

Intensional Polymorphism in Type-Erasure Semantics*

Karl Cray Stephanie Weirich Greg Morrisett

Cornell University

Abstract

Intensional polymorphism, the ability to dispatch to different routines based on types at run time, enables a variety of advanced implementation techniques for polymorphic languages, including tag-free garbage collection, unboxed function arguments, polymorphic marshalling, and flattened data structures. To date, languages that support intensional polymorphism have required a type-passing (as opposed to type-erasure) interpretation where types are constructed and passed to polymorphic functions at run time. Unfortunately, type-passing suffers from a number of drawbacks: it requires duplication of constructs at the term and type levels, it prevents abstraction, and it severely complicates polymorphic closure conversion.

We present a type-theoretic framework that supports intensional polymorphism, but avoids many of the disadvantages of type passing. In our approach, run-time type information is represented by ordinary terms. This avoids the duplication problem, allows us to recover abstraction, and avoids complications with closure conversion. In addition, our type system provides another improvement in expressiveness; it allows unknown types to be refined in place thereby avoiding certain beta-expansions required by other frameworks.

1 Introduction

Type-directed compilers use type information to enable optimizations and transformations that are impossible (or prohibitively difficult) without such information [13,

11, 17, 2, 23, 25, etc.]. However, type-directed compilers for some languages such as Modula-3 and ML face the difficulty that some type information cannot be known at compile time. For example, polymorphic code in ML may operate on inputs of type α where α is not only unknown, but may in fact be instantiated by a variety of different types.

In order to use type information in contexts where it cannot be provided statically, a number of advanced implementation techniques process type information at run time [11, 17, 28, 19, 25]. Such type information is used in two ways: behind the scenes, typically by tag-free garbage collectors [28, 1], and explicitly in program code, for a variety of purposes such as efficient data representation and marshalling [17, 11, 26]. In this paper we focus on the latter area of applications.

To lay a solid foundation for programs that analyze types at run time, Harper and Morrisett devised an internal language, called λ_i^{ML} , that supports the first-class *intensional analysis*¹ of types (following earlier work by Constable [3, 4]). The λ_i^{ML} language and its derivatives were then used extensively in the high-performance ML compilers TIL/ML [27, 20] and FLINT [26]. The primary novelty of λ_i^{ML} is the presence of “typecase” operators at the level of terms and types, that allow computations and type expressions to depend upon the values of other type expressions at run time.

Supporting intensional type analysis (and the use of type information at run time in general) seems to require semantics where types are constructed and passed to polymorphic functions during computation. However, there are a number of practical and theoretical reasons why type-passing is unattractive:

- A type-passing language such as λ_i^{ML} requires that type information *always* be constructed and passed to polymorphic functions. This framework can result in considerable overhead if types are rarely examined at run time, and, as we discuss later, it makes abstraction impossible.
- Type passing results in considerable complexity in

*This material is based upon work supported in part by AFOSR grant F49620-97-1-001, ARPA/RADC grant FC30602-96-1-0317, NSF grant CCR-9244739, and ONR grant N00014-92-J-1764. The second author is also supported under a National Science Foundation Graduate Fellowship. Any opinions, findings, and conclusions or recommendations in this publication are those of the authors and do not reflect the views of these agencies.

¹“Intensional” since types are analyzed by the structure of their names, rather than by what terms they contain. This approach is critical for practicality.

language semantics, due in large part to the number of semantic devices that must be duplicated for both terms and types. For example, in semantics that make memory allocation explicit [18, 19] a central device is a formal heap in which data is stored; in a type-erasure framework one such heap suffices, but when types are passed it is necessary to add a second heap (and all the attendant machinery) for type data.

- Type passing also greatly complicates low-level intermediate languages, again due in large part to the duplication of computational devices at the type level, and also to the need to support mixed-phase devices (constructs with both type and term level components). This can pose a serious problem for typed intermediate languages, because these devices can disrupt the essential symmetries on which elegant type systems depend. For example, a type-passing semantics for Typed Assembly Language [21] requires additional instructions for allocating and initializing types, which in turn requires the typing machinery for allocation and initialization to be lifted an additional level into the kind structure.
- As a particularly important example of the preceding issue, type passing severely complicates typed closure conversion (compare the type-passing system of Minamide *et al.* [15] to the type-erasure system of Morrisett *et al.* [21]). In a type-erasure framework, the partial application of a polymorphic function to a type may still be considered a value (since the application has no run-time significance), which means that closed code may simply be instantiated with its type environment when a closure is created. In a type-passing framework, the instantiation with a type environment can have some run-time effect, so it must be delayed until the function is invoked. Consequently, closures must include a type environment, necessitating complicated mechanisms including abstract kinds and translucent types [15].

In this paper we propose a typed calculus, called λ_R , that ameliorates the problems of type passing without sacrificing intensional type analysis. If run-time type dispatch is supported, then clearly on some level types must be passed. The fundamental idea behind our approach is to construct and pass *terms* that represent types instead of the types themselves. The connection between a type τ and its term representation e is made in the static semantics by assigning e the special type $R(\tau)$. Semantically, we may interpret $R(\tau)$ as a singleton type that contains only the representation of τ .

This framework resolves the difficulties with type passing semantics discussed above. In particular, as representations of types are simply terms, we can use the pre-existing term operations to deal with run-time type information in languages and their semantics. Furthermore, we can eliminate the difficulties associated with polymorphic closure conversion, as we show in Section 4. Finally, our approach enables the choice *not* to pass

representations. In turn, this choice allows us to eliminate the overhead of constructing and passing representations of types where it is not necessary.²

Perhaps more importantly, the ability not to pass types allows abstraction and parametricity to be recovered. In most type systems, abstraction may be achieved by hiding the identity of types either through parametric polymorphism [22] or through existential types [16]. However, when all types are passed and may be analyzed (as in λ_i^{ML}), the identity of types cannot be hidden and consequently abstraction is impossible. In contrast, a λ_R type can be analyzed only when its representation is available at run time, so abstraction can be achieved simply by not supplying type representations.

For example, consider the type $\exists\alpha.\alpha$. When all types may be analyzed, this type implements a *dynamic* type; an expression of this type provides an object of some unknown type, and that unknown type’s identity can be determined at run time by analyzing α . In λ_R , as in most other type systems, $\exists\alpha.\alpha$ implements an abstract type (in this particular example, a useless abstract type), because no representation of α is provided. Dynamic types are implemented in λ_R by including a representation of the unknown type, as in $\exists\alpha. R(\alpha) \times \alpha$.

1.1 Expressiveness

In the interest of clarity of presentation, we express λ_R as an extension of Harper and Morrisett’s λ_i^{ML} and focus on their differences. The principal difference is the restriction of type analysis to those types for which representations are provided. This change does not diminish the expressiveness of our calculus; λ_i^{ML} may be translated in a straightforward syntax-directed manner into λ_R .

Moreover, the λ_R calculus incorporates an additional improvement in expressiveness over λ_i^{ML} that is independent of explicit type passing: In λ_i^{ML} , information gained by analyzing a type is not propagated to other variables having that type. Consequently, when analyzing a type α with the interest of processing an object of type α , it is necessary to create a function with argument type α and then apply that function to the object of interest. In other words, the type system of λ_i^{ML} requires the use of beta-expansions that are not operationally necessary. In λ_R we resolve this shortcoming by strengthening the typing rule for typecase so that it refines types in place.

1.2 Overview

The remainder of this paper is organized as follows: In Section 2 we review the λ_i^{ML} calculus. We then present, in Section 3, our λ_R calculus and discuss its formal semantics, including representation terms, R -types, and the strengthened typecase rule. In Section 4, we discuss

²In fact, the TIL/ML compiler already finds it necessary to use annotations that mark whether a type must be passed at run-time. Our system provides a formal basis for that mechanism.

the simplification of polymorphic closure conversion by explicit type passing. We end with related work and conclusions in Sections 5 and 6. In the appendices we relate our typed semantics to an untyped one through type erasure (Appendix A), discuss the analysis of quantified types (Appendix B), and provide the formal operational and static semantics (Appendices C and D.)

2 Intensional Type Analysis

Suppose we wanted to efficiently store an array of boolean values. Most computer architectures require that memory accesses are a word at a time, but it is a terrible waste of space to store booleans as integers. The solution is to pack thirty two booleans into one word and use bit manipulations to retrieve the correct value. To subscript from a packed boolean array, we might use the following function (with \ll for shift left, $\&$ for bitwise and, and $\langle \rangle$ for inequality):

```
val bitsub : array[int] * int -> bool =
  fn (a,i) =>
    sub(a,i div 32) & (1<<(i mod 32)) <> 0
```

This function is fine when we know a given array contains boolean values, but we would like code polymorphic over all arrays to be able to use this mechanism. Below we define a new array constructor, `PackedArray`, which will produce an array of integers to hold booleans, and an ordinary array for other types. We also define an associated subscript operation, `packedsub`, which calls `bitsub` on arrays of booleans and the ordinary subscript operator on arrays of other types. These constructs can be created with intensional type analysis, where in both cases an argument type is examined with a “typecase” form:

```
type PackedArray[α] =
  Typecase α of
    bool => array[int]
    | _ => array[α]

val packedsub : ∀α. PackedArray[α] * int -> α =
  Fn [α] =>
    typecase α of
      bool => bitsub
      | _ => sub
```

2.1 The λ_i^{ML} calculus

To formalize the tools of intensional type analysis, we begin by summarizing Harper and Morrisett’s λ_i^{ML} calculus [11]. The λ_i^{ML} calculus provides these tools in a form that is relatively simple, but already quite powerful.

The syntax of λ_i^{ML} appears below (modified slightly for presentation). The backbone is a predicative variant of Girard’s F_ω [9, 8] in which the quantified type $\forall\alpha:\kappa.\sigma$ ranges only over type constructors and “small” types (*i.e.*, monotypes), which do not include the quantified

types. The type analysis operators are `Typerec` and `typecase` at the constructor and term levels respectively.

```
(kinds)  κ ::= Type | κ1 → κ2
(con's)  c ::= α | int | c1 → c2 | c1 × c2 |
          λα:κ.c | c1c2 |
          Typerec c (cint, c→, c×)
(types)  σ ::= c | σ1 → σ2 | σ1 × σ2 | ∀α:κ.σ
(terms)  e ::= x | i | λx:σ.e | fix f:σ.v | e1e2 |
          ⟨e1, e2⟩ | π1e | π2e |
          Λα:κ.v | e[c] |
          typecase[α.σ] c of
            int ⇒ eint
            β → γ ⇒ e→
            β × γ ⇒ e×

(values) v ::= i | λx:σ.e | fix x:σ.v |
          ⟨v1, v2⟩ | Λα:κ.v
```

Occasionally, for brevity, we will write `typecase` terms as `typecase[α.σ] c (eint, βγ.e→, βγ.e×)`.

```
fix tostring : (∀α:Type. α → string).
Λα:Type.
  typecase[δ.δ → string] α of
    int ⇒ int2string
    string ⇒ λobj:string.obj
    β → γ ⇒
      λobj:(β → γ)."function"
    β × γ ⇒
      λobj:(β × γ).
        "<~(tostring[β](π1 obj))~"
        ", ~"(tostring[γ](π2 obj))~">
```

Figure 1: The function `tostring`

As an example of the use of type analysis in λ_i^{ML} (with the addition of another base type, `string`), consider the function `tostring`, presented in Figure 1. This function uses `typecase` to produce a string representation of a data object. For example, the call `tostring[int]3` returns the string “3”. As we cannot provide any information about the implementation of functions, we just return the word “function” when one is encountered, as in the call:

$$\text{tostring}[(\text{int} \rightarrow \text{int}) \times \text{int}] (\lambda x:\text{int}. x + 1, 3)$$

which returns

“⟨function, 3⟩”.

When the argument to `tostring` is a product type, the function calls itself recursively. In this branch, the type variables β and γ are bound to the types of the first and second components of the tuple, so that the recursive call can be instantiated with the correct type.

The `typecase` form has a type annotation for type checking without type inference; the annotation `[α.σ]` indicates that given a type constructor argument `c`, the

`typecase` computes a value with type $\sigma[c/\alpha]$ (where this syntax denotes the capture-avoiding substitution of c for α in σ). In this example, each arm returns a function from δ to *string*, where δ is replaced by the appropriate type, such as *int* in the *int* branch, and $\beta \times \gamma$ in the product branch.

With this intuition, the typing rule for `typecase` is the natural one (but we will see that this rule is somewhat restrictive):

$$\frac{\begin{array}{l} \Gamma \vdash c : \mathbf{Type} \quad \Gamma, \alpha : \mathbf{Type} \vdash \sigma \text{ type} \\ \Gamma \vdash e_{int} : \sigma[int/\alpha] \\ \Gamma, \beta : \mathbf{Type}, \gamma : \mathbf{Type} \vdash e_{\rightarrow} : \sigma[\beta \rightarrow \gamma/\alpha] \\ \Gamma, \beta : \mathbf{Type}, \gamma : \mathbf{Type} \vdash e_{\times} : \sigma[\beta \times \gamma/\alpha] \end{array}}{\Gamma \vdash (\mathbf{typecase}[\alpha.\sigma] c \text{ of}) : \sigma[c/\alpha]} \quad \begin{array}{l} int \Rightarrow e_{int} \\ \beta \rightarrow \gamma \Rightarrow e_{\rightarrow} \\ \beta \times \gamma \Rightarrow e_{\times} \end{array}$$

Often, to compute the result type σ of a `typecase` expression the constructor-level `Typerec` on the argument α will be required. `Typerec` allows the creation of new types by similar intensional analysis. Several examples of its use appear in Harper and Morrisett [11], including type-directed data layout, marshalling and unboxing.

While recursion in the term-level `typecase` is handled by `fix`, at the the constructor level there is no such mechanism. For this reason, `Typerec` is essentially a “fold” operation (or catamorphism) over inductively defined types. It provides primitive recursion by calling itself recursively on all of the components of the argument type. Also unlike `typecase`, where the branches explicitly bind arguments for the components of the type, the c_{\rightarrow} and c_{\times} branches of `Typerec` are constructor functions. For example, if the argument of a `Typerec` operation is $c_1 \times c_2$, then that operation reduces to its c_{\times} branch (a constructor function of four arguments) applied to the components c_1 and c_2 , and to the result of recursively computing the `Typerec` operation on those components.

$$\mathbf{Typerec}(c_1 \times c_2)(c_{int}, c_{\rightarrow}, c_{\times}) = \begin{array}{l} c_{\times} c_1 c_2 \\ (\mathbf{Typerec} c_1 (c_{int}, c_{\rightarrow}, c_{\times})) \\ (\mathbf{Typerec} c_2 (c_{int}, c_{\rightarrow}, c_{\times})) \end{array}$$

The kinding rule for `Typerec` is again the natural one. To compute a constructor of kind κ , present a type argument and three branches returning κ constructors:

$$\frac{\begin{array}{l} \Gamma \vdash c : \mathbf{Type} \quad \Gamma \vdash c_{int} : \kappa \\ \Gamma \vdash c_{\rightarrow} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \\ \Gamma \vdash c_{\times} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \end{array}}{\Gamma \vdash \mathbf{Typerec} c (c_{int}, c_{\rightarrow}, c_{\times}) : \kappa}$$

3 The λ_R calculus

Figure 2 presents the syntax of λ_R , which we describe in detail in the following section.

$$\begin{array}{ll} (\mathit{kinds}) & \kappa ::= \mathbf{Type} \mid \kappa \rightarrow \kappa \\ (\mathit{con's}) & c ::= \mathit{int} \mid \alpha \mid \lambda\alpha:\kappa.c \mid c_1 c_2 \mid \\ & c_1 \rightarrow c_2 \mid c_1 \times c_2 \mid R(c) \mid \\ & \mathbf{Typerec} c (c_{int}, c_{\times}, c_{\rightarrow}, c_R) \\ (\mathit{types}) & \sigma ::= c \mid \sigma_1 \rightarrow \sigma_2 \mid \sigma_1 \times \sigma_2 \mid \\ & \forall\alpha:\kappa.\sigma \mid \exists\alpha:\kappa.\sigma \\ (\mathit{terms}) & e ::= x \mid i \mid \langle e_1, e_2 \rangle \mid \pi_1 e \mid \pi_2 e \mid \lambda x:\sigma.e \mid \\ & e_1 e_2 \mid \Lambda\alpha:\kappa.v \mid e[c] \mid \mathbf{fix} f:\sigma.v \mid \\ & \mathbf{pack} e \text{ as } \exists\alpha.\sigma_1 \mathbf{hiding} \sigma_2 \mid \\ & \mathbf{unpack}(\alpha, x) = e_1 \mathbf{in} e_2 \mid R_{int} \mid \\ & R_{\rightarrow}(e_1, e_2) \mid R_{\times}(e_1, e_2) \mid R_R(e) \mid \\ & \mathbf{typecase}[\delta.c] e \text{ of} \\ & \quad R_{int} \Rightarrow e_{int} \\ & \quad R_{\rightarrow}(x, y) \text{ as } (\beta \rightarrow \gamma) \Rightarrow e_{\rightarrow} \\ & \quad R_{\times}(x, y) \text{ as } (\beta \times \gamma) \Rightarrow e_{\times} \\ & \quad R_R(x) \text{ as } R(\beta) \Rightarrow e_R \\ (\mathit{values}) & v ::= i \mid \langle v_1, v_2 \rangle \mid \lambda x:\sigma.e \mid \\ & \Lambda\alpha:\kappa.v \mid v[c] \mid \mathbf{fix} f:\sigma.v \mid \\ & \mathbf{pack} v \text{ as } \exists\alpha.\sigma_1 \mathbf{hiding} \sigma_2 \mid R_{int} \mid \\ & R_{\rightarrow}(v_1, v_2) \mid R_{\times}(v_1, v_2) \mid R_R(v) \end{array}$$

Figure 2: Syntax of λ_R

3.1 Term Representations of Types

The key feature we add to the term language of λ_R is the representations of types as terms, which remain when the types themselves are ultimately erased. The base type, *int*, has a corresponding representation constant R_{int} . Likewise, inductive types have inductively defined representations; the type $int \rightarrow int$ is represented by the term $R_{\rightarrow}(R_{int}, R_{int})$.

Accordingly, the argument to the term level `typecase` is the representation of a type, instead of a type. For example, if the argument e is of the form $R_{\rightarrow}(e_1, e_2)$, the arrow branch (e_{\rightarrow}) is taken. The type variables β and γ are still bound to the types that e_1 and e_2 represent, but, because we need not only the component types but also their representations, x and y are bound to e_1 and e_2 . This notion is reflected in the following rule of the operational semantics:

$$\mathbf{typecase}[\delta.c] (R_{\rightarrow}(e_1, e_2)) (e_{int}, \beta\gamma xy.e_{\rightarrow}, \dots) \mapsto e_{\rightarrow}[\mathcal{D}(e_1), \mathcal{D}(e_2), e_1, e_2/\beta, \gamma, x, y]$$

The operation $\mathcal{D}(\cdot)$ in this rule converts a representation to the type that it denotes (Figure 3). The rest of our dynamic semantics is formalized in Appendix C. It is presented as a call-by-value, small step operational semantics.

In order to assign a type to these representations of types, we have extended the type constructor level of λ_R with the R construct, where the representation of a type τ is given the type $R(\tau)$, and extended the static semantics accordingly. For example, the formation rule

$$\begin{aligned}
\mathcal{D}(R_{int}) &= int \\
\mathcal{D}(R_{\times}(e_1, e_2)) &= \mathcal{D}(e_1) \times \mathcal{D}(e_2) \\
\mathcal{D}(R_{\rightarrow}(e_1, e_2)) &= \mathcal{D}(e_1) \rightarrow \mathcal{D}(e_2) \\
\mathcal{D}(R_R(e)) &= R(\mathcal{D}(e))
\end{aligned}$$

Figure 3: Translating Representations to Types

for the representation of function types is

$$\frac{\Gamma \vdash e_1 : R(\tau_1) \quad \Gamma \vdash e_2 : R(\tau_2)}{\Gamma \vdash R_{\rightarrow}(e_1, e_2) : R(\tau_1 \rightarrow \tau_2)} \text{ (rep}_{\rightarrow}\text{)}$$

which says that if the two subterms, e_1 and e_2 , are type representations of τ_1 and τ_2 , then $R_{\rightarrow}(e_1, e_2)$ will be a representation of $\tau_1 \rightarrow \tau_2$.

As an example of the use of λ_R , the *tostring* function from the previous section can be transliterated into λ_R by requiring it to take an additional term argument, x_α for the representation of the argument type:

```

fix tostring : ( $\forall \alpha$ :Type.  $R(\alpha) \rightarrow \alpha \rightarrow string$ ).
 $\Lambda \alpha$ :Type.  $\lambda x_\alpha$ : $R(\alpha)$ .
  typecase[ $\delta$ . $\alpha \rightarrow string$ ]  $x_\alpha$  of
     $R_{int} \Rightarrow int2string$ 
     $R_{string} \Rightarrow \lambda obj:string. obj$ 
     $R_{\rightarrow}(x, y) \text{ as } \beta \rightarrow \gamma \Rightarrow$ 
       $\lambda obj:\beta \rightarrow \gamma.$ "function"
     $R_{\times}(x, y) \text{ as } \beta \times \gamma \Rightarrow$ 
       $\lambda obj:\beta \times \gamma.$ 
      " $<$ " ^ (tostring [ $\beta$ ]  $x$  ( $\pi_1$   $obj$ )) ^
      " $,$ " ^ (tostring [ $\gamma$ ]  $y$  ( $\pi_2$   $obj$ )) ^ " $>$ "

```

The static semantics we have defined ensures that these R -types are singleton types; for each one there is exactly one value which inhabits it. This fact allows us to express constraints between types and their representations at a very fine level. For instance, in the *tostring* example, the representation argument must be the representation of the type of the object.

Furthermore, as we have added a new way to form types to the constructor language, we must add another term construct, $R_R(\cdot)$, to form the representation of representation types. We also extend **typecase** with an extra branch to handle these terms and **Typerec** to handle R -types.

3.2 In-place Refinement of Types

The typing rules of λ_i^{ML} often force an inefficient use of **typecase**. In the *tostring* example in Section 2, we were required to create closures in each of the branches of the **typecase**. It would be more efficient if we could lift the lambdas outside of the **typecase** and have each branch of the **typecase** return a *string*. We could then write

this function as:

```

fix tostring : ( $\forall \alpha$ :Type.  $R(\alpha) \rightarrow \alpha \rightarrow string$ ).
 $\Lambda \alpha$ :Type.  $\lambda x_\alpha$ : $R(\alpha)$ .  $\lambda obj$ : $\alpha$ .
  typecase[ $\delta$ .string]  $x_\alpha$  of
     $R_{int} \Rightarrow int2string$   $obj$ 
     $R_{string} \Rightarrow obj$ 
     $R_{\rightarrow}(x, y) \text{ as } \beta \rightarrow \gamma \Rightarrow$ 
      "function"
     $R_{\times}(x, y) \text{ as } \beta \times \gamma \Rightarrow$ 
      " $<$ " ^ (tostring [ $\beta$ ]  $x$  ( $\pi_1$   $obj$ )) ^
      " $,$ " ^ (tostring [ $\gamma$ ]  $y$  ( $\pi_2$   $obj$ )) ^ " $>$ "

```

The reason we could not write this function in λ_i^{ML} is that it requires the type of *obj* to change based upon which branch of the **typecase** is selected. In λ_i^{ML} , all that is known in the product branch is that *obj* is of type α , not a tuple. In order to project from it in the recursive calls, the typing rules must update the type of *obj* to reflect the fact that we know that α is $\beta \times \gamma$ in the product branch.

With the right enhancement to the static semantics this optimization is possible. We have held off discussion of the λ_R 's **typecase** formation rule in order to emphasize this point. The basic idea is that in some cases **typecase** increases our knowledge of the argument type. We separate the formation rule into situations where **typecase** gives us new information, such as when the argument is of type $R(\alpha)$, and when it does not, such as when the argument is of type $R(\tau_1 \rightarrow \tau_2)$. In the inference rule for type checking a **typecase** term, when the argument is of type $R(\alpha)$, we refine types containing α to reflect the gain in information, as shown below. For simplicity, only some of the rule is given here (the complete rule appears in Appendix D.3):

$$\frac{\begin{array}{l}
\Gamma, \alpha : \text{Type}, \Gamma' \vdash e : R(\alpha) \\
\Gamma, \Gamma' [int/\alpha] \vdash e_{int} [int/\alpha] : c[int, int/\alpha, \delta] \\
\Gamma, \beta : \text{Type}, \gamma : \text{Type}, \Gamma' [\beta \rightarrow \gamma/\alpha], x : R(\beta), \\
y : R(\gamma) \vdash e_{\rightarrow} [\beta \rightarrow \gamma/\alpha] : c[\beta \rightarrow \gamma, \beta \rightarrow \gamma/\alpha, \delta] \\
\vdots
\end{array}}{\Gamma, \alpha : \text{Type}, \Gamma' \vdash \text{typecase}[\delta.c] e (e_{int}, \beta \gamma xy.e_{\rightarrow}, \dots) : c[\alpha/\delta]}$$

For example, to typecheck the e_{\rightarrow} branch, we substitute $\beta \rightarrow \gamma$ for α everywhere, including the surrounding context.³ Consequently the types of the variables bound in the context will be refined by that substitution. Because λ_i^{ML} only makes this substitution in the return type of the branch, and not in the context, in order to propagate this information one must abstract over all variables of interest.

When we know more about the argument because of the singularity of the R -types, we can deduce statically which branch of the **typecase** will be taken. Therefore we do not need to typecheck the other branches at all, leading to a much simpler rule. For example, if we know

³The substitution for α is applied within the branches themselves in order to avoid creating a hole in the scope of α . In practice, a typechecker would implement this operation by a local type definition, rather than by substitution.

Judgment	Meaning
$\Gamma \vdash c : \kappa$	c is a valid constructor of kind κ
$\Gamma \vdash \sigma$	σ is a valid type
$\Gamma \vdash c_1 = c_2 : \kappa$	c_1 and c_2 are equal constructors
$\Gamma \vdash \sigma_1 = \sigma_2$	σ_1 and σ_2 are equal types
$\Gamma \vdash e : \sigma$	e is a term of type σ

Figure 4: Judgments of λ_R

the argument is of type $R(\tau_1 \rightarrow \tau_2)$, we only need to examine the e_{\rightarrow} branch, as in the rule:

$$\frac{\Gamma \vdash e : R(\tau_1 \rightarrow \tau_2) \quad \Gamma, x:R(\tau_1), y:R(\tau_2) \vdash e_{\rightarrow}[\tau_1, \tau_2/\beta, \gamma] : c[\tau_1 \rightarrow \tau_2/\delta]}{\Gamma \vdash \text{typecase}[\delta.c] e (e_{int}, \beta\gamma xy.e_{\rightarrow}, \dots) : c[\tau_1 \rightarrow \tau_2/\delta]}$$

3.3 Properties of the Formal Semantics

Formally, the static semantics of λ_R consists of a collection of rules for deriving judgments of the forms shown in Figure 4. In these judgments, Γ is a unified type and kind context, mapping constructor variables (α, β, \dots) to kinds and term variables (x, y, \dots) to types. The formal operational and static semantics of λ_R appear in Appendices C and D, and from them we can prove several useful properties about λ_R .

First, we would like to prove the decidability of λ_R type-checking. The only mildly difficult part is equivalence checking for constructors. Based upon the equivalence rules in Appendix D.2 we can define a notion of constructor reduction to a normal form in an obvious manner. This reduction relation can be proved strongly normalizing and confluent (in a manner similar to Morrisett [17]) from which it follows that constructor equivalence is decidable. Therefore we can state the following theorem:

Theorem 3.1 (Decidability) *It is decidable whether or not $\Gamma \vdash e : \tau$ is derivable in λ_R .*

Next, we would like to show that the static semantics guarantees safety; that is, if a term typechecks, then the operational semantics will not get *stuck* (where a term that is not a value, and for which no rule of our operational semantics applies, is *stuck*):

Theorem 3.2 (Type Safety) *If $\emptyset \vdash e : \sigma$ and $e \mapsto^* e'$ then e' is not stuck.*

The proof of this theorem is standard, relying on the usual progress, subject reduction and substitution lemmas.

4 Polymorphic Typed Closure Conversion

As a final example, we consider typed closure conversion in a λ_R -like framework. The key idea behind closure conversion is to shift from a substitution-based model of execution to an environment-based model via a source-to-source translation. In particular, all functions are replaced with explicit closures which are represented within the language as pairs consisting of a λ -abstraction (the code of the closure), and a tuple (the environment of the closure). The environment contains values for the free variables of the function. The code abstracts the environment as well as the arguments of the function and is thus closed. Hence, the code may be hoisted to the top-level, allocated at compile time, and shared among all substitution instances. Application is rewritten so that the code of a closure is first applied to its environment and then to its arguments.

In the monomorphic case no discrepancy arises between type-passing [15] and type-erasure [21] closure conversion. An existential type is used to hold the type of the closure's environment abstract, so a closure for a $\tau_1 \rightarrow \tau_2$ function is given the type $\exists\alpha.((\tau_1 \times \alpha) \rightarrow \tau_2) \times \alpha$.

However, with the introduction of polymorphism, significant differences arise between type-passing and type-erasure. The issue stems from the fact that functions may contain free type variables as well as free value variables, and closed code must abstract both. Closures must then provide some way to apply such code to the appropriate type variables. In a type-erasure setting, type application has no run-time effect, so the partial application of code to the appropriate type variables may be performed when closures are created. Consequently, these type variables do not appear in the type of a closure. In fact, closures have the same type ($\exists\alpha.((\tau_1 \times \alpha) \rightarrow \tau_2) \times \alpha$) as before.

However, in a type-passing semantics, the application to type arguments is a run-time operation and so such applications must be suspended until the closure is called. Thus, it is necessary for the closure to include a type environment as well as a value environment. The kind of the type environment must be hidden (as did the type of the value environment in the monomorphic case), and the closure's type must enforce the requirement that the code be applied only to the proper type environment (see Minamide *et al.* [15] for detailed explanations of why). The former requires the use of abstract kinds and the latter requires the use of translucent types [10]. This approach results in a closure having the considerably more complicated type (again, see Minamide *et al.* [15] for a formalization of the necessary type theory):

$$\exists k_{\text{tenv}}:\text{Kind}. \exists \alpha_{\text{venv}}:\text{Type}. \exists \beta_{\text{tenv}}:k_{\text{tenv}}. (\forall \gamma:k_{\text{tenv}}=\beta_{\text{tenv}}. (\tau_1 \times \alpha_{\text{venv}}) \rightarrow \tau_2) \times \alpha_{\text{venv}}$$

In the above type, k_{tenv} abstracts the kind of the type environment, α_{venv} abstracts the type of the value of the value environment, and β_{tenv} provides the type environment. The code type then takes a type environment γ of kind k_{tenv} as an argument, but γ is constrained (using translucent types) to be the appropriate environment, β_{tenv} .

Since our framework is one of type-erasure, type environments may be resolved by partial application, resulting in the simpler type for closures. However, it is instructive to examine the details. Suppose the function to be closure-converted is the function $f = \lambda x:\tau_1.e$ with type $\tau_1 \rightarrow \tau_2$ and suppose further that the function contains free occurrences of the type variable α and its representation $x_\alpha:R(\alpha)$.

First the function is rewritten in closed form as:

$$\begin{aligned} f' &: \forall \alpha. (\tau_1 \times R(\alpha)) \rightarrow \tau_2 \\ &= \Lambda \alpha. \lambda y: (\tau_1 \times R(\alpha)). e[\pi_1 y, \pi_2 y/x, x_\alpha] \end{aligned}$$

Then (at run time) f' is instantiated with the type environment (that is, α):

$$f'' : (\tau_1 \times R(\alpha)) \rightarrow \tau_2 = f'[\alpha]$$

Finally, a closure is created:

$$f''' = \text{pack} \langle f'', x_\alpha \rangle \text{ as } \exists \beta. ((\tau_1 \times \beta) \rightarrow \tau_2) \times \beta \\ \text{hiding } R(\alpha)$$

Consider what has become of the mechanisms for type-passing closure conversion: The type of f'' requires that it be applied (for its second argument) only to the representation of α . So the translucency mechanism appears again, suggesting that translucency is inherent in type-passing closure conversion. However, this version of translucency has two advantages; the necessary type theory is simpler, and the translucency is completely hidden by the existential packaging in the eventual closure. On the other hand, abstract kinds do not appear in the process, suggesting them to be an artifact of true type-passing.

5 Related Work

Closely related to our work is the work of Minamide on lifting of type parameters for tag-free garbage collection [14]. Minamide was interested in lifting type parameters out of code so they could be preallocated at compile time. His lifting procedure required the maintenance of interrelated constraints between type parameters to retain type soundness, and he used a system similar to ours that makes explicit the passing of type parameters in order to simplify the expression of such constraints. The principal difference between Minamide's system and ours is that Minamide did not consider intensional type analysis or first-class polymorphism. Minamide's system also makes a distinction between type representations (which he calls *evidence*, following Jones [12]) and ordinary terms, while λ_R type representations are fully first-class.

The issue of type parameter lifting is an important one for compilers based on λ_R . The construction of type representations at run time would likely lead to significant cost and, in practice, should be lifted out to compile time whenever possible. (Unfortunately, in the presence of polymorphic recursion, which λ_R supports, it is not always possible.) Mechanisms for such lifting have been

developed by Minamide (in the work discussed above) and by Saha and Shao [24].

Dubois *et al.* [6] also pass explicit type representations to polymorphic functions when compiling ad-hoc polymorphism. However, their system differs from ours and Minamide's in that no mechanism is provided for connecting representations to the types they denote, and consequently, information gained by analyzing type representations does not propagate into the type system.

Duggan [7] proposes another typed framework for intensional type analysis that is similar in some ways to λ_i^{ML} . Like λ_i^{ML} , Duggan's system passes types implicitly and allows for the intensional analysis of types at the term level. Duggan's system does not support intensional type analysis at the constructor level, as λ_i^{ML} and λ_R do, but it adds a facility for defining type classes (using union and recursive kinds) and allows type analysis to be restricted to members of such classes.

6 Conclusions and Future Directions

We have presented a type-theoretic framework that supports the passing and analysis of type information at run time, but that avoids the shortcomings associated with previous such frameworks (*e.g.*, duplication of constructs, lack of abstraction, and complication of closure conversion). This new framework makes it feasible to use intensional type analysis in settings where the shortcomings previously made it impractical.

For example, Morrisett *et al.* [21] developed typing mechanisms for low-level intermediate and target languages that allow type information to be used all the way to the end of compilation. It would be desirable, in a system based on those mechanisms, to be able to exploit that type information using intensional type analysis. Unfortunately, the shortcomings of type-passing semantics made it incompatible with some of those low-level typing mechanisms. This unfortunate incompatibility has made it infeasible to use the mechanisms of Morrisett *et al.* in type-analyzing compilers such as TIL/ML [27, 20] and FLINT [26], and has made it infeasible to use intensional type analysis in the end-to-end typed compiler TALC [21]. The framework in this paper makes it possible to unify these two lines of work for the first time.

In pursuance of this aim, an important direction for future work is to extend the mechanisms of λ_R into lower-level typed intermediate languages such as typed assembly language [21]. Among the issues to be explored in such research is how to analyze the more complicated types used in typed assembly language, and how to perform type-directed dispatch without an atomic **typecase** construct. Another issue to explore is analysis of quantified types (an initial proposal appears in Appendix B), and whether such mechanisms are useful in practice.

Another important question is whether a parametricity

theorem like that of Reynolds [22] can be shown for λ_R . Polymorphism is clearly non-parametric in λ_i^{ML} , but the lowering of type analysis to explicit term-level representatives makes it plausible that some sort of parametricity could be shown for λ_R . In other words, we discussed at an intuitive level in Section 1 how the explicit passing of types restores the ability to abstract types that was discarded by λ_i^{ML} ; it would be interesting to explore how that intuition may be formalized.

7 Acknowledgments

We would like to thank Neal Glew, Fred Smith, Chris Stone, Dave Walker, Steve Zdancewic, and the anonymous referees for comments and suggestions.

References

- [1] S. Aditya and A. Caro. Compiler-directed type reconstruction for polymorphic languages. In *Conference on Functional Programming Languages and Computer Architecture*, pages 74–82, Copenhagen, June 1993.
- [2] L. Birkedal, M. Tofte, and M. Vejlstrup. From region inference to von Neumann machines via region representation inference. In *Twenty-Third ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 171–183, St. Petersburg, Jan. 1996.
- [3] R. L. Constable. Intensional analysis of functions and types. Technical Report CSR-118-82, Department of Computer Science, University of Edinburgh, June 1982.
- [4] R. L. Constable and D. R. Zlatin. The type theory of PL/CV3. *ACM Transactions on Programming Languages and Systems*, 6(1):94–117, Jan. 1984.
- [5] J. Despeyroux, F. Pfenning, and Carsten Schürmann. Primitive recursion for higher-order abstract syntax. In *Third International Conference on Typed Lambda Calculi and Applications*, volume 1210 of *Lecture Notes in Computer Science*, pages 147–163, Nancy, France, Apr. 1997. Springer-Verlag. Extended version published as Carnegie Mellon University technical report CMU-CS-96-172.
- [6] C. Dubois, F. Rouaix, and P. Weis. Extensional polymorphism. In *Twenty-Second ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 118–129, San Francisco, Jan. 1995.
- [7] D. Duggan. A type-based semantics for user-defined marshalling in polymorphic languages. In *Second Workshop on Types in Compilation*, Mar. 1998.
- [8] J.-Y. Girard. Une extension de l’interprétation de Gödel à l’analyse, et son application à l’élimination de coupures dans l’analyse et la théorie des types. In J. E. Fenstad, editor, *Proceedings of the Second Scandinavian Logic Symposium*, pages 63–92. North-Holland Publishing Co., 1971.
- [9] J.-Y. Girard. *Interprétation fonctionnelle et élimination des coupures de l’arithmétique d’ordre supérieur*. PhD thesis, Université Paris VII, 1972.
- [10] R. Harper and M. Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *Twenty-First ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 123–137, Portland, Oregon, Jan. 1994.
- [11] R. Harper and G. Morrisett. Compiling polymorphism using intensional type analysis. In *Twenty-Second ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 130–141, San Francisco, Jan. 1995.
- [12] M. P. Jones. A theory of qualified types. In *Fourth European Symposium on Programming*, volume 582 of *Lecture Notes in Computer Science*, Rennes, France, 1992. Springer-Verlag.
- [13] X. Leroy. Unboxed objects and polymorphic typing. In *Nineteenth ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 177–188, 1992.
- [14] Y. Minamide. Full lifting of type parameters. Submitted for publication. Earlier version published as “Compilation Based on a Calculus for Explicit Type-Passing” in the *Second Fuji International Workshop on Functional and Logic Programming*, 1996.
- [15] Y. Minamide, G. Morrisett, and R. Harper. Typed closure conversion. In *Twenty-Third ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 271–283, St. Petersburg, Florida, Jan. 1996.
- [16] J. C. Mitchell and G. D. Plotkin. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems*, 10(3):470–502, July 1988.
- [17] G. Morrisett. *Compiling with Types*. PhD thesis, Carnegie Mellon University, School of Computer Science, Pittsburgh, Pennsylvania, Dec. 1995.
- [18] G. Morrisett, M. Felleisen, and R. Harper. Abstract models of memory management. In *Conference on Functional Programming Languages and Computer Architecture*, 1995.
- [19] G. Morrisett and R. Harper. Semantics of memory management for polymorphic languages. In A. D. Gordon and A. M. Pitts, editors, *Higher Order Operational Techniques in Semantics*. Cambridge University Press, 1997.
- [20] G. Morrisett, D. Tarditi, P. Cheng, C. Stone, R. Harper, and P. Lee. The TIL/ML compiler: Performance and safety through types. In *Workshop on Compiler Support for Systems Software*, Tucson, Feb. 1996.

- [21] G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly language. In *Twenty-Fifth ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 85–97, San Diego, Jan. 1998. Extended version published as Cornell University technical report TR97-1651.
- [22] J. C. Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing '83*, pages 513–523. North-Holland, 1983. Proceedings of the IFIP 9th World Computer Congress.
- [23] E. Ruf. Partitioning dataflow analyses using types. In *Twenty-Fourth ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 15–26, Paris, Jan. 1997.
- [24] B. Saha and Z. Shao. Optimal type lifting. In *Second Workshop on Types in Compilation*, Mar. 1998.
- [25] Z. Shao. Flexible representation analysis. In *1997 ACM SIGPLAN International Conference on Functional Programming*, pages 85–98, Amsterdam, June 1997.
- [26] Z. Shao. An overview of the FLINT/ML compiler. In *1997 Workshop on Types in Compilation*, Amsterdam, June 1997. ACM SIGPLAN. Published as Boston College Computer Science Department Technical Report BCCS-97-03.
- [27] D. Tarditi, G. Morrisett, P. Cheng, C. Stone, R. Harper, and P. Lee. TIL: A type-directed optimizing compiler for ML. In *1996 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 181–192, May 1996.
- [28] A. Tolmach. Tag-free garbage collection using explicit type parameters. In *ACM Conference on Lisp and Functional Programming*, pages 1–11, Orlando, June 1994.

A Untyped Calculus

Although the formal static and operational semantics for λ_R are for a typed language, we would like to emphasize the point that types are unnecessary for computation and can safely be erased. Accordingly, we exhibit an untyped language, λ_R° , a translation of λ_R to this language through type erasure, and the following theorem, which states that execution in the untyped language mirrors execution in the typed language:

- Theorem A.1** 1. If $e_1 \mapsto^* e_2$ then $e_1^\circ \mapsto^* e_2^\circ$.
2. If $\emptyset \vdash e_1 : \tau$ and $e_1^\circ \mapsto^* u$ then there exists e_2 such that $e_1 \mapsto^* e_2$ and $e_2^\circ = u$.

From this theorem and type safety for λ_R it follows that our untyped semantics is safe.

Corollary A.2 If $\emptyset \vdash e : \tau$ and $e^\circ \mapsto^* u$ then u is not stuck.

A.1 Syntax of Untyped Calculus

(terms) $u ::= x \mid i \mid \lambda x.u \mid \mathbf{fix} f.w \mid u_1 u_2 \mid \langle u_1, u_2 \rangle \mid \pi_1 u \mid \pi_2 u \mid \mathbf{R}_{int} \mid \mathbf{R}_\rightarrow(u_1, u_2) \mid \mathbf{R}_\times(u_1, u_2) \mid \mathbf{R}_R(u) \mid \mathbf{typecase} u \mathbf{of}$
 $\mathbf{R}_{int} \Rightarrow u_{int}$
 $\mathbf{R}_\rightarrow(x, y) \Rightarrow u_\rightarrow$
 $\mathbf{R}_\times(x, y) \Rightarrow u_\times$
 $\mathbf{R}_R(x) \Rightarrow u_R$

(values) $w ::= i \mid \langle w_1, w_2 \rangle \mid \lambda x.u \mid \mathbf{fix} f.w \mid \mathbf{R}_{int} \mid \mathbf{R}_\times(w_1, w_2) \mid \mathbf{R}_\rightarrow(w_1, w_2) \mid \mathbf{R}_R(w)$

A.2 Type Erasure

$$\begin{aligned} x^\circ &= x \\ i^\circ &= i \\ \langle e_1, e_2 \rangle^\circ &= \langle e_1^\circ, e_2^\circ \rangle \\ (\pi_i e)^\circ &= \pi_i e^\circ \\ (\lambda x:c.e)^\circ &= \lambda x.e^\circ \\ (\Lambda\alpha:\kappa.v)^\circ &= v^\circ \\ (\mathbf{fix} f:c.v)^\circ &= \mathbf{fix} f.v^\circ \\ (e_1 e_2)^\circ &= e_1^\circ e_2^\circ \\ e[c]^\circ &= e^\circ \\ \mathbf{pack} e \mathbf{as} c \mathbf{hiding} c^\circ &= e^\circ \\ \mathbf{unpack} \langle \alpha, x \rangle = e_1 \mathbf{in} e_2^\circ &= (\lambda x.e_2^\circ) e_1^\circ \\ \mathbf{R}_{int}^\circ &= \mathbf{R}_{int} \\ \mathbf{R}_\rightarrow(e_1, e_2)^\circ &= \mathbf{R}_\rightarrow(e_1^\circ, e_2^\circ) \\ \mathbf{R}_\times(e_1, e_2)^\circ &= \mathbf{R}_\times(e_1^\circ, e_2^\circ) \\ \mathbf{R}_R(e_1)^\circ &= \mathbf{R}_R(e_1^\circ) \\ (\mathbf{typecase}[\delta.c] e \mathbf{of} &= \mathbf{typecase} e^\circ \mathbf{of} \\ \mathbf{R}_{int} \Rightarrow e_{int} &\mathbf{R}_{int} \Rightarrow e_{int}^\circ \\ \mathbf{R}_\rightarrow(x, y) \mathbf{as} &\mathbf{R}_\rightarrow(x, y) \Rightarrow e_\rightarrow \\ (\beta \rightarrow \gamma) \Rightarrow e_\rightarrow &\mathbf{R}_\times(x, y) \Rightarrow e_\times^\circ \\ \mathbf{R}_\times(x, y) \mathbf{as} &\mathbf{R}_R(x) \Rightarrow e_R^\circ \\ (\beta \times \gamma) \Rightarrow e_\times & \\ \mathbf{R}_R(x) \mathbf{as} & \\ R(\beta) \Rightarrow e_R & \end{aligned}$$

A.3 Operational Semantics of λ_R°

$$\begin{aligned} (\lambda x.u)w &\mapsto u[w/x] \\ (\mathbf{fix} f.w)w' &\mapsto (w[\mathbf{fix} f.w/f])w' \\ \pi_1 \langle w_1, w_2 \rangle &\mapsto w_1 \quad \pi_2 \langle w_1, w_2 \rangle \mapsto w_2 \\ \mathbf{typecase} \mathbf{R}_{int} (u_{int}, xy.u_\rightarrow, &xy.u_\times, x.u_R) \mapsto u_{int} \\ \mathbf{typecase} (\mathbf{R}_\times(w_1, w_2)) (u_{int}, xy.u_\rightarrow, &xy.u_\times, x.u_R) \mapsto u_\times[w_1, w_2/x, y] \\ \mathbf{typecase} (\mathbf{R}_\rightarrow(w_1, w_2)) (u_{int}, xy.u_\rightarrow, &xy.u_\times, x.u_R) \mapsto u_\rightarrow[w_1, w_2/x, y] \end{aligned}$$

$$\begin{array}{c}
\text{typecase}(\mathbf{R}_R(w)) (u_{int}, xy.u_{\rightarrow}, \\
xy.u_{\times}, x.u_R) \mapsto u_R[w/x] \\
\\
\frac{u_1 \mapsto u'_1}{u_1 u_2 \mapsto u'_1 u_2} \quad \frac{u \mapsto u'}{wu \mapsto wu'} \\
\frac{u_1 \mapsto u'_1}{\langle u_1, u_2 \rangle \mapsto \langle u'_1, u_2 \rangle} \quad \frac{u \mapsto u'}{\langle w, u \rangle \mapsto \langle w, u' \rangle} \\
\frac{u \mapsto u'}{\pi_1 u \mapsto \pi_1 u'} \quad \frac{u \mapsto u'}{\pi_2 u \mapsto \pi_2 u'} \\
\frac{u_1 \mapsto u'_1}{\mathbf{R}_{\rightarrow}(u_1, u_2) \mapsto \mathbf{R}_{\rightarrow}(u'_1, u_2)} \quad \frac{u \mapsto u'}{\mathbf{R}_{\rightarrow}(w, u) \mapsto \mathbf{R}_{\rightarrow}(w, u')} \\
\frac{u_1 \mapsto u'_1}{\mathbf{R}_{\times}(u_1, u_2) \mapsto \mathbf{R}_{\times}(u'_1, u_2)} \quad \frac{u \mapsto u}{\mathbf{R}_{\times}(w, u) \mapsto \mathbf{R}_{\times}(w, u')} \\
\frac{u \mapsto u'}{\mathbf{R}_R(u) \mapsto \mathbf{R}_R(u')} \\
\frac{u \mapsto u'}{\text{typecase } u (u_{int}, xy.u_{\rightarrow}, xy.u_{\times}, x.u_R) \mapsto \text{typecase } u' (u_{int}, xy.u_{\rightarrow}, xy.u_{\times}, x.u_R)}
\end{array}$$

B Analysis of Quantified Types

In an impredicative or recursive variant of λ_R it would be desirable to analyze quantified types. This modification can be done easily by limiting the analysis to the outermost operator that is, a **Typerec** or **typecase** will determine that a type is a \forall or \exists type, but will not provide any information about the body of the type. Experience from the TIL/ML compiler suggests that this solution might be enough in practice. However, in this section we briefly explore what would be required to extend λ_R to support full analysis of quantified types.

Adding term-level representations for quantified types is straightforward. The representation for $\forall\alpha:\kappa.c$ must carry information expressing how to produce the appropriate c given an appropriate α of kind κ . For example, a member of $R(\forall\alpha:\text{Type}.c)$ would be built by a constructor $R_{\forall\text{Type}}$ that takes a function with type $\forall\alpha:\text{Type}.R(\alpha) \rightarrow R(c)$. Representations of polymorphic types at any other kind κ would be built by an analogous constructor $R_{\forall\kappa}$ or $R_{\exists\kappa}$. Note that this solution requires an infinite collection of such constructors, two for each kind.

Analysis of types is also straightforward if quantification is restricted to **Type**, or to any finite set of kinds. In that case, the appropriate branches can be added to the **Typerec** and **typecase** operations, and each branch would return the appropriate representation function discussed above. General quantification could be handled, at the expense of additional complexity, by adding kind variables, term representations of kinds, and a **Kindrec** facility for analyzing such representations. With such additions, the quantifier branches would return a kind representation and the appropriate representation function.

However, with such a mechanism in place, the amount of useful analysis that can be performed is still quite limited. Type analyzing code may apply the representation function to an argument and analyze its output, but it could not analyze the function itself. This fact makes it impossible, for example, to print quantified types. The modal type theory of Despeyroux *et al.* [5] is intended for precisely this sort of application; it provides mechanisms for primitive recursion on higher-order syntax and should provide a solution to this problem.

Also, allowing any analysis of quantified types by **Typerec** sacrifices strong normalization of type expressions. (A correspondence between **Type** and **Type** \rightarrow **Type** can be built, permitting the encoding of the untyped lambda calculus.) A modal type discipline appears to solve this problem as well.

C Operational Semantics

$$\begin{array}{c}
(\lambda x:c.e)v \mapsto e[v/x] \\
(\Lambda\alpha:\kappa.v)[c] \mapsto v[c/\alpha] \\
\pi_1 \langle v_1, v_2 \rangle \mapsto v_1 \quad \pi_2 \langle v_1, v_2 \rangle \mapsto v_2 \\
(\mathbf{fix } f:c.v)v' \mapsto (v[\mathbf{fix } f:c.v/f])v' \\
(\mathbf{fix } f:c.v)[c] \mapsto (v[\mathbf{fix } f:c.v/f])[c] \\
\\
\text{unpack } \langle \alpha, x \rangle = (\text{pack } v \text{ as } \exists\beta.c_1 \text{ hiding } c_2) \\
\text{in } e_2 \mapsto e_2[c_2, v/\alpha, x] \\
\\
\text{typecase}[\delta.c] \mathbf{R}_{int} (e_{int}, \beta\gamma xy.e_{\rightarrow}, \\
\beta\gamma xy.e_{\times}, \beta x.e_R) \mapsto e_{int} \\
\\
\text{typecase}[\delta.c] (\mathbf{R}_{\rightarrow}(v_1, v_2)) (e_{int}, \beta\gamma xy.e_{\rightarrow}, \\
\beta\gamma xy.e_{\times}, \beta x.e_R) \mapsto e_{\rightarrow}[\mathcal{D}(v_1), \mathcal{D}(v_2), v_1, v_2/\beta, \gamma, x, y] \\
\\
\text{typecase}[\delta.c] (\mathbf{R}_{\times}(v_1, v_2)) (e_{int}, \beta\gamma xy.e_{\rightarrow}, \\
\beta\gamma xy.e_{\times}, \beta x.e_R) \mapsto e_{\times}[\mathcal{D}(v_2), \mathcal{D}(v_2), v_1, v_2/\beta, \gamma, x, y] \\
\\
\text{typecase}[\delta.c] (\mathbf{R}_R(v)) (e_{int}, \beta\gamma xy.e_{\rightarrow}, \\
\beta\gamma xy.e_{\times}, \beta x.e_R) \mapsto e_R[\mathcal{D}(v), v/\beta, x] \\
\\
\frac{e_1 \mapsto e'_1}{e_1 e_2 \mapsto e'_1 e_2} \quad \frac{e \mapsto e'}{ve \mapsto ve'} \quad \frac{e \mapsto e'}{e[c] \mapsto e'[c]} \\
\\
\frac{e \mapsto e'}{\pi_i e \mapsto \pi_i e'} \quad \frac{e_1 \mapsto e'_1}{\langle e_1, e_2 \rangle \mapsto \langle e'_1, e_2 \rangle} \quad \frac{e \mapsto e'}{\langle v, e \rangle \mapsto \langle v, e' \rangle} \\
\\
\frac{e \mapsto e'}{\text{pack } e \text{ as } \exists\beta.c_1 \text{ hiding } c_2 \mapsto \text{pack } e' \text{ as } \exists\beta.c_1 \text{ hiding } c_2} \\
\\
\frac{e \mapsto e'}{\text{unpack } \langle \alpha, x \rangle = e \text{ in } e_2 \mapsto \text{unpack } \langle \alpha, x \rangle = e' \text{ in } e_2} \\
\\
\frac{e \mapsto e'}{\text{typecase}[\delta.\sigma] e (e_{int}, \beta\gamma xy.e_{\rightarrow}, \beta\gamma xy.e_{\times}, \beta x.e_R) \mapsto \text{typecase}[\delta.\sigma] e' (e_{int}, \beta\gamma xy.e_{\rightarrow}, \beta\gamma xy.e_{\times}, \beta x.e_R)}
\end{array}$$

$$\frac{e_1 \mapsto e'_1}{\mathbf{R}_{\rightarrow}(e_1, e_2) \mapsto \mathbf{R}_{\rightarrow}(e'_1, e_2)}$$

$$\frac{e \mapsto e'}{\mathbf{R}_{\rightarrow}(v, e) \mapsto \mathbf{R}_{\rightarrow}(v, e')} \quad \frac{e_1 \mapsto e'_1}{\mathbf{R}_{\times}(e_1, e_2) \mapsto \mathbf{R}_{\times}(e'_1, e_2)}$$

$$\frac{e \mapsto e'}{\mathbf{R}_{\times}(v, e) \mapsto \mathbf{R}_{\times}(v, e')} \quad \frac{e \mapsto e'}{\mathbf{R}_R(e) \mapsto \mathbf{R}_R(e')}$$

D Static Semantics

D.1 Type formation

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

$$\frac{}{\Gamma \vdash \mathit{int} : \mathbf{Type}} \quad \frac{}{\Gamma \vdash \alpha : \kappa} \quad (\Gamma(\alpha) = \kappa)$$

$$\frac{\Gamma \vdash c_1 : \mathbf{Type} \quad \Gamma \vdash c_2 : \mathbf{Type}}{\Gamma \vdash c_1 \rightarrow c_2 : \mathbf{Type}}$$

$$\frac{\Gamma \vdash c_1 : \mathbf{Type} \quad \Gamma \vdash c_2 : \mathbf{Type}}{\Gamma \vdash c_1 \times c_2 : \mathbf{Type}}$$

$$\frac{\Gamma, \alpha : \kappa_1 \vdash c : \kappa_2}{\Gamma \vdash \lambda \alpha : \kappa_1. c : \kappa_1 \rightarrow \kappa_2}$$

$$\frac{\Gamma \vdash c_1 : \kappa_1 \rightarrow \kappa_2 \quad \Gamma \vdash c_2 : \kappa_1}{\Gamma \vdash c_1 c_2 : \kappa_2} \quad (\alpha \notin \mathit{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c : \mathbf{Type}}{\Gamma \vdash R(c) : \mathbf{Type}}$$

$$\frac{\Gamma \vdash c : \mathbf{Type} \quad \Gamma \vdash c_{\mathit{int}} : \kappa \quad \Gamma \vdash c_{\rightarrow} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \quad \Gamma \vdash c_{\times} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \quad \Gamma \vdash c_R : \mathbf{Type} \rightarrow \kappa \rightarrow \kappa}{\Gamma \vdash \mathbf{Type} \mathit{rec} \ c(c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R) : \kappa}$$

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

$$\frac{\Gamma \vdash c : \mathbf{Type}}{\Gamma \vdash c}$$

$$\frac{\Gamma \vdash \sigma_1 \quad \Gamma \vdash \sigma_2}{\Gamma \vdash \sigma_1 \times \sigma_2} \quad \frac{\Gamma \vdash \sigma_1 \quad \Gamma \vdash \sigma_2}{\Gamma \vdash \sigma_1 \rightarrow \sigma_2}$$

$$\frac{\Gamma, \alpha : \kappa \vdash \sigma}{\Gamma \vdash \forall \alpha : \kappa. \sigma} \quad (\alpha \notin \mathit{Dom}(\Gamma)) \quad \frac{\Gamma, \alpha : \kappa \vdash \sigma}{\Gamma \vdash \exists \alpha : \kappa. \sigma} \quad (\alpha \notin \mathit{Dom}(\Gamma))$$

D.2 Type Equivalence (selected rules)

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

$$\frac{\Gamma, \alpha : \kappa' \vdash c_1 : \kappa \quad \Gamma \vdash c_2 : \kappa'}{\Gamma \vdash (\lambda \alpha : \kappa'. c_1) c_2 = c_1 [c_2 / \alpha] : \kappa} \quad (\alpha \notin \mathit{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c : \kappa_1 \rightarrow \kappa_2}{\Gamma \vdash \lambda \alpha : \kappa_1. c \alpha = c : \kappa_1 \rightarrow \kappa_2} \quad (\alpha \notin \mathit{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash c_{\mathit{int}} : \kappa \quad \Gamma \vdash c_{\rightarrow} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \quad \Gamma \vdash c_{\times} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \quad \Gamma \vdash c_R : \mathbf{Type} \rightarrow \kappa \rightarrow \kappa}{\Gamma \vdash \mathbf{Type} \mathit{rec}(c_{\mathit{int}})(c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R) = c_{\mathit{int}} : \kappa}$$

$$\frac{\Gamma \vdash c_1 : \mathbf{Type} \quad \Gamma \vdash c_2 : \mathbf{Type} \quad \Gamma \vdash c_{\mathit{int}} : \kappa \quad \Gamma \vdash c_{\rightarrow} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \quad \Gamma \vdash c_{\times} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \quad \Gamma \vdash c_R : \mathbf{Type} \rightarrow \kappa \rightarrow \kappa}{\Gamma \vdash \mathbf{Type} \mathit{rec}(c_1 \rightarrow c_2)(c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R) = c_{\rightarrow} c_1 c_2 (\mathbf{Type} \mathit{rec} c_1 (c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R)) (\mathbf{Type} \mathit{rec} c_2 (c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R)) : \kappa}$$

$$\left\{ \begin{array}{l} \Gamma \vdash \mathbf{Type} \mathit{rec}(c_1 \times c_2)(c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R) = c_{\times} c_1 c_2 (\mathbf{Type} \mathit{rec} c_1 (c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R)) (\mathbf{Type} \mathit{rec} c_2 (c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R)) : \kappa \\ \Gamma \vdash \mathbf{Type} \mathit{rec}(c_1 \times c_2)(c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R) = c_{\times} c_1 c_2 (\mathbf{Type} \mathit{rec} c_1 (c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R)) (\mathbf{Type} \mathit{rec} c_2 (c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R)) : \kappa \end{array} \right.$$

$$\frac{\Gamma \vdash c : \mathbf{Type} \quad \Gamma \vdash c_{\mathit{int}} : \kappa \quad \Gamma \vdash c_{\rightarrow} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \quad \Gamma \vdash c_{\times} : \mathbf{Type} \rightarrow \mathbf{Type} \rightarrow \kappa \rightarrow \kappa \rightarrow \kappa \quad \Gamma \vdash c_R : \mathbf{Type} \rightarrow \kappa \rightarrow \kappa}{\Gamma \vdash \mathbf{Type} \mathit{rec}(R(c))(c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R) = c_R c (\mathbf{Type} \mathit{rec} c (c_{\mathit{int}}, c_{\rightarrow}, c_{\times}, c_R)) : \kappa}$$

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

$$\frac{\Gamma \vdash \sigma_1 = \sigma_2 : \kappa}{\Gamma \vdash \sigma_1 = \sigma_2}$$

D.3 Term Formation

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

$$\frac{}{\Gamma \vdash i : \mathit{int}} \quad \frac{}{\Gamma \vdash x : \sigma} \quad (\Gamma(x) = \sigma)$$

$$\frac{\Gamma \vdash e_1 : \sigma_1 \quad \Gamma \vdash e_2 : \sigma_2}{\Gamma \vdash \langle e_1, e_2 \rangle : \sigma_1 \times \sigma_2}$$

$$\frac{\Gamma \vdash e : \sigma_1 \times \sigma_2}{\Gamma \vdash \pi_1 e : \sigma_1} \quad \frac{\Gamma \vdash e : \sigma_1 \times \sigma_2}{\Gamma \vdash \pi_2 e : \sigma_2}$$

$$\frac{\Gamma \vdash e_1 : \sigma_2 \rightarrow \sigma_1 \quad \Gamma \vdash e_2 : \sigma_2}{\Gamma \vdash e_1 e_2 : \sigma_1}$$

Term Formation (continued)

$$\frac{\Gamma \vdash e : \forall \alpha : \kappa . \sigma \quad \Gamma \vdash c : \kappa}{\Gamma \vdash e[c] : \sigma[c/\alpha]} \quad \frac{\Gamma, x : \sigma_2 \vdash e : \sigma_1 \quad \Gamma \vdash \sigma_2}{\Gamma \vdash \lambda x : \sigma_2 . e : \sigma_2 \rightarrow \sigma_1} \quad (x \notin \text{Dom}(\Gamma)) \quad \frac{\Gamma, \alpha : \kappa \vdash e : \sigma}{\Gamma \vdash \Lambda \alpha : \kappa . e : \forall \alpha : \kappa . \sigma} \quad (x \notin \text{Dom}(\Gamma))$$

$$\frac{\Gamma, \alpha : \kappa \vdash \sigma_1 : \text{Type} \quad \Gamma \vdash \sigma_2 : \kappa \quad \Gamma \vdash e : \sigma_1[\sigma_2/\alpha]}{\Gamma \vdash \text{pack } e \text{ as } \exists \alpha : \kappa . \sigma \text{ hiding } \sigma_2 : \exists \alpha : \kappa . \sigma_1} \quad (\alpha \notin \text{Dom}(\Gamma))$$

$$\frac{\Gamma \vdash e_1 : \exists \alpha : \kappa . \sigma_2 \quad \Gamma, \alpha : \kappa, x : \sigma_2 \vdash e_2 : \sigma_1}{\Gamma \vdash \text{unpack } \langle \alpha, x \rangle = e_1 \text{ in } e_2 : \sigma_1} \quad (\alpha, x \notin \text{Dom}(\Gamma))$$

$$\frac{\Gamma, f : \sigma \vdash e : \sigma \quad \Gamma \vdash \sigma}{\Gamma \vdash \text{fix } f . e : \sigma} \quad (f \notin \text{Dom}(\Gamma), \sigma = \forall \alpha_1 : \kappa_1 \cdots \alpha_n : \kappa_n . \sigma_1 \rightarrow \sigma_2, n \geq 0)$$

$$\frac{}{\Gamma \vdash R_{\text{int}} : R(\text{int})} \quad \frac{\Gamma \vdash e_1 : R(c_1) \quad \Gamma \vdash e_2 : R(c_2)}{\Gamma \vdash R_{\rightarrow}(e_1, e_2) : R(c_1 \rightarrow c_2)} \quad \frac{\Gamma \vdash e_1 : R(c_1) \quad \Gamma \vdash e_2 : R(c_2)}{\Gamma \vdash R_{\times}(e_1, e_2) : R(c_1 \times c_2)}$$

$$\frac{\Gamma \vdash e : R(c)}{\Gamma \vdash R_R(e) : R(R(c))} \quad \frac{\Gamma \vdash e : \sigma_2 \quad \Gamma \vdash \sigma_1 = \sigma_2}{\Gamma \vdash e : \sigma_1}$$

$$\frac{\Gamma \vdash e : R(\text{int}) \quad \Gamma \vdash e_{\text{int}} : \sigma[\text{int}/\delta]}{\Gamma \vdash \text{typecase}[\delta . \sigma] e (e_{\text{int}}, \beta \gamma xy . e_{\rightarrow}, \beta \gamma xy . e_{\times}, \beta x . e_R) : \sigma[\text{int}/\delta]}$$

$$\frac{\Gamma \vdash e : R(c_1 \rightarrow c_2) \quad \Gamma, x : R(c_1), y : R(c_2) \vdash e_{\rightarrow}[c_1, c_2/\beta, \gamma] : \sigma[c_1 \rightarrow c_2/\delta]}{\Gamma \vdash \text{typecase}[\delta . \sigma] e (e_{\text{int}}, \beta \gamma xy . e_{\rightarrow}, \beta \gamma xy . e_{\times}, \beta x . e_R) : \sigma[c_1 \rightarrow c_2/\delta]}$$

$$\frac{\Gamma \vdash e : R(c_1 \times c_2) \quad \Gamma, x : R(c_1), y : R(c_2) \vdash e_{\times}[c_1, c_2/\beta, \gamma] : \sigma[c_1 \times c_2/\delta]}{\Gamma \vdash \text{typecase}[\delta . \sigma] e (e_{\text{int}}, \beta \gamma xy . e_{\rightarrow}, \beta \gamma xy . e_{\times}, \beta x . e_R) : \sigma[c_1 \times c_2/\delta]}$$

$$\frac{\Gamma \vdash e : R(R(c)) \quad \Gamma, x : R(c) \vdash e_R[c/\beta] : \sigma[R(c)/\delta]}{\Gamma \vdash \text{typecase}[\delta . \sigma] e (e_{\text{int}}, \beta \gamma xy . e_{\rightarrow}, \beta \gamma xy . e_{\times}, \beta x . e_R) : \sigma[R(c)/\delta]}$$

$$\frac{\begin{array}{l} \Gamma, \alpha : \text{Type}, \Gamma' \vdash e : R(\alpha) \quad \Gamma(\Gamma'[\text{int}/\alpha]) \vdash e_{\text{int}}[\text{int}/\alpha] : \sigma[\text{int}, \text{int}/\alpha, \delta] \\ \Gamma, \beta : \text{Type}, \gamma : \text{Type}, (\Gamma'[\beta \rightarrow \gamma/\alpha]), x : R(\beta), y : R(\gamma) \vdash e_{\rightarrow}[\beta \rightarrow \gamma/\alpha] : \sigma[\beta \rightarrow \gamma, \beta \rightarrow \gamma/\alpha, \delta] \\ \Gamma, \beta : \text{Type}, \gamma : \text{Type}, (\Gamma'[\beta \times \gamma/\alpha]), x : R(\beta), y : R(\gamma) \vdash e_{\times}[\beta \times \gamma/\alpha] : \sigma[\beta \times \gamma, \beta \times \gamma/\alpha, \delta] \\ \Gamma, \beta : \text{Type}, (\Gamma'[R(\beta)/\alpha]), x : R(\beta) \vdash e_R[R(\beta)/\alpha] : \sigma[R(\beta), R(\beta)/\alpha, \delta] \end{array}}{\Gamma, \alpha : \text{Type}, \Gamma' \vdash \text{typecase}[\delta . \sigma] e (e_{\text{int}}, \beta \gamma xy . e_{\rightarrow}, \beta \gamma xy . e_{\times}, \beta x . e_R) : \sigma[\alpha/\delta]} \quad (\alpha, \beta, \gamma \notin \text{Dom}(\Gamma, \Gamma'))$$

$$\frac{\begin{array}{l} \Gamma \vdash e : R(c) \quad \Gamma \vdash e_{\text{int}} : \sigma[\text{int}/\delta] \\ \Gamma, \beta : \text{Type}, \gamma : \text{Type}, x : R(\beta), y : R(\gamma) \vdash e_{\rightarrow} : \sigma[\beta \rightarrow \gamma/\delta] \\ \Gamma, \beta : \text{Type}, \gamma : \text{Type}, x : R(\beta), y : R(\gamma) \vdash e_{\times} : \sigma[\beta \times \gamma/\delta] \\ \Gamma, \beta : \text{Type}, x : R(\beta) \vdash e_R : \sigma[R(\beta)/\delta] \end{array}}{\Gamma \vdash \text{typecase}[\delta . \sigma] e (e_{\text{int}}, \beta \gamma xy . e_{\rightarrow}, \beta \gamma xy . e_{\times}, \beta x . e_R) : \sigma[c/\delta]} \quad (\beta, \gamma \notin \text{Dom}(\Gamma, \Gamma'))$$