

A Separate Compilation Extension to Standard ML

David Swasey Tom Murphy VII Karl Crary Robert Harper

Carnegie Mellon

{swasey,tom7,crary,rwh}@cs.cmu.edu

Abstract

We present an extension to Standard ML, called SMLSC, to support separate compilation. The system gives meaning to individual program fragments, called units. Units may depend on one another in a way specified by the programmer. A dependency may be mediated by an interface (the type of a unit); if so, the units can be compiled separately. Otherwise, they must be compiled in sequence. We also propose a methodology for programming in SMLSC that reflects code development practice and avoids syntactic repetition of interfaces. The language is given a formal semantics, and we argue that this semantics is implementable in a variety of compilers.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features—Modules, packages; D.3.1 [Programming Languages]: Formal Definitions and Theory—Syntax

General Terms Languages

Keywords standard ml, separate compilation, incremental compilation, types

Introduction

We propose an extension to Standard ML called SMLSC. SMLSC supports separate compilation in the sense that it gives a static semantics to individual program fragments, which we call *units*. A unit may depend on other units, and can be type-checked independently of those units by specifying what it expects of them. These expectations are given in the form of *interfaces* for those other units. When unit A is checked against another unit B via a mediating interface, we need not have access to B at all. Therefore we say that A is separately compiled (SC) against B.

It is also useful to allow unit A to depend on another unit B without specifying an interface for B. In this case, the only way to derive the context necessary to check A is to first check B and read off its actual interface. In this scenario we say that A is incrementally compiled (IC) against B.

Units may be compiled and then linked together to satisfy dependencies. The compiled form of a unit or set of linked units is called a *linkset*. A linkset may be further linked with other linksets. If a linkset has no remaining dependencies, then it can be transformed into an executable program.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ML'06 September 16, 2006, Portland, Oregon, USA.

Copyright © 2006 ACM 1-59593-483-9/06/0009...\$5.00.

The goal of this work is to consolidate and synthesize previous work on compilation management for ML into a formally defined extension to the Standard ML language. The extension itself is syntactically and conceptually very simple. A unit is a series of Standard ML top-level declarations, given a name. To the current top-level declarations such as `structure` and `functor` we add an `import` declaration that is to units what the `open` declaration is to structures. An `import` declaration may optionally specify an interface for the unit, in which case we are able to compile separately against that dependency; with no interface we must compile incrementally. Compatibility with existing compilers, including whole-program compilers, is assured by making no commitment to the precise meaning of “compile” and “link”—a compiler is free to limit compilation to elaboration and type checking, and to perform code generation as part of linking.

Sections 1 and 2 summarize our main design principles, and provide an overview of the system. In Section 2 we give a small example of its use. The semantics, formulated in the framework of the Harper-Stone semantics of ML [12, 13], is given in Section 3.¹ Some implementation issues are discussed in Section 4. We conclude with a discussion of related work in Section 6.

1. Design Principles

A language, not a tool. We propose an extension to the Standard ML language to support separate compilation, rather than a tool to implement it. The extension is defined by a semantics that extends the semantics of Standard ML to provide a declarative description of the meanings of the language constructs. The semantics provides a clear correctness criterion for implementations to ensure source-level compatibility among them.

Flexibility. A compilation unit consists of any sequence of top-level declarations, including signature and functor declarations.² However, since Standard ML lacks syntactically expressible signatures, some units cannot be separately compiled from one another. We therefore support incremental, as well as separate, compilation for any unit. This means that the interface of a unit can either be inferred from its source (incremental compilation) or explicitly specified (separate compilation) at the programmer’s discretion.

Simplicity. The design provides only the minimum functionality of a separate compilation system. It omits any form of compilation parameters, conditional compilation directives, or compiler directives. We leave for future work the specification of such additional machinery.

Conservativity. The semantics of Standard ML should not be changed by the introduction of separate compilation. In particular, we do not permit “circular dependencies” or similar concepts

¹We give a semantics in the framework of *The Definition of Standard ML* [16] in the companion technical report [21].

²Consequently, units cannot be identified with Standard ML structures.

<code>srcunit</code>	<code>::= unit unitid = top topdec end</code>	unit declaration
<code>topdec</code>	<code>::= import impexp strdec sigdec fundec local topdec₁ in topdec₂ end topdec₁ topdec₂</code>	open units structure-level declaration signature declaration functor declaration local declaration
<code>impexp</code>	<code>::= unitid ⟨: intexp⟩ impexp₁ impexp₂</code>	open unit <i>unitid</i>
<code>intexp</code>	<code>::= intf topspec end</code>	interface expression
<code>topspec</code>	<code>::= spec functor funspec topspec₁ topspec₂</code>	structure-level specification functor specification
<code>funspec</code>	<code>::= funid(strid : sigexp) : sigexp' ⟨and funspec⟩</code>	

Figure 1. SMLSC concrete syntax

that are not otherwise expressible in Standard ML. This ensures that compilers should not be disturbed by the proposed extension beyond what is required to implement the extension itself.

Explicit dependencies. The dependencies among units are explicitly specified, not inferred. The chief reason for this is that dependencies among units may not be syntactically evident—for example, the side effects of one unit may influence the behavior of another. There are in general many ways to order effects consistently with syntactic dependencies, and these orderings need not be equivalent. A lesser reason is that supporting dependency inference requires restrictions on compilation units that are not semantically necessary, reducing flexibility.

Environment independence. The separate compilation system is defined independently of any environment in which it might be implemented. The design speaks in terms of linguistic and semantic entities, rather than implementation-specific concepts such as files or directories.

Separation of units from modules. The separate compilation system is designed as a proper extension to Standard ML so as to ensure backward compatibility of source code. It is tempting to identify compilation units with modules, but to do so would require that functors, signatures, and fixity declarations be permitted as components of modules. Permitting such an extension is not entirely straightforward; for example, permitting signature declarations in modules and their types can lead to undecidability of type checking [10].

2. Overview

Units and Interfaces

The SMLSC extension is organized around the concept of a *unit*. A unit consists of top-level declarations, which include declarations of signatures, structures, and functors. Each unit is given a name by which the unit is known throughout the program. One unit may refer to the components of another using an `import` declaration, which records the dependency of the importing unit on the imported unit, and opens it for use within the importing unit. This is the only means by which one unit may refer to another; we do not support “dot notation” for accessing the components of a unit. An `import` declaration is a new form of top-level declaration. (This is the only modification that we make to an existing syntactic category of Standard ML.)

The compilation context for a unit is entirely determined by its imports. That is, all dependencies of a unit on another unit must be explicitly indicated using `import` declarations. The dependency of

one unit on another is mediated by an *interface*, the type of a unit. The interface of an imported unit can be specified in one of two ways, either *implicitly* or *explicitly*, corresponding to *incremental* or *separate* compilation.

An `import` declaration of the form `import unitid : intexp` specifies an explicit interface for the imported unit. This permits the importing unit to be compiled independently of the implementation of *unitid*, relying only on the specified interface. This is called *separate compilation*, or *SC* for short. An `import` declaration of the form `import unitid` specifies that the interface for *unitid* is to be inferred from its source code. This is called *incremental compilation*, or *IC* for short.

The concrete syntax for units and interfaces in SMLSC is given in Figure 1. We extend *topdec*s to add `import` and `local`. The `import` declaration (like `open`) allows multiple units to be simultaneously imported. Interfaces are *topspec*s; this is the syntactic class *spec* of Standard ML, with the addition of a specification form for functors. The `local` declaration limits the scope of imports, just as the structure declaration of the same name.

Projects and Linksets

A *linkset* consists of several compiled units, called its exports, together with the names and interfaces of its imports, the units on which it depends (following Cardelli [5]). A *project* consists of a linearly ordered sequence of source units and linksets. The ordering of the components in a project is significant, both because it specifies the order of identifier resolution, and because it specifies the order of computational effects when executed. Compilation of a project consists of processing the source units in the specified order to obtain linksets, and then knitting them together to resolve dependencies.

Linking consists of resolving inter-unit dependencies by binding exports to imports among linksets. When all references have been resolved, the resulting linkset can be completed to form an executable.

We do not give a concrete syntax for linksets, as we do not intend for programmers to write them, nor do we expect compatibility of linksets across implementations. Rather, they are left as implementation-specific concepts (such as object files), which are modeled here by the abstract semantic objects described in Section 3.

Examples

We begin with a few simple examples to illustrate the features of the system.

Suppose that we have a library of data structures whose name is `Collections`. It is natural to place this library in a unit. Let's assume it contains only the queue data structure:

```
unit Collections =
top
  signature QUEUE =
  sig
    type 'a queue
    val empty : 'a queue
    val push : 'a * 'a queue -> 'a
    val pop : 'a queue -> 'a * 'a queue
  end
  structure Queue :> QUEUE =
  struct (* ... *) end
end
```

A client of the `Collections` library can import it using IC easily:

```
unit Scheduler =
top
  import Collections
  structure Sched =
  struct
    type job = (* ... *)
    val readyqueue =
      ref Queue.empty : job Queue.queue ref
    (* ... *)
  end
end
```

In these examples we use *link* to stand for the semantic operation of compiling and linking a list of source units and linksets. We can compile and link this program as

$$L = \text{link}(\text{Collections}, \text{Scheduler})$$

or we can compile the library and then the client

$$L_0 = \text{link}(\text{Collections})$$

$$L_1 = \text{link}(L_0, \text{Scheduler})$$

but the `Scheduler` unit may not be compiled on its own.

Incremental compilation is convenient when we have source or a compiled linkset for the `Collections` unit. We may prefer to use separate compilation, or may be forced to because the implementation for `Collections` is not available. A client with an SC import looks like this:

```
unit Scheduler2 =
top
  import Collections :
  intf
    structure Queue :
    sig
      type 'a queue
      val empty : 'a queue
      val push : 'a * 'a queue -> 'a
      val pop : 'a queue -> 'a * 'a queue
    end
  end

  structure Sched =
  struct
    type job = (* ... *)
    val readyqueue =
      ref Queue.empty : job Queue.queue ref
    (* ... *)
  end
end
```

This allows `Scheduler2` and `Collections` to be compiled separately:

$$L_0 = \text{link}(\text{Scheduler2})$$

$$L_1 = \text{link}(\text{Collections})$$

$$L_2 = \text{link}(L_1, L_0)$$

However, writing the SC import this way forces an undesirable repetition of code. If more than one client uses `Collections`—which we would expect—each client repeats the interface for its import of the unit. A further problem is that this style asks the client to supply the interface of the library, but the interface of a library is usually provided by the library author, not the client. Fortunately, a combination of SC and IC allows us to use the system in a much cleaner way.

Handoff Units

A programmer who wishes his code to be available for separate compilation can provide a *handoff unit* which supplies the interface. Starting from scratch, the handoff unit contains an SC import:

```
unit Collections =
top
  signature QUEUE =
  sig
    type 'a queue
    val empty : 'a queue
    val push : 'a * 'a queue -> 'a
    val pop : 'a queue -> 'a * 'a queue
  end
  import CollectionsImpl :
  intf
    structure Queue : QUEUE
  end
end
```

The implementation of the collections library is moved to the unit `CollectionsImpl`. Because the `import` declaration opens the imported unit, all of the contents of `CollectionsImpl` are available in the `Collections` unit. (Note that signature declarations do not appear in interfaces, but we can write these before the `import` declaration and then use the signature bindings in the interface for the same effect.) Clients wishing to make use of the library simply import the handoff unit using IC:

```
unit Scheduler3 =
top
  import Collections
  structure Sched = (* ... *)
end
```

This additionally has the benefit that the clients only need to know the name of the handoff unit, not the implementation unit. A few such clients can be linked with the handoff unit:

$$L_0 = \text{link}(\text{Collections}, \text{Scheduler3}, \text{OtherClient})$$

The result can later be linked with the implementation of the `Collections` library:

$$L_1 = \text{link}(\text{CollectionsImpl}, L_0)$$

Definite References

In the terminology of Harper and Pierce [11] an import of one unit in another is interpreted as a *definite reference*—that is, as a free variable that refers to single, specific unit through an interface for it, either inferred or specified. This ensures that if two separate units import a common unit, such as a well-known library, these units share a common understanding of the abstract types exported by that unit. No additional sharing specifications are required to use the separate compilation system.

This is in sharp contrast to the so-called *fully functorized style* of use of the ML modules system, in which a functor is λ -abstracted over the modules on which it depends. In this formulation references are interpreted as *indefinite*, in that the functors involved may be applied to *any* modules satisfying those interfaces, not necessarily in a coherent manner. Different functors with a parameter referring to “the” shared library might be applied to different instances of it. To ensure coherence in the presence of indefinite references, one must resort to explicit type sharing specifications, which are potentially burdensome.

The handoff methodology in SMLSC facilitates programming with definite references. Two pieces of code that SC import the same unit may only be linked if they import that unit at equivalent interfaces. The imports are then consolidated into a single import, ensuring that type equations hold. When the same handoff unit is used to create the two imports, these interfaces will always be equivalent. In corner cases such as skew between versions of a library’s handoff unit, the programmer may manually consolidate two imports. We discuss this further in Section 5.

3. Semantics

We give a semantics to SMLSC by extending Harper and Stone’s Typed Semantics (TS) for Standard ML [13].³ At a high level the typed semantics consists of an elaboration relation from an *external language*, called TSEL, into an *internal language*, called TSIL. The external language is a slight extension of the abstract syntax of Standard ML. The internal language is a typed λ -calculus based on the Harper-Lillibridge type theory for modules [10]. Elaboration comprises type inference, pattern compilation, equality compilation, identifier resolution, and insertions of coercions for signature matching. The result of elaboration is a well-formed program in the TSIL, to which a dynamic semantics is given to provide an execution model. The semantics of SMLSC is an extension of the Harper-Stone semantics that elaborates units into linksets that can be completed for execution.

The TSIL. We begin with a brief review of the structure of the TSIL. The TSIL consists of a core level and a module level. The core level includes expressions exp , constructors con , and kinds knd . Kinds classify constructors. Constructors of kind Ω are types; they classify expressions. The module level includes modules mod and signatures sig , which classify modules. We write $\{\}$ to denote the empty record, and $mod.lab$ to denote the projection of a component named lab from the structure mod . The semantics works mainly with modules, ultimately elaborating units to TSIL structures.

Declaration lists serve as contexts in the TSIL static semantics. A declaration list $decs = dec_1, \dots, dec_n$ declares expression ($var:con$), constructor ($var:knd(=con)$), and module ($var:sig$) variables. A structure declaration list $sdecs$ has the form

$$lab_1 \triangleright dec_1, \dots, lab_n \triangleright dec_n$$

associating a label with each declaration. The structure declaration list $lab \triangleright dec, sdecs$ binds the variable declared by dec with scope $sdecs$. We write $[sdecs]$ to denote the signature of a structure containing fields described by $sdecs$. Variables express dependencies between components in a structure signature and may be freely alpha-varied. Labels name components for external reference and may not be renamed without changing the meaning of the signature. Consider the declaration of a structure m containing an opaque type component T and value component X of that type:

$$m : [T \triangleright t:\Omega, X \triangleright x:t].$$

³In the companion technical report [21], we also give a semantics by extending *The Definition of Standard ML* [16].

Judgement...	Meaning...
$\vdash decs \text{ ok}$	$decs$ is well-formed
$decs \vdash sdecs \text{ ok}$	$sdecs$ is well-formed
$decs \vdash sig : \text{Sig}$	sig is well-formed
$decs \vdash sig \equiv sig' : \text{Sig}$	signature equivalence
$decs \vdash sbnds : sdecs$	$sbnds$ has declaration list $sdecs$
$decs \vdash mod : sig$	mod has signature sig

Figure 2. TSIL judgements (summary)

We can systematically rename the bound variables t and x . A *path* is a module variable followed by a list of labels, serving a role similar to SML long identifiers. The paths $m.T$ (a constructor) and $m.X$ (an expression) refer to m ’s components.

A *bind* binds a variable to an expression ($var=exp$), constructor ($var=con$), or module ($var=mod$). A structure binding list $sbnds$ has the form

$$lab_1 \triangleright bind_1, \dots, lab_n \triangleright bind_n.$$

A structure is written $[sbnds]$. The module syntax is closed under the formation of functors: dependently typed functions from modules to modules.

We shall use the TSIL judgements given in Figure 2. These judgements have the following meaning.

- $decs \vdash sdecs \text{ ok}$. No label is used twice and every declaration is well-formed. For example,

$$\vdash T \triangleright t:\Omega, X \triangleright x:t \text{ ok.}$$

- $decs \vdash sig : \text{Sig}$. The signature sig is well-formed.
- $decs \vdash sig \equiv sig' : \text{Sig}$. The signatures sig and sig' declare the same components, in the same order, with the same labels, and corresponding type components are equivalent.
- $decs \vdash sbnds : sdecs$. The structure binding list $sbnds$ matches the structure declaration list $sdecs$. Corresponding labels must agree and each bound expression, constructor, or module in $sbnds$ must match its declaration in $sdecs$. For example, the judgement

$$decs \vdash (lab \triangleright var=mod, sbnds) : (lab \triangleright var:sig, sdecs)$$

holds if $decs \vdash mod : sig$ and $decs, var:sig \vdash sbnds : sdecs$.

- $decs \vdash mod : sig$. The module mod has signature sig . The signature sig may or may not be fully transparent. For example, we may derive both

$$m : [T \triangleright t:\Omega, X \triangleright x:t] \vdash m : [T \triangleright t:\Omega=m.T, X \triangleright x:t]$$

and

$$m : [T \triangleright t:\Omega, X \triangleright x:t] \vdash m : [T \triangleright t:\Omega, X \triangleright x:t].$$

The former signature is said to be *selfified* with respect to the variable m .

TS elaboration. Harper and Stone give a semantics to Standard ML by elaboration of TSEL into TSIL. Elaboration is performed in a context Γ consisting of a structure declaration list ($sdecs$) that, due to shadowing, may have duplicate labels. We shall use the TS elaboration judgements given in Figure 3. These judgements have the following meaning:

- $\Gamma \vdash strdec \rightsquigarrow sbnds : sdecs$. Elaborate the TSEL structure declaration $strdec$ to the structure binding list $sbnds : sdecs$. Since the TSEL permits functors within structures, this includes elaboration of functor declarations.
- $\Gamma \vdash sigexp \rightsquigarrow sig : \text{Sig}$. Elaborate the TSEL signature expression $sigexp$ to the signature sig . The TSEL does not

Judgement...	Meaning...
$\Gamma \vdash \text{strdec} \rightsquigarrow \text{sbnds} : \text{sdecs}$	structure declaration elaboration
$\Gamma \vdash \text{sigexp} \rightsquigarrow \text{sig} : \text{Sig}$	signature elaboration
$\Gamma \vdash \text{spec} \rightsquigarrow \text{sdecs}$	specification elaboration
$\Gamma \vdash_{\text{ctx}} \text{labs} \rightsquigarrow \text{path} : \text{class}$	context lookup
$\text{decs} \vdash_{\text{sub}} \text{path} : \text{sig}_0 \preceq \text{sig} \rightsquigarrow \text{mod} : \text{sig}'$	coercion compilation

Figure 3. TS elaboration judgements (summary)

include signature declarations; we treat them as abbreviations for TSIL signatures, recording them in linksets and expanding them during elaboration.

- $\Gamma \vdash \text{spec} \rightsquigarrow \text{sdecs}$. Elaborate the TSEL specification spec to the structure declaration list sdecs . This includes elaboration of functor specifications.
- $\Gamma \vdash_{\text{ctx}} \text{labs} \rightsquigarrow \text{path} : \text{class}$. Perform identifier resolution in the context Γ . The input is a list of labels, which is derived from an SML long identifier; the output is a path classified by the type, kind, or signature class . Some labels in the context are annotated with a star, indicating that they are “open” (in the sense of the SML `open` declaration). Identifier resolution searches Γ from right to left, descending into structures with starred labels. For example, we may derive

$$T \triangleright t_1 : \Omega, T \triangleright t_2 : \Omega = \{ \} \vdash T \rightsquigarrow t_2 : \Omega = \{ \}$$

and

$$X \triangleright x_1 : \{ \}, 1^* \triangleright m : [T \triangleright t : \Omega, X \triangleright x_2 : t] \vdash X \rightsquigarrow m.X : m.T.$$

- $\text{decs} \vdash_{\text{sub}} \text{path} : \text{sig}_0 \preceq \text{sig} \rightsquigarrow \text{mod} : \text{sig}'$. Perform transparent signature ascription. The inputs are a signature sig_0 , a path having that signature, and a target signature sig . The output is a module $\text{mod} : \text{sig}'$, where sig' has the same shape as sig but is fully transparent relative to path .

Elaboration maps TSEL identifiers to TSIL labels using a function $\bar{\cdot}$. To implement identifier “shadowing,” elaboration employs a function $\text{sbnds} ++ \text{sbnds}' : \text{sdecs} ++ \text{sdecs}'$ that concatenates $\text{sbnds} : \text{sdecs}$ and $\text{sbnds}' : \text{sdecs}'$, renaming labels in the left hand sides that appear in the right hand sides. The function chooses fresh labels that do not correspond to TSEL identifiers. For example, if

$$\begin{aligned} \text{sbnds} : \text{sdecs} &= \bar{T} \triangleright t_1 = \{ \} : \bar{T} \triangleright t_1 : \Omega = \{ \} \\ \text{sbnds}' : \text{sdecs}' &= \bar{T} \triangleright t_2 = \text{Int} : \bar{T} \triangleright t_2 : \Omega = \text{Int}, \end{aligned}$$

then $\text{sbnds} ++ \text{sbnds}' : \text{sdecs} ++ \text{sdecs}'$ might be

$$(\text{lab} \triangleright t_1 = \{ \}, \bar{T} \triangleright t_2 = \text{Int}) : (\text{lab} \triangleright t_1 : \Omega = \{ \}, \bar{T} \triangleright t_2 : \Omega = \text{Int})$$

where lab is not in the range of the $\bar{\cdot}$ function.

3.1 Linking

We define linking for the TSIL by giving rules for deriving the linking judgements in Figure 4. A linkset

$$\text{sdecs}_0 \rightarrow \text{sbnds} : \text{sdecs}; S$$

comprises imports sdecs_0 , exports $\text{sbnds} : \text{sdecs}$, and signature abbreviations S .

- The imports sdecs_0 describe the TSIL structures on which the linkset depends; they must be well-formed in the ambient context. For example, the imports

$$\text{sdecs}_{AB} = \begin{aligned} &A \triangleright a : [T \triangleright t : \Omega, X \triangleright x : t], \\ &B \triangleright b : [Y \triangleright y : a.T] \end{aligned}$$

express dependency on structures labelled A and B .

Judgement...	Meaning...
$\text{decs} \vdash L \text{ ok}$	L is well-formed
$\text{decs} \vdash S \text{ ok}$	S is well-formed
$L \rightsquigarrow \text{exp} : \{ \}$	L completes to exp
$\text{decs} \vdash L ++ L' \rightsquigarrow L''$	L and L' merge to L''

Figure 4. Linking judgements

$L ::= \text{sdecs}_0 \rightarrow \text{sbnds} : \text{sdecs}; S$	linkset
$S ::= \cdot$	
$S, \text{sigid} = \text{sig}$	top-level
$S, \text{unitid} = S'$	declared by unitid

Figure 5. Linkset syntax

Imports specify assumptions to be satisfied by linking. A linkset with imports sdecs_{AB} assumes structure B binds (at least) a value $Y : a.T$ but can be linked with (a linkset exporting) a structure B providing more components.

- The exports $\text{sbnds} : \text{sdecs}$ are the TSIL code associated with the linkset. They may make reference to the linkset’s imports. Continuing our example, the exports

$$\text{sbnds}_{ZR} : \text{sdecs}_{ZR} = \begin{aligned} &Z \triangleright z = b.Y, R \triangleright r = a.T : \\ &Z \triangleright z : a.T, R \triangleright r : \Omega = a.T \end{aligned}$$

reference the imports sdecs_{AB} to bind an expression Z of the imported type and an equivalent type R .

- The signature abbreviations S are used during elaboration. They may make reference to the linkset’s imports and exports. Continuing our example, the signature abbreviations

$$S_{\text{SIG}} = \text{SIG} = [M \triangleright m : \Omega = r]$$

specify that elaboration should treat the signature identifier SIG as an abbreviation for a TSIL signature referencing the exported type R .

The dynamic semantics for SMLSC is very simple. The completion judgment $L \rightsquigarrow \text{exp} : \{ \}$ translates a linkset

$$\cdot \rightarrow \text{sbnds} : \text{sdecs}; S$$

with no imports to a TSIL expression

$$[\text{sbnds}, \text{lab} \triangleright \text{var} = \{ \}]. \text{lab} : \{ \}$$

where lab and var are fresh. Under the TSIL dynamic semantics, the resulting expression evaluates the linkset’s exports from left to right for their side-effects. Evaluation terminates when an uncaught exception is raised or when every export has been evaluated.

We give the full syntax for linksets in Figure 5 and the rules in Appendix A. The remainder of this section explains the rules for linkset merge.

Notation. We write $\text{decs}, \text{sdecs}$ to extend a context decs , implicitly dropping the labels in sdecs . We define the domain of a structure declaration list, $\text{dom}(\text{sdecs})$, by

$$\text{dom}(\text{lab}_1 \triangleright \text{dec}_1, \dots, \text{lab}_n \triangleright \text{dec}_n) = \{ \text{lab}_1, \dots, \text{lab}_n \}.$$

We write $\{ \text{var} / \text{var}' \} L$ for the capture-free substitution of var for free occurrences of var' in L .⁴

Linkset merge. The rules for linkset merge $\text{decs} \vdash L_1 ++ L_2 \rightsquigarrow L_3$ combine L_1 and L_2 to produce L_3 . The rules presuppose that L_1 is well-formed with respect to decs but permit L_2 to make

⁴Linkset bound variables and scopes are discussed in Appendix A.

reference (via free TSIL variables) not only to *decs* but to the imports and exports of L_1 . Formally, the rules satisfy the following property.⁵

If $L_1 = sdecs_1 \rightarrow sbnds : sdecs$
and $decs \vdash L_1 \text{ ok}$
and $decs, sdecs_1, sdecs \vdash L_2 \text{ ok}$,
and $decs \vdash L_1 ++ L_2 \rightsquigarrow L_3$,
then $decs \vdash L_3 \text{ ok}$.

If a linkset is well-formed, then it neither imports nor exports the same label twice (although it may both import and export a particular label).

The rules process the imports in L_2 from left to right. If L_2 has no imports, then the following rule applies.

$$\frac{L = sdecs_0 \rightarrow sbnds : sdecs}{\frac{decs \vdash L ++ (\cdot \rightarrow sbnds' : sdecs') \rightsquigarrow}{sdecs_0 \rightarrow sbnds ++ sbnds' : sdecs ++ sdecs'}}$$

L_3 imports what L_1 does and exports what L_1 and L_2 do. To ensure that L_3 is well-formed, the rule uses the TS function $++$ to concatenate the exports in L_1 with the exports in L_2 , renaming labels exported by L_1 that are also exported by L_2 .

Otherwise, the rules examine the first import $lab \triangleright var : sig$ in L_2 and distinguish three mutually exclusive cases:

- L_1 exports *lab*.

$$\frac{\begin{array}{l} sdecs = sdecs'', lab \triangleright var' : sig', sdecs''' \\ decs, sdecs_0, sdecs \vdash_{\text{sub}} var' : sig' \preceq sig \rightsquigarrow mod : sig'' \\ sbnd := 1 \triangleright var = mod \quad sdec := 1 \triangleright var : sig'' \\ L := sdecs_0 \rightarrow sbnds ++ sbnd : sdecs ++ sdec \\ decs \vdash L ++ (sdecs_1 \rightarrow sbnds' : sdecs') \rightsquigarrow L'' \end{array}}{decs \vdash (sdecs_0 \rightarrow sbnds : sdecs) ++ (lab \triangleright var : sig, sdecs_1 \rightarrow sbnds' : sdecs') \rightsquigarrow L''}$$

The first premise picks out the L_1 export $lab \triangleright var' : sig'$ for *lab*; there can be at most one since L_1 is well-formed. The second premise calls the TS coercion compiler to match the export $var' : sig'$ to the import signature *sig*. Linking fails if no match is possible; otherwise, sig'' has the same “shape” as *sig*, but is fully transparent relative to the variable var' . The structure binding $sbnd : sdec$ is constructed using the coercion module *mod* at the signature sig'' , maximizing type sharing. The linkset L has the same imports as L_1 , and exports those of L_1 plus the result of the preceding coercion. To ensure that L is well-formed—in particular, that it exports nothing more than once—the rule uses $++$ to construct its exports.

- L_1 imports *lab* but does not export it.

$$\frac{\begin{array}{l} lab \notin \text{dom}(sdecs) \\ sdecs_0 = sdecs'', lab \triangleright var' : sig', sdecs''' \\ decs, sdecs_0, sdecs \vdash sig \equiv sig' : \text{Sig} \\ L' := \{var' / var\}(sdecs_1 \rightarrow sbnds' : sdecs') \\ decs \vdash (sdecs_0 \rightarrow sbnds : sdecs) ++ L' \rightsquigarrow L'' \end{array}}{decs \vdash (sdecs_0 \rightarrow sbnds : sdecs) ++ (lab \triangleright var : sig, sdecs_1 \rightarrow sbnds' : sdecs') \rightsquigarrow L''}$$

The first premise ensures L_1 does not export *lab*. The second premise picks out the L_1 import $lab \triangleright var' : sig'$. Linking fails if *sig* and sig' are not equivalent; otherwise, L' is constructed by changing references in the remainder of L_2 to use the import in L_1 .

⁵In this description of linkset merge, we suppress all details related to signature abbreviations.

Judgement...	Meaning...
$project \rightsquigarrow L$	project elaboration
$\Gamma \vdash srcunit \rightsquigarrow L$	unit elaboration
$\Gamma \vdash topdec \rightsquigarrow L$	top-level declaration elaboration
$\Gamma \vdash impexp \rightsquigarrow L$	import expression elaboration
$\Gamma \vdash sigbind \rightsquigarrow S$	signature binding elaboration
$\Gamma \vdash_{\text{ctx}} sigid \rightsquigarrow sig : \text{Sig}$	signature lookup
$\Gamma \vdash_{\text{ctx}} unitid \rightsquigarrow S$	
$\Gamma \text{ ok}$	Γ is well-formed

Figure 6. Elaboration judgements

<i>project</i>	::= \cdot	empty
	<i>project, srcunit</i>	source unit
	<i>project, L</i>	compiled unit(s)
<i>srcunit</i>	::= unit <i>unitid</i> = <i>topdec</i>	unit declaration
<i>topdec</i>	::= import <i>impexp</i>	open units
	<i>strdec</i>	
	signature <i>sigbind</i>	
	local <i>topdec</i> ₁ in <i>topdec</i> ₂ end	
	<i>topdec</i> ₁ <i>topdec</i> ₂	
<i>impexp</i>	::= <i>unitid</i> \langle intf spec end \rangle	open <i>unitid</i>
	<i>impexp</i> ₁ <i>impexp</i> ₂	
<i>sigbind</i>	::= <i>sigid</i> = <i>sigexp</i> \langle and <i>sigbind</i> \rangle	

Figure 7. SMLSC abstract syntax

- L_1 neither imports nor exports *lab*.

$$\frac{\begin{array}{l} lab \notin \text{dom}(sdecs) \cup \text{dom}(sdecs_0) \\ decs, sdecs_0, sdecs \vdash sig \equiv sig' : \text{Sig} \\ decs, sdecs_0 \vdash sig' : \text{Sig} \\ L := sdecs_0, lab \triangleright var : sig' \rightarrow sbnds : sdecs \\ decs \vdash L ++ (sdecs_1 \rightarrow sbnds' : sdecs') \rightsquigarrow L'' \end{array}}{decs \vdash (sdecs_0 \rightarrow sbnds : sdecs) ++ (lab \triangleright var : sig, sdecs_1 \rightarrow sbnds' : sdecs') \rightsquigarrow L''}$$

The first premise ensures that L_1 neither imports nor exports *lab*. The next two premises choose a signature sig' equivalent to *sig* but well-formed without reference to the exports of L_1 . Linking fails if no such signature exists—when opaque types exported by L_1 occur in *sig*. Otherwise, L is constructed by adding a new import to the imports in L_1 .

3.2 Elaboration

We define a semantics for SMLSC by giving rules for the elaboration judgements in Figure 6. We give the abstract syntax for SMLSC in Figure 7. The elaboration rules appear in Appendix B. These judgements have the following meaning.

- $project \rightsquigarrow L$. Elaborate *project*, using linkset merge to accumulate a resulting linkset L . A source unit is elaborated in a context Γ that declares the imports and exports in L .
- $\Gamma \vdash srcunit \rightsquigarrow L$. Elaborate the *topdec* in *srcunit* to the linkset

$$sdecs_0 \rightarrow sbnds : sdecs; S.$$

The imports $sdecs_0$ arise from the import declarations in *topdec*. The exports $sbnds : sdecs$ arise from the structure declarations in *topdec*. The signature abbreviations S arise from the signature declarations in *topdec*. The result, L , exports a single module

$$\overline{unitid} \triangleright var = [sbnds] : \overline{unitid} \triangleright var : [sdecs].$$

- $\Gamma \vdash \text{topdec} \rightsquigarrow L$. Elaborate *topdec* using linkset merge and identifier resolution.
- $\Gamma \vdash \text{impexp} \rightsquigarrow L$. Elaborate *impexp* using identifier resolution and *spec* elaboration.
- $\Gamma \vdash \text{sigbind} \rightsquigarrow S$. Elaborate *sigbind* using signature elaboration.

4. Implementation

The semantics of SMLSC avoids commitment to the meaning of “compilation,” “linking,” and “completion” to ensure compatibility with various implementation strategies. These phases may be implemented using classical methods (code generation during compilation, object code weaving during linking, and writing an executable for completion), or in other, more novel, ways (such as type checking during compilation, and code generation during linking). The design is, as far as we know, implementable in all current Standard ML compilers without requiring radical changes to their infrastructure.

Parallel Build. The purpose of separate compilation is to permit a client unit to be compiled independently of its implementation. A compiler can exploit this by permitting clients of a separately compiled unit to be compiled in parallel with one another in order to speed up system build times. The TILT compiler, which implements an earlier version of the present extension, implements such a strategy. Moreover, it also implements cut-off incremental recompilation [1], where it is able to interrupt the normal cascade of recompilation when a source change does not cause a unit’s interface to change.

Parsing. This presentation of SMLSC provides concrete and abstract syntax, but does not formalize parsing. The only issue that entangles separate compilation and parsing is fixity declarations. To support fixity declarations at parse-time, we include a parsing context in the concrete representation of linksets (object files). A source unit that is incrementally compiled against a linkset is parsed using that linkset’s included parsing context. We do not permit fixity specifications in user-specified interfaces, and therefore they do not affect interface matching or any other part of the semantics.

Note that a library may specify fixity information by placing appropriate declarations in the handoff unit. For example, to describe a matrix library that supplies an infix `**` operator for multiplication, we may write the following handoff unit:

```
unit Matrices =
top
  import MatricesImpl :
  intf
    type matrix
    val ** : matrix * matrix -> matrix
    (* ... *)
  end
  infix **
end
```

5. Multiple Interfaces for the Same Import

In Section 2 we presented the programming methodology of handoff units. As long as two linksets that import the same unit identifier do so by using the same handoff unit, they will always agree on the interface for that unit and so can be linked together. However, in some situations it may be useful to permit two clients to import the same unit, each with a different interface. Since interface matching, like signature matching, is coercive, this complicates the method-

ology of definite references by introducing “views” of the same underlying unit.

For example, suppose that two linksets L_1 and L_2 import the same unit `MathLib` at disparate interfaces I_1 and I_2 . This may happen because the developers of L_1 and L_2 compiled using different versions of the handoff unit for `MathLib`, or because the developers wrote their import interfaces by hand. The link

$$\text{link}(L_1, L_2)$$

fails because the linksets are required to agree on the interfaces of their common imports. Aside from recompiling the two linksets to use the same interface, the programmer has several options for resolving this situation. First, she can satisfy the imports by providing the implementation of `MathLib`:

$$\begin{aligned} L'_1 &= \text{link}(\text{MathLib}, L_1) \\ L &= \text{link}(L'_1, L_2) \end{aligned}$$

The first step satisfies the SC import of `MathLib` in L_1 , as long as the actual interface of `MathLib` matches the import interface I_1 . The result L'_1 does not import `MathLib`, so it does not conflict with the import of `MathLib` in L_2 . L'_1 does export `MathLib`, so if the actual interface of `MathLib` matches I_2 , then the second link succeeds. Because linking is left-associative, $L = \text{link}(\text{MathLib}, L_1, L_2)$ accomplishes the same thing.

Any implementation of `MathLib` that satisfies both I_1 and I_2 will suffice. Because we do not require unit names to be globally unique, this implementation of `MathLib` might even import `MathLib` (again) and then contain some glue code to make it compatible with the two given interfaces I_1 and I_2 (Figure 8). We expect such cases to be uncommon, the preferred methodology being to use a single handoff unit for all clients.

6. Related Work

There are several closely related systems that influenced the design of SMLSC.

The notion of linkset in SMLSC comes from Cardelli’s investigation of separate compilation and type-safe linking in the simply-typed λ -calculus [5]. Our formalization of linking extends these ideas to support the Standard ML module system including signature subtyping, abstract types, and module and type definitions in structures.

Harper and Pierce [11] discuss language design for module systems, including separate compilation. Particularly relevant to the current work is their discussion of sharing of abstract types. They describe the use of definite references to avoid the coherence problems (and excess sharing specifications) that arise from aliasing.

The notion of a handoff unit bears some resemblance to the use of `.h` files in C. The presence of function prototypes in a `.h` file provides an interface for application code that includes that header file. Code that references a prototyped function triggers a link-time demand for that function. The degree of link-time type-checking varies across C implementations. Usually, type correctness is assured by programming conventions.




Figure 8. The linkset L_{glue} imports `MathLib` at interface I and then exports it to satisfy the imports in L_1 and L_2 at disparate interfaces I_1 and I_2 .

Glew and Morrisett [8] describe separate compilation for Typed Assembly Language [19]. Their language, MTAL, permits type definitions, abstract types, and polymorphic types in interfaces and supports recursive linking.

Jim [14] describes a λ -calculus P_2 with rank 2 intersection types that has *principal typings*. The principal typings property means that from a term M , one can infer both Γ and τ such that any typing derivation $\Gamma' \vdash M : \tau'$ is an instance of $\Gamma \vdash M : \tau$. In a system with principal typings, program fragments can be separately compiled without context information, meaning that SC imports need not even specify interfaces. Standard ML, however, does not have principal typings. It remains an open problem to design a type system that supports principal types for features such as abstract and recursive types, and type definitions in modules.

Objective Caml. The separate compilation system of Objective Caml (O'Caml) [15] is similar in many regards to SMLSC. The declaration of a unit U is an O'Caml module stored within a file called $U.m1$. The interface for U may optionally be given in a file called $U.m1i$. If the interface is present, other units depending on U can compile even if the implementation is not available, just as in SMLSC. Because the filename of an interface indicates the unit that it describes, O'Caml interfaces play the role of handoff units in SMLSC. Additionally, O'Caml's use of the filesystem to provide a canonical location for each unit and interface means that all unit references are definite.

On the other hand, O'Caml's dependence on the filesystem means that the language is not independent from its environment. For instance, unit names are limited to valid filenames on the host system, and restructuring a project on disk may force changes to the code. Another significant difference is that O'Caml conflates the notions of units and modules. This earns O'Caml some conceptual economy, but it makes it impossible to separate the notions of top-level declarations and structure components. This makes it necessary to support signature and functor definitions within structures, so such a choice would not be compatible with our design principle of conservativity over Standard ML. Finally, unlike SMLSC, O'Caml and its separate compilation system are defined informally in terms of their implementation.

Moscow ML. The Moscow ML [20] compiler for Standard ML supports a separate compilation system nearly identical to Objective Caml's. Moscow ML extends the Standard ML module system to allow (among other things) `functor` and `signature` declarations in structures and specifications for them in signatures. Then, like O'Caml, units are structures. In contrast, SMLSC does not require any changes to the Standard ML module language.

Other Standard ML implementations include mechanisms for breaking programs up into compilation units. None support separate compilation in the sense we use it here; they use the term to mean cut-off incremental recompilation (recall Section 4).

SML/NJ CM. The Compilation Manager for Standard ML of New Jersey (CM) [3] is a tool for compiling Standard ML programs spread across many source files. CM permits a program to be divided into a hierarchy of libraries [4]. A library comprises a list of imported libraries, Standard ML source files, and a list of symbols exported by the library. Dependencies between libraries are explicit but dependencies among the source files in a library are inferred [2, 9]. CM provides control over the identifiers visible to a source file, and supports conditional compilation, parallel compilation, and cut-off incremental recompilation. CM provides no way for the programmer to write interfaces nor to compile against unimplemented units. SMLSC is not a replacement for CM; we believe that dependency analysis and recompilation tools are useful, and that SMLSC provides a good linguistic target for such tools.

ML Basis. The MLton compiler [18] and ML Kit [17] implement a language called ML Basis. A "basis" in their terminology is what we call a unit. An ML Basis program is a series of declarations, including a binding construct for bases and an open construct for basis identifiers. These are analogous to SMLSC's unit declaration and `IC import` declaration. Like SMLSC, the order of compilation entities is explicit, and thus each program has unambiguous meaning. ML Basis is given a formal semantics [6] in terms of *The Definition of Standard ML*. The implementation of ML Basis in the ML Kit supports cut-off incremental recompilation based on Elsmann's thesis work [7]. Like CM, ML Basis does not provide a way for programmers to write down interfaces or separately compile against unimplemented bases.

7. Conclusion

We have presented an extension to Standard ML for separate compilation called SMLSC. Its focus is the *unit*, a program fragment that can depend on other program fragments through either separate or incremental compilation. Via the programming idiom of handoff units—that uses both separate and incremental compilation—we limit the number and complexity of linguistic mechanisms while supporting a convenient programming style. Our formal and abstract definition of the language ensures that it is unambiguously specified, and admits a variety of implementation strategies.

Acknowledgments

We thank the anonymous referees for their constructive comments.

References

- [1] Rolf Adams, Walter Tichy, and Annette Weinert. The cost of selective recompilation and environment processing. *ACM Transactions on Software Engineering and Methodology*, 3(1):3–28, January 1994.
- [2] Matthias Blume. Dependency analysis for Standard ML. *ACM Transactions on Programming Languages and Systems*, 21(4):790–812, 1999.
- [3] Matthias Blume. CM: The SML/NJ compilation and library manager (for SML/NJ version 110.40 and later) user manual, 2002. <http://www.smlnj.org/doc/CM/new.pdf>.
- [4] Matthias Blume and Andrew W. Appel. Hierarchical modularity. *ACM Transactions on Programming Languages and Systems*, 21(4):813–847, 1999.
- [5] Luca Cardelli. Program fragments, linking, and modularization. In *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 266–277. ACM Press, 1997.
- [6] Henry Cejtin, Matthew Fluet, Suresh Jagannathan, and Stephen Weeks. Formal specification of the ML Basis system, January 2005. <http://mlton.org/MLBasis>.
- [7] Martin Elsmann. *Program Modules, Separate Compilation, and Intermodule Optimisation*. PhD thesis, Department of Computer Science, University of Copenhagen, January 1999.
- [8] Neal Glew and Greg Morrisett. Type-safe linking and modular assembly language. In *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 250–261. ACM Press, 1999.
- [9] Robert Harper, Peter Lee, Frank Pfenning, and Eugene Rollins. Incremental recompilation for Standard ML of New Jersey. Technical Report CMU-CS-94-116, School of Computer Science, Carnegie Mellon University, February 1994.
- [10] Robert Harper and Mark Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 123–137. ACM Press, 1994.

- [11] Robert Harper and Benjamin C. Pierce. Design considerations for ML-style module systems. In Benjamin C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 8, pages 293–346. MIT Press, 2005.
- [12] Robert Harper and Christopher Stone. An interpretation of Standard ML in type theory. Technical Report CMU-CS-97-147, School of Computer Science, Carnegie Mellon University, June 1997.
- [13] Robert Harper and Christopher Stone. A type-theoretic interpretation of Standard ML. In Gordon Plotkin, Colin Stirling, and Mads Tofte, editors, *Proof, Language, and Interaction: Essays in Honour of Robin Milner*. MIT Press, 2000.
- [14] Trevor Jim. What are principal typings and what are they good for? In *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press, 1996.
- [15] Xavier Leroy, Damien Doligez, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. The Objective Caml system release 3.09: Documentation and user’s manual, 2005. <http://caml.inria.fr/pub/docs/manual-ocaml/index.html>.
- [16] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.
- [17] MLKit web site. http://www.itu.dk/research/mlkit/index.php/Main_Page.
- [18] MLton web site. <http://mlton.org/>.
- [19] Greg Morrisett, David Walker, Karl Cray, and Neal Glew. From system F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):527–568, 1999.
- [20] Sergei Romanenko, Claudio Russo, and Peter Sestoft. Moscow ML owner’s manual version 2.00, June 2000. <http://www.dina.kvl.dk/~sestoft/mosml/manual.pdf>.
- [21] David Swasey, Tom Murphy, VII, Karl Cray, and Robert Harper. A separate compilation extension to Standard ML (extended technical report). Technical Report CMU-CS-06-133, School of Computer Science, Carnegie Mellon University, 2006.

A. Linking Rules

The typed semantics defines a closed structure $mod_{basis}:sig_{basis}$ serving as an initial basis for the TSIL. The elaborator assumes Γ declares $basis:sig_{basis}$, which includes components such as the built-in `Match` exception. This basis structure is introduced in Rule 6 for completion and Rule 12 for elaboration of source units in projects.

We use the following definitions and notation.

- Writing $BV(dec)$ for the variable declared by dec , we define the bound variables of a structure declaration list, $BV(sdecs)$, by

$$BV(lab_1 \triangleright dec_1, \dots, lab_n \triangleright dec_n) = \{BV(dec_1), \dots, BV(dec_n)\}.$$

A linkset

$$L = sdecs_0 \rightarrow sbnds : sdecs; S$$

binds variables $BV(sdecs_0)$ with scope $sbnds : sdecs; S$ and variables $BV(sdecs)$ with scope S . We write $BV(L)$ for $BV(sdecs_0) \cup BV(sdecs)$.

- For readability, we sometimes elide variables in structure bindings and declarations. It should be immediately obvious how to consistently restore these with fresh variables.
- We assume that unit identifiers are disjoint from all other identifier classes.
- We assume that the TS overbar injection $\bar{\cdot}$ maps identifiers of different classes to different labels and that there are infinitely many labels not in its range.

We assume that its range includes neither the distinguished label basis, nor the labels chosen fresh in the rules.

- Structure declaration lists $sdecs$, signature abbreviations S , and so on specify lists of elements. We adopt the following notation for lists.

- We denote by (\cdot, \cdot) the operation of syntactic concatenation; for example, S, S' .

- We sometimes use pattern matching at the left end of a list, writing $sigid=sig, S$ to match the first binding in the list.

- We usually omit the initial $;$; for example,

$$sigid_1=sig_1, \dots, sigid_n=sig_n.$$

- We define the domain of a signature abbreviation, $\text{dom}(S)$, by

$$\begin{aligned} \text{dom}(\cdot) &= \emptyset \\ \text{dom}(S, sigid=sig) &= \text{dom}(S) \cup \{sigid\} \\ \text{dom}(S, unitid=S_u) &= \text{dom}(S) \cup \{unitid\}. \end{aligned}$$

- We define the function $S++S'$ by

$$\begin{aligned} (\cdot ++ S') &= S' \\ ((sigid=sig, S) ++ S') &= \begin{cases} sigid=sig, S'' & \text{if } sigid \notin \text{dom}(S'') \\ S'' & \text{otherwise} \end{cases} \\ &\text{where } S'' = S++S' \\ ((unitid=S_u, S) ++ S') &= \begin{cases} unitid=S_u, S'' & \text{if } unitid \notin \text{dom}(S'') \\ S'' & \text{otherwise} \end{cases} \\ &\text{where } S'' = S++S'. \end{aligned}$$

It concatenates S and S' , making the result well-formed by dropping signature abbreviations if $\text{dom}(S) \cap \text{dom}(S') \neq \emptyset$.

- We write both “=” and “:=” in side-conditions. Interpreting the rules algorithmically, the former pattern-matches inputs, and the latter specifies an output.

$$\boxed{decs \vdash L \text{ ok}}$$

$$\frac{\begin{array}{l} decs \vdash sdecs_0 \text{ ok} \\ decs, sdecs_0 \vdash sbnds : sdecs \\ decs, sdecs_0, sdecs \vdash S \text{ ok} \\ sdecs_0 = lab_1 \triangleright var_1 : [sdecs_1], \dots, lab_n \triangleright var_n : [sdecs_n] \end{array}}{decs \vdash sdecs_0 \rightarrow sbnds : sdecs; S \text{ ok}} \quad (1)$$

Rule 1: Imports are restricted to structures. The elaborator in Appendix B needs nothing else.

$$\boxed{decs \vdash S \text{ ok}}$$

$$\frac{\vdash decs \text{ ok}}{decs \vdash \cdot \text{ ok}} \quad (2)$$

$$\frac{decs \vdash sig : \text{Sig} \quad decs \vdash S \text{ ok} \quad sigid \notin \text{dom}(S)}{decs \vdash sigid=sig, S \text{ ok}} \quad (3)$$

$$\frac{\begin{array}{l} decs \vdash S' \text{ ok} \quad decs \vdash S \text{ ok} \quad unitid \notin \text{dom}(S) \\ S' = (sigid_1=sig_1, \dots, sigid_n=sig_n) \end{array}}{decs \vdash unitid=S', S \text{ ok}} \quad (4)$$

$$\boxed{L \rightsquigarrow exp : \{\}}$$

$$\frac{lab \notin \text{dom}(sdecs)}{\cdot \rightarrow sbnds : sdecs; S \rightsquigarrow [sbnds, lab=\{\}], lab : \{}} \quad (5)$$

$$\frac{L_{basis} := \cdot \rightarrow \mathbf{basis} = \mathit{mod}_{basis} : \mathbf{basis} : \mathit{sig}_{basis} ; \cdot}{\frac{\vdash L_{basis} ++ L \rightsquigarrow L' \quad L' \rightsquigarrow \mathit{exp} : \{\}}{L \rightsquigarrow \mathit{exp} : \{\}}} \quad (6)$$

$$\boxed{\mathit{decs} \vdash L ++ L' \rightsquigarrow L''}$$

$$\frac{L = \mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S}{\mathit{decs} \vdash L ++ (\cdot \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow \mathit{sdecs}_0 \rightarrow \mathit{sbnds} ++ \mathit{sbnds}' : \mathit{sdecs} ++ \mathit{sdecs}' ; S ++ S'} \quad (7)$$

$$\frac{\begin{array}{l} \mathit{sdecs} = \mathit{sdecs}'' , \mathit{lab} \triangleright \mathit{var}' : \mathit{sig}' , \mathit{sdecs}''' \\ \mathit{decs} , \mathit{sdecs}_0 , \mathit{sdecs} \vdash_{\text{sub}} \mathit{var}' : \mathit{sig}' \preceq \mathit{sig} \rightsquigarrow \mathit{mod} : \mathit{sig}'' \\ \mathit{sbnd} := 1 \triangleright \mathit{var} = \mathit{mod} \quad \mathit{sdec} := 1 \triangleright \mathit{var} : \mathit{sig}'' \\ L := \mathit{sdecs}_0 \rightarrow \mathit{sbnds} ++ \mathit{sbnd} : \mathit{sdecs} ++ \mathit{sdec} ; S \\ \mathit{decs} \vdash L ++ (\mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L'' \end{array}}{\mathit{decs} \vdash L ++ (\mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''} \quad (8)$$

$$\frac{\mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ (\mathit{lab} \triangleright \mathit{var} : \mathit{sig} , \mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''}{\mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ (\mathit{lab} \triangleright \mathit{var} : \mathit{sig} , \mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''}$$

$$\frac{\begin{array}{l} \mathit{lab} \notin \text{dom}(\mathit{sdecs}) \\ \mathit{sdecs}_0 = \mathit{sdecs}'' , \mathit{lab} \triangleright \mathit{var}' : \mathit{sig}' , \mathit{sdecs}''' \\ \mathit{decs} , \mathit{sdecs}_0 , \mathit{sdecs} \vdash \mathit{sig} \equiv \mathit{sig}' : \mathbf{Sig} \\ L := \{ \mathit{var}' / \mathit{var} \} (\mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \\ \mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ L' \rightsquigarrow L'' \end{array}}{\mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ (\mathit{lab} \triangleright \mathit{var} : \mathit{sig} , \mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''} \quad (9)$$

$$\frac{\mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ (\mathit{lab} \triangleright \mathit{var} : \mathit{sig} , \mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''}{\mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ (\mathit{lab} \triangleright \mathit{var} : \mathit{sig} , \mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''}$$

$$\frac{\begin{array}{l} \mathit{lab} \notin \text{dom}(\mathit{sdecs}) \cup \text{dom}(\mathit{sdecs}_0) \\ \mathit{decs} , \mathit{sdecs}_0 , \mathit{sdecs} \vdash \mathit{sig} \equiv \mathit{sig}' : \mathbf{Sig} \\ \mathit{decs} , \mathit{sdecs}_0 \vdash \mathit{sig}' : \mathbf{Sig} \\ L := \mathit{sdecs}_0 , \mathit{lab} \triangleright \mathit{var} : \mathit{sig}' \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S \\ \mathit{decs} \vdash L ++ (\mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L'' \end{array}}{\mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ (\mathit{lab} \triangleright \mathit{var} : \mathit{sig} , \mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''} \quad (10)$$

$$\frac{\mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ (\mathit{lab} \triangleright \mathit{var} : \mathit{sig} , \mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''}{\mathit{decs} \vdash (\mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S) ++ (\mathit{lab} \triangleright \mathit{var} : \mathit{sig} , \mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L''}$$

B. Elaboration Rules

We change the TS elaborator to expand signature abbreviations. First, we modify every TS elaboration judgement and rule using a TS elaboration context sdecs to use $\mathit{sdecs} ; S$. A context $\mathit{sdec} , \mathit{sdecs} ; S$ binds the variable $\text{BV}(\mathit{sdec})$ with scope $\mathit{sdecs} ; S$ and a context $\mathit{sdecs} ; S$ binds variables $\text{BV}(\mathit{sdecs})$ with scope S . We define $\text{BV}(\Gamma)$ by $\text{BV}(\mathit{sdecs})$. Second, we extend the syntax for TSEL signature expressions:

$$\mathit{sigexp} ::= \dots \quad \mathit{sigid} \quad \text{signature identifier}$$

Finally, we extend the TS judgement $\Gamma \vdash \mathit{sigexp} \rightsquigarrow \mathit{sig} : \mathbf{Sig}$, adding the rule

$$\frac{\Gamma \vdash_{\text{ctx}} \mathit{sigid} \rightsquigarrow \mathit{sig} : \mathbf{Sig}}{\Gamma \vdash \mathit{sigid} \rightsquigarrow \mathit{sig} : \mathbf{Sig}}$$

to elaborate signature identifiers.

We use the following definitions and notation.

- To extend an elaboration context $\Gamma = \mathit{sdecs} ; S$, we write

$$\begin{array}{l} \Gamma , \mathit{dec} \text{ for } \mathit{sdecs} , 1 \triangleright \mathit{dec} ; S , \\ \Gamma , \mathit{sdecs}' \text{ for } \mathit{sdecs} , \mathit{sdecs}' ; S , \text{ and} \\ \Gamma , S' \text{ for } \mathit{sdecs} ; S , S' . \end{array}$$

We also define a function $R(\mathit{sdecs})$ that renames the labels in sdecs to make them inaccessible to identifier resolution:

$$R(\mathit{lab}_1 \triangleright \mathit{dec}_1 , \dots , \mathit{lab}_n \triangleright \mathit{dec}_n) = 1 \triangleright \mathit{dec}_1 , \dots , 1 \triangleright \mathit{dec}_n .$$

- We define a function $U(\mathit{sdecs})$ that drops the labels in sdecs :

$$U(\mathit{lab}_1 \triangleright \mathit{dec}_1 , \dots , \mathit{lab}_n \triangleright \mathit{dec}_n) = \mathit{dec}_1 , \dots , \mathit{dec}_n .$$

- When an elaboration context $\Gamma = \mathit{sdecs} ; S$ appears in a judgement requiring an IL context decs , we implicitly coerce Γ to $U(\mathit{sdecs})$.

- We define the substitution function $\sigma(\mathit{var} , \mathit{sdecs} , S)$ by

$$\begin{array}{l} \sigma(\mathit{var} , \cdot , S) = S \\ \sigma(\mathit{var} , (\mathit{lab} \triangleright \mathit{dec} , \mathit{sdecs}) , S) = \\ \quad \{ \mathit{var} . \mathit{lab} / \text{BV}(\mathit{dec}) \} \sigma(\mathit{var} , \mathit{sdecs} , S) \end{array}$$

where $\{ \mathit{path} / \mathit{var} \} S$ denotes the capture-free substitution of path for free occurrences of var in S . Rule 14 uses σ to elaborate source units.

$$\boxed{\mathit{project} \rightsquigarrow L}$$

$$\frac{\cdot \rightsquigarrow (\cdot \rightarrow \cdot : \cdot ; \cdot)}{\cdot \rightsquigarrow (\cdot \rightarrow \cdot : \cdot ; \cdot)} \quad (11)$$

$$\frac{\begin{array}{l} \mathit{project} \rightsquigarrow L \quad \mathit{basis} \notin \text{BV}(L) \\ L = \mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S \\ \Gamma := \mathit{basis} : \mathit{sig}_{\mathit{basis}} , R(\mathit{sdecs}_0) , \mathit{sdecs} ; S \\ \Gamma \vdash \mathit{srcunit} \rightsquigarrow L' \quad \mathit{var} \notin \text{BV}(L') \\ \mathit{sdecs}'_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S' := \{ \mathit{var} / \mathit{basis} \} L' \\ \mathit{sdecs}_1 := \mathit{basis} \triangleright \mathit{var} : \mathit{sig}_{\mathit{basis}} , \mathit{sdecs}'_1 \\ \vdash L ++ (\mathit{sdecs}_1 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S') \rightsquigarrow L'' \end{array}}{\mathit{project} , \mathit{srcunit} \rightsquigarrow L''} \quad (12)$$

Rule 12: The side-condition $\mathit{basis} \notin \text{BV}(L)$ can always be achieved by renaming bound variables in L .

$$\frac{\mathit{project} \rightsquigarrow L \quad \text{BV}(L) \cap \text{BV}(L') = \emptyset \quad \vdash L' \text{ ok} \quad \vdash L ++ L' \rightsquigarrow L''}{\mathit{project} , L' \rightsquigarrow L''} \quad (13)$$

$$\boxed{\Gamma \vdash \mathit{srcunit} \rightsquigarrow L}$$

$$\frac{\begin{array}{l} \Gamma \vdash \mathit{topdec} \rightsquigarrow L \\ L = \mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S \\ \mathit{var} \notin \text{BV}(\Gamma) \cup \text{BV}(L) \\ \mathit{sbnds}' := \overline{\mathit{unitid}} \triangleright \mathit{var} = [\mathit{sbnds}] \\ \mathit{sdecs}' := \mathit{unitid} \triangleright \mathit{var} : [\mathit{sdecs}] \\ S' := \mathit{unitid} = \sigma(\mathit{var} , \mathit{sdecs} , S) \end{array}}{\Gamma \vdash \mathbf{unit} \mathit{unitid} = \mathit{topdec} \rightsquigarrow \mathit{sdecs}_0 \rightarrow \mathit{sbnds}' : \mathit{sdecs}' ; S'} \quad (14)$$

$$\boxed{\Gamma \vdash \mathit{topdec} \rightsquigarrow L}$$

$$\frac{\Gamma \vdash \mathit{impexp} \rightsquigarrow L}{\Gamma \vdash \mathbf{import} \mathit{impexp} \rightsquigarrow L} \quad (15)$$

$$\frac{\Gamma \vdash \mathit{strdec} \rightsquigarrow \mathit{sbnds} : \mathit{sdecs}}{\Gamma \vdash \mathit{strdec} \rightsquigarrow \cdot \rightarrow \mathit{sbnds} : \mathit{sdecs} ; \cdot} \quad (16)$$

$$\frac{\Gamma \vdash \mathit{sigbind} \rightsquigarrow S}{\Gamma \vdash \mathbf{signature} \mathit{sigbind} \rightsquigarrow \cdot \rightarrow \cdot : \cdot ; S} \quad (17)$$

$$\frac{\begin{array}{l} \Gamma \vdash \mathit{topdec} \rightsquigarrow \mathit{sdecs}_0 \rightarrow \mathit{sbnds} : \mathit{sdecs} ; S \\ \mathit{var} \notin \text{BV}(\Gamma) \cup \text{BV}(\mathit{sdecs}_0) \\ \Gamma , R(\mathit{sdecs}_0) , 1^* \triangleright \mathit{var} : [\mathit{sdecs}] , S \vdash \mathit{topdec}' \rightsquigarrow L' \\ L := \mathit{sdecs}_0 \rightarrow 1 \triangleright \mathit{var} = [\mathit{sbnds}] : 1 \triangleright \mathit{var} : [\mathit{sdecs}] ; \cdot \\ \Gamma \vdash L ++ L' \rightsquigarrow L'' \end{array}}{\Gamma \vdash \mathbf{local} \mathit{topdec} \text{ in } \mathit{topdec}' \text{ end} \rightsquigarrow L''} \quad (18)$$

$$\frac{\begin{array}{l} \Gamma \vdash \text{topdec} \rightsquigarrow L \\ L = \text{sdecs}_0 \rightarrow \text{sbnds} : \text{sdecs}; S \\ \Gamma, R(\text{sdecs}_0), \text{sdecs}, S \vdash \text{topdec}' \rightsquigarrow L' \\ \Gamma \vdash L++L' \rightsquigarrow L'' \end{array}}{\Gamma \vdash \text{topdec topdec}' \rightsquigarrow L''} \quad (19)$$

$$\frac{\text{sdecs}; S \text{ ok} \quad \text{sdecs} \vdash \text{sig} : \text{Sig}}{\text{sdecs}; S, \text{sigid}=\text{sig} \text{ ok}} \quad (31)$$

$$\frac{\text{sdecs}; S \text{ ok} \quad \text{sdecs} \vdash S' \text{ ok}}{S' = (\text{sigid}_1=\text{sig}_1, \dots, \text{sigid}_n=\text{sig}_n)} \quad (32)$$

$$\boxed{\Gamma \vdash \text{impexp} \rightsquigarrow L}$$

$$\frac{\begin{array}{l} \Gamma \vdash_{\text{ctx}} \overline{\text{unitid}} \rightsquigarrow \text{var} : \text{sig} \\ \Gamma \vdash_{\text{ctx}} \text{unitid} \rightsquigarrow S \quad \text{var}' \notin \text{BV}(\Gamma) \\ L := \overline{\text{unitid}} \triangleright \text{var}' : \text{sig} \rightarrow 1^* = \text{var}' : 1^* : \text{sig}; S \end{array}}{\Gamma \vdash \text{unitid} \rightsquigarrow L} \quad (20)$$

Rule 20: Rules 12, 18, and 19 use $R(\cdot)$ to hide imported units from IR imports. The signature sig should be fully selfified.

$$\frac{\begin{array}{l} \Gamma \vdash \text{spec} \rightsquigarrow \text{sdecs} \quad \text{var}' \notin \text{BV}(\Gamma) \\ \Gamma, \text{var}' : [\text{sdecs}] \vdash \text{var}' : \text{sig} \\ L := \overline{\text{unitid}} \triangleright \text{var}' : [\text{sdecs}] \rightarrow 1^* = \text{var}' : 1^* : \text{sig}; \cdot \end{array}}{\Gamma \vdash \text{unitid} : \text{intf spec end} \rightsquigarrow L} \quad (21)$$

Rule 21: The signature sig should be fully selfified.

$$\frac{\begin{array}{l} \Gamma \vdash \text{impexp} \rightsquigarrow L \quad \Gamma \vdash \text{impexp}' \rightsquigarrow L' \\ \text{BV}(L) \cap \text{BV}(L') = \emptyset \quad \Gamma \vdash L++L' \rightsquigarrow L'' \end{array}}{\Gamma \vdash \text{impexp impexp}' \rightsquigarrow L''} \quad (22)$$

$$\boxed{\Gamma \vdash \text{sigbind} \rightsquigarrow S}$$

$$\frac{\begin{array}{l} \Gamma \vdash \text{sigexp} \rightsquigarrow \text{sig} : \text{Sig} \quad S := \text{sigid}=\text{sig} \\ \langle \Gamma \vdash \text{sigbind} \rightsquigarrow S' \quad \text{sigid} \notin \text{dom}(S') \rangle \end{array}}{\Gamma \vdash \text{sigid} = \text{sigexp} \langle \text{and sigbind} \rangle \rightsquigarrow S \langle S' \rangle} \quad (23)$$

Rule 23: Either all optional elements or none must be present.

$$\boxed{\Gamma \vdash_{\text{ctx}} \text{sigid} \rightsquigarrow \text{sig} : \text{Sig}}$$

$$\frac{}{\text{sdecs}; S, \text{sigid}=\text{sig} \vdash \text{sigid} \rightsquigarrow \text{sig} : \text{Sig}} \quad (24)$$

$$\frac{\text{sigid}' \neq \text{sigid}}{\text{sdecs}; S \vdash \text{sigid} \rightsquigarrow \text{sig} : \text{Sig}} \quad (25)$$

$$\frac{\text{sdecs}; S, \text{sigid}'=\text{sig}' \vdash \text{sigid} \rightsquigarrow \text{sig} : \text{Sig}}{\text{sdecs}; S \vdash \text{sigid} \rightsquigarrow \text{sig} : \text{Sig}} \quad (26)$$

$$\frac{}{\text{sdecs}; S, \text{unitid}=S' \vdash \text{sigid} \rightsquigarrow \text{sig} : \text{Sig}}$$

$$\boxed{\Gamma \vdash_{\text{ctx}} \text{unitid} \rightsquigarrow S}$$

$$\frac{}{\text{sdecs}; S, \text{unitid}=S' \vdash \text{unitid} \rightsquigarrow S'} \quad (27)$$

$$\frac{\text{unitid}' \neq \text{unitid}}{\text{sdecs}; S \vdash \text{unitid} \rightsquigarrow S''} \quad (28)$$

$$\frac{\text{sdecs}; S, \text{unitid}'=S' \vdash \text{unitid} \rightsquigarrow S''}{\text{sdecs}; S \vdash \text{unitid} \rightsquigarrow S'} \quad (29)$$

$$\frac{}{\text{sdecs}; S, \text{sigid}=\text{sig} \vdash \text{unitid} \rightsquigarrow S'}$$

$$\boxed{\Gamma \text{ ok}}$$

$$\frac{\vdash U(\text{sdecs}) \text{ ok}}{\text{sdecs}; \cdot \text{ ok}} \quad (30)$$