What is a Recursive Module?*

Karl Crary
Carnegie Mellon University

Robert Harper
Carnegie Mellon University

Sidd Puri
Microsoft Corporation

Abstract

A hierarchical module system is an effective tool for structuring large programs. Strictly hierarchical module systems impose an acyclic ordering on import dependencies among program units. This can impede modular programming by forcing mutually-dependent components to be consolidated into a single module. Recently there have been several proposals for module systems that admit cyclic dependencies, but it is not clear how these proposals relate to one another, nor how one might integrate them into an expressive module system such as that of ML.

To address this question we provide a type-theoretic analysis of the notion of a recursive module in the context of a "phase-distinction" formalism for higher-order module systems. We extend this calculus with a recursive module mechanism and a new form of signature, called a *recursively dependent signature*, to support the definition of recursive modules. These extensions are justified by an interpretation in terms of more primitive language constructs. This interpretation may also serve as a guide for implementation.

1 Introduction

Hierarchical decomposition is a fundamental design principle for controlling the complexity of large programs. According to this principle a software system is to be decomposed into a collection of modules whose dependency relationships form a directed, acyclic graph. Most modern programming languages include module systems that support hierarchical decomposition. Many, such as Standard ML [16] and O'Caml [14], also support parameterized, or generic, modules to better support code re-use.

There is no question that hierarchical design is an important tool for structuring large systems. It has often been noted, however, that strict adherence to a hierarchical architecture can preclude the decomposition of a system into "mind-sized" components. In some situations the natural decomposition of a system into modules introduces cyclic dependencies, which cannot be expressed in a purely hierarchical formalism. The only solution is to consolidate mutually-dependent fragments into a single module, which partially undermines the very idea of modular organization.

In response several authors have proposed linguistic mechanisms to support non-hierarchical modular decomposition. Recent examples include: Sirer, *et al.'s* extension of Modula-3 with a "cross-linking" mechanism [21]; Flatt and Felleisen's extension of their MzScheme language with cyclically-dependent "units" [8]; Duggan and Sourelis's "mixin modules" that extend the Standard ML module system with a special "mixlink" construct for integrating mutually-dependent structures [6, 7]; and Ancona and Zucca's algebraic formalism for mixin modules [2]. Each of these proposals seeks to address the problem of cyclic dependencies in a module system, but each does so in a slightly different way. For example, Flatt and Felleisen's formalism does not address the critical issue of controlling propagation of type information across module boundaries. Duggan and Sourelis's framework relies on a syntactic transformation that, in effect, coalesces the code of mutually-dependent modules into a single module. It is not clear what are the fundamental ideas, nor is it clear how to integrate the various aspects of these proposals into a full-featured module system.

It is natural to ask: what is a recursive module? We propose to address this question in the framework of type theory, which has proved to be a powerful tool for both the design and implementation of module systems. We conduct our analysis in the context of the "phase distinction" module formalism introduced by Harper, Mitchell, and Moggi [11] (hereafter, HMM), augmented to support recursive types and functions, and to support type definitions in signatures [9, 13]. The phase distinction calculus provides a rigorous account of higher-order modules (supporting hierarchy and parameterization) in a framework that makes explicit the critical distinction between the static, or compile-time, part of a module

*This research was sponsored by the Advanced Research Projects Agency CSTO under the title "The Fox Project: Advanced Languages for Systems Software", ARPA Order No. C533, issued by ESC/ENS under Contract No. F19628-95-C-0050. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

To appear at the 1999 Conference on Programming Language Design and Implementation.

and the dynamic, or run-time, part. This calculus has proved to be of fundamental importance to the implementation of higher-order modules, as evidenced by its use in Shao's FLINT formalism used in the SML/NJ compiler [18, 20] and in the TIL/ML compiler [23].

Our analysis proceeds in two stages. First we consider a straightforward extension of the phase distinction calculus with a notion of recursive (self-referential) module. An interpretation of this new construct is provided by an interpretation of it into the primitive module formalism of the phase distinction calculus. This interpretation renders the compile-time part as a recursive type and the run-time part as a recursive function, as might be expected. In essence a recursive module is just a convenient way of introducing recursive types and functions.

Unfortunately this simple-minded extension does not go far enough to be of much practical use. As Duggan and Sourelis have observed [7], it is of critical importance for most practical examples that the type equations that hold of a recursive module be propagated into the definition of the recursive module itself. In essence the definitions of the type components of a recursive module must be taken to be the types that they will eventually turn out to be once the recursive declaration has been processed. Accounting for this "forward reference" is the core contribution of our work. We introduce a new form of signature (interface) for recursive modules, called a *recursively dependent signature*, that allows us to capture the required type identities during type checking of a recursive module binding. This significantly increases the expressive power of the recursive module formalism, and is, we assert, of fundamental importance to the very idea of recursive modules.

In this paper we aim to focus on the core issues lying at the center of a recursive module system, so we study recursive modules in the framework of a small *internal* language that is sufficient to bring out the main issues and that could be used by a type-directed compiler to implement recursive modules. Therefore, we make no specific proposals as to what form an *external* language supporting recursive modules should take, although we do present most of our examples in a hypothetical external language. Indeed, some important questions regarding the design of an external language remain open, such as the practicality of typechecking. In Section 5 we make some observations and preliminary proposals regarding the design of an external language.

2 Type-Theoretic Framework

We begin by presenting the framework in which we conduct our analysis. We will conduct our examples using an informal external language closely modeled after the syntax of Standard ML. The external language is then elaborated into the type-theoretic internal language that we describe below. We will treat the elaboration process informally, illustrating it by examples. Details of how elaboration may be formalized in a general setting appear in Harper and Stone [12].

Our internal language is an extension of the *phase distinction calculus* of Harper, Mitchell, and Moggi [11]. The language consists of two main components: a *core calculus*, a predicative variant of Girard's F_ω, and a *structure calculus*, extending the core language with a primitive module construct without explicit mechanisms for hierarchy (*e.g.*, substructures) or parameterization (*e.g.*, functors). Primitive modules consist of a static, or compile-time, part containing the type constructors of the module, together with a dynamic, or run-time, part containing the executable code of the module. This separation is known as the *phase distinction*. An important property of the formalism is that the phase distinction is maintained, even in the presence of higher-order (and, as we shall see, recursive) module constructs.

The main result of HMM is that higher-order module constructs are a definitional extension of the primitive structure calculus. In other words higher-order constructs are *already present* in the primitive structure calculus in the sense that they may be defined in terms of existing constructs. (This interpretation may be thought of as a compilation strategy for higher-order modules, and indeed this fact has been exploited in the FLINT [20] and TIL [23] compilers.) This means that we need not explicitly discuss higher-order module constructs in this paper, but rather appeal to HMM for a detailed discussion of their implicit presence.

To support the extension with recursive modules we enrich the core phase distinction calculus with these additional constructs:

1. Singleton and dependent kinds to allow expression of type sharing information in signatures. Related formalisms for expressing type sharing information are given by Harper and Lillibridge [9] and Leroy [13].

2. A fixed point operation for building collections of mutually-recursive type constructors. These recursive constructors are definitionally equal to their unrollings. We term such constructors *equi-recursive*, to distinguish them from the more conventional *iso-recursive* constructors, for which conversions between the constructors and their unrollings must be mediated by the explicit use of an isomorphism. We discuss the interplay of equi- and iso-recursive constructors in Section 5.3.

3. A fixed point operation for building collections of mutually-recursive functions. As will become apparent later on, we cannot (as in SML) limit this operation to collections of explicit lambda abstractions. Instead we formalize a notion of *valuability* (indicating terminating expressions) and a corresponding notion of *total function*, essentially as in Harper and Stone [12], but with the additional idea that recursively defined variables are not considered valuable within the body of their definitions, but are considered valuable in their subsequent scope.

In subsequent sections of this paper, we will further augment our structure calculus with various constructs for recursive modules, and then show how those constructs can be reduced to the elementary constructs discussed in this section.

$$
\begin{array}{lrcl}
\textit{kinds} & \kappa & ::= & T \mid 1 \mid \mathfrak{S}(c) \mid \Pi\alpha{:}\kappa_1.\kappa_2 \mid \Sigma\alpha{:}\kappa_1.\kappa_2 \\
\textit{constructors} & c & ::= & \alpha \mid \star \mid \lambda\alpha{:}\kappa.c \mid c_1 c_2 \mid \langle c_1, c_2\rangle \mid \\
 & & & \pi_i(c) \mid 1 \mid c_1 \rightharpoonup c_2 \mid c_1 \times c_2 \mid \mu\alpha{:}\kappa.c \\
\textit{types} & \sigma & ::= & c \mid \sigma_1 \to \sigma_2 \mid \sigma_1 \rightharpoonup \sigma_2 \mid \sigma_1 \times \sigma_2 \mid \\
 & & & \forall\alpha{:}\kappa.\sigma \\
\textit{terms} & e & ::= & x \mid \star \mid \lambda x{:}\sigma.e \mid e_1 e_2 \mid \langle e_1, e_2\rangle \mid \\
 & & & \pi_i(e) \mid \Lambda\alpha{:}\kappa.e \mid e[c] \mid \mathit{fix}(x{:}\sigma.e) \\
\textit{contexts} & \Gamma & ::= & \epsilon \mid \Gamma[\alpha : \kappa] \mid \Gamma[x : \sigma] \mid \\
 & & & \Gamma[\alpha \uparrow \kappa] \mid \Gamma[x \uparrow \sigma]
\end{array}
$$

Figure 1: The Core Calculus

$$
\begin{array}{rcl}
\mathfrak{S}(c : T) & \stackrel{\text{def}}{=} & \mathfrak{S}(c) \\
\mathfrak{S}(c : \Pi\alpha{:}\kappa_1.\kappa_2) & \stackrel{\text{def}}{=} & \Pi\alpha{:}\kappa_1.\mathfrak{S}(c\,\alpha : \kappa_2) \\
 & & (\text{for } \alpha \text{ not free in } c) \\
\mathfrak{S}(c : \kappa_1 \times \kappa_2) & \stackrel{\text{def}}{=} & \mathfrak{S}(\pi_1(c) : \kappa_1) \times \mathfrak{S}(\pi_2(c) : \kappa_2) \\
\mathfrak{S}(c : 1) & \stackrel{\text{def}}{=} & 1
\end{array}
$$

Figure 2: Higher-Order Singletons

2.1 The Core Calculus

The core phase distinction calculus contains four syntactic classes: kinds, type constructors (or just "constructors"), types, and terms. As usual, types classify terms and kinds classify constructors. The constructors provide a lambda calculus for constructing types. The syntax of the core calculus appears in Figure 1. We shall consider expressions that differ only in the names of bound variables to be identical, and write capture-avoiding substitution of E for X in E' as $E'[E/X]$.

The kinds include the kind T of all monotypes; the trivial kind 1, containing only the constructor $\star$; dependent products $\Pi\alpha{:}\kappa_1.\kappa$, containing constructor functions from κ_1 to κ_2 where α stands for the argument and may appear free in κ_2; and dependent sums $\Sigma\alpha{:}\kappa_1.\kappa_2$, containing constructor pairs built from κ_1 and κ_2 (respectively) where α stands for the left-hand member and may appear free in κ_2. As usual, if α does not appear free in κ_2, we write $\kappa_1 \to \kappa_2$ for $\Pi\alpha{:}\kappa_1.\kappa_2$ and $\kappa_1 \times \kappa_2$ for $\Sigma\alpha{:}\kappa_1.\kappa_2$.

Finally, for any constructor c having kind T, the *singleton kind* $\mathfrak{S}(c)$ contains monotypes definitionally equal to c. Thus, if c has kind $\mathfrak{S}(c')$, the calculus permits the deduction of the equation $c = c' : T$. Singleton kinds provide a mechanism for expressing type sharing information [9, 13]. Although singleton kinds exist only for monotypes, they may be used in conjunction with dependent kinds to express higher-order sharing information. For instance, if c has kind $\Pi\alpha{:}T.\mathfrak{S}(\mathit{list}(\alpha))$, it follows that $c = \mathit{list} : T \to T$. The definition in Figure 2 generalizes this idea.[1]

[1]Note that higher-order singletons are not defined for kinds containing strictly positive singleton or dependent sum kinds; this would eliminate the useful property that κ is a kind whenever $\mathfrak{S}(c : \kappa)$ is. This does not reduce the expressive power of the construct; any constructor whose kind contains strictly positive singletons or dependent sums can be given a kind (exactly as coarse) without them.

The type constructors are largely standard. The trivial type 1 contains the trivial term $\star$. The types $c_1 \to c_2$ and $c_1 \rightharpoonup c_2$ are the types of total and partial functions from c_1 to c_2 and are discussed in more detail below. The *equi-recursive* constructor $\mu\alpha{:}\kappa.c[\alpha]$ is a fixed point of the equation $\alpha = c[\alpha]$. Thus $\mu\alpha{:}\kappa.c[\alpha]$ is equal to its unrolling $c[\mu\alpha{:}\kappa.c[\alpha]]$. This is in contrast to the somewhat more conventional *iso-recursive* formulation, where conversions between the two must be mediated by explicit operations. In Section 5.3 we discuss how to simplify the type theory to use only iso-recursive constructors.

The final construct, $\mathit{fix}(x{:}\sigma.e)$ at the term level, allows the definition of recursive values. However, we wish to prevent the definition of cyclic data structures such as $\mathit{fix}(x : \texttt{int list}.\texttt{1::}x)$, which cannot be defined in ML. We do this by imposing a value restriction on the bodies of recursive definitions. The calculus contains judgements $\Gamma \vdash e \downarrow \sigma$ asserting that e has type σ and terminates without computational effects. (In the present setting, the only computational effect is nontermination.) With this so-called value restriction in place, the formation rule for recursive values is:

$$\frac{\Gamma[x \uparrow \sigma] \vdash e \downarrow \sigma \quad \Gamma \vdash \sigma \ \mathit{type}}{\Gamma \vdash \mathit{fix}(x{:}\sigma.e) : \sigma} \ (x \notin \mathrm{Dom}(\Gamma))$$

This rule is read: $\mathit{fix}(x{:}\sigma.e)$ has type σ if e terminates with type σ under the assumption that x has type σ *but cannot be taken as valuable.* This rules out the cyclic list proposed above, since $\texttt{1::}x$ is not valuable unless x is valuable. The value restriction implies that all appearances of x must be guarded by (*i.e.*, within the body of) a lambda abstraction; lambda abstractions are always valuable, regardless of the valuability status of their free variables. As in Harper and Stone [12], the collection of valuable expressions is enlarged by including a type for total (pure) functions, such as cons (`::`). The application of a valuable total function to a valuable argument is considered valuable. Total functions are considered to be types, but not type constructors, in order to prevent their erroneous use in conjunction with recursive types.

A similar restriction is made on the formation of equi-recursive type constructors. In order to show that $\mu\alpha{:}\kappa.c$ is well-formed, one must show that c is *contractive* in α [1]. This is in contrast to iso-recursive types, which require no such condition. Informally, contractiveness means that the infinite tree specified by iterating the body c is actually infinite, or, even more informally, that iterating the body "goes somewhere." Thus, $\mu\alpha{:}T.\texttt{int} \times \alpha$ is legal, but $\mu\alpha{:}T.\alpha$ is not. Contractiveness is formalized in the type theory by a judgement $[\alpha \uparrow \kappa] \vdash c \downarrow \kappa$, which is read: c is contractive in kind κ under the assumption that α has kind κ *but cannot be taken as contractive.* The rules for contractiveness of constructors are similar to those for valuability of terms; they ensure that all occurrences of the recursive variable are guarded by a type construction operation (such as $\texttt{int} \times \alpha$ in the above example).

2.2 The Structure Calculus

Atop the core calculus we erect a structure calculus, exactly as in HMM. To review, we add two syntactic classes, one for flat signatures and one for flat structures

$$
\begin{array}{lrcl}
\textit{constructors} & c & ::= & \cdots \mid \mathit{Fst}\, s \\
\textit{terms} & e & ::= & \cdots \mid \mathit{Snd}\, s \\
\textit{signatures} & S & ::= & [\alpha{:}\kappa, \sigma] \\
\textit{modules} & M & ::= & [c, e] \\
\textit{contexts} & \Gamma & ::= & \cdots \mid \Gamma[s : S] \mid \Gamma[s \uparrow S]
\end{array}
$$

Figure 3: The Structure Calculus

$$
\frac{\begin{array}{cc} \Gamma \vdash \kappa\ \mathit{kind} & \Gamma[\alpha : \kappa] \vdash \sigma\ \mathit{type} \\ \Gamma[\alpha \uparrow \kappa] \vdash c \downarrow \kappa & \Gamma[\alpha : \kappa][x \uparrow \sigma] \vdash e \downarrow \sigma \end{array}}{\begin{array}{r} \Gamma \vdash \mathit{fix}(s{:}[\alpha{:}\kappa.\sigma].[c[\mathit{Fst}\, s/\alpha], e[\mathit{Fst}\, s, \mathit{Snd}\, s/\alpha, x]]) = \\ [\alpha = \mu\alpha{:}\kappa.c, \mathit{fix}(x{:}\sigma.e)] : [\alpha{:}\kappa.\sigma] \\ (\alpha, x, s \notin \mathrm{Dom}(\Gamma)) \end{array}}
$$

Figure 4: Phase-Splitting Recursive Modules

(Figure 3). Structures are pairs $[c, e]$ of constructors and terms. The left-hand component is referred to as the compile-time (or, static) component, and the right-hand component is referred to as the run-time (or, dynamic) component. Signatures, which classify structures, have the form $[\alpha{:}\kappa, \sigma]$, where α stands for the compile-time component and may appear free in σ. The structure, $[c, e]$ has kind $[\alpha{:}\kappa, \sigma]$ if c has kind κ and e has type $\sigma[c/\alpha]$. Often we will write $[\alpha = c, e]$ as shorthand for $[c, e[c/\alpha]]$. We also add constructor and term constructs $\mathit{Fst}\, s$ and $\mathit{Snd}\, s$ for extracting the first and second components out of structures named by variables. We will occasionally treat these constructs as variables and allow substitution for them.

The structure calculus shows an explicit *phase distinction* between compile-time and run-time expressions [11, 4]. Static expressions may be separated from dynamic ones, and static ones will never depend on dynamic ones. This ensures that programs may be type-checked without the need to execute any run-time code.

HMM show that higher-order modules can be reduced to the simple structure calculus given here. Therefore we will omit explicit discussion of higher-order modules, without any loss of generality. In this paper, we show how recursive modules may similarly be reduced to the structure calculus given here. In so doing, we will show that despite the apparent intertwining of static and dynamic expressions in recursive modules, that the phase distinction can be preserved, just as HMM showed for higher-order modules.

3 Opaque Recursive Modules

We begin our examination by considering what we call "opaque" recursive modules. These will prove to insufficiently expressive for most applications, but they will serve to illustrate the main ideas and motivate the more complex machinery in the next section.

In the (informal) external language, we write an opaque recursive module definition as:

```
structure rec S :> SIG = struct ... end
```

The structure variable `S` is, of course, permitted to appear free within the structure's body. The signature `SIG` then expresses all the information that is known about `S` in the body or in the subsequent code. (We borrow the "`:>`" symbol from Standard ML 1997 [16] to suggest this opacity.) In particular, the opaque signature obscures the fact that the types in `S` are recursively defined.

This declaration construct corresponds to a module fixed point operation in the internal language, written $\mathit{fix}(s{:}S.M)$. For reasons similar to those in the previous section, we must impose a value restriction on M, resulting in the following typing rule:

$$
\frac{\Gamma[s \uparrow S] \vdash M \downarrow S \quad \Gamma \vdash S\ \mathit{sig}}{\Gamma \vdash \mathit{fix}(s{:}S.M) : S}\ (s \notin \mathrm{Dom}(\Gamma))
$$

Thus, a recursive module is valid if its body (M) is valuable without assuming the recursive variable (s) to be valuable. If a module M is $[c, e]$, then M will be valuable exactly when e is valuable (*i.e.*, constructors are always valuable).

Following HMM, we wish to reduce recursive modules to the primitive structure formalism by defining $\mathit{fix}(s{:}S.M)$ in terms of primitive constructs. We will do this by phase-splitting recursive modules into run-time and compile-time components. Suppose S is the signature $[\alpha{:}\kappa.\sigma]$ and M is the structure $[c(\mathit{Fst}\, s), e(\mathit{Fst}\, s, \mathit{Snd}\, s)]$. Then we can interpret $\mathit{fix}(s{:}S.M)$ by wrapping the static and dynamic components in fixed point expressions:

$$
\mathit{fix}(s{:}S.M) = [\alpha = \mu\alpha{:}\kappa.c(\alpha), \mathit{fix}(x : \sigma.e(\alpha, x))]
$$

This definition is formalized in the type theory by the equational rule in Figure 4. This rule parallels the non-standard equational rules from HMM, and illustrates that recursive modules are *already present* in the underlying calculus. In particular, the formation rule for recursive modules given above follows from the definition and need not appear as a primitive rule.

3.1 Trouble with Opacity

The opaque interpretation of recursive modules is pleasantly simple, but unfortunately, it is not sufficiently expressive to support some desired programming idioms. One common application of recursive modules is to break up mutually recursive data types. As a particularly simple (though somewhat contrived) example, consider an implementation of integer lists as a recursive module that defers recursively to itself for an implementation of the tail:

```
signature LIST =
  sig
    type t
    val nil : t
    val null : t -> bool
    val cons : int * t -> t
    val uncons : t -> int * t
  end

structure rec List :> LIST =
  struct
    datatype t = NIL | CONS of int * List.t
```

```
    val nil = NIL

    fun null NIL = true
      | null (CONS _) = false

    fun cons (n : int, l : t) =
       case l of
         NIL => CONS (n, List.nil)
       | CONS (n' : int, l' : List.t) =>
           CONS (n, List.cons (n', l'))

    fun uncons NIL = raise Fail
      | uncons (CONS (n : int, l : List.t)) =
          if List.null l then
            (n, NIL)
          else
            (n, CONS (List.uncons l))
  end
```

This implementation typechecks properly, and it is observationally equivalent to a conventional implementation. However, intensionally it is very different, because each use of `cons` and `uncons` must traverse the entire list, leading to poor behavior in practice. A more direct implementation is impossible because the opacity of `List.t` precludes any knowledge that `List.t` is the same as `t`.

Some other examples cannot be written in the opaque case at all (but see Section 4.3). For example, consider an implementation of abstract syntax trees using mutually dependent modules for expressions and declarations. These modules interact with each other through the `let` expression, which contains a declaration, and the `val` declaration, which contains an expression. To optimize a common case, the expression code includes a function for `let val` expressions that defers to the declaration code to build a declaration:

```
signature EXPR =
  sig
    type exp
    type dec
    val make_let : dec * exp -> exp
                   (* let DEC in EXP end *)
    val make_let_val :
        identifier * exp * exp -> exp
        (* let val ID = EXP in EXP end *)
    ⋮
  end

signature DECL =
  sig
    type dec
    type exp
    val make_val : identifier * exp -> dec
                   (* val ID = EXP *)
    ⋮
  end

structure rec Expr :> EXPR =
  struct
    datatype exp = LET of Decl.dec * exp | ...
    type dec = Decl.dec

    fun make_let (d : dec, e : exp) =
      LET (d, e)

    fun make_let_val (id : identifier,
                      e1 : exp, e2 : exp) =
      let val d = Decl.make_val (id, e1)
        (* type error! e1 : exp ≠ Decl.exp *)
      in
        LET (d, e2)
      end
    ⋮
  end
and Decl :> DECL =
  struct
    datatype dec =
      VAL of identifier * Expr.exp | ...
    type exp = Expr.exp
    ⋮
  end
```

Unfortunately, this code does not typecheck. The call to `make_val` within `make_let_val` expects an argument with type `Decl.exp`, which, because of the opacity of `Decl`, is not known to be the same type as `exp`, the type of its actual argument `e1`. The type error occurs because the type system cannot tell that `exp` is equal to `Decl.exp`, even though an examination of the recursive definition reveals that it is actually true.

4 Transparent Recursive Modules

The difficulties described in the previous section can be traced to the inability to track sufficient type information in the context of a recursive structure binding. In the abstract syntax example the proposed binding fails to typecheck because within the definition of `Expr` it is not apparent that the type `exp` is equivalent to the type `Decl.exp`, *even though* this equation will be valid once the recursive binding is in force. Similarly, within the definition of `Decl` it is not apparent that the type `dec` is equivalent to the type `Expr.dec`, which will turn out to be true once the binding is in force. Were this equation available while the definitions of `Expr` and `Decl` are being typechecked, the entire declaration would be seen to be valid, and these very equations would hold true afterwards. Similarly, the inefficiency of the suggested implementation of lists may be traced to the failure to identify the types `List.t` and `t` inside the definition of `List`.

What is needed is a means of propagating the type equations that will turn out to be true of the recursively defined structures into the scope of the recursive definition itself. This makes it possible to exploit the recursive definitions of the types involved during typechecking of the dynamic part of the recursively defined modules, leading to a much more flexible and useful notion of recursive module. In effect we are exploiting the phase distinction by solving the *static* recursion equations prior to checking the *dynamic* typing conditions of the module.

How is this additional type sharing information to be propagated? The obvious solution is to add the ap-

propriate equations to the signatures of the modules involved. For example, in the list example we may propagate the required information as follows:

```
structure rec List :>
  sig
    datatype t = NIL | CONS of int * List.t
    ⋮
    val cons : int * t -> t
    val uncons : t -> int * t
  end = ...
```

The boxed phrase highlights the occurrence of the structure variable introduced by the recursive structure binding. Since the signatures of the recursively defined structure variables depend on the structures themselves, we call these signatures *recursively dependent signatures*, or *rds's* for short. In Section 4.2 we shall see how recursively dependent signatures are formalized.

The purpose of a recursively dependent signature is to express the sorts of recursive type equations that are required to recover the ill-formed examples of the preceding section. Let us now revisit those examples to see how rds's are used to resolve the difficulties those examples raise. Using a recursively dependent signature it is possible to give an implementation of lists with constant-time primitive operations as follows:

```
structure rec List :>
  sig
    datatype t = NIL | CONS of int * List.t
    ⋮
    val cons : int * t -> t
    val uncons : t -> int * t
  end =
  struct
    datatype t = NIL | CONS of int * List.t
    ⋮
    fun cons (n : int, l : t) = CONS (n, l)
    fun uncons NIL = raise Fail
      | uncons (CONS (n, l)) = (n, l)
  end
```

The effect of the recursively dependent signature in this example is to ensure that the implementation type of the recursive datatype `List.t` coincides with the implementation type of the type `t` within the body of the definition.

This example also raises a important point about recursive datatypes in the context of a recursive structure binding. We must impose a *structural*, or *transparent*, interpretation of recursive datatypes within the scope of a recursive structure binding, rather than the more familiar *nominal*, or *opaque*, interpretation used in Standard ML. In type-theoretic terms the rds ascribed to `List` is tantamount to a signature that *transparently* defines the type `t` to be the underlying iso-recursive type of the recursive datatype. We note, however, that this interpretation can be limited to the recursive structure binding itself, and need not propagated into the subsequent scope of the binding: the elaborator may "seal" the structure with an opaque signature hiding the implementation type of `List.t` after the binding has been processed.

The abstract syntax example may be handled similarly to the list example, as shown below. The recursively dependent signatures ascribed to `Expr` and `Decl` allow `Decl.make_val`'s second argument to be given the type `Expr.exp`, and under the transparent interpretation of datatypes, `exp` = `Expr.exp` holds within the scope of `exp`. Consequently, the call to `Decl.make_val` is type correct.

```
structure rec Expr :>
  sig
    datatype exp =
      LET of Decl.dec * exp | ...
    val make_let : Decl.dec * exp -> exp
    val make_let_val :
        identifier * exp * exp -> exp
    ⋮
  end =
  struct
    datatype exp = LET of Decl.dec * exp | ...

    fun make_let (d : Decl.dec, e : exp) =
      LET (d, e)

    fun make_let_val (id, e1 : exp, e2 : exp) =
      let val d = Decl.VAL (id, e1)
        (* typechecks, since exp = Expr.exp *)
      in
        LET (d, e2)
      end
    ⋮
  end
and Decl :>
  sig
    datatype dec =
      VAL of identifier * Expr.exp | ...
    val make_val :
        identifier * Expr.exp -> dec
    ⋮
  end = struct ... end
```

In each of these examples, the recursively dependent signatures used were fully transparent, in the sense that every type component was given by an explicit type definition. This was necessary in order to make the given examples typecheck. More generally, fully transparent rds's provide optimal propagation of type information, thereby maximizing the set of programs that can be typechecked. Moreover, we will see in Section 4.2 that in order to phase-split recursively dependent signatures, it is necessary to *require* full transparency of all rds's, and to require a contractiveness condition of them as well. In Section 5.1 we illustrate how the elaborator can ensure that these conditions are satisfied.

4.1 Functors and Separate Compilation

In the preceding abstract syntax example, the mutually recursive modules `Expr` and `Decl` were compiled simultaneously in a single recursive definition. In practice, however, it is important for it to be possible to compile each mutually recursive module separately [7]. In the

absence of separate compilation, the structuring of code as mutually recursive module definitions would often be a largely cosmetic exercise.

One may separately compile mutually recursive modules by rewriting them as closed functors and then gluing those functors together by instantiating them in a recursive structure binding. Each closed functor may then be separately compiled. However, it is instructive to examine the details, as a naive attempt to do so runs afoul of the opacity problem once again. This problem is demonstrated by the following functorized version of the abstract syntax example:

```
structure rec Expr :> sig ... end =
  ExprFun (Expr, Decl)
and Decl :> sig ... end =
  DeclFun (Expr, Decl)

functor ExprFun (structure Expr : EXPR
                 structure Decl : DECL) =
  ... as above ...

functor DeclFun (structure Expr : EXPR
                 structure Decl : DECL) =
  ... as above ...
```

When defined in this manner, `ExprFun` does not typecheck because the arguments `Expr` and `Decl` are given opaque signatures, causing exactly the same problem as in Section 3.1. To make this work, we must use recursively dependent signatures for the functor's parameters:

```
functor ExprFun
        (structure rec Expr :
           sig
             datatype exp =
                LET of Decl.Dec * exp | ...
             :
             :
           end
         and Decl : sig ... end) =
  ... as above ...
```

This example now reveals an important limitation on the degree of separate compilation that is possible. This version of the functor typechecks and may be compiled independently, but in order to make it typecheck, we have been forced to provide a recursively dependent signature for both `Expr` and `Decl`, thereby specifying all the type components of the *other* module `Decl`. Hence we observe that the code may be independently compiled, but in some sense the types may not, since they must be kept consistent among both mutually recursive components.

This is not a frivolous restriction; it is a simple consequence of supplying enough type information to allow each module to typecheck. However, the restriction is stronger than necessary in one regard: we require that rds's be fully transparent, but it is not always necessary to know the definitions of *all* the type components of a mutually recursive module (though it was in the abstract syntax example). Therefore we may gain some additional expressiveness by relaxing the full transparency requirement of rds's. We discuss how to do this in Section 5.1.

4.2 Formalization of Recursively Dependent Signatures

The addition of recursively dependent signatures to the phase distinction calculus is performed in two stages. First, we extend the syntax of signatures with the recursively dependent form, which we write $\rho s.S$, and extend the signature formation and equivalence rules with rules governing this new form. We also extend the module formation rules to include introductory and eliminatory rules for recursively dependent signatures. Second, we show that this enrichment of the structure calculus may be interpreted into the original structure calculus (over the extended core language described in Section 2) by exhibiting an equation between rds's and ordinary signatures.

Informally, the recursively dependent signature $\rho s.S$ contains those modules M that belong to S where s may appear free in S and stands for M. In other words, M belongs to $\rho s.S$ when M belongs to $S[M/s]$. Formally, rds's adhere to the following introductory and eliminatory rules:

$$\frac{\Gamma \vdash M : S[M/s] \quad \Gamma \vdash \rho s.S \ \ sig}{\Gamma \vdash M : \rho s.S} \qquad \frac{\Gamma \vdash M : \rho s.S}{\Gamma \vdash M : S[M/s]}$$

As discussed previously, in the rds $\rho s.S$ we require that the static component of S be fully transparent, that is, that it completely specify the identity of its static component using singleton kinds. Thus, in order for an rds $\rho s.S$ to be well-formed, S must be fully transparent and well-formed under the assumption that s has signature S', where S' is obtained from S by stripping out the singleton kinds specifying the identity of the static component. Formally, rds's have the following formation rule:[2]

$$\frac{\begin{array}{c}\Gamma \vdash S \ \ sig \qquad \Gamma[s \uparrow S] \vdash c \downarrow \kappa \\ \Gamma[s:S] \vdash [\alpha{:}\mathfrak{S}(c:\kappa), \sigma] \ \ sig\end{array}}{\Gamma \vdash \rho s.[\alpha{:}\mathfrak{S}(c:\kappa), \sigma] \ \ sig} \quad \left(\begin{array}{l} s \notin \mathrm{Dom}(\Gamma) \ \ and \\ S \ \ is \\ [\alpha{:}\kappa, \sigma[\alpha/Fst\, s]]\end{array}\right)$$

The third subgoal is the contractiveness condition alluded to previously. This may be seen as part of the full transparency requirement, since if c is noncontractive in s, then it is just retelling s, and provides no useful information. Both the full transparency and contractiveness conditions are necessary in order to interpret rds's into the basic structure calculus, as we will see in a moment.

As with the recursive modules of Section 3, we wish to reduce recursively dependent signatures to primitive constructs of the structure formalism. We do this by wrapping the compile-time component of the rds in a fixed point expression, and by redirecting recursive references in the run-time component:

$$\begin{array}{c}\rho s.[\alpha{:}\mathfrak{S}(c(Fst\, s):\kappa), \sigma(\alpha, Fst\, s)] \\ = \\ [\alpha{:}\mathfrak{S}(\boxed{\mu\beta{:}\kappa.c(\beta)} : \kappa), \sigma(\alpha, \boxed{\alpha})]\end{array}$$

In the second highlighted fragment, recursive references using $Fst\, s$ are redirected to use α. The interesting part is the first highlighted fragment: Suppose $[c', e]$

[2] Recall that we consider $Fst\, s$ to be a variable and allow substitution for it.

$$\frac{\Gamma \vdash \kappa\ \mathit{kind} \quad \Gamma[\beta:\kappa] \vdash \mathfrak{S}(c:\kappa)\ \mathit{kind} \quad \Gamma[\beta \uparrow \kappa] \vdash c \downarrow \kappa \quad \Gamma[\alpha:\kappa] \vdash \sigma[\alpha/\mathit{Fst}\, s]\ \mathit{type}}{\Gamma \vdash \rho s.[\alpha{:}\mathfrak{S}(c[\mathit{Fst}\, s/\beta]:\kappa), \sigma] = [\alpha{:}\mathfrak{S}(\mu\beta{:}\kappa.c\ :\ \kappa), \sigma[\alpha/\mathit{Fst}\, s]]\ \mathit{sig}} \ (\alpha, \beta, s \notin \mathrm{Dom}(\Gamma))$$

Figure 5: Phase-Splitting Recursively Dependent Signatures

is a prospective member of the rds on the left. Since $\mathit{Fst}\,[c', e]$ (so to speak) is c', the rds dictates that c' have kind $\mathfrak{S}(c(c') : \kappa)$, and consequently that $c' = c(c') : \kappa$. Therefore, c' may be taken to be $\mu\beta{:}\kappa.c(\beta)$, as provided by the first highlighted fragment.

This definition makes clear the need for full transparency and contractiveness of rds's. When translated into the structure calculus, an rds specifies its static component to be $\mu\beta{:}\kappa.c(\beta)$, as above. To extract the necessary constructor c, the rds must be fully transparent. Furthermore, $c(\beta)$ must be contractive in β, or else the specified static component is ill-formed.

The definition is formalized in the type theory by the equational rule in Figure 5. As in Section 3, this rule illustrates that recursively dependent signatures are already present in the underlying calculus. In particular, the introductory and eliminatory rules given above follow from the definition and need not appear as primitive rules.

4.3 Opacity Revisited

The phase-splitting rule for recursively dependent signatures reveals an interesting fact: of the two forms of recursive dependency, static-on-static (*i.e.,* types on types) and dynamic-on-static (*i.e.,* terms on types), only static-on-static dependencies are *essentially* recursive (as shown in the first highlighted fragment above). In contrast, dynamic-on-static dependencies were resolved without recursion (in the second fragment), simply by redirecting them from the recursive variable ($\mathit{Fst}\, s$) to the variable standing for the signature's static component (α). Thus, dynamic-on-static dependencies are not truly recursive at all. (In fact, such dependencies would never arise when programming in a fully phase-split style [11].)

Is it possible to program with only dynamic-on-static dependencies and thereby largely dispense with recursively dependent signatures? To some degree, yes. (This is the expressiveness provided by Flatt and Felleisen's "units" [8].) Recall the abstract syntax example as corrected using a recursively dependent signature. That example used static-on-static dependencies in the specifications of `exp` and `dec`, and used dynamic-on-static dependencies in the specifications of `make_let` and `make_val`. The dependency of `dec` on `exp` was used in `make_let_val` when constructing a `dec` (named `d`) from an identifier and an `exp`. However, the need to know `dec`'s specification in terms of `exp` can be eliminated by instead using the function `make_val`, which is specified by a dynamic-on-static dependency. The cost of this is that one must incur a function call whenever going between `exp` and `dec`, and such function calls *cannot* be safely inlined without using static-on-static dependency information.

This is an instance of the well-known fact that one can always program opaquely (*i.e.,* without any type sharing information) if one is willing to incur mediating function calls between types that are actually equal. Tools for supplying type equality information, such as sharing [15], translucent sums [9, 13], and recursively dependent signatures, serve only to improve performance.

5 External Language Issues

5.1 Elaboration of Recursively Dependent Signatures

As discussed previously, the internal language requires that all recursively dependent signatures be fully transparent and contractive in their static component. In an external language, however, this requirement can be burdensome to satisfy. Therefore, we would like to remove that requirement from the external language, and instead have the compiler satisfy the requirement as it elaborate external code into the internal language.

In the case of a recursive module definition, it is easy for the elaborator to supply any omitted type definitions, because it may inspect the actual module to discover the omitted information. However, as we saw in Section 4.1, it is important to allow rds's as signatures for functor arguments. In such cases there is no particular module to inspect for the missing information, so if the elaborator is to rewrite such an rds to be transparent, it must do so without supplying additional information.

Given a prospective rds that fails to be fully transparent, the elaborator may transform it to an isomorphic signature that is permitted by naming any abstract types within the rds and hoisting them out. (A similar device is used by the generative stamps in the Definition of Standard ML [16].) For example, the signature

```
rec S : sig
          type t
          type u = S.u -> t
        end
```

can be made permissible by introducing a type definition for `t` setting it equal to an abstract type that is defined outside the rds. The resulting signature is permissible because the rds, which now lies within an outer signature, is fully transparent:

```
sig
  type t'
  structure rec S :
    sig
      type t = t'
      type u = S.u -> t
    end
end
```

Recall that we also require that recursively dependent signatures be contractive. In most cases, the elab-

orator can transform signatures to satisfy this requirement in the same manner as to satisfy transparency. For example, consider the noncontractive signature:

```
rec S : sig
          type t = S.t
          type u = S.u -> t
        end
```

Note that the signature does not provide any information as to the identity of `t` (except that it is equal to itself). In this case, `t` is essentially abstract, as it was in the previous example, so the signature can be transformed to a permissible one exactly as before.

It should be noted that this technique can fail in the presence of unknown type constructors. Suppose `f` is an unknown constructor with kind $T \to T$ and consider the signature:

```
rec S : sig
          type t = S.t f
          type u = S.u -> t
        end
```

In this signature it is not clear whether or not `t` is abstract. If `f` were the identity, then this signature would reduce to the previous example, and `t` could be hoisted. On the other hand, if `f` were, say, $\lambda\alpha{:}T.\texttt{int} + \alpha$, then `t` would not be abstract and it would incorrect to hoist `t`, and indeed the signature would be properly contractive. The fact that `f` is unknown stymies any attempt to resolve this situation.

5.2 Typechecking

An important problem in a practical implementation is typechecking of recursive modules. Suppose a typechecker is presented with a recursive module definition:

```
structure rec A : ASIG(A) =
  struct ... body ... end
```

In terms of the external language, an appealing typechecking strategy is to check first that the rds $\rho s.\texttt{ASIG}(s)$ is well-formed, and second that the body of the structure definition has signature `ASIG(A)` under the assumption that `A` does. However, it is not immediately clear that this strategy is sound or complete, since the type theory requires (instead of the second condition above) that the body have signature $\rho s.\texttt{ASIG}(s)$ under the assumption that `A` does.

We wish to show that these two typechecking strategies are equivalent. The conditions on the recursive variable `A` are certainly equivalent, using the introductory and eliminatory rules for rds's. To show the conditions on the body to be equivalent, we observe that the signatures `ASIG(A)` and $\rho s.\texttt{ASIG}(s)$ are equal whenever `A` has signature $\rho s.\texttt{ASIG}(s)$.

Suppose $\rho s.S(s)$ is a well-formed rds, and suppose that the variable s is given that signature. Then (for some c, κ and σ):

$$\begin{array}{rcl} S(s) & = & [\alpha{:}\mathbb{S}(c(Fst\,s) : \kappa), \sigma(Fst\,s)] \\ & = & [\alpha{:}\mathbb{S}(c(\mu\beta{:}\kappa.c(\beta)) : \kappa), \sigma(\mu\beta{:}\kappa.c(\beta))] \\ & = & [\alpha{:}\mathbb{S}(\mu\beta{:}\kappa.c(\beta) : \kappa), \sigma(\mu\beta{:}\kappa.c(\beta))] \\ & = & [\alpha{:}\mathbb{S}(\mu\beta{:}\kappa.c(\beta) : \kappa), \sigma(\alpha)] \\ & = & \rho s.S(s) \end{array}$$

$$\frac{\Gamma \vdash c = c'[c/\alpha] : \kappa \quad \Gamma[\alpha \uparrow \kappa] \vdash c' \downarrow \kappa}{\Gamma \vdash c = \mu\alpha{:}\kappa.c' : \kappa} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

Figure 6: Bisimilarity

That $S(s)$ has the form given in the first line follows from the well-formedness of $\rho s.S(s)$; the second line follows since s's signature and phase-splitting dictate that $Fst\,s = \mu\beta{:}\kappa.c(\beta)$; the third and fourth lines follow by equational reasoning using recursive constructors and singleton kinds; and the last follows by the phase-splitting rule.

Constructor equality With or without this external-level typechecking strategy, one significant problem for typechecking remains: the typechecker must be able to determine whether two constructors are equal. This problem is made difficult by singleton kinds and by equi-recursive constructors. At this time neither problem is known to be decidable, although algorithms exist for singleton kinds that work well in practice.

For equi-recursive constructors, Amadio and Cardelli [1] give an algorithm for checking equality at kind type, but this algorithm does not extend to higher kinds. Some recent work suggests that the problem may be decidable at higher kinds as well: Solomon [22] showed in 1978 that type equality with a somewhat similar notion of recursive type could be reduced to equivalence of deterministic pushdown automata, which was recently shown decidable by Sénizergues [17], though not by a very practical algorithm.

These two problems remain the main outstanding issues confronting a practical implementation of recursive modules.

5.3 Equi- versus Iso-recursive Constructors

Given the difficulty of typechecking in the presence of equi-recursive constructors, a natural question is whether the reliance on equi-recursive constructors is essential for supporting recursive modules. (For example, Duggan and Sourelis's formalism does not rely on this form of recursive types.) We conjecture that it is not essential, based on the following observations. Under the standard type-theoretic interpretation of ML (for example, Harper and Mitchell [10]), the implementation of a recursive datatype is an iso-recursive type. (We will write iso-recursive types as $\mu_{\simeq}\alpha.c$.) If we restrict recursive modules to datatypes (as in Duggan and Sourelis' formalism), and adopt the "transparent" interpretation outlined in Section 4, then equi-recursive types are completely eliminable by the translation into the underlying structure calculus, provided that we adopt *Shao's equation* for iso-recursive types:

$$\mu_{\simeq}\alpha.c(\alpha) = \mu_{\simeq}\alpha.c(\mu_{\simeq}\alpha.c(\alpha))$$

This equation was introduced by Shao [19] in his FLINT formalism in order to support the compilation of Standard ML. It is also discussed by Crary *et al.* [5], who argue that this equation is already essential for efficient

$$\dfrac{\dfrac{}{\Gamma \vdash \mu_{\cong}\beta.c(\boxed{\beta},\beta) = \mu_{\cong}\beta.c(\boxed{\mu_{\cong}\beta.c(\beta,\beta)},\beta)}\ \text{Shao's equation} \quad \dfrac{}{\Gamma[\alpha \uparrow T] \vdash \mu_{\cong}\beta.c(\alpha,\beta)\downarrow}\ \mu_{\cong}\text{ contractiveness}}{\Gamma \vdash \mu_{\cong}\beta.c(\beta,\beta) = \mu\alpha{:}T.\mu_{\cong}\beta.c(\alpha,\beta) : T}\ \text{(bisimilarity)}$$

Figure 7: Elimination of Equi-Recursive Types

compilation of Standard ML's opaque datatypes, even in the absence of recursive modules.

The relevance of Shao's equation to the elimination of equi-recursive types is based on the following observation. After translation into the pure structure calculus, datatypes in the body of a recursive module definition have implementation types of the form

$$\mu\alpha.\mu_{\cong}\beta.c(\alpha,\beta)$$

for some constructor c, where α results from recursion via recursive modules and β results from ordinary datatype recursion. By invoking a bisimilarity rule for equi-recursive types (Figure 6) and applying Shao's equation, as shown in Figure 7, we may prove that this type is equivalent to the type

$$\mu_{\cong}\beta.c(\beta,\beta)$$

which is a purely iso-recursive type.

This observation sheds light on the nature of Duggan and Sourelis's restriction on the recursively defined type components of a mixin module to datatypes, which are implicitly iso-recursive. Strictly speaking, this restriction is not necessary, but if it were to be adopted, it would, by the observation above, allow the elimination of equi-recursive types from the internal language of a type-based compiler for ML.

5.4 Transparent Signature Ascription

Transparent signature ascription raises significant issues for the design of an external language with recursive modules. In Standard ML it is common practice to write structure declarations in the form

```
structure S : SIG =
  struct type t=int ... end
```

where `SIG` specifies, but does not define, a type component `t`. The elaborator processes the right-hand side, extracting the binding of `t`, and implicitly propagating it to the signature `SIG`. In effect, the binding is made opaque, but with an augmented signature containing additional defining equations for types, as follows:

```
structure S :> SIG where type t=int =
  struct type t=int ... end
```

It is natural to consider whether transparent ascription may be extended to recursive module bindings. The difficulty is that the right-hand side of the binding can involve references to `S` itself. As we have seen, the code on the right cannot be type-checked in the absence of equations for the type components of the module. One approach to this problem is to adopt a two-pass elaboration process. In the first pass the right-hand side is processed to extract the type information that is to be added to the ascribed signature. In the second pass, the augmented, fully-transparent signature can be used to elaborate the run-time parts of the right-hand to complete elaboration.

A related question is whether an ascribed signature may be omitted entirely in the recursive case. Here we face the difficulty that it is not even apparent what are the type components of the recursively-defined structure, let alone what are their definitions. To recover this information would seem to require the kind of pre-elaboration used in the SML/NJ Compilation Manager [3], whereby the defined components of a module are determined before processing begins.

6 Conclusions

Purely hierarchical module systems, such as the Standard ML module system, may be criticized on the grounds that they lack adequate support for cyclic dependencies among components. Several authors (including Duggan and Sourelis [6, 7] and Flatt and Felleisen [8]) have proposed module systems that better support such cyclic dependencies among units. With at least two different proposals for recursive modules in hand, it is natural to ask "what is a recursive module?" We provide an answer to this question in the form of a type-theoretic analysis of recursive modules based on the "phase distinction" calculus of higher-order modules [9].

We propose an extension of the phase distinction calculus with a new form of recursive module and a new form of signature, called a recursively dependent signature. Following the paradigm of the phase distinction interpretation of higher-order modules, we demonstrate the sensibility of this extension by giving an interpretation of it into a pure calculus of structures (without explicit recursive module constructs). This interpretation demonstrates that in a precise sense, recursive modules are already present in the pure structure calculus.

To make these ideas practical more work remains to be done. It is important to demonstrate that typechecking is decidable in this framework. The central issue is decidability of type equality in the presence of singleton kinds and equi-recursive constructors of higher kind. It is also important to consider a dynamic semantics for the extended language and to demonstrate the soundness of the type system for this dynamic semantics.

Acknowledgements

We would like to thank Matthias Felleisen, Matthew Flatt, Xavier Leroy, David MacQueen, Zhong Shao, and the anonymous referees for many helpful comments.

References

[1] Roberto Amadio and Luca Cardelli. Subtyping recursive types. *ACM TOPLAS*, 15(4):575–631, 1993.

[2] Davide Ancona and Elena Zucca. An algebra of mixin modules. In F. Parisi-Presicce, editor, *WADT '97 12th Workshop on Algebraic Development Techniques – Selected Papers*, volume 1376 of *Lecture Notes in Computer Science*, pages 92–106, Berlin, 1997. Springer Verlag.

[3] Matthias Blume. *Hierarchical Modularity and Intermodule Optimization*. PhD thesis, Princeton University, Department of Computer Science, Princeton, New Jersey, November 1997.

[4] Luca Cardelli. Phase distinctions in type theory. Unpublished manuscript.

[5] Karl Crary, Robert Harper, Perry Cheng, Leaf Petersen, and Chris Stone. Transparent and opaque interpretations of datatypes. Technical Report CMU-CS-98-177, Carnegie Mellon University, School of Computer Science, November 1998.

[6] Dominic Duggan and Constantinos Sourelis. Mixin modules. In *1996 ACM SIGPLAN International Conference on Functional Programming*, pages 262–273, Philadelphia, Pennsylvania, June 1996.

[7] Dominic Duggan and Constantinos Sourelis. Parameterized modules, recursive modules, and mixin modules. In *1998 ACM SIGPLAN Workshop on ML*, pages 87–96, Baltimore, Maryland, September 1998.

[8] Matthew Flatt and Matthias Felleisen. Units: Cool modules for HOT languages. In *1998 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 236–248, Montreal, Canada, June 1998.

[9] Robert Harper and Mark Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *Twenty-First ACM Symposium on Principles of Programming Languages*, pages 123–137, Portland, Oregon, January 1994.

[10] Robert Harper and John C. Mitchell. On the type structure of Standard ML. *ACM Transactions on Programming Languages and Systems*, 15(2):211–252, April 1993.

[11] Robert Harper, John C. Mitchell, and Eugenio Moggi. Higher-order modules and the phase distinction. In *Seventeenth ACM Symposium on Principles of Programming Languages*, San Francisco, California, January 1990.

[12] Robert Harper and Chris Stone. A type-theoretic interpretation of Standard ML. In *Proof, Language and Interaction: Essays in Honour of Robin Milner*. The MIT Press, 1998. To appear.

[13] Xavier Leroy. Manifest types, modules, and separate compilation. In *Proceedings of the Twenty-first Annual ACM Symposium on Principles of Programming Languages*, pages 109–122, Portland, Oregon, January 1994.

[14] Xavier Leroy. The Objective Caml system: Documentation and user's guide. Available at `http://pauillac.inria.fr/ocaml/htmlman/`., 1996.

[15] David MacQueen. Modules for Standard ML. In *1984 ACM Conference on Lisp and Functional Programming*, pages 198–207, Austin, Texas, August 1984.

[16] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.

[17] Géraud Sénizergues. The equivalence problem for deterministic pushdown automata is decidable. In *Twenty-Fourth International Colloquium on Automata, Languages, and Programming*, volume 1256 of *Lecture Notes in Computer Science*, pages 671–681, Bologna, Italy, July 1997. Springer-Verlag.

[18] Zhong Shao. An overview of the FLINT/ML compiler. In *Proceedings of the 1997 ACM SIGPLAN Workshop on Types in Compilation*, Kyoto, Japan, June 1997.

[19] Zhong Shao. Equality of recursive types. (Private communication), September 1998.

[20] Zhong Shao. Typed cross-module compilation. In *1998 ACM SIGPLAN International Conference on Functional Programming*, pages 141–152, Baltimore, Maryland, September 1998.

[21] Emin Gün Sirer, Marc E. Fiucynski, Przemysław Pardyak, and Brian N. Bershad. Safe dynamic linking in an extensible operating system. In *Workshop on Compiler Support for System Software*, Tucson, Arizona, February 1996.

[22] Marvin Solomon. Type definitions with parameters (extended abstract). In *Fifth ACM Symposium on Principles of Programming Languages*, pages 31–38, Tucson, Arizona, January 1978.

[23] David Tarditi, Greg Morrisett, Perry Cheng, Chris Stone, Robert Harper, and Peter Lee. TIL: A type-directed optimizing compiler for ML. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 181–192, Philadelphia, Pennsylvania, May 1996.

A Type Theory

A.1 Core Calculus

$\boxed{\Gamma \vdash \kappa \ kind}$

$$\frac{}{\Gamma \vdash T \ kind} \qquad \frac{}{\Gamma \vdash 1 \ kind} \qquad \frac{\Gamma \vdash c : Type}{\Gamma \vdash \mathfrak{S}(c) \ kind}$$

$$\frac{\Gamma \vdash \kappa_1 \ kind \quad \Gamma[\alpha : \kappa_1] \vdash \kappa_2 \ kind}{\Gamma \vdash \Pi\alpha{:}\kappa_1.\kappa_2 \ kind} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash \kappa_1 \ kind \quad \Gamma[\alpha : \kappa_1] \vdash \kappa_2 \ kind}{\Gamma \vdash \Sigma\alpha{:}\kappa_1.\kappa_2 \ kind} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash \kappa_1 = \kappa_2 \ kind}$

$$\frac{}{\Gamma \vdash T = T \ kind} \qquad \frac{}{\Gamma \vdash 1 = 1 \ kind}$$

$$\frac{\Gamma \vdash c_1 = c_2 : Type}{\Gamma \vdash \mathfrak{S}(c_1) = \mathfrak{S}(c_2) \ kind}$$

$$\frac{\Gamma \vdash \kappa_1 = \kappa_1' \ kind \quad \Gamma[\alpha : \kappa_1] \vdash \kappa_2 = \kappa_2' \ kind}{\Gamma \vdash \Pi\alpha{:}\kappa_1.\kappa_2 = \Pi\alpha{:}\kappa_1'.\kappa_2' \ kind} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash \kappa_1 = \kappa_1' \ kind \quad \Gamma[\alpha : \kappa_1] \vdash \kappa_2 = \kappa_2' \ kind}{\Gamma \vdash \Sigma\alpha{:}\kappa_1.\kappa_2 = \Sigma\alpha{:}\kappa_1'.\kappa_2' \ kind} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash \kappa_1 \leq \kappa_2 \ \mathit{kind}}$

$$\frac{}{\Gamma \vdash T \leq T \ \mathit{kind}} \qquad \frac{}{\Gamma \vdash 1 \leq 1 \ \mathit{kind}}$$

$$\frac{\Gamma \vdash c_1 = c_2 : \mathit{Type}}{\Gamma \vdash \mathfrak{S}(c_1) \leq \mathfrak{S}(c_2) \ \mathit{kind}} \qquad \frac{\Gamma \vdash c : T}{\Gamma \vdash \mathfrak{S}(c) \leq T \ \mathit{kind}}$$

$$\frac{\begin{array}{c}\Gamma \vdash \kappa_1' \leq \kappa_1 \ \mathit{kind} \\ \Gamma[\alpha : \kappa_1'] \vdash \kappa_2 \leq \kappa_2' \ \mathit{kind} \\ \Gamma[\alpha : \kappa_1] \vdash \kappa_2 \ \mathit{kind}\end{array}}{\Gamma \vdash \Pi\alpha{:}\kappa_1.\kappa_2 \leq \Pi\alpha{:}\kappa_1'.\kappa_2' \ \mathit{kind}} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\begin{array}{c}\Gamma \vdash \kappa_1 \leq \kappa_1' \ \mathit{kind} \\ \Gamma[\alpha : \kappa_1] \vdash \kappa_2 \leq \kappa_2' \ \mathit{kind} \\ \Gamma[\alpha : \kappa_1'] \vdash \kappa_2' \ \mathit{kind}\end{array}}{\Gamma \vdash \Sigma\alpha{:}\kappa_1.\kappa_2 \leq \Sigma\alpha{:}\kappa_1'.\kappa_2' \ \mathit{kind}} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash c : \kappa}$

$$\frac{}{\Gamma \vdash \alpha : \kappa} \ (\alpha : \kappa \in \Gamma) \qquad \frac{}{\Gamma \vdash \alpha : \kappa} \ (\alpha \uparrow \kappa \in \Gamma) \qquad \frac{}{\Gamma \vdash \star : 1}$$

$$\frac{\Gamma \vdash \kappa_1 \ \mathit{kind} \quad \Gamma[\alpha : \kappa_1] \vdash c : \kappa_2}{\Gamma \vdash \lambda\alpha{:}\kappa_1.c : \Pi\alpha{:}\kappa_1.\kappa_2} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c_1 : \Pi\alpha{:}\kappa_1.\kappa_2 \quad \Gamma \vdash c_2 : \kappa_1}{\Gamma \vdash c_1 c_2 : \kappa_2[c_2/\alpha]}$$

$$\frac{\begin{array}{c}\Gamma \vdash c_1 : \kappa_1 \\ \Gamma \vdash c_2 : \kappa_2[c_1/\alpha] \\ \Gamma[\alpha : \kappa_1] \vdash \kappa_2 \ \mathit{kind}\end{array}}{\Gamma \vdash \langle c_1, c_2 \rangle : \Sigma\alpha{:}\kappa_1.\kappa_2} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c : \Sigma\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash \pi_1(c) : \kappa_1} \qquad \frac{\Gamma \vdash c : \Sigma\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash \pi_2(c) : \kappa_2[\pi_1(c)/\alpha]}$$

$$\frac{\Gamma \vdash c_1 : T \quad \Gamma \vdash c_2 : T}{\Gamma \vdash c_1 \rightharpoonup c_2 : T} \qquad \frac{\Gamma \vdash c_1 : T \quad \Gamma \vdash c_2 : T}{\Gamma \vdash c_1 \times c_2 : T}$$

$$\frac{}{\Gamma \vdash 1 : T} \qquad \frac{\Gamma \vdash \kappa \ \mathit{kind} \quad \Gamma[\alpha \uparrow \kappa] \vdash c \downarrow \kappa}{\Gamma \vdash \mu\alpha{:}\kappa.c : \kappa} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c : \kappa' \quad \Gamma \vdash \kappa' \leq \kappa \ \mathit{kind}}{\Gamma \vdash c : \kappa}$$

$$\frac{\Gamma \vdash \lambda\alpha{:}\kappa_1.c\,\alpha : \Pi\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash c : \Pi\alpha{:}\kappa_1.\kappa_2} \ (\alpha \text{ not free in } c)$$

$$\frac{\Gamma \vdash \langle \pi_1(c), \pi_2(c) \rangle : \Sigma\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash c : \Sigma\alpha{:}\kappa_1.\kappa_2}$$

$\boxed{\Gamma \vdash c \downarrow \kappa}$

$$\overline{\Gamma[B]} \stackrel{\text{def}}{=} \begin{cases} \overline{\Gamma}[\alpha : \kappa] & \text{if } [B] \text{ is } [\alpha \uparrow \kappa] \\ \overline{\Gamma}[B] & \text{otherwise} \end{cases}$$

$$\frac{}{\Gamma \vdash \alpha \downarrow \kappa} \ (\alpha : \kappa \in \Gamma) \qquad \frac{}{\Gamma \vdash \star \downarrow 1}$$

$$\frac{\Gamma \vdash \kappa_1 \ \mathit{kind} \quad \Gamma[\alpha : \kappa_1] \vdash c \downarrow \kappa_2}{\Gamma \vdash \lambda\alpha{:}\kappa_1.c \downarrow \Pi\alpha{:}\kappa_1.\kappa_2} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c_1 \downarrow \Pi\alpha{:}\kappa_1.\kappa_2 \quad \Gamma \vdash c_2 \downarrow \kappa_1}{\Gamma \vdash c_1 c_2 \downarrow \kappa_2[c_1/\alpha]}$$

$$\frac{\begin{array}{c}\Gamma \vdash c_1 \downarrow \kappa_1 \\ \Gamma \vdash c_2 \downarrow \kappa_2[c_1/\alpha] \\ \Gamma[\alpha : \kappa_1] \vdash \kappa_2 \ \mathit{kind}\end{array}}{\Gamma \vdash \langle c_1, c_2 \rangle \downarrow \Sigma\alpha{:}\kappa_1.\kappa_2} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c \downarrow \Sigma\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash \pi_1(c) \downarrow \kappa_1} \qquad \frac{\Gamma \vdash c \downarrow \Sigma\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash \pi_2(c) \downarrow \kappa_2[\pi_1(c)/\alpha]}$$

$$\frac{\overline{\Gamma} \vdash c_1 : T \quad \overline{\Gamma} \vdash c_2 : T}{\Gamma \vdash c_1 \rightharpoonup c_2 \downarrow T} \qquad \frac{\overline{\Gamma} \vdash c_1 : T \quad \overline{\Gamma} \vdash c_2 : T}{\Gamma \vdash c_1 \times c_2 \downarrow T}$$

$$\frac{}{\Gamma \vdash 1 \downarrow T} \qquad \frac{\Gamma \vdash \kappa \ \mathit{kind} \quad \Gamma[\alpha \uparrow \kappa] \vdash c \downarrow \kappa}{\Gamma \vdash \mu\alpha{:}\kappa.c \downarrow \kappa} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c \downarrow \kappa' \quad \Gamma \vdash \kappa' \leq \kappa \ \mathit{kind}}{\Gamma \vdash c \downarrow \kappa} \qquad \frac{\Gamma \vdash c = c' : \kappa \quad \Gamma \vdash c' \downarrow \kappa}{\Gamma \vdash c \downarrow \kappa}$$

$\boxed{\Gamma \vdash c_1 = c_2 : \kappa}$

$$\frac{\Gamma \vdash c : \kappa}{\Gamma \vdash c = c : \kappa} \qquad \frac{\Gamma \vdash c_2 = c_1 : \kappa}{\Gamma \vdash c_1 = c_2 : \kappa}$$

$$\frac{\Gamma \vdash c_1 = c_2 : \kappa \quad \Gamma \vdash c_2 = c_3 : \kappa}{\Gamma \vdash c_1 = c_3 : \kappa}$$

$$\frac{\Gamma \vdash \kappa_1 = \kappa_1' \ \mathit{kind} \quad \Gamma[\alpha : \kappa_1] \vdash c = c' : \kappa_2}{\Gamma \vdash \lambda\alpha{:}\kappa_1.c = \lambda\alpha{:}\kappa_1'.c' : \Pi\alpha{:}\kappa_1.\kappa_2} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c_1 = c_1' : \Pi\alpha{:}\kappa_1.\kappa_2 \quad \Gamma \vdash c_2 = c_2' : \kappa_1}{\Gamma \vdash c_1 c_2 = c_1' c_2' : \kappa_2[c_2/\alpha]}$$

$$\frac{\begin{array}{c}\Gamma \vdash c_1 = c_1' : \kappa_1 \\ \Gamma \vdash c_2 = c_2' : \kappa_2[c_1/\alpha] \\ \Gamma[\alpha : \kappa_1] \vdash \kappa_2 \ \mathit{kind}\end{array}}{\Gamma \vdash \langle c_1, c_2 \rangle = \langle c_1', c_2' \rangle : \Sigma\alpha{:}\kappa_1.\kappa_2} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c = c' : \Sigma\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash \pi_1(c) = \pi_1(c') : \kappa_1}$$

$$\frac{\Gamma \vdash c = c' : \Sigma\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash \pi_2(c) = \pi_2(c') : \kappa_2[\pi_1(c)/\alpha]}$$

$$\frac{\Gamma \vdash c_1 = c_1' : T \quad \Gamma \vdash c_2 = c_2' : T}{\Gamma \vdash c_1 \rightharpoonup c_2 = c_1' \rightharpoonup c_2' : T}$$

$$\frac{\Gamma \vdash c_1 = c_1' : T \quad \Gamma \vdash c_2 = c_2' : T}{\Gamma \vdash c_1 \times c_2 = c_1' \times c_2' : T}$$

$$\frac{\Gamma \vdash \kappa = \kappa' \quad \Gamma[\alpha:\kappa] \vdash c = c' : \kappa \quad \Gamma[\alpha \uparrow \kappa] \vdash c \downarrow \kappa \quad \Gamma[\alpha \uparrow \kappa] \vdash c' \downarrow \kappa}{\Gamma \vdash \mu\alpha{:}\kappa.c = \mu\alpha{:}\kappa'.c' : \kappa} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c_1 = c_2 : \kappa' \quad \Gamma \vdash \kappa' \leq \kappa \ \ kind}{\Gamma \vdash c_1 = c_2 : \kappa}$$

$$\frac{\Gamma \vdash c : \mathfrak{S}(c')}{\Gamma \vdash c = c' : T} \qquad \frac{\Gamma \vdash c : 1}{\Gamma \vdash c = \star : 1}$$

$$\frac{\Gamma \vdash c_1 : \kappa_1 \quad \Gamma[\alpha:\kappa_1] \vdash c_2 : \kappa_2}{\Gamma \vdash (\lambda\alpha{:}\kappa_1.c_2)c_1 = c_2[c_1/\alpha] : \kappa_2[c_1/\alpha]} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c : \Pi\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash (\lambda\alpha{:}\kappa_1.c\,\alpha) = c : \Pi\alpha{:}\kappa_1.\kappa_2} \ (\alpha \text{ not free in } c)$$

$$\frac{\Gamma \vdash c_1 : \kappa_1 \quad \Gamma \vdash c_2 : \kappa_2}{\Gamma \vdash \pi_i(\langle c_1, c_2\rangle) = c_i : \kappa_i} \ (i = 1,2)$$

$$\frac{\Gamma \vdash c : \Sigma\alpha{:}\kappa_1.\kappa_2}{\Gamma \vdash \langle \pi_1(c), \pi_2(c)\rangle = c : \Sigma\alpha{:}\kappa_1.\kappa_2}$$

$$\frac{\Gamma \vdash \kappa \ \ kind \quad \Gamma[\alpha \uparrow \kappa] \vdash c \downarrow \kappa}{\Gamma \vdash \mu\alpha{:}\kappa.c = c[(\mu\alpha{:}\kappa.c)/\alpha] : \kappa} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c = c'[c/\alpha] : \kappa \quad \Gamma[\alpha \uparrow \kappa] \vdash c' \downarrow \kappa}{\Gamma \vdash c = \mu\alpha{:}\kappa.c' : \kappa} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash \sigma \ \ type}$

$$\frac{\Gamma \vdash c : T}{\Gamma \vdash c \ \ type} \qquad \frac{\Gamma \vdash \sigma_1 \ \ type \quad \Gamma \vdash \sigma_2 \ \ type}{\Gamma \vdash \sigma_1 \to \sigma_2 \ \ type}$$

$$\frac{\Gamma \vdash \sigma_1 \ \ type \quad \Gamma \vdash \sigma_2 \ \ type}{\Gamma \vdash \sigma_1 \rightharpoonup \sigma_2 \ \ type} \qquad \frac{\Gamma \vdash \sigma_1 \ \ type \quad \Gamma \vdash \sigma_2 \ \ type}{\Gamma \vdash \sigma_1 \times \sigma_2 \ \ type}$$

$$\frac{\Gamma \vdash \kappa \ \ kind \quad \Gamma[\alpha:\kappa] \vdash \sigma \ \ type}{\Gamma \vdash \forall\alpha{:}\kappa.\sigma \ \ type} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash \sigma_1 = \sigma_2 \ \ type}$

$$\frac{\Gamma \vdash \sigma \ \ type}{\Gamma \vdash \sigma = \sigma \ \ type} \qquad \frac{\Gamma \vdash \sigma_2 = \sigma_1 \ \ type}{\Gamma \vdash \sigma_1 = \sigma_2 \ \ type}$$

$$\frac{\Gamma \vdash \sigma_1 = \sigma_2 \ \ type \quad \Gamma \vdash \sigma_2 = \sigma_3 \ \ type}{\Gamma \vdash \sigma_1 = \sigma_3 \ \ type}$$

$$\frac{\Gamma \vdash c = c' : T}{\Gamma \vdash c = c' \ \ type}$$

$$\frac{\Gamma \vdash \sigma_1 = \sigma_1' \ \ type \quad \Gamma \vdash \sigma_2 = \sigma_2' \ \ type}{\Gamma \vdash \sigma_1 \to \sigma_2 = \sigma_1' \to \sigma_2' \ \ type}$$

$$\frac{\Gamma \vdash \sigma_1 = \sigma_1' \ \ type \quad \Gamma \vdash \sigma_2 = \sigma_2' \ \ type}{\Gamma \vdash \sigma_1 \rightharpoonup \sigma_2 = \sigma_1' \rightharpoonup \sigma_2' \ \ type}$$

$$\frac{\Gamma \vdash \sigma_1 = \sigma_1' \ \ type \quad \Gamma \vdash \sigma_2 = \sigma_2' \ \ type}{\Gamma \vdash \sigma_1 \times \sigma_2 = \sigma_1' \times \sigma_2' \ \ type}$$

$$\frac{\Gamma \vdash \kappa = \kappa' \quad \Gamma[\alpha:\kappa] \vdash \sigma = \sigma' \ \ type}{\Gamma \vdash \forall\alpha{:}\kappa.\sigma = \forall\alpha{:}\kappa'.\sigma' \ \ type} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash e : \sigma}$

$$\frac{}{\Gamma \vdash x : \sigma} \ (x : \sigma \in \Gamma) \qquad \frac{}{\Gamma \vdash x : \sigma} \ (x \uparrow \sigma \in \Gamma) \qquad \frac{}{\Gamma \vdash \star : 1}$$

$$\frac{\Gamma \vdash \sigma \ \ type \quad \Gamma[x:\sigma] \vdash e \downarrow \sigma}{\Gamma \vdash \lambda x{:}\sigma.e : \sigma \to \sigma'} \ (x \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash \sigma \ \ type \quad \Gamma[x:\sigma] \vdash e : \sigma}{\Gamma \vdash \lambda x{:}\sigma.e : \sigma \rightharpoonup \sigma'} \ (x \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash e_1 : \sigma \to \sigma' \quad \Gamma \vdash e_2 : \sigma}{\Gamma \vdash e_1 e_2 : \sigma'}$$

$$\frac{\Gamma \vdash e_1 : \sigma \rightharpoonup \sigma' \quad \Gamma \vdash e_2 : \sigma}{\Gamma \vdash e_1 e_2 : \sigma'}$$

$$\frac{\Gamma \vdash e_1 : \sigma_1 \quad \Gamma \vdash e_2 : \sigma_2}{\Gamma \vdash \langle e_1, e_2\rangle : \sigma_1 \times \sigma_2} \qquad \frac{\Gamma \vdash e : \sigma_1 \times \sigma_2}{\Gamma \vdash \pi_i(e) : \sigma_i} \ (i = 1,2)$$

$$\frac{\Gamma \vdash \kappa \ \ kind \quad \Gamma[\alpha:\kappa] \vdash e \downarrow \sigma}{\Gamma \vdash \Lambda\alpha{:}\kappa.e : \forall\alpha{:}\kappa.\sigma} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash e : \forall\alpha{:}\kappa.\sigma \quad \Gamma \vdash c : \kappa}{\Gamma \vdash e[c] : \sigma[c/\alpha]}$$

$$\frac{\Gamma \vdash \sigma \ \ type \quad \Gamma[x \uparrow \sigma] \vdash e \downarrow \sigma}{\Gamma \vdash \mathit{fix}(x{:}\sigma.e) : \sigma} \ (x \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash e : \sigma' \quad \Gamma \vdash \sigma = \sigma' \ \ type}{\Gamma \vdash e : \sigma}$$

$\boxed{\Gamma \vdash e \downarrow \sigma}$

$$\frac{}{\Gamma \vdash x \downarrow \sigma} \ (x : \sigma \in \Gamma) \qquad \frac{}{\Gamma \vdash \star \downarrow 1}$$

$$\frac{\Gamma \vdash \sigma \ \ type \quad \Gamma[x:\sigma] \vdash e \downarrow \sigma}{\Gamma \vdash \lambda x{:}\sigma.e \downarrow \sigma \to \sigma'} \ (x \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash \sigma \ \ type \quad \Gamma[x:\sigma] \vdash e : \sigma}{\Gamma \vdash \lambda x{:}\sigma.e \downarrow \sigma \rightharpoonup \sigma'} \ (x \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash e_1 \downarrow \sigma \to \sigma' \quad \Gamma \vdash e_2 \downarrow \sigma}{\Gamma \vdash e_1 e_2 \downarrow \sigma'}$$

$$\frac{\Gamma \vdash e_1 \downarrow \sigma_1 \quad \Gamma \vdash e_2 \downarrow \sigma_2}{\Gamma \vdash \langle e_1, e_2\rangle \downarrow \sigma_1 \times \sigma_2} \qquad \frac{\Gamma \vdash e \downarrow \sigma_1 \times \sigma_2}{\Gamma \vdash \pi_i(e) \downarrow \sigma_i} \ (i = 1,2)$$

$$\frac{\Gamma \vdash \kappa \ \ kind \quad \Gamma[\alpha:\kappa] \vdash e \downarrow \sigma}{\Gamma \vdash \Lambda\alpha{:}\kappa.e \downarrow \forall\alpha{:}\kappa.\sigma} \ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash e \downarrow \forall\alpha{:}\kappa.\sigma \quad \Gamma \vdash c : \kappa}{\Gamma \vdash e[c] \downarrow \sigma[c/\alpha]}$$

$$\frac{\Gamma \vdash \sigma \ \ type \quad \Gamma[x \uparrow \sigma] \vdash e \downarrow \sigma}{\Gamma \vdash \mathit{fix}(x{:}\sigma.e) \downarrow \sigma} \ (x \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash e \downarrow \sigma' \quad \Gamma \vdash \sigma = \sigma' \ \ type}{\Gamma \vdash e \downarrow \sigma}$$

A.2 Structure Calculus

$\boxed{\Gamma \vdash c : \kappa}$

$$\frac{\Gamma \vdash S = [\alpha{:}\kappa, \sigma]\ \ sig}{\Gamma \vdash Fst\, s : \kappa}\ (s : S \in \Gamma)$$

$$\frac{\Gamma \vdash S = [\alpha{:}\kappa, \sigma]\ \ sig}{\Gamma \vdash Fst\, s : \kappa}\ (s \uparrow S \in \Gamma)$$

$\boxed{\Gamma \vdash c \downarrow \kappa}$

$$\frac{\Gamma \vdash S = [\alpha{:}\kappa, \sigma]\ \ sig}{\Gamma \vdash Fst\, s \downarrow \kappa}\ (s : S \in \Gamma)$$

$\boxed{\Gamma \vdash e : \sigma}$

$$\frac{\Gamma \vdash S = [\alpha{:}\kappa, \sigma]\ \ sig}{\Gamma \vdash Snd\, s : \sigma[Fst\, s/\alpha]}\ (s : S \in \Gamma)$$

$$\frac{\Gamma \vdash S = [\alpha{:}\kappa, \sigma]\ \ sig}{\Gamma \vdash Snd\, s : \sigma[Fst\, s/\alpha]}\ (s \uparrow S \in \Gamma)$$

$\boxed{\Gamma \vdash e \downarrow \sigma}$

$$\frac{\Gamma \vdash S = [\alpha{:}\kappa, \sigma]\ \ sig}{\Gamma \vdash Snd\, s \downarrow \sigma[Fst\, s/\alpha]}\ (s : S \in \Gamma)$$

$\boxed{\Gamma \vdash S\ \ sig}$

$$\frac{\Gamma \vdash \kappa\ \ kind \quad \Gamma[\alpha : \kappa] \vdash \sigma\ \ type}{\Gamma \vdash [\alpha{:}\kappa, \sigma]\ \ sig}\ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash S_1 = S_2\ \ sig}$

$$\frac{\Gamma \vdash S\ \ sig}{\Gamma \vdash S = S\ \ sig} \qquad \frac{\Gamma \vdash S_2 = S_1\ \ sig}{\Gamma \vdash S_1 = S_2\ \ sig}$$

$$\frac{\Gamma \vdash S_1 = S_2\ \ sig \quad \Gamma \vdash S_2 = S_3\ \ sig}{\Gamma \vdash S_1 = S_3\ \ sig}$$

$$\frac{\Gamma \vdash \kappa = \kappa'\ \ kind \quad \Gamma[\alpha : \kappa] \vdash \sigma = \sigma'\ \ type}{\Gamma \vdash [\alpha{:}\kappa, \sigma] = [\alpha{:}\kappa', \sigma']\ \ sig}\ (\alpha \notin \mathrm{Dom}(\Gamma)$$

$\boxed{\Gamma \vdash S_1 \leq S_2\ \ sig}$

$$\frac{\Gamma \vdash S_1 = S_2\ \ sig}{\Gamma \vdash S_1 \leq S_2\ \ sig}$$

$$\frac{\Gamma \vdash S_1 \leq S_2\ \ sig \quad \Gamma \vdash S_2 \leq S_3\ \ sig}{\Gamma \vdash S_1 \leq S_3\ \ sig}$$

$$\frac{\begin{array}{c}\Gamma \vdash \kappa \leq \kappa'\ \ kind \\ \Gamma[\alpha : \kappa] \vdash \sigma = \sigma'\ \ type \\ \Gamma[\alpha : \kappa'] \vdash \sigma'\ \ type\end{array}}{\Gamma \vdash [\alpha{:}\kappa, \sigma] \leq [\alpha{:}\kappa', \sigma']\ \ sig}\ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash M : S}$

$$\frac{\begin{array}{c}\Gamma \vdash c : \kappa \\ \Gamma \vdash e : \sigma[c/\alpha] \\ \Gamma[\alpha : \kappa] \vdash \sigma\ \ type\end{array}}{\Gamma \vdash [c, e] : [\alpha{:}\kappa, \sigma]}\ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash M : S' \quad \Gamma \vdash S' \leq S\ \ sig}{\Gamma \vdash M : S}$$

$\boxed{\Gamma \vdash M \downarrow S}$

$$\frac{\begin{array}{c}\Gamma \vdash c \downarrow \kappa \\ \Gamma \vdash e \downarrow \sigma[c/\alpha] \\ \Gamma[\alpha : \kappa] \vdash \sigma\ \ type\end{array}}{\Gamma \vdash [c, e] \downarrow [\alpha{:}\kappa, \sigma]}\ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash M \downarrow S' \quad \Gamma \vdash S' \leq S\ \ sig}{\Gamma \vdash M \downarrow S}$$

$$\frac{\Gamma \vdash M' \downarrow S \quad \Gamma \vdash M = M' : S}{\Gamma \vdash M \downarrow S}$$

$\boxed{\Gamma \vdash M_1 = M_2 : S}$

$$\frac{\Gamma \vdash M : S}{\Gamma \vdash M = M : S} \qquad \frac{\Gamma \vdash M_2 = M_1 : S}{\Gamma \vdash M_1 = M_2 : S}$$

$$\frac{\Gamma \vdash M_1 = M_2 : S \quad \Gamma \vdash M_2 = M_3 : S}{\Gamma \vdash M_1 = M_3 : S}$$

$$\frac{\begin{array}{c}\Gamma \vdash c = c' : \kappa \\ \Gamma \vdash e : \sigma[c/\alpha] \\ \Gamma[\alpha : \kappa] \vdash \sigma\ \ type\end{array}}{\Gamma \vdash [c, e] = [c', e] : [\alpha{:}\kappa, \sigma]}\ (\alpha \notin \mathrm{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash M_1 = M_2 : S' \quad \Gamma \vdash S \leq S'\ \ sig}{\Gamma \vdash M_1 = M_2 : S}$$

A.3 Recursive Module Calculus

$\boxed{\Gamma \vdash M : S}$

$$\frac{\Gamma[s \uparrow S] \vdash M \downarrow S \quad \Gamma \vdash S\ \ sig}{\Gamma \vdash fix(s{:}S.M) : S}\ (s \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash M_1 = M_2 : S}$

$$\frac{\begin{array}{cc}\Gamma \vdash \kappa\ \ kind & \Gamma[\alpha : \kappa] \vdash \sigma\ \ type \\ \Gamma[\alpha \uparrow \kappa] \vdash c \downarrow \kappa & \Gamma[\alpha : \kappa][x \uparrow \sigma] \vdash e \downarrow \sigma\end{array}}{\begin{array}{l}\Gamma \vdash fix\,(s{:}[\alpha{:}\kappa.\sigma].[c[Fst\, s/\alpha], e[Fst\, s, Snd\, s/\alpha, x]]) = \\ \quad [\alpha = \mu\alpha{:}\kappa.c, fix\,(x{:}\sigma.e)] : [\alpha{:}\kappa, \sigma]\end{array}}$$
$$(\alpha, x, s \notin \mathrm{Dom}(\Gamma))$$

$\boxed{\Gamma \vdash S\ \ sig}$

$$\frac{\begin{array}{cc}\Gamma \vdash S\ \ sig & \Gamma[s \uparrow S] \vdash c \downarrow \kappa \\ \Gamma[s : S] \vdash [\alpha{:}\mathfrak{S}(c : \kappa), \sigma]\ \ sig & \end{array}}{\Gamma \vdash \rho s.[\alpha{:}\mathfrak{S}(c : \kappa), \sigma]\ \ sig}\ \left(\begin{array}{l} s \notin \mathrm{Dom}(\Gamma)\ and \\ S\ is \\ [\alpha{:}\kappa, \sigma[\alpha/Fst\, s]] \end{array}\right)$$

$\boxed{\Gamma \vdash S_1 = S_2\ \ sig}$

$$\frac{\begin{array}{cc}\Gamma \vdash \kappa\ \ kind & \Gamma[\beta : \kappa] \vdash \mathfrak{S}(c : \kappa)\ \ kind \\ \Gamma[\beta \uparrow \kappa] \vdash c \downarrow \kappa & \Gamma[\alpha : \kappa] \vdash \sigma[\alpha/Fst\, s]\ \ type\end{array}}{\begin{array}{l}\Gamma \vdash \rho s.[\alpha{:}\mathfrak{S}(c[Fst\, s/\beta] : \kappa), \sigma] = \\ \quad [\alpha{:}\mathfrak{S}(\mu\beta{:}\kappa.c : \kappa), \sigma[\alpha/Fst\, s]]\ \ sig\end{array}}$$
$$(\alpha, \beta, s \notin \mathrm{Dom}(\Gamma))$$