps both modules' exports, so it can be the provider of a further link.

**Reusing a Compiled Client.** Rule **CL-LINK** lets one compiled client serve providers it was never compiled against. Take  $n'$ , the compiled Double of Section 2.3, whose interface is  $\{\text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}\} \rightarrow \top$  with  $\top = \{\text{double} : \text{Int} \rightarrow \text{Int}\}$ , and a provider compiled with another interface:

$$e_1 = \{\text{unit} = 1\}, \{\text{op} = \lambda \text{Int}. \lambda \text{Int}. \underline{1} \times \underline{0}\} \quad I_1 = \{\text{unit} : \text{Int}\} \& \{\text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}\}$$

Since  $\text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \in I_1$  holds, rule **CL-LINK** derives

$$e_1 : I_1 \bowtie n' : \{\text{op} : \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}\} \rightarrow \top \rightsquigarrow_{\text{link}} e_\ell$$

Here  $\text{op}$  multiplies, so  $e_\ell.\text{double } 5$  evaluates to 25. Type equality is structural, so  $I_1$  differs from the interface  $P$  of the compiled Add, which lists  $\text{op}$  first. The client  $n'$  is still reused unchanged, because rule **CL-LINK** checks only that  $I_1$  contains  $\text{op}$ .

**Separate Compilation.** Separate compilation correctness is a corollary of semantics and behavior preservation, because the core linker emits the elaboration of the source link. The provider and the client may depend on the ambient context  $\Gamma$ , for example when a function receives its provider as an argument. The linked core term then runs under the elaboration of the source environment.

In the corollaries, the provider's interface  $I_1$ , the import type  $I_l$  and the client's exports  $I_2$  are source types. The core linker receives their translations  $\llbracket I_1 \rrbracket$ ,  $\llbracket I_l \rrbracket$  and  $\llbracket I_2 \rrbracket$ .

**COROLLARY 4.12 (DYNAMIC SEPARATE COMPILATION).** *If*

- (1)  $\Gamma \vdash e_1 : I_1 \rightsquigarrow e'_1 \wedge l : I_l \in I_1$
- (2)  $\Gamma \vdash e_2 : \{l : I_l\} \rightarrow_m I_2 \rightsquigarrow e'_2$
- (3)  $e'_1 : \llbracket I_1 \rrbracket \bowtie e'_2 : \{l : \llbracket I_l \rrbracket\} \rightarrow \llbracket I_2 \rrbracket \rightsquigarrow_{\text{link}} e_\ell$
- (4)  $\varepsilon \vdash v : \Gamma \rightsquigarrow u$

*Then  $e_1 \bowtie e_2$  under  $v$  and  $e_\ell$  under  $u$  have the same behavior:*

- (a)  $v \vdash e_1 \bowtie e_2 \Rightarrow v'$  implies  $u \vdash e_\ell \Rightarrow u'$  and  $\varepsilon \vdash v' : I_1 \& I_2 \rightsquigarrow u'$  for some  $u'$ ;
- (b)  $u \vdash e_\ell \Rightarrow u'$  implies  $v \vdash e_1 \bowtie e_2 \Rightarrow v'$  and  $\varepsilon \vdash v' : I_1 \& I_2 \rightsquigarrow u'$  for some  $v'$ ; and
- (c)  $v \vdash e_1 \bowtie e_2 \Uparrow$  if and only if  $u \vdash e_\ell \Uparrow$ .

**PROOF SKETCH.** With (1)–(3) and rules **CL-LINK** and **ELAB-LINK**, we have  $\Gamma \vdash e_1 \bowtie e_2 : I_1 \& I_2 \rightsquigarrow e_\ell$ . Theorem 4.8, applied to this elaboration with (4), yields (a), and Theorem 4.9 yields (b) and (c).  $\square$

| Property                                                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Context independence<br>(Lemma 4.1)                         | A sandboxed functor $f = \text{functor}_{\blacksquare}(A)$ $e$ has the same type and elaborates to the same core term in every context.<br>$\Gamma_1 \vdash f : C \rightsquigarrow e' \implies \Gamma_2 \vdash f : C \rightsquigarrow e'$ , for all $\Gamma_1, \Gamma_2$ .                                                                                                                                                   |
| Interface abstraction<br>(Theorem 4.14)                     | Fix a client $e$ of type $\{l : A\} \rightarrow_m B$ under $\Gamma$ , and an environment $v$ of type $\Gamma$ . Provider values $p_1, p_2$ that contain $\{l : A\}$ in their types and select the same import value give<br>$v \vdash p_1 \bowtie e \Rightarrow p_1 \circ w \iff v \vdash p_2 \bowtie e \Rightarrow p_2 \circ w$ .                                                                                           |
| Separate compilation<br>(Corollary 4.12;<br>Corollary 4.13) | Linking independently compiled modules preserves the result of their source link. For $\Gamma \vdash e_1 \bowtie e_2 : I_1 \& I_2$ , let $e_\ell$ link their separately elaborated core terms. If $\varepsilon \vdash v : \Gamma \rightsquigarrow u$ , then<br>$v \vdash e_1 \bowtie e_2 \Rightarrow v' \implies \exists u', u \vdash e_\ell \Rightarrow u' \wedge \varepsilon \vdash v' : I_1 \& I_2 \rightsquigarrow u'$ . |

Fig. 11. Properties of linking and separate compilation.

A sandboxed client is compiled once, under  $\varepsilon$ , before any provider or scope is chosen: by Lemma 4.1, its compiled term satisfies (2) for every  $\Gamma$ . An open client's compiled term may refer to  $\Gamma$ , so it is valid only under the context  $\Gamma$  in which it was compiled. The corollary also covers such clients, each in its own context, as well as clients that are variables bound to loaded modules. Loading itself lies outside the calculus: through premise (4), the corollary assumes that a loaded module's value elaborates to its compiled form. Static linking instead compiles both fragments under  $\varepsilon$  and runs them under the empty environment.

**COROLLARY 4.13 (STATIC SEPARATE COMPILATION).** *Assume (1)–(3) of Corollary 4.12 with  $\Gamma = \varepsilon$ . Then (a)–(c) hold with  $v = u = \varepsilon$ .*

**PROOF.** Apply Corollary 4.12 with  $v = u = \varepsilon$ . □

In words, the core linker's output evaluates to the elaboration of the source link's result exactly when the source link terminates, and diverges exactly when the source link does. Both corollaries give this for a client compiled once and a separately compiled provider.

The corollaries cover a single link. An SCE-Lang program is a sequence of files, each linked against the earlier ones, or merged with them if it has no imports. Theorem C.1 extends Corollary 4.13 to such a sequence of binary links and merges with import-free files (Appendix C.1).

**Interface Abstraction.** A new provider preserves the client's result if the selected import value is unchanged. We fix the client, ambient environment, and import value, with its captured environments.

**THEOREM 4.14 (INTERFACE ABSTRACTION).** *Let  $p_1$  and  $p_2$  be values. Suppose  $\Gamma \vdash e : \{l : A\} \rightarrow_m B$ ,  $\varepsilon \vdash v : \Gamma$ , and, for  $i = 1, 2$ ,  $\varepsilon \vdash p_i : P_i$ ,  $l : A \in P_i$ , and  $p_i.l \rightsquigarrow v_l$ . Then, for every value  $w$ ,*

$$v \vdash p_1 \bowtie e \Rightarrow p_1 \circ w \iff v \vdash p_2 \bowtie e \Rightarrow p_2 \circ w$$

*and  $v \vdash p_1 \bowtie e \Uparrow$  if and only if  $v \vdash p_2 \bowtie e \Uparrow$ .*

**PROOF SKETCH.** Typing makes selection of  $l$  unique, so rule EVAL-LINK supplies the same record  $\{l = v_l\}$  to the client in both links. Reusing the client and body derivations proves each direction. A well-typed link diverges exactly when it does not terminate, which gives the divergence part. □

Figure 11 summarizes context independence, interface abstraction, and separate compilation correctness. Its entries omit divergence, and the separate compilation entry shows only clause (a).

|                   |                                                                                                                                                                                                                    |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Types             | $A, B, \Gamma ::= \dots \mid A \vee B \mid \alpha \mid \mu\alpha. A$                                                                                                                                               |
| Expressions       | $e ::= \dots \mid \mathbf{inl}_B e \mid \mathbf{inr}_A e \mid \mathbf{case } e \mathbf{ of } e_1 \mid e_2$<br>$\mid \mathbf{fix}(A \rightarrow B). e \mid \mathbf{fold } [\mu\alpha. A] e \mid \mathbf{unfold } e$ |
| Import interfaces | $D ::= \{l : A\} \mid D \& \{l : A\}$                                                                                                                                                                              |
| Modules           | $m ::= \dots \mid e_1 \bowtie e_2$                                                                                                                                                                                 |

Fig. 12. Unions and recursion extend both calculi, and multi-import linking extends  $\lambda_{\text{SCE}}$ .

**Sandbox and Link Confinement.** A link never passes its own ambient environment to the client. The client runs under the environment it captured, extended with the selected import. A sandboxed client captures the empty environment (**S-FUN**), even inside a function, as the box in its elaboration (**ELAB-FUN**) does. Its body therefore sees only that import: no binding in scope where the client is written or linked, and no other export of the provider. Since typing makes every selection unique, source evaluation is deterministic, and the agreement of results follows from determinism, both for a sandboxed functor applied directly and for links. For divergence, both runs take the same first steps, then run the same body under the same environment, so they diverge together.

**THEOREM 4.15 (DETERMINISM OF EVALUATION).** *If  $\Gamma \vdash e : A$ ,  $\varepsilon \vdash v : \Gamma$ ,  $v \vdash e \Rightarrow v_1$ , and  $v \vdash e \Rightarrow v_2$ , then  $v_1 = v_2$ .*

**THEOREM 4.16 (SANDBOX CONFINEMENT).** *Let  $\varepsilon \& A \vdash e : B$ ,  $\varepsilon \vdash v_l : A$ , and  $f = \mathbf{functor}_{\blacksquare}(A) e$ . For values  $\rho_1, \rho_2$ , if  $\rho_i \vdash f \bullet v_l \Rightarrow v_i$  for  $i = 1, 2$ , then  $v_1 = v_2$ . Moreover,  $\rho_1 \vdash f \bullet v_l \Uparrow$  if and only if  $\rho_2 \vdash f \bullet v_l \Uparrow$ .*

**THEOREM 4.17 (LINK CONFINEMENT).** *Let  $c = \langle \varepsilon, \mathbf{functor}(\{l : A\}) e \rangle$  with  $\varepsilon \vdash c : \{l : A\} \rightarrow_m B$ , and  $\varepsilon \vdash v_l : A$ . For  $i = 1, 2$ , let  $p_i$  be a value with  $p_i.l \rightsquigarrow v'$  exactly when  $v' = v_l$ . For values  $\rho_1, \rho_2$ , if  $\rho_i \vdash p_i \bowtie c \Rightarrow p_i \circ w_i$  for  $i = 1, 2$ , then  $w_1 = w_2$ . Moreover,  $\rho_1 \vdash p_1 \bowtie c \Uparrow$  if and only if  $\rho_2 \vdash p_2 \bowtie c \Uparrow$ .*

The client's exports  $w_i$  thus depend neither on the ambient environments  $\rho_i$  nor on the rest of the providers. The closure  $c$  captures the empty environment, which excludes a client that captured a nonempty one, such as an inline open functor. A sandboxed client evaluates to a closure over the empty environment wherever it is written (Section 2.3). For terminating runs, the client may be any expression that evaluates only to this closure, such as a variable bound to a loaded module.

## 5 Extensions, Implementation and Discussion

We extend the calculus and its metatheory (Section 5.1), present the WebAssembly compiler (Section 5.2), and discuss limitations and further extensions (Section 5.3).

### 5.1 Extensions

We add unions, recursive types, and fixpoints to both calculi, and multi-import linking to  $\lambda_{\text{SCE}}$  (Figure 12). Every theorem so far extends to these constructs, and Lean checks the new proof cases.

**Unions and Recursion.** Union types let the loader of Section 2.4 return either a plugin or an error. Matching the **Plugin** case binds the loaded plugin at the functor type required for linking. Iso-recursive types describe recursive data, and fixpoints define recursive functions. These constructs elaborate directly to their counterparts in  $\lambda_E$ ; Appendix B.1 gives the rules.

**Multi-Import Linking.** Multi-import linking  $e_1 \bowtie e_2$  selects a unique provider field of the required type for each import in  $D = \{l_1 : A_1\} \& \dots \& \{l_n : A_n\}$ . Elaboration replaces the single-field record in term (1) with  $\{l_1 = p.l_1\} \circ \dots \circ \{l_n = p.l_n\}$ , where  $p$  names the provider. Core linking constructs the same term from separately compiled modules (**CL-LINKN**): source and core links terminate with results related by elaboration or diverge together, for both static and dynamic links. For a

fixed well-typed client and ambient environment, links with typed provider values supplying the same import package yield the same client exports and diverge together. Separate compilation correctness also covers files combined in dependency order by linking against accumulated exports or merging import-free files. Rules and proofs appear in Appendix B.2 (Theorems B.4 and B.6) and Appendix C.1 (Theorems C.1 and C.2).

## 5.2 Implementation

SCE-Lang implements  $\lambda_{\text{SCE}}$  in about 4,500 lines of OCaml. The compiler parses and desugars source programs into  $\lambda_{\text{SCE}}$  terms, then typechecks and elaborates them into  $\lambda_E$ . It provides interpreters for both calculi and a WasmGC backend, and implements every extension above.

*Units and Linking.* SCE-Lang represents an importing file as a sandboxed functor over its imports, and an import-free file as a structure boxed under the empty environment. Each unit elaborates under the empty context to a core term, stored with the interfaces the loader checks. The linker chains units, saving each result as a unit for staged linking. Theorem C.2 proves chains correct.

*Compilation to WebAssembly.* The WasmGC backend compiles these units without a dedicated linking instruction. Figure 13 shows the compilation of  $\lambda_E$ 's environment operations, omitting tags and casts. At run time, a local  $\rho$  holds an environment reference, and `struct.new` allocates garbage-collected pairs and closures. The merge code pairs  $\rho$  with the value  $x$  of  $e_1$  to evaluate  $e_2$ , then pairs  $x$  with the result. The core linking term (1) needs only these operations, index lookup, records, and field selection. For multiple files, the backend emits a WebAssembly link module that imports unit entry points and links their results using the core construction of Theorem C.2. Saved linked units can likewise be compiled and linked again. Appendix C.2 gives the full rules; Appendix D traces the compilation of Add and Double.

## 5.3 Discussion and Limitations

The calculus is simply typed: signatures describe value exports, without ML-style abstract type components, sharing constraints, generativity, or polymorphism. These features would let separately compiled units hide representation choices and expose reusable interfaces over abstract types.

Subtyping offers a concrete next step. We could add an interface restriction  $e : A$ : when the type of  $e$  is a subtype of  $A$ , a cast restricts its value to  $A$ , discarding components outside that interface. This suggests the following module-application rule:

$$\frac{\Gamma \vdash e_f : D \rightarrow_m B \rightsquigarrow e'_f \quad \Gamma \vdash e_p : P \rightsquigarrow e'_p \quad \llbracket P \rrbracket <: \llbracket D \rrbracket}{\Gamma \vdash e_f \bullet e_p : B \rightsquigarrow e'_f (e'_p : \llbracket D \rrbracket)} \text{APP-SUB}$$

Here the inserted restriction supplies the functor's import interface. At the level of interfaces, this would allow applying Double directly to Add: restriction keeps only `op`. Application would thus perform import matching but return only the client's result: a form of *linking as application*. Disjointness at merges and context extensions would keep casts unambiguous. Such a discipline could draw on compositional programming [53] and type-directed operational semantics [21].

Sealing and polymorphism call for further extensions. One direction is to enrich the target with polymorphism and higher-order type operators, as in System  $F_\omega$ , while retaining first-class environments. We could then investigate an elaboration in the style of F-ing modules [40], with the aim of supporting ML-style modules alongside *first-class* linking and environments.

```

1030  C[env]ρ = local.get ρ
1031  C[e1 ▸ e2]ρ = C[e1]ρ local.set ρ' C[e2]ρ'
1032  C[e1 , e2]ρ = C[e1]ρ local.set x
1033                  local.get ρ local.get x struct.new $Pair local.set ρ' ;; ρ' = (ρ, x)
1034                  local.get x C[e2]ρ' struct.new $Pair
1035  C[λA. e]ρ = ref.func $f local.get ρ struct.new $Clos
1036              where the lifted function $f(self, x) is
1037                  local.get self struct.get $Clos env local.get x
1038                  struct.new $Pair local.set ρ' ;; ρ' = (self.env, x)
1039                  C[e]ρ'
1040  C[e1 e2]ρ = C[e1]ρ local.tee c C[e2]ρ
1041                  local.get c struct.get $Clos fn call_ref $Fn
1042
1043
1044
1045

```

Fig. 13. Compiling the environment operations of  $\lambda_E$  to WasmGC.

## 6 Related Work

*Program fragments and separate compilation.* Cardelli’s linksets (Section 2.2) formalize simply typed separate compilation, preserving typing under combination and linking [7]. SMLSC brings linksets to Standard ML [48]. Sewell proves that two ML-style structures build in sequence exactly when the second builds as a functor on the first [42]. Manifest types [26], Backpack [24] and MTAL [14] check units, packages and object files against imports. All link outside the language; none proves that linked code matches the merged sources. Corollary 4.13 proves this property.

*Linking as a language construct.* In Units, a compound expression links first-class units even inside functions [12]. Owens and Flatt type units [32]. Alice ML models components as functions of imports and static linking as an ordinary term [37]. MixML links mixins as a module construct [39]. Hirschowitz and Leroy compile a call-by-value CMS, a mixin calculus with mutual recursion, to a  $\lambda$ -calculus with records and let rec [3, 19]. These second-class mixins cannot be function arguments or results. Units and MixML support recursive linking, Owens and Flatt’s units have type components, and Alice ML loads components at run time. These capabilities exceed  $\lambda_{\text{SCE}}$ ; none of these typed designs proves that linking compiled code agrees with the source link.

Flatt and Felleisen state no theorem. Owens and Flatt prove type soundness, and Rossberg states it unproved. MixML and Hirschowitz and Leroy prove only type safety, via typed translations. CMS and MM, a call-by-value mixin calculus, have reduction semantics [3, 20], but no compiler is proved against them. Lee and Friedman give two untyped calculi,  $\lambda C$  and  $\lambda CD$  [25]. Source programs in  $\lambda C$  are contexts (terms with holes), with no linking construct. They prove that primitive linking of compiled code agrees with compiling source contexts after meta-level hole filling.  $\lambda CD$  adds definitions, combined with seq and rec, whose compiled forms link compiled components. For  $\lambda CD$ , they prove no such theorem. In  $\lambda_{\text{SCE}}$ , a link is a source expression, reduced by rule S-LINKV, and Corollary 4.12 covers it, even at run time.

Units’ implementation and Alice ML’s semantics represent clients as functions of their imports. Units’ formal model merges source text and contains no compiled code; Alice ML’s static linking is only sketched. In  $\lambda_{\text{SCE}}$ , this form is the link’s elaboration. Flatt and Felleisen argue informally for separate compilation of closed units; Lemma 4.1 proves that a sandboxed client compiles identically in every scope. Units and MixML check a component’s body under its enclosing scope, as our open functors do [12, 39]. Alice ML instead uses a custom component manager to restrict imports at run time [37]. A sandboxed  $\lambda_{\text{SCE}}$  client is closed by typing and, once compiled, by its box, as in  $\lambda_E$  [50], wherever it is written, and Theorem 4.17 proves its confinement. Corollary 4.12 proves that the

core linker agrees with source linking for every well-typed link, even under binders. We know of no other typed calculus with linking as an expression that states such a theorem.

*Modules by elaboration.* Module systems are often defined by elaboration [10, 38, 40]. Crary proves in Coq that his elaboration of ML modules is contextually correct and fully abstract [8]. These systems are richer than  $\lambda_{\text{SCE}}$ , with type abstraction, sharing and generativity, and full abstraction is stronger than our behavior preservation, which relates runs, not contextual equivalence. Yet none defines a linker for compiled fragments or a linking construct of its own. Harper and Lillibridge model linking as functor application and claim, without proof, that separate and whole-program compilation agree [16]. Corollary 4.13 proves a semantic counterpart for simple types.

*Dynamic linking.* Ancona et al. combine static and dynamic linking in one calculus, separating configuration and command languages [2]. They leave first-class linkset values as future work. In  $\lambda_{\text{SCE}}$ , links are ordinary expressions; providers and clients are first-class values. Links under binders realize Cardelli’s suggestion of imports not restricted to the top level [6]. Drossopoulou et al., Assemblages, Acute, dynamic rebinding and TAL/Load model run-time code loading [11, 18, 27, 43, 44], covering lazy resolution, re-linking and versioning, which  $\lambda_{\text{SCE}}$  does not. They prove soundness (conjectured for Acute), and Drossopoulou et al. and dynamic rebinding also prove strategies equivalent. None proves that run-time linking of compiled code agrees with source linking.

*Verified separate compilation.* Verified compilers treat linking outside the source language. For SepCompCert, Kang et al. prove in Coq that linked, separately compiled C files refine the linked source, using whole-program correctness for passes that commute with linking [23]. CertiCoq’s linker is a term template that plugs a library into an open client, and not every CertiCoq pass commutes with it [33]. Bowman and Ahmed link by substitution and prove that closure conversion commutes with it up to equivalence [5]. CakeML compiles code at run time but does not support partial compilation and linking [45]. SepCompCert and CertiCoq verify optimizing pipelines, far beyond our compiler. SepCompCert also relates I/O traces.  $\lambda_{\text{SCE}}$  has no I/O, and relates results and divergence. Still, none has a linking construct in its source, so commutation needs its own proof or fails. Nor do their theorems cover dynamic links. In  $\lambda_{\text{SCE}}$ , rule CL-LINK wraps compiled fragments in a fixed  $\lambda_E$  template, exactly the term that ELAB-LINK assigns to the source link. Commutation is thus a design property rather than a proof obligation, and Theorems 4.8 and 4.9 give correctness, static and dynamic.

Patterson and Ahmed distinguish linking with code from the same compiler, other compilers and other languages [34]. Ours is of the first kind, whereas Compositional CompCert, Pilsner, Perconti and Ahmed, and CompCertM cover the other kinds [30, 35, 46, 47]. Patterson and Ahmed warn that a theorem whose linker compiles source code is meaningless, as readers must trust the compiler. CL-LINK sees only compiled fragments and interfaces and emits a fixed  $\lambda_E$  template.

*First-class environments.* First-class environments turn bindings into values [13]; untyped languages use them for modules and run-time linking [22]. Queinnec and De Roure represent modules as functions of environments supplying their free variables [36]. Their import evaluates code in a supplied environment chained to the lexical one, but no theorem relates it to compiled code. The typed calculi of Nishizaki and Sato et al. prove subject reduction and confluence [31, 41], and Hashimoto and Otori suggest contexts for run-time linking [17], but none states a linking theorem.  $\lambda_E$  encodes sandboxed modules as boxed functions, with no module types and no link [49, 50]. Like definitional interpreters [1], its metatheory avoids term substitution. Closure conversion also pairs closed code with environments, but translates a substitution-based source [28]. Over  $\lambda_E$ ,  $\lambda_{\text{SCE}}$  adds the module language, links, the core linker and the separate compilation corollary. Commutation itself needs no first-class environments: any compositional translation of the link would do. They determine which links a program can write. A link can take the current scope `env` as its provider,

at a static type, and Corollary 4.12 covers it. A linked library can be the scope of a computation, as in  $\text{box } \text{lib in } e$ . Because a closure captures its whole environment, the  $\text{box } \varepsilon \triangleright (\lambda[A]. e')$  empties the environment of a compiled sandboxed client, even under a binder.

## 7 Conclusion

Separate compilation correctness needs an account of linking. Linking as evaluation makes that account the source semantics. The core linker emits the source link's elaboration, so compilation commutes with linking. Semantics and behavior preservation give links the same behavior at top level, under binders or at run time. Three ingredients shape the design: source reduction rules for linking, elaboration using only compiled operands and their interfaces, and contexts and interfaces as ordinary types, as in  $\lambda_E$  [49, 50]. Earlier work we know of has at most two. Cardelli's linksets, SMLSC and verified compilers lack source-level linking, and MixML and Alice ML define it only by translation. Units has the construct but merges source text. The first two yield separate compilation correctness, and the third decides which links a program can write.

Lean checks every proof, and the WebAssembly backend needs no linking instruction. Our results cover simple types. A concrete next step is subtyping, through an interface restriction. Sealing and polymorphism could build on a target that keeps first-class environments and adds the type operators of System  $F_\omega$ , with an elaboration in the style of F-ing modules [40].

## References

- [1] Nada Amin and Tiark Rompf. 2017. Type soundness proofs with definitional interpreters. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL '17)*. ACM, 666–679. doi:10.1145/3009837.3009866
- [2] Davide Ancona, Sonia Fagorzi, and Elena Zucca. 2003. A Calculus for Dynamic Linking. In *Theoretical Computer Science*. Springer, 284–301. doi:10.1007/978-3-540-45208-9\_23
- [3] Davide Ancona and Elena Zucca. 2002. A calculus of module systems. *Journal of Functional Programming* 12, 2 (March 2002), 91–132. doi:10.1017/s0956796801004257
- [4] Andrew W. Appel and David B. MacQueen. 1994. Separate compilation for Standard ML. In *Proceedings of the ACM SIGPLAN 1994 Conference on Programming Language Design and Implementation (Orlando, Florida, USA) (PLDI '94)*. Association for Computing Machinery, New York, NY, USA, 13–23. doi:10.1145/178243.178245
- [5] William J. Bowman and Amal Ahmed. 2018. Typed closure conversion for the calculus of constructions. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '18)*. ACM, 797–811. doi:10.1145/3192366.3192372
- [6] Luca Cardelli. 1989. *Typeful Programming*. SRC Research Report 45. Digital Equipment Corporation, Systems Research Center, Palo Alto, CA. Revised January 1, 1993.
- [7] Luca Cardelli. 1997. Program Fragments, Linking, and Modularization. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '97)*. ACM, 266–277. doi:10.1145/263699.263735
- [8] Karl Crary. 2019. Fully abstract module compilation. *Proceedings of the ACM on Programming Languages* 3, POPL (Jan. 2019), 1–29. doi:10.1145/3290323
- [9] N.G de Bruijn. 1972. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae (Proceedings)* 75, 5 (1972), 381–392. doi:10.1016/1385-7258(72)90034-0
- [10] Derek Dreyer, Karl Crary, and Robert Harper. 2003. A Type System for Higher-Order Modules. In *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2003)*. ACM, 236–249. doi:10.1145/604131.604151
- [11] Sophia Drossopoulou, Giovanni Lagorio, and Susan Eisenbach. 2003. Flexible Models for Dynamic Linking. In *Programming Languages and Systems*. Springer, 38–53. doi:10.1007/3-540-36575-3\_4
- [12] Matthew Flatt and Matthias Felleisen. 1998. Units: Cool Modules for HOT Languages. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '98)*. ACM, 236–248.
- [13] David Gelernter, Suresh Jagannathan, and Thomas London. 1987. Environments as First Class Objects. In *Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL)*. ACM, 98–110. doi:10.1145/41625.41634
- [14] Neal Glew and Greg Morrisett. 1999. Type-Safe Linking and Modular Assembly Language. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '99)*. ACM, 250–261. doi:10.1145/

- 292540.292563
- [15] Andreas Haas, Andreas Rossberg, Derek L. Schuff, Ben L. Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon Zakai, and JF Bastien. 2017. Bringing the web up to speed with WebAssembly. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Barcelona, Spain) (*PLDI 2017*). Association for Computing Machinery, New York, NY, USA, 185–200. doi:10.1145/3062341.3062363
  - [16] Robert Harper and Mark Lillibridge. 1994. A Type-Theoretic Approach to Higher-Order Modules with Sharing. In *Proceedings of the 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '94)*. ACM, 123–137. doi:10.1145/174675.176927
  - [17] Masatomo Hashimoto and Atsushi Ohori. 2001. A typed context calculus. *Theoretical Computer Science* 266, 1-2 (Sept. 2001), 249–272. doi:10.1016/s0304-3975(00)00174-2
  - [18] Michael Hicks, Stephanie Weirich, and Karl Cray. 2001. Safe and Flexible Dynamic Linking of Native Code. In *Types in Compilation*. Springer Berlin Heidelberg, 147–176. doi:10.1007/3-540-45332-6\_6
  - [19] Tom Hirschowitz and Xavier Leroy. 2005. Mixin modules in a call-by-value setting. *ACM Transactions on Programming Languages and Systems* 27, 5 (Sept. 2005), 857–881. doi:10.1145/1086642.1086644
  - [20] Tom Hirschowitz, Xavier Leroy, and J. B. Wells. 2004. Call-by-Value Mixin Modules: Reduction Semantics, Side Effects, Types. In *Programming Languages and Systems, 13th European Symposium on Programming, ESOP 2004 (Lecture Notes in Computer Science)*, David Schmidt (Ed.). Springer Berlin Heidelberg, 64–78. doi:10.1007/978-3-540-24725-8\_6
  - [21] Xuejing Huang and Bruno C. d. S. Oliveira. 2020. A Type-Directed Operational Semantics For a Calculus with a Merge Operator. In *34th European Conference on Object-Oriented Programming (ECOOP 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 166)*, Robert Hirschfeld and Tobias Pape (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 26:1–26:32. doi:10.4230/LIPIcs.ECOOP.2020.26
  - [22] Suresh Jagannathan. 1994. Metalevel building blocks for modular systems. *ACM Transactions on Programming Languages and Systems* 16, 3 (May 1994), 456–492. doi:10.1145/177492.177578
  - [23] Jeehoon Kang, Yoonseung Kim, Chung-Kil Hur, Derek Dreyer, and Viktor Vafeiadis. 2016. Lightweight Verification of Separate Compilation. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '16)*. ACM, 178–190. doi:10.1145/2837614.2837642
  - [24] Scott Kilpatrick, Derek Dreyer, Simon Peyton Jones, and Simon Marlow. 2014. Backpack: retrofitting Haskell with interfaces. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (*POPL '14*). Association for Computing Machinery, New York, NY, USA, 19–31. doi:10.1145/2535838.2535884
  - [25] Shinn-Der Lee and Daniel P. Friedman. 1996. Enriching the lambda calculus with contexts: toward a theory of incremental program construction. In *Proceedings of the first ACM SIGPLAN international conference on Functional programming (ICFP96)*. ACM, 239–250. doi:10.1145/232627.232652
  - [26] Xavier Leroy. 1994. Manifest types, modules, and separate compilation. In *Proceedings of the 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Portland, Oregon, USA) (*POPL '94*). Association for Computing Machinery, New York, NY, USA, 109–122. doi:10.1145/174675.176926
  - [27] Yu David Liu and Scott F. Smith. 2004. Modules with Interfaces for Dynamic Linking and Communication. In *ECOOP 2004 – Object-Oriented Programming*. Springer Berlin Heidelberg, 415–439. doi:10.1007/978-3-540-24851-4\_19
  - [28] Yasuhiko Minamide, Greg Morrisett, and Robert Harper. 1996. Typed closure conversion. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT symposium on Principles of programming languages - POPL '96 (POPL '96)*. ACM Press, 271–283. doi:10.1145/237721.237791
  - [29] Leonardo de Moura and Sebastian Ullrich. 2021. The Lean 4 Theorem Prover and Programming Language. In *Automated Deduction – CADE 28: 28th International Conference on Automated Deduction, Virtual Event, July 12–15, 2021, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 625–635. doi:10.1007/978-3-030-79876-5\_37
  - [30] Georg Neis, Chung-Kil Hur, Jan-Oliver Kaiser, Craig McLaughlin, Derek Dreyer, and Viktor Vafeiadis. 2015. Pilsner: a compositionally verified compiler for a higher-order imperative language. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming* (Vancouver, BC, Canada) (*ICFP 2015*). Association for Computing Machinery, New York, NY, USA, 166–178. doi:10.1145/2784731.2784764
  - [31] Shin-ya Nishizaki. 1994. Simply Typed Lambda Calculus with First-Class Environments. *Publications of the Research Institute for Mathematical Sciences* 30, 6 (Dec. 1994), 1055–1121. doi:10.2977/prims/1195164948
  - [32] Scott Owens and Matthew Flatt. 2006. From structures and functors to modules and units. In *Proceedings of the eleventh ACM SIGPLAN international conference on Functional programming (ICFP06)*. ACM, 87–98. doi:10.1145/1159803.1159815
  - [33] Zoe Paraskevopoulou, John M. Li, and Andrew W. Appel. 2021. Compositional optimizations for CertiCoq. *Proceedings of the ACM on Programming Languages* 5, ICFP (Aug. 2021), 1–30. doi:10.1145/3473591
  - [34] Daniel Patterson and Amal Ahmed. 2019. The Next 700 Compiler Correctness Theorems (Functional Pearl). *Proc. ACM Program. Lang.* 3, ICFP, Article 85 (Aug. 2019), 29 pages. doi:10.1145/3341689

- [35] James T. Perconti and Amal Ahmed. 2014. Verifying an Open Compiler Using Multi-language Semantics. In *Programming Languages and Systems*, Zhong Shao (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 128–148.
- [36] Christian Queinnec and David De Roure. 1996. Sharing Code through First-class Environments. In *Proceedings of the First ACM SIGPLAN International Conference on Functional Programming (ICFP)*. ACM, 251–261. doi:10.1145/232627.232653
- [37] Andreas Rossberg. 2006. The missing link: dynamic components for ML. In *Proceedings of the eleventh ACM SIGPLAN international conference on Functional programming (ICFP06)*. ACM, 99–110. doi:10.1145/1159803.1159816
- [38] Andreas Rossberg. 2018. 1ML – Core and Modules United. *Journal of Functional Programming* 28 (2018). doi:10.1017/S0956796818000205
- [39] Andreas Rossberg and Derek Dreyer. 2013. Mixin’ Up the ML Module System. *ACM Transactions on Programming Languages and Systems* 35, 1 (2013), 1–84. doi:10.1145/2450136.2450137
- [40] Andreas Rossberg, Claudio Russo, and Derek Dreyer. 2014. F-ing Modules. *Journal of Functional Programming* 24, 5 (2014), 529–607. doi:10.1017/S0956796814000264
- [41] Masahiko Sato, Takafumi Sakurai, and Yukiyooshi Kameyama. 2001. A simply typed context calculus with first-class environments. In *Functional and Logic Programming*, Herbert Kuchen and Kazunori Ueda (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 359–374.
- [42] Peter Sewell. 2001. Modules, abstract types, and distributed versioning. In *Proceedings of the 28th ACM SIGPLAN-SIGACT symposium on Principles of programming languages (POPL01)*. ACM, 236–247. doi:10.1145/360204.360225
- [43] Peter Sewell, James J. Leifer, Keith Wansbrough, Francesco Zappa Nardelli, Mair Allen-Williams, Pierre Habouzit, and Viktor Vafeiadis. 2007. Acute: High-level programming language design for distributed computation. *Journal of Functional Programming* 17, 4-5 (July 2007), 547–612. doi:10.1017/s0956796807006442
- [44] Peter Sewell, Gareth Stoye, Michael Hicks, Gavin Bierman, and Keith Wansbrough. 2008. Dynamic rebinding for marshalling and update, via redex-time and destruct-time reduction. *Journal of Functional Programming* 18, 4 (2008), 437–502. doi:10.1017/s0956796807006600
- [45] Thomas Sewell, Magnus O. Myreen, Yong Kiam Tan, Ramana Kumar, Alexander Mihajlovic, Oskar Abrahamsson, and Scott Owens. 2023. Cakes That Bake Cakes: Dynamic Computation in CakeML. *Proceedings of the ACM on Programming Languages* 7, PLDI (June 2023), 1121–1144. doi:10.1145/3591266
- [46] Youngju Song, Minki Cho, Dongjoo Kim, Yonghyun Kim, Jecheon Kang, and Chung-Kil Hur. 2020. CompCertM: CompCert with C-assembly linking and lightweight modular verification. *Proceedings of the ACM on Programming Languages* 4, POPL (Jan. 2020), 1–31. doi:10.1145/3371091
- [47] Gordon Stewart, Lennart Beringer, Santiago Cuellar, and Andrew W. Appel. 2015. Compositional CompCert. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL ’15)*. ACM, 275–287. doi:10.1145/2676726.2676985
- [48] David Swasey, Tom Murphy VII, Karl Cray, and Robert Harper. 2006. A Separate Compilation Extension to Standard ML. In *Proceedings of the 2006 Workshop on ML*. ACM, 32–42. doi:10.1145/1159876.1159883
- [49] Jinhao Tan and Bruno C. d. S. Oliveira. 2023. Dependent Merges and First-Class Environments. In *37th European Conference on Object-Oriented Programming (ECOOP 2023) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 263)*, Karim Ali and Guido Salvaneschi (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 34:1–34:32. doi:10.4230/LIPIcs.ECOOP.2023.34
- [50] Jinhao Tan and Bruno C. d. S. Oliveira. 2024. A Case for First-Class Environments. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 360 (Oct. 2024), 30 pages. doi:10.1145/3689800
- [51] Amin Timany, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2024. A Logical Approach to Type Soundness. *J. ACM* 71, 6 (2024), 1–75. doi:10.1145/3676954
- [52] A. K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *Information and Computation* 115, 1 (1994), 38–94. doi:10.1006/INCO.1994.1093
- [53] Weixin Zhang, Yaozhu Sun, and Bruno C. D. S. Oliveira. 2021. Compositional Programming. *ACM Trans. Program. Lang. Syst.* 43, 3, Article 9 (Sept. 2021), 61 pages. doi:10.1145/3460228

## A Big-Step Semantics of $\lambda_{\text{SCE}}$

Figure 14 gives the big-step semantics used in Section 4.3. The lemmas below establish its equivalence with small-step evaluation (Theorem 4.4) and show that evaluation produces values.

(Big-step evaluation)

|                                                                                                                                                                                                                                                                                                                        |                                                                                                |                                                                                                                                                                                                             |                                                                                                                                                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <div style="border: 1px solid black; padding: 2px; margin-bottom: 10px;"><math>v \vdash e \Rightarrow v'</math></div> <div style="text-align: center;">EVAL-VAL</div> $\frac{}{v \vdash v' \Rightarrow v'}$                                                                                                            | <div style="text-align: center;">EVAL-ENV</div> $\frac{}{v \vdash \mathbf{env} \Rightarrow v}$ | <div style="text-align: center;">EVAL-IDX</div> $\frac{v \vdash e \Rightarrow v_1 \quad \text{lookupV}(v_1, n) = v_2}{v \vdash e.n \Rightarrow v_2}$                                                        | <div style="text-align: center;">EVAL-LAM</div> $\frac{}{v \vdash \lambda A. e \Rightarrow \langle v, \lambda A. e \rangle}$                       |
| <div style="text-align: center;">EVAL-APP</div> $\frac{v \vdash e_1 \Rightarrow \langle v_1, \lambda A. e \rangle \quad v \vdash e_2 \Rightarrow v_2 \quad v_1, v_2 \vdash e \Rightarrow v'}{v \vdash e_1 e_2 \Rightarrow v'}$                                                                                         |                                                                                                | <div style="text-align: center;">EVAL-BOX</div> $\frac{v \vdash e_1 \Rightarrow v_1 \quad v_1 \vdash e_2 \Rightarrow v_2}{v \vdash e_1 \triangleright e_2 \Rightarrow v_2}$                                 |                                                                                                                                                    |
| <div style="text-align: center;">EVAL-MRG</div> $\frac{v \vdash e_1 \Rightarrow v_1 \quad v, v_1 \vdash e_2 \Rightarrow v_2}{v \vdash e_1, e_2 \Rightarrow v_1, v_2}$                                                                                                                                                  | <div style="text-align: center;">EVAL-RCD</div> $\frac{}{v \vdash e \Rightarrow v_1}$          |                                                                                                                                                                                                             | <div style="text-align: center;">EVAL-SEL</div> $\frac{v \vdash e \Rightarrow v_1 \quad v_1.l \rightsquigarrow v_2}{v \vdash e.l \Rightarrow v_2}$ |
| <div style="text-align: center;">EVAL-LET</div> $\frac{v \vdash e_1 \Rightarrow v_1 \quad v, v_1 \vdash e_2 \Rightarrow v_2}{v \vdash \mathbf{let} \, e_1 \, \mathbf{in} \, e_2 \Rightarrow v_2}$                                                                                                                      |                                                                                                | <div style="text-align: center;">EVAL-OPEN</div> $\frac{v \vdash e_1 \Rightarrow \{l = v_1\} \quad v, v_1 \vdash e_2 \Rightarrow v_2}{v \vdash \mathbf{open} \, e_1 \, \mathbf{in} \, e_2 \Rightarrow v_2}$ |                                                                                                                                                    |
| <div style="text-align: center;">EVAL-STR</div> $\frac{}{v \vdash e \Rightarrow v'}$                                                                                                                                                                                                                                   |                                                                                                | <div style="text-align: center;">EVAL-FUN</div> $\frac{}{v \vdash \mathbf{functor}_s(A) \, e \Rightarrow \langle \text{amb}_s(v), \mathbf{functor}(A) \, e \rangle}$                                        |                                                                                                                                                    |
| <div style="text-align: center;">EVAL-MAPP</div> $\frac{v \vdash e_1 \Rightarrow \langle v_1, \mathbf{functor}(A) \, e \rangle \quad v \vdash e_2 \Rightarrow v_2 \quad v_1, v_2 \vdash e \Rightarrow v'}{v \vdash e_1 \bullet e_2 \Rightarrow v'}$                                                                    |                                                                                                |                                                                                                                                                                                                             |                                                                                                                                                    |
| <div style="text-align: center;">EVAL-LINK</div> $\frac{v \vdash e_1 \Rightarrow v_1 \quad v \vdash e_2 \Rightarrow \langle v_2, \mathbf{functor}(\{l : A\}) \, e \rangle \quad v_1.l \rightsquigarrow v_l \quad v_2, v_3 \vdash \{l = v_l\} \vdash e \Rightarrow v_3}{v \vdash e_1 \bowtie e_2 \Rightarrow v_1, v_3}$ |                                                                                                |                                                                                                                                                                                                             |                                                                                                                                                    |

Fig. 14. Big-step semantics of  $\lambda_{\text{SCE}}$ .

LEMMA A.1 (SOUNDNESS OF BIG-STEP). If  $v \vdash e \Rightarrow v'$  then  $v \vdash e \longrightarrow^* v'$ .

LEMMA A.2 (COMPLETENESS OF BIG-STEP). If  $v \vdash e \longrightarrow^* e'$  and  $e'$  is a value, then  $v \vdash e \Rightarrow e'$ .

LEMMA A.3 (EVALUATION PRODUCES VALUES). If  $v \vdash e \Rightarrow e'$  then  $e'$  is a value.

## B Extensions

The following rules complete the extensions of Section 5.1.

### B.1 Unions, Recursive Functions, and Recursive Types

Unions and recursion extend both calculi with the syntax of Figure 12 and the following runtime forms:

|             |                                                                                                                                                                                                    |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Expressions | $e ::= \dots \mid \langle v, \mathbf{fix}(A \rightarrow B). e \rangle_f$                                                                                                                           |
| Values      | $v ::= \dots \mid \mathbf{inl}_B v \mid \mathbf{inr}_A v \mid \langle v, \mathbf{fix}(A \rightarrow B). e \rangle_f \mid \mathbf{fold} [\mu\alpha.A] v$                                            |
| Frames      | $F ::= \dots \mid \mathbf{inl}_B \square \mid \mathbf{inr}_A \square \mid \mathbf{case} \square \mathbf{of} \, e_1 \mid e_2 \mid \mathbf{fold} [\mu\alpha.A] \square \mid \mathbf{unfold} \square$ |

In the paper, types in judgments and annotations are closed: every type variable is bound by an enclosing  $\mu$ . We write  $A[B/\alpha]$  for capture-avoiding type substitution. The mechanization represents type variables by de Bruijn indices, and its fold constructor stores the body of the recursive-type annotation. The lemma `elab_substTyp` proves that type translation commutes with the mechanized substitution at any depth.

|                                                                                                                                                                                                          |                                                                                                                                                               |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| $\Gamma \vdash e : A \rightsquigarrow e'$                                                                                                                                                                | (Elaboration: extensions)                                                                                                                                     |
| ELAB-INJL                                                                                                                                                                                                | ELAB-INJR                                                                                                                                                     |
| $\Gamma \vdash e : A \rightsquigarrow e'$                                                                                                                                                                | $\Gamma \vdash e : B \rightsquigarrow e'$                                                                                                                     |
| $\Gamma \vdash \text{inl}_B e : A \vee B \rightsquigarrow \text{inl}_{\llbracket B \rrbracket} e'$                                                                                                       | $\Gamma \vdash \text{inr}_A e : A \vee B \rightsquigarrow \text{inr}_{\llbracket A \rrbracket} e'$                                                            |
| ELAB-CASE                                                                                                                                                                                                | ELAB-FIX                                                                                                                                                      |
| $\Gamma \vdash e : A \vee B \rightsquigarrow e'$                                                                                                                                                         | $\Gamma \& (A \rightarrow B) \& A \vdash e : B \rightsquigarrow e'$                                                                                           |
| $\Gamma \& A \vdash e_1 : C \rightsquigarrow e'_1 \quad \Gamma \& B \vdash e_2 : C \rightsquigarrow e'_2$                                                                                                | $\Gamma \vdash \text{fix}(A \rightarrow B). e : A \rightarrow B \rightsquigarrow \text{fix}(\llbracket A \rrbracket \rightarrow \llbracket B \rrbracket). e'$ |
| $\Gamma \vdash \text{case } e \text{ of } e_1 \mid e_2 : C \rightsquigarrow \text{case } e' \text{ of } e'_1 \mid e'_2$                                                                                  |                                                                                                                                                               |
| ELAB-FIXCLOS                                                                                                                                                                                             |                                                                                                                                                               |
| $\varepsilon \vdash v : \Gamma_1 \rightsquigarrow v' \quad \Gamma_1 \& (A \rightarrow B) \& A \vdash e : B \rightsquigarrow e'$                                                                          |                                                                                                                                                               |
| $\Gamma \vdash \langle v, \text{fix}(A \rightarrow B). e \rangle_f : A \rightarrow B \rightsquigarrow \langle v', \text{fix}(\llbracket A \rrbracket \rightarrow \llbracket B \rrbracket). e' \rangle_f$ |                                                                                                                                                               |
| ELAB-FOLD                                                                                                                                                                                                | ELAB-UNFOLD                                                                                                                                                   |
| $\Gamma \vdash e : A[\mu\alpha.A/\alpha] \rightsquigarrow e'$                                                                                                                                            | $\Gamma \vdash e : \mu\alpha.A \rightsquigarrow e'$                                                                                                           |
| $\Gamma \vdash \text{fold } [\mu\alpha.A] e : \mu\alpha.A \rightsquigarrow \text{fold } [\llbracket \mu\alpha.A \rrbracket] e'$                                                                          | $\Gamma \vdash \text{unfold } e : A[\mu\alpha.A/\alpha] \rightsquigarrow \text{unfold } e'$                                                                   |

The core typing rules are obtained by replacing each judgment  $\Gamma \vdash e : A \rightsquigarrow e'$  with  $\Gamma \vdash e : A$  and restricting all metavariables to the core syntax. Type preservation for fold and unfold uses the following lemma.

LEMMA B.1 (TRANSLATION COMMUTES WITH SUBSTITUTION). *If  $\mu\alpha.A$  is closed, then*

$$\llbracket A[\mu\alpha.A/\alpha] \rrbracket = \llbracket A \rrbracket [\llbracket \mu\alpha.A \rrbracket / \alpha].$$

|                                                                                                                                                                                   |                                                                                                                    |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| $v \vdash e \longrightarrow e'$                                                                                                                                                   | (Small-step reduction: extensions)                                                                                 |
| S-CASEL                                                                                                                                                                           | S-CASER                                                                                                            |
| $v \vdash \text{case } \text{inl}_B v_1 \text{ of } e_1 \mid e_2 \longrightarrow (v \circ v_1) \triangleright e_1$                                                                | $v \vdash \text{case } \text{inr}_A v_1 \text{ of } e_1 \mid e_2 \longrightarrow (v \circ v_1) \triangleright e_2$ |
| S-FIX                                                                                                                                                                             | S-UNFOLDV                                                                                                          |
| $v \vdash \text{fix}(A \rightarrow B). e \longrightarrow \langle v, \text{fix}(A \rightarrow B). e \rangle_f$                                                                     | $v \vdash \text{unfold } (\text{fold } [\mu\alpha.A] v_1) \longrightarrow v_1$                                     |
| S-FIXBETA                                                                                                                                                                         |                                                                                                                    |
| $v \vdash \langle v_2, \text{fix}(A \rightarrow B). e \rangle_f v_1 \longrightarrow (v_2 \circ \langle v_2, \text{fix}(A \rightarrow B). e \rangle_f \circ v_1) \triangleright e$ |                                                                                                                    |

The same frames and rules extend core reduction, with  $\vdash$  in place of  $\vdash$  throughout.

|                             |                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                  |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| $v \vdash e \Rightarrow v'$ |                                                                                                                                                                                                                          | (Big-step evaluation: extensions)                                                                                                                                                                                                                |
|                             | $\frac{\text{EVAL-INJL} \quad v \vdash e \Rightarrow v_1}{v \vdash \text{inl}_B e \Rightarrow \text{inl}_B v_1}$                                                                                                         | $\frac{\text{EVAL-INJR} \quad v \vdash e \Rightarrow v_1}{v \vdash \text{inr}_A e \Rightarrow \text{inr}_A v_1}$                                                                                                                                 |
|                             | $\frac{\text{EVAL-CASEL} \quad v \vdash e \Rightarrow \text{inl}_B v_1 \quad v, v_1 \vdash e_1 \Rightarrow v'}{v \vdash \text{case } e \text{ of } e_1 \mid e_2 \Rightarrow v'}$                                         | $\frac{\text{EVAL-CASER} \quad v \vdash e \Rightarrow \text{inr}_A v_1 \quad v, v_1 \vdash e_2 \Rightarrow v'}{v \vdash \text{case } e \text{ of } e_1 \mid e_2 \Rightarrow v'}$                                                                 |
|                             | $\frac{\text{EVAL-FIX} \quad v \vdash \text{fix}(A \rightarrow B). e \Rightarrow \langle v, \text{fix}(A \rightarrow B). e \rangle_f}{v \vdash \text{fold } [\mu\alpha.A] e \Rightarrow \text{fold } [\mu\alpha.A] v_1}$ | $\frac{\text{EVAL-FIXAPP} \quad v_f = \langle v_1, \text{fix}(A \rightarrow B). e \rangle_f \quad v \vdash e_1 \Rightarrow v_f \quad v \vdash e_2 \Rightarrow v_2 \quad v_1, v_f, v_2 \vdash e \Rightarrow v'}{v \vdash e_1 e_2 \Rightarrow v'}$ |
|                             | $\frac{\text{EVAL-FOLD} \quad v \vdash e \Rightarrow v_1}{v \vdash \text{fold } [\mu\alpha.A] e \Rightarrow \text{fold } [\mu\alpha.A] v_1}$                                                                             | $\frac{\text{EVAL-UNFOLD} \quad v \vdash e \Rightarrow \text{fold } [\mu\alpha.A] v_1}{v \vdash \text{unfold } e \Rightarrow v_1}$                                                                                                               |

Rule **VAL-VAL** covers the extended values. Replacing source judgments with core judgments gives the corresponding big-step rules.

## B.2 Multi-Import Linking

Import interfaces  $D$  are defined in Figure 12.

|                |                                                                  |                                                                                               |
|----------------|------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| $I_1 \Vdash D$ |                                                                  | (Interface satisfaction)                                                                      |
|                | $\frac{\text{OK-ONE} \quad l : A \in I_1}{I_1 \Vdash \{l : A\}}$ | $\frac{\text{OK-MORE} \quad I_1 \Vdash D \quad l : A \in I_1}{I_1 \Vdash D \ \& \ \{l : A\}}$ |

**Elaboration and Core Linking.** Figure 15 generalizes (1):  $\text{wire}_k(D)$  builds the import package from a provider at index  $k$ . Theorem B.3 establishes link typing under fragment typing and interface satisfaction. Unlike **CL-LINK**, rule **CL-LINKN** has no premise: Theorems B.3 and B.4 assume the satisfaction  $I_1 \Vdash D$  instead.

|                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                              |                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| $\Gamma \vdash e : A \rightsquigarrow e', \quad e_1 : \llbracket I_1 \rrbracket \rightsquigarrow e'_1, \quad e_2 : \llbracket D \rrbracket \rightarrow \llbracket I_2 \rrbracket \rightsquigarrow_{\text{link}} e'_2$ |                                                                                                                                                                                                                                                                                                                                                                              | (Multi-import linking) |
|                                                                                                                                                                                                                       | $\frac{\text{ELAB-LINKN} \quad \Gamma \vdash e_1 : I_1 \rightsquigarrow e'_1 \quad \Gamma \vdash e_2 : D \rightarrow_m I_2 \rightsquigarrow e'_2 \quad I_1 \Vdash D}{\Gamma \vdash e_1 \rightsquigarrow e'_1 \quad e_2 : I_1 \ \& \ I_2 \rightsquigarrow \text{link}_c(\llbracket I_1 \rrbracket, \llbracket D \rrbracket, \llbracket I_2 \rrbracket, e'_1, e'_2)}$          |                        |
|                                                                                                                                                                                                                       | $\frac{\text{CL-LINKN} \quad e_1 : \llbracket I_1 \rrbracket \rightsquigarrow e'_1 \quad e_2 : \llbracket D \rrbracket \rightarrow \llbracket I_2 \rrbracket \rightsquigarrow_{\text{link}} e'_2 \quad \text{link}_c(\llbracket I_1 \rrbracket, \llbracket D \rrbracket, \llbracket I_2 \rrbracket, e_1, e_2)}{\text{wire}_k(\{l : A\}) \triangleq \{l = \underline{k}.l\}}$ |                        |
|                                                                                                                                                                                                                       | $\text{wire}_k(D \ \& \ \{l : A\}) \triangleq \text{wire}_k(D) \ , \ \{l = \underline{k+1}.l\}$                                                                                                                                                                                                                                                                              |                        |
|                                                                                                                                                                                                                       | $\text{link}_c(I_1, D, I_2, e_1, e_2) \triangleq (\lambda I_1. \lambda(D \rightarrow I_2). \underline{1} \ , \ (\underline{1} \ \text{wire}_0(D))) \ e_1 \ e_2$                                                                                                                                                                                                              |                        |

Fig. 15. Multi-import linking: elaboration, core linking, and the linking term for an import interface  $D$ . The term of (1) is the instance  $D = \{l : A\}$ .

The right operand of each dependent merge uses index  $k + 1$ , since the left operand adds one environment entry. For  $D = \{x : \text{Int}\} \& \{y : \text{Int}\}$ , the wire is  $\text{wire}_0(D) = \{x = \underline{0}.x\} \circ \{y = \underline{1}.y\}$ , and the provider  $\{y = 2\} \circ \{x = 1\} \circ \{z = 3\}$  supplies the package  $\{x = 1\} \circ \{y = 2\}$ .

*Semantics.* The judgment  $\text{selpkg}(v, D) = p$  states that selecting the imports of  $D$  from provider  $v$  gives the package  $p$ .

$$\begin{array}{c}
 \boxed{\text{selpkg}(v, D) = p} \quad \text{(Package extraction)} \\
 \text{SP-ONE} \quad \frac{v.l \rightsquigarrow v_l}{\text{selpkg}(v, \{l : A\}) = \{l = v_l\}} \quad \text{SP-MORE} \quad \frac{\text{selpkg}(v, D) = p \quad v.l \rightsquigarrow v_l}{\text{selpkg}(v, D \& \{l : A\}) = p \circ \{l = v_l\}}
 \end{array}$$

Two additional frames evaluate the provider and then the client under the ambient environment:

$$\text{Frames} \quad F ::= \dots \mid \square \bowtie_e e \mid v \bowtie_e \square.$$

$$\begin{array}{c}
 \boxed{v \vdash e \longrightarrow e'} \quad \text{(Multi-import linking)} \\
 \text{S-LINKN} \quad \frac{\text{selpkg}(v_1, D) = p}{v \vdash v_1 \bowtie_e \langle v_2, \text{functor}(D) e \rangle \longrightarrow v_1 \circ ((v_2 \circ p) \triangleright e)}
 \end{array}$$

$$\begin{array}{c}
 \boxed{v \vdash e \Rightarrow e'} \quad \text{(Multi-import linking)} \\
 \text{EVAL-LINKN} \quad \frac{v \vdash e_1 \Rightarrow v_1 \quad v \vdash e_2 \Rightarrow \langle v_2, \text{functor}(D) e \rangle \quad \text{selpkg}(v_1, D) = p \quad v_2 \circ p \vdash e \Rightarrow v_3}{v \vdash e_1 \bowtie_e e_2 \Rightarrow v_1 \circ v_3}
 \end{array}$$

*Properties.*

LEMMA B.2 (WIRE TYPING). *If  $I_1 \Vdash D$  and  $\text{lookup}(\Gamma, k) = \llbracket I_1 \rrbracket$ , then  $\Gamma \vdash \text{wire}_k(\llbracket D \rrbracket) : \llbracket D \rrbracket$ .*

THEOREM B.3 (MULTI-IMPORT LINK TYPING). *If  $e_1 : \llbracket I_1 \rrbracket \bowtie_e e_2 : \llbracket D \rrbracket \rightarrow \llbracket I_2 \rrbracket \rightsquigarrow_{\text{link}} e_\ell$ ,  $I_1 \Vdash D$ ,  $\Gamma_c \vdash e_1 : \llbracket I_1 \rrbracket$ , and  $\Gamma_c \vdash e_2 : \llbracket D \rrbracket \rightarrow \llbracket I_2 \rrbracket$ , then  $\Gamma_c \vdash e_\ell : \llbracket I_1 \rrbracket \& \llbracket I_2 \rrbracket$ .*

THEOREM B.4 (MULTI-IMPORT STATIC SEPARATE COMPILATION). *Let  $e_2 = \text{functor}_\blacksquare(D) e$  with  $\varepsilon \vdash e_2 : D \rightarrow_m I_2 \rightsquigarrow e'_2$ . Suppose  $\varepsilon \vdash e_1 : I_1 \rightsquigarrow e'_1$ ,  $I_1 \Vdash D$ , and  $e'_1 : \llbracket I_1 \rrbracket \bowtie_e e'_2 : \llbracket D \rrbracket \rightarrow \llbracket I_2 \rrbracket \rightsquigarrow_{\text{link}} e_\ell$ . Then clauses (a)–(c) of Corollary 4.12 hold for  $e_1 \bowtie_e e_2$  and  $e_\ell$ , with  $v = u = \varepsilon$ .*

PROOF SKETCH OF THEOREM B.4. Rules ELAB-LINKN and CL-LINKN produce the same core term under the theorem's premises. Theorems 4.8 and 4.9 then give clauses (a)–(c).  $\square$

As in Corollary 4.12, the result generalizes to a provider elaborated under  $\Gamma$  and evaluation under a corresponding environment. Lemma 4.1 reuses the client's elaboration under  $\varepsilon$  unchanged.

For multi-import linking, the mechanized replacement law requires the replacement provider to preserve every extractable package, over all import interfaces.

**THEOREM B.5 (MULTI-IMPORT INTERFACE REPLACEMENT).** *Suppose  $v \vdash e_1 \Downarrow e \Rightarrow p_1 \circ w, v \vdash e_2 \Rightarrow p_2$ , and, for every  $D$  and  $q$ ,  $\text{selpkg}(p_1, D) = q$  implies  $\text{selpkg}(p_2, D) = q$ . Then*

$$v \vdash e_2 \Downarrow e \Rightarrow p_2 \circ w.$$

*If also  $\Gamma \vdash e_2 \Downarrow e : A, \varepsilon \vdash v : \Gamma$ , and  $v \vdash e_2 \Downarrow e \Rightarrow p_2 \circ w'$ , then  $w = w'$ .*

**PROOF SKETCH.** Reuse the client and body evaluations in rule **Eval-LinkN**, with the replacement provider and the same import package. The equality follows from determinism (Theorem 4.15).  $\square$

The first part of Theorem B.5 needs no typing, and the theorem concerns only terminating runs. For typed provider values, a two-way law holds, the multi-import counterpart of Theorem 4.14. Providers that supply the same import package give the client the same exports, whatever they supply at other interfaces, and one link diverges exactly when the other does.

**THEOREM B.6 (MULTI-IMPORT INTERFACE ABSTRACTION).** *Let  $p_1$  and  $p_2$  be values. Suppose  $\Gamma \vdash e : D \rightarrow_m B, \varepsilon \vdash v : \Gamma$ , and, for  $i = 1, 2$ ,  $\varepsilon \vdash p_i : P_i, P_i \Vdash D$ , and  $\text{selpkg}(p_i, D) = q$ . Then, for every value  $w$ ,*

$$v \vdash p_1 \Downarrow e \Rightarrow p_1 \circ w \iff v \vdash p_2 \Downarrow e \Rightarrow p_2 \circ w$$

*and  $v \vdash p_1 \Downarrow e \Uparrow$  if and only if  $v \vdash p_2 \Downarrow e \Uparrow$ .*

## C Implementation Details

We give the unit-linking construction and WasmGC compilation used by SCE-Lang (Section 5.2).

### C.1 Units and Linking

An importing file is represented by the sandboxed functor **functor** $\blacksquare(\{M_1 : A_1\} \& \cdots \& \{M_k : A_k\}) e$ , with each  $A_i$  read from  $M_i$ 's interface file. An import-free file is a structure boxed under the empty environment.

Let  $e_1, \dots, e_n$  be units in dependency order, with  $n \geq 1$  and  $e_1$  import-free. Each  $e_i$  elaborates independently under  $\varepsilon$  to  $c_i$ , at type  $E_i$  if import-free and  $D_i \rightarrow_m E_i$  otherwise. Write  $S_i$ ,  $L_i$ , and  $I_i$  for the source composition, core composition, and accumulated export type. Initially,

$$S_1 = e_1 \quad L_1 = c_1 \quad I_1 = E_1.$$

For an importing unit  $e_{i+1}$ , require  $I_i \Vdash D_{i+1}$  and set

$$S_{i+1} = S_i \Downarrow e_{i+1},$$

$$L_{i+1} = \text{link}_c(\llbracket I_i \rrbracket, \llbracket D_{i+1} \rrbracket, \llbracket E_{i+1} \rrbracket, L_i, c_{i+1}).$$

For an import-free unit, set  $S_{i+1} = S_i \circ e_{i+1}$  and construct  $L_{i+1}$  from  $L_i$  and  $c_{i+1}$  using the encoding of rule **ELAB-NMRG**. In either case,  $I_{i+1} = I_i \& E_{i+1}$ . For binary chains, each importing unit has a singleton interface  $D_{i+1} = \{l : A\}$ ; require  $l : A \in I_i$  and use  $S_{i+1} = S_i \Downarrow e_{i+1}$ . The core construction is unchanged.

**THEOREM C.1 (UNIT-CHAIN CORRECTNESS).** *For a binary chain constructed above,  $\varepsilon \vdash S_n : I_n \rightsquigarrow L_n$ , and clauses (a)–(c) of Corollary 4.10 hold for  $S_n$  and  $L_n$  at type  $I_n$ .*

**THEOREM C.2 (MULTI-IMPORT UNIT-CHAIN CORRECTNESS).** *For a multi-import chain constructed above,  $\varepsilon \vdash S_n : I_n \rightsquigarrow L_n$ , and clauses (a)–(c) of Corollary 4.10 hold for  $S_n$  and  $L_n$  at type  $I_n$ .*

**PROOF SKETCH.** Induct on the chain, starting from the first unit's elaboration. Each extension uses **ELAB-LINK** or **ELAB-LINKN** for an importing unit, and **ELAB-NMRG** for an import-free unit. Corollary 4.10 then gives clauses (a)–(c).  $\square$

For WebAssembly, a link module imports unit entry points and compiles the same composition with environment projections for their results (Appendix D.3).

## C.2 Compilation to WebAssembly

Figure 16 compiles well-typed program terms of extended  $\lambda_E$  (Appendix B); runtime closures are not compilation inputs. Field selection uses the inferred type of its operand. The Wasm backend represents integers and performs arithmetic with i32, so the literal rule is restricted to signed 32-bit integers.

|                                  |                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| $C[\text{env}]_\rho$             | <code>local.get <math>\rho</math></code>                                                                                                                                                                                                                                                                                                                                                |
| $C[\varepsilon]_\rho$            | <code>global.get \$unit</code>                                                                                                                                                                                                                                                                                                                                                          |
| $C[i]_\rho$                      | <code>i32.const <math>\tau_{\text{int}}</math> i32.const <math>i</math> struct.new \$I32Box</code>                                                                                                                                                                                                                                                                                      |
| $C[e.n]_\rho$                    | <code><math>C[e]_\rho</math> <math>\mathcal{L}^n</math> <math>\mathcal{R}</math></code>                                                                                                                                                                                                                                                                                                 |
| $C[e.l]_\rho$                    | <code><math>C[e]_\rho</math> <math>\mathcal{P}(A, l)</math> ref.cast (ref \$Lrec)<br/>struct.get \$Lrec value</code>                                                                                                                                                                                                                                                                    |
| $C[\{l = e\}]_\rho$              | <code>i32.const <math>\tau_{\text{rec}}</math> <math>\mathcal{B}(l)</math> <math>C[e]_\rho</math> struct.new \$Lrec</code>                                                                                                                                                                                                                                                              |
| $C[e_1 \triangleright e_2]_\rho$ | <code><math>C[e_1]_\rho</math> local.set <math>\rho'</math> <math>C[e_2]_{\rho'}</math></code>                                                                                                                                                                                                                                                                                          |
| $C[e_1, e_2]_\rho$               | <code><math>C[e_1]_\rho</math> local.set <math>x</math><br/>i32.const <math>\tau_{\text{mrg}}</math> local.get <math>\rho</math> local.get <math>x</math><br/>struct.new \$Pair local.set <math>\rho'</math><br/>i32.const <math>\tau_{\text{mrg}}</math> local.get <math>x</math> <math>C[e_2]_{\rho'}</math> struct.new \$Pair</code>                                                 |
| $C[\lambda A. e]_\rho$           | <code>i32.const <math>\tau_{\text{clos}}</math> ref.func \$f local.get <math>\rho</math> struct.new \$Clos<br/>where the fresh lifted function \$f(self, x) is<br/>i32.const <math>\tau_{\text{mrg}}</math> local.get self ref.cast (ref \$Clos)<br/>struct.get \$Clos env local.get <math>x</math> struct.new \$Pair<br/>local.set <math>\rho'</math> <math>C[e]_{\rho'}</math></code> |
| $C[e_1 e_2]_\rho$                | <code><math>C[e_1]_\rho</math> local.tee <math>c</math> <math>C[e_2]_\rho</math> local.get <math>c</math><br/>ref.cast (ref \$Clos) struct.get \$Clos fn call_ref \$Fn</code>                                                                                                                                                                                                           |

|                             |                                                                                                                                       |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| $\mathcal{L}$               | <code>ref.cast (ref \$Pair) struct.get \$Pair left</code>                                                                             |
| $\mathcal{R}$               | <code>ref.cast (ref \$Pair) struct.get \$Pair right</code>                                                                            |
| $\mathcal{B}(l)$            | <code>i32.const <math>\text{off}_l</math> i32.const <math>\text{len}_l</math> array.new_data \$Bytes \$strings</code>                 |
| $\mathcal{P}(\{l : A\}, l)$ | <code>[]</code>                                                                                                                       |
| $\mathcal{P}(A \& B, l)$    | <code><math>\mathcal{L} \mathcal{P}(A, l)</math> if <math>l \in \text{label}(A)</math>, <math>l \notin \text{label}(B)</math></code>  |
| $\mathcal{P}(A \& B, l)$    | <code><math>\mathcal{R} \mathcal{P}(B, l)</math> if <math>l \notin \text{label}(A)</math>, <math>l \in \text{label}(B)</math>.</code> |

Fig. 16. WasmGC compilation of extended  $\lambda_E$  (1/2): the base calculus. The local  $\rho$  holds the environment, and  $\tau_k$  is the tag for constructor  $k$ . Introduced locals are fresh; each compiled term leaves one value reference on the stack. In selection,  $A$  is the inferred type of  $e$ . Juxtaposition concatenates instruction sequences;  $\mathcal{L}^n$  repeats  $\mathcal{L}$   $n$  times. The pair  $(\text{off}_l, \text{len}_l)$  locates  $l$  in the passive data segment \$strings.

```

1569  C[[fix(A→B). e]]ρ      =  i32.const τfclos  ref.func $f  local.get ρ
1570                               struct.new $Clos
1571                               where the fresh lifted function $f(self, x) is
1572                               i32.const τmrg  i32.const τmrg
1573                               local.get self  ref.cast(ref $Clos)  struct.get $Clos env
1574                               local.get self  struct.new $Pair  local.get x
1575                               struct.new $Pair  local.set ρ'  C[[e]]ρ'
1576
1577  C[[inlB e]]ρ          =  i32.const τinl  C[[e]]ρ  struct.new $Wrap
1578
1579  C[[inrA e]]ρ          =  i32.const τinr  C[[e]]ρ  struct.new $Wrap
1580
1581  C[[case e of e1 | e2]]ρ =  C[[e]]ρ  local.set u
1582                               local.get u  struct.get $Val tag
1583                               i32.const τinl  i32.eq
1584                               if (result (ref null $Val))
1585                               E(ρ, u, ρ1)  C[[e1]]ρ1
1586                               else
1587                               E(ρ, u, ρ2)  C[[e2]]ρ2
1588                               end
1589
1590  C[[fold [μα.A] e]]ρ    =  i32.const τfold  C[[e]]ρ  struct.new $Wrap
1591
1592  C[[unfold e]]ρ        =  C[[e]]ρ  ref.cast(ref $Wrap)  struct.get $Wrap value
1593
1594  E(ρ, u, ρ')            =  i32.const τmrg  local.get ρ  local.get u
1595                               ref.cast(ref $Wrap)  struct.get $Wrap value
1596                               struct.new $Pair  local.set ρ'
1597
1598  $Val    tag:i32                $I32Box  tag:i32, n:i32
1599  $Pair    tag:i32, left:V, right:V  $Lrec   tag:i32, label:B, value:V
1600  $Clos    tag:i32, fn:F, env:V      $Wrap   tag:i32, value:V
1601
1602  V = (ref null $Val),    B = (ref null $Bytes),    F = (ref null $Fn)
1603  $Bytes = (array (mut i8)),  $Fn : (V, V) → V,    $UnitFn : () → V.

```

Fig. 16. WasmGC compilation of extended  $\lambda_E$  (2/2): unions and recursion, with runtime layouts in field-index order. Struct fields are immutable; each struct except \$Val is a final subtype of \$Val. Value locals and results have type V. The global \$unit is a \$Val with tag 3. Tags for integers, merges, records, closures, recursive closures, left and right injections, and folds are 0, 4, 5, 6, 7, 8, 9, 10, respectively.

## D The Double Example Through Compilation

The files of Figure 1 appear below in the paper's surface notation.

Add.scei

```
sig
  op : Int -> Int -> Int
  unit : Int
end
```

Double.scei

```
sig double : Int -> Int end
```

Add.sce

```
module Add = struct
  let op (x y : Int) : Int = x + y
  let unit : Int = 0
end
```

Double.sce

```
import Add
module Double = struct
  let double (x : Int) : Int = Add.op x x
end
```

### D.1 Link, Then Compile

The unit wrappers give the following top-level link  $S$ :

```
link (sandbox struct
  module Add = struct
    let op (x y : Int) : Int = x + y
    let unit : Int = 0
  end
end) with sandbox functor
(M : { Add : sig op : Int -> Int -> Int; unit : Int end }) ->
open M in struct
  module Double = struct
    let double (x : Int) : Int = Add.op x x
  end
end
```

We abbreviate the core terms below, with  $\underline{n} = \text{env}.n$ . Identity applications and the internal label %open are retained from the compiler output.

$$O = \text{Int} \rightarrow \text{Int} \rightarrow \text{Int}, \quad P = \{\text{op} : O\} \ \& \ \{\text{unit} : \text{Int}\},$$

$$T = \{\text{double} : \text{Int} \rightarrow \text{Int}\}, \quad I = \{\text{Add} : P\}, \quad J = \{\text{Double} : T\}.$$

$$\begin{aligned} id_A &= \lambda A. \underline{0}, \\ a_+ &= \lambda \text{Int}. \lambda \text{Int}. \underline{1} + \underline{0}, \\ b &= \lambda \text{Int}. (\underline{1}.\text{Add.op}) \underline{0} \underline{0}, \\ r_A &= \{\text{op} = id_O \ a_+\} \ , \ \{\text{unit} = id_{\text{Int}} \ 0\}, \\ r_D &= \{\text{Double} = \{\text{double} = id_{\text{Int} \rightarrow \text{Int}} \ b\}\}, \\ a &= \varepsilon \triangleright \{\text{Add} = r_A\}, \\ d &= \varepsilon \triangleright \lambda I. (\lambda I. r_D) \ (\{\%open = \underline{0}\}.\%open). \end{aligned}$$

Compiling  $S$  gives

$$L = \underbrace{(\lambda I. \lambda (I \rightarrow J). \underline{1} \ , \ (\underline{1} \ \{\text{Add} = \underline{0}.\text{Add}\}))}_K \ a \ d.$$

The linked library exports Add and Double, and  $S.\text{Double}.\text{double} \ 5$  evaluates to 10.

## D.2 Compile, Then Link in the Core

Compiling the files independently produces:

$$\begin{aligned} \text{Add.sceo} &\mapsto a : I \\ \text{Double.sceo} &\mapsto d : I \rightarrow J. \end{aligned}$$

The core linker selects Add for the client's import and produces

$$\begin{aligned} L_{\text{core}} &= (\lambda I. \lambda (I \rightarrow J). \underline{1} \text{ } _9 (\underline{1} \{ \text{Add} = \underline{0}. \text{Add} \})) \text{ } a \text{ } d \\ &= L. \end{aligned}$$

Both routes produce the same core term, with

$$\varepsilon \vdash L_{\text{core}}. \text{Double.double } 5 \Rightarrow 10.$$

## D.3 Compile, Then Link in WebAssembly

The core units compile to separate modules, each exporting `main : () -> (ref null $Val)`:

$$a \mapsto \text{Add.wasm} \quad d \mapsto \text{Double.wasm}.$$

If their entry points return  $p$  and  $f$ , the link module constructs  $\rho = (\varepsilon \text{ } _9 p) \text{ } _9 f$  and executes the compiled term  $K \text{ } \underline{1} \text{ } \underline{0}$ . The listing uses symbolic type and function names and omits shared type declarations and rendering support.

Lib.wasm

```
(module
  ;; Shared types and rendering support omitted.
  (import "u0" "main" (func $Add.main (type $UnitFn)))
  (import "u1" "main" (func $Double.main (type $UnitFn)))
  (global $unit (ref null $Val) (struct.new $Val (i32.const 3)))
  (data $labels "Add")
  (elem declare func $link0 $link1)

  (func $main (export "main") (type $UnitFn)
    (local (ref null $Val) (ref null $Val) (ref null $Val))
    global.get $unit local.set 0
    i32.const 4 local.get 0 call $Add.main struct.new $Pair local.set 0
    i32.const 4 local.get 0 call $Double.main struct.new $Pair local.set 0
    i32.const 6 ref.func $link0 local.get 0 struct.new $Clos local.tee 1
    local.get 0 ref.cast (ref $Pair) struct.get $Pair 1 ref.cast (ref $Pair)
    struct.get $Pair 2 local.get 1 ref.cast (ref $Clos) struct.get $Clos 1
    call_ref $Fn
    local.tee 2
    local.get 0 ref.cast (ref $Pair) struct.get $Pair 2 local.get 2
    ref.cast (ref $Clos) struct.get $Clos 1 call_ref $Fn
  )
)
```

Lib.wasm (continued)

```

1716 (func $link0 (type $Fn)
1717   (local (ref null $Val))
1718   i32.const 4 local.get 0 ref.cast (ref $Clos) struct.get $Clos 2
1719   local.get 1 struct.new $Pair local.set 2
1720   i32.const 6 ref.func $link1 local.get 2 struct.new $Clos
1721 )
1722
1723 (func $link1 (type $Fn)
1724   (local (ref null $Val) (ref null $Val) (ref null $Val))
1725   (local (ref null $Val))
1726   i32.const 4 local.get 0 ref.cast (ref $Clos) struct.get $Clos 2
1727   local.get 1 struct.new $Pair local.set 2
1728   local.get 2 ref.cast (ref $Pair) struct.get $Pair 1 ref.cast (ref $Pair)
1729   struct.get $Pair 2 local.set 3
1730   i32.const 4 local.get 2 local.get 3 struct.new $Pair local.set 4
1731   i32.const 4 local.get 3 local.get 4 ref.cast (ref $Pair)
1732   struct.get $Pair 1 ref.cast (ref $Pair) struct.get $Pair 2 local.tee 5
1733   i32.const 5 i32.const 0 i32.const 3 array.new_data $Bytes $labels
1734   local.get 4 ref.cast (ref $Pair) struct.get $Pair 2 ref.cast (ref $Lrec)
1735   struct.get $Lrec 2 struct.new $Lrec local.get 5 ref.cast (ref $Clos)
1736   struct.get $Clos 1 call_ref $Fn
1737   struct.new $Pair
1738 )
1739 ;; sce.units: Add, Double
1740 )

```

Each unit and the link module declares the same recursive group of ABI types (Figure 16), allowing values to pass between instances. The host supplies the unit entry points and calls the linker:

```

1743 const imports = {
1744   u0: { main: add.exports.main },
1745   u1: { main: double.exports.main }
1746 };
1747 const lib = new WebAssembly.Instance(linkModule, imports);
1748 const result = lib.exports.main();

```

The linked result, rendered by the implementation, is

```
{ Add = { op = <fun>, unit = 0 }, Double = { double = <fun> } }
```

Calling its `Double.double` function on 5 returns 10, as in Sections D.1 and D.2.