

1 Introduction

Functional languages are gaining support throughout the academic world due to their strong mathematical basis and amenity to proof. Many different languages are available – for example, Miranda¹, Hope, Ponder, Pebble, ML, SML, and Haskell – all based on the lambda calculus. Given their common basis, an intermediate language can be designed to meet the following objectives:

- Span a large portion of the spectrum between a users program and the machine it is running on thus relieving a large part of the compilation burden for new (experimental) languages
- Provide a standard intermediate representation for portability and to facilitate programming environments which handle several languages.
- Provide an intermediate compiler which may be proven correct with the view to future “total system” correctness. (The focus of current research in the U of C’s VLSI group is towards proving the correctness of microchip designs. In conjunction with a proven compiler, this would be a major step towards a system in which both hardware and software are verified).
- Provide a simple yet expressive language with which to do research into functional language issues. It should provide mechanisms to assist research without the complexities of a high level language.

In order to define this intermediate functional language, the complete semantics must be worked out so that the meaning of each statement is well defined. As well, it is possible to find semantics for the translation of the intermediate language, which will tie in well with current research in complete system correctness.

2 Background

Following are some justifications for using functional languages:

- Functional languages are all based on the lambda calculus. By definition there exists an intermediate language for them – the lambda calculus. Further, the expressive power of the lambda calculus is well known.
- The semantics of the lambda calculus have been completely defined and are relatively easy. Thus, the semantics of an intermediate language (which will be based on lambda) will benefit.
- High level functional languages and the lambda calculus are amenable to proof. In particular, the pattern matching explained below matches the form of proofs done by *structural induction*. As well, since no expression in a functional language can cause a side effect, these types of complications are removed from proofs.
- Large amounts of recent research has looked at functional languages and machines for them. The direct correlation between functional languages and functional architectures is mathematically nice in that the properties of the lambda calculus span the boundary. Current architectures for functional languages tend to be slow, but faster machines are on the horizon. The Church-Rosser property holds for the lambda calculus, and thus parallel implementations are also possible.

2.1 Intermediate Functional Languages

Current intermediate functional languages (for, example [6],[16]) can be characterized as:

- They exist at the level of the typed lambda calculus with lots of builtin functions and very few “high-level” structures. This means:

¹Miranda is a trademark of Research Software Ltd.

- The languages cannot themselves be typechecked but they rely on well-typed expressions. That is, the only expressions which have meanings in these languages are those which have been type-checked by some external agent.
- They are difficult source languages to program in. For a researcher who needs access to low level structures but does not want to be burdened by “assembly-like” programming these languages are insufficient.
- They ignore the fact that high level functional languages are interactive. This means that environment issues must be addressed elsewhere.

But, these languages form an important part of a system in that:

- They provide the lowest level representation from which a designer would like to generate machine code. This is because the “higher” level structures they contain allow optimizations for machine code.
- They contain a complete set of primitive functions for implementing high level functional languages.
- They allow many compilation functions to be performed as source to source transformations. Performing dependency analysis, strictness analysis, lazy lambda lifting and other functions can be performed directly on the intermediate language. (These issues are discussed in [15]).

This paper deals with a design for an intermediate language at a higher level than those discussed above. This language is designed to:

- Provide a wide array of structures and primitives so that high level languages are easily translated. The composition of these structures and primitives will support high-level structures with a minimum of effort, but will not limit the generality of the high-level structures.
- Provide a type system which will be sufficient for most applications and which can be sidestepped for those applications which require a different system.
- Handle the interactive features of high level languages so that binding in the environment is well modelled.
- Use existing intermediate languages (or a slight variation of them) as a target language. In this way, the primitives from the existing languages will help to implement the structures in the new intermediate language.

The most important consideration in the design of this language is: “Which structures can be contained in the intermediate language without a loss of generality to the structures in high level languages?” The next section works through an example of a high level expression in order to motivate an answer to this question.

2.2 Introduction to Functional Expressions

Consider the following expression (in a generic high level functional language syntax):

```
datatype list *      := nil | cons * (list *)
length nil           == 0
length (cons a l)    == 1 + (length l)
```

This is the declaration of the type `list *` and the function `length` on lists. A list is either `nil` or it is the `cons` of an item of type `*` onto a list of type `list *`. The `*` is a type variable, meaning that the type `list` is polymorphic. We may have items of type `list int`, `list (list *)`, `list char`,....

`length` is defined using *pattern matching* – each constructor in the type of `list` is given a case in the definition of the function (this corresponds well with proofs done by *structural induction* for the correctness of programs). Thus, if the argument to `length` is `nil`, the first expression will be evaluated. Otherwise, the argument must be a `cons`, so the second expression will be evaluated.

If the argument is not `nil` it must be `cons` since, in general, high level functional languages are statically typechecked – the type of an expression is computed at compile-time. The type of `length` is a function from `list *` to `int`, usually written `length: list * → int`. Thus, if the first defining equation for `length` has a parameter of type `list *`, the following equations must have a type at least as specific as `list *` (a more specific type, such as `list int` is also acceptable).

Programming occurs in an *interactive* environment. After each function is typed in, it is compiled or interpreted. It is important to note that several interactive commands may be given before a function is defined (as in the case of `length`).

In order to compile the `length` example, we must have some representation of the type `list *`. It is sufficient to have a representation for both `nil` and `cons` since together they totally define `list`:

```
nil          ==> (list,0)
cons * (list *) ==> (list,1,*,list *)
```

Thus, `cons 3 nil` would be represented as: `(list,1,3,(list,0))`. In general, types can be *packed* in this manner, and then specific parts of the type may be *selected* at will.

Compiling the `length` function will require a function to select an expression based on which parameter matches the argument. Given the representation of `list` this would look like:

```
length n == if n = (list,0) then 1 else
              if n = (list,1,a,l) then 1 + (length l)
              else error
```

In this case `error` will never be encountered (since `length` was typechecked), but incomplete pattern matches are possible. This definition contains no pattern matching – the matching has been replaced by equality checks on the representation of lists. This process is known as *pattern matching compilation*.

The next step in compiling `length` is to compile it into the lambda calculus, from which the back end is generated. These steps are completely covered in [15].

2.3 Characteristics of Functional Languages

A review of the ideas touched on above and an introduction to some other characteristics of functional languages is given below:

- *Data types and pattern matching* as introduced above.
- *Strongly typed* expressions are the only ones accepted. Types are *polymorphic* and may be *explicit* (supplied by the user) or *implicit* (inferred by the compiler). Explicit systems allow more general types than implicit ones, but lack the user friendliness that inferring types provides. Most widely used languages (ie. Miranda and SML) have implicit type systems although considerable recent research has been into explicit systems (for example, [4]).
- *List comprehensions* are akin to the declaration of mathematical sets. For example,

```
[n * n | n < 100]
```

is the list containing the squares of all the integers less than 100. These structures allow for very succinct clear programming. Consider this quicksort from [15]:

```
qs nil == nil
qs (cons a l) == cons (qs [b | mem b l; b <= a]) (cons a (qs [b | mem b l; b > a]))
```

where `mem` is a function which returns `true` if its first argument is included in the list which is its second argument.

- *Modules* are user friendly and useful when programming large systems. SML [9] and Pebble [1] are two functional languages which take radically different approaches to modules. Pebble treats modules as values and thus types become first class objects. SML maintains the ability to infer types by raising a strict distinction between functions and modules. Both these systems can be viewed as experimental since neither has been widely accepted yet. Other functional languages, such as Miranda and Haskell do not contain modules.
- *Arrays* are desirable in a functional language if sharing analysis is done [11], since they provide constant access time. With no sharing analysis each update of an array would require making a copy of the entire array.
- *Exceptions* allow errors to be flagged so that functions can recover. For example:

```
exception div0;
div x 0 == raise div0
div x y == x / y;
```

The exception `div0` can then be *handled* by functions which wish to provide a default value or print out an error message.

- *Lazy evaluation* ensures that expressions which can be evaluated are. Strictness analysis allows some functions to be called by value, and thus allows parallelism. It also increases the speed of non-parallel machines [22].

3 An Intermediate Functional Language (IFL)

My goal is to define an intermediate functional language, IFL, which incorporates as many of the above structures as possible without being too specialized.

The following priorities for designing and implementing a language (given in order of importance) reflect my views on the importance of having a well defined language before it is implemented:

- *Design*: The language must be designed to meet specific applications. A useful subset of programming languages must be defined. In this case I have limited myself to lazy functional languages and have outlined the major features which are necessary in the design. This subset is mathematically clean and the implementation is not overly complex. An abstract syntax must be given for the language, listing the components of the language and their relation to each other. It does not include syntactic information or precedence information. The abstract syntax is sufficient for defining the semantics of the language.
- *Semantics*: Before any implementation is done the semantics of the language must be totally defined. If this is done, all implementations should give identical results for any program. Semantics also give important feedback on the design of a language – if the semantics are complicated this is an indication that the language may be too complex. At this stage it is also worth designing a concrete syntax which specifies syntactic and precedence constraints on programs. If this is done, the language will be totally defined before it is implemented.
- *Implementation*: Once the language is totally defined, a correct implementation, one which maintains the semantics of expressions, can be devised. A proof of the correctness would be advantageous.
- *Optimization*: Any optimization which improves performance without changing the semantics can be added once the implementation is known to be correct.

The following design decisions were incorporated into IFL:

- Data types and pattern matching are the most pervasive feature of recent functional languages. A general type definition and pattern matching system, akin to that of Miranda, ML, SML, Hope, and others has been included.

- Typechecking will be the most general implicit system based on the one developed by Milner [14] and outlined in [15]. Explicit types will be allowed through annotation so that the user has control of types, if he so desires. This also means that more powerful explicit types can be given to IFL expressions, sidestepping the inferencing algorithm.
- The interactive environment will be explicit in IFL. There is a distinction between commands which cause bindings and those which evaluate to some value.
- Modules are not yet included due to the lack of agreement as to their definition. A semantic description of modules is being worked on, and primitives for implementing them may be added later. If, in the meantime, some agreement as to module definition is arrived at, this will be incorporated. (Semantically, a module is simply a set of bindings to be added or saved from the environment. The problem is in enforcing types between modules in a general manner. So far the two solutions to this problem are exemplified by SML and Pebble as discussed earlier).
- List comprehensions can be transformed into iteration over lists. With the appropriate functions defined in the original environment for IFL, transformations for list comprehensions will be easy. This will allow a very general class of list comprehensions, whereas defining a list comprehension structure directly in IFL may lead to some limitations.
- Exceptions are included in IFL by using the builtin operator `fail` and lazy evaluation. (The failures which are simulated are not as general as those in SML, but SML is not purely functional and uses some imperative features in the implementation of its exceptions. The method of translating exceptions into a functional language is given in [21]).
- Arrays will not be included in the first version of IFL although it is expected that they will be included in the final version. (This is because sharing analysis is a difficult topic which would simply complicate the original version).
- Lazy evaluation with strictness analysis will be the evaluation mechanisms for IFL.

3.1 Abstract Syntax

The BNF for IFL has been organized in order to highlight the difference between commands which evaluate, `exp`, and those which cause bindings in the environment, `bind`. In the definition in Figure 1, x^+ means one or more occurrences of x , x^* means zero or more occurrences of x .

3.2 Informal Description of IFL

An ifl term can be an expression, `exp`, which evaluates to a value in a given environment, or a binding, `bind`, which enriches the environment for subsequent evaluations, or a series of these two, `ifl*`.

A **pattern** is a match for a builtin type. Examples of patterns will be shown as we progress. Constants include integer operations, real operations (prefixed by `r`), boolean operations, and `fail`.

A binding is one of the following:

- A *let* binding causes each variable in the pattern on the left-hand side to be bound to the corresponding value on the right-hand side. Multiple bindings are done in parallel. For example:

`let x = 3, y = 4`

binds 3 to `x` and 4 to `y`. But,

`let x = 3, y = (+ x 1)`

binds 3 to `x` and produces an error (if `x` was not previously defined) for the `y` binding. If an error occurs in any part of a binding, none of the variables become bound.

ifl	::= bind exp ifl ⁺
bind	::= letbind <i>lettype</i> tyname = type,...,tyname = type <i>abstype</i> tyname = type,...,tyname = type <i>in</i> letbind
letbind	::= <i>let</i> pattern = exp,...,pattern = exp <i>letrec</i> pattern = exp,...,pattern = exp
exp	::= var const exp exp annot exp number string λ pattern ⁺ .exp <i>let</i> pattern = exp,...,pattern = exp <i>in</i> exp <i>letrec</i> pattern = exp,...,pattern = exp <i>in</i> exp <i>case</i> var <i>of</i> pattern => exp,...,pattern => exp \square exp exp (exp)
pattern	::= const pattern* var constructor pattern*
tyname	::= var constructor tyname*
type	::= tyname knowntype var constructor type*
constructor	::= <i>sum</i> <i>prod</i> <i>fun</i> var
knowntype	::= <i>int</i> <i>bool</i> <i>char</i> <i>string</i> <i>real</i> tyname
annot	::= var var => exp ... var => exp type
const	::= + - * / r+ r- r* r/ <i>if</i> <i>or</i> <i>and</i> <i>xor</i> <i>true</i> <i>false</i> <i>fail</i> num string
var	::= <i>alphabetic alphanumeric</i> *
num	::= digit ⁺ -num num.num num e num
digit	::= 0..9
string	::= " <i>alphanumeric</i> *

Figure 1: Abstract Syntax for IFL

- A *letrec* binding works in the same manner as the *let* except that definitions may be recursive or mutually recursive. For example:

$$\text{letrec } f = \lambda x. \text{ if } (= x 0) 1 (+ x (f (- x 1)))$$

- A *lettype* binding declares a new type. An enumerated type is simply the sum of the possible values.

$$\text{lettype colors} = \text{sum } (\text{sum } (\text{green}, \text{blue}), \text{red})$$

A data type can be more complicated and may contain type variables.

$$\begin{aligned} \text{lettype list } a &= \text{sum } (\text{nil}, \text{cons } a (\text{list } a)) \\ \text{lettype pair} &= \text{prod } (a, b) \end{aligned}$$

When more than one type definition is given they are assumed to be dependent on each other.

- A *abstype* binding binds the operators for an abstract type. For example, a stack could be defined as:

$$\begin{aligned} \text{abstype stack } a &= \text{sum } (\text{nils}, \text{element } a (\text{stack } a)) \text{ in} \\ \text{let push} &= \lambda a \ s. \text{ element } a \ s, \text{ pop} = \lambda (\text{element } a \ s). a \end{aligned}$$

An exp can be a variable, a constant, an application, a number, a string or one of the following:

- A lambda expression is a function from one or more patterns to the result of the body.

$$\begin{aligned} \text{let head} &= \lambda (\text{cons } x \ y). x \\ \text{let tail} &= \lambda (\text{cons } x \ y). y \end{aligned}$$

Here are the first examples of patterns in IFL. The $\text{cons } x \ y$ matches our builtin type list. Thus, *head* could be applied to any list which is not nil.

- A *let* expression executes the body of the *let* in the environment given by the *let* binding. The difference between the binding and the expression is that the binding is permanent. This also holds for the *letrec* expression.
- The *case* expression allows pattern matching on sum types.

$$\begin{aligned} \text{let head} &= \lambda x. \text{ case } x \text{ of } (\text{cons } a \ b) => a, \text{ nil} => \text{fail} \\ \text{let tail} &= \lambda x. \text{ case } x \text{ of } (\text{cons } a \ b) => b, \text{ nil} => \text{fail} \end{aligned}$$

These would be more accurate definitions of *head* and *tail*.

- The \square expression executes the first expression. If this evaluates, the second expression is ignored. If the first expression fails, the result of the second is returned.

$$\text{let tail2} = \lambda x. \square (\text{tail } x) \text{ nil}$$

- Annotated expressions allow extra information to be given to the compiler. At the present time annotations may contain strictness and typing information. In the future, sharing, concurrency, and other annotations may also be added.

3.3 Types

The number of different type systems which have been developed is immense (see [18] or [3] for a review of those of interest here). The Milner type system is one of the most widely used since there is a well known algorithm for inferring the types of expressions. The type system I develop here is an extension of the Milner system to IFL with many ideas from the papers cited above.

The usual method for specifying type inference is through a series of inference rules. The *premisses* consist of a *type assignment* which maps (possibly free) identifiers onto types. The following variables are replaced by expressions or types as appropriate to create an instance of a rule.

Variable	Meaning
m	type assignment
$v \ v_1 \dots v_n$	variables
$c \ c_1 \dots c_n$	constructors
$k \ k_1 \dots k_n$	constants
$e \ e_1 \dots e_n$	expressions
$p \ p_1 \dots p_n$	patterns
n	integers ($n \geq 0$)
$t \ t_1 \dots t_n$	types
$tv_1 \dots tv_n$	type variables

A *type scheme* is a type containing type variables. A *principle typing* of an expression e is a type scheme which can be instantiated to any valid type for e .

A substitution is given by $\delta, \delta_1, \dots, \delta_n$. $\delta(t)$ is the result of applying the substitution δ to the type t . $\delta(m)$ is the result of applying the substitution δ to all types assigned by m . Unification, ϑ , finds the most general substitution for two types. That is, $\vartheta(\delta, t_1, t_2) = \delta_1$ where $\delta_1(t_1) = \delta_1(t_2)$. If unification fails, the typing will fail. In Figure 2 the inference rules are given for typing IFL expressions. δm is used as a shorthand for $\delta(m)$.

Figure 3 gives two examples of how the typechecking rules are used to infer the types of expressions.

VAR	$\overline{m \vdash v : t}$ when m assigns t to v
CST	$\overline{\vdash k : t}$ when t is the type of k
APP	$\frac{\delta m \vdash e_1 : \delta(t \rightarrow t_1) \quad \delta m \vdash e_2 : \delta t'}{\delta_1 m \vdash e_1 e_2 : \delta_1 t_1} \text{ where } \delta_1 = \vartheta(\delta, t, t')$
FUN	$\frac{m, p : t, v_1 : t_1, \dots, v_n : t_n \vdash e : t_{n+1}}{m \vdash \lambda p. e : t \rightarrow t_{n+1}}$ <p>when $v_1 \dots v_n$ are the variables in the pattern p</p>
LET	$\frac{m \vdash e_1 : t \quad m, p : t \vdash e_2 : t_1}{m \vdash \text{let } p = e_1 \text{ in } e_2 : t_1}$
CSE	$\frac{\delta m, v : \delta t, p_1 : \delta t \vdash e_1 : \delta t_1 \quad \dots \quad \delta m, v : \delta t, p_n : \delta t \vdash e_n : \delta t_n}{\delta_1 m \vdash \text{case } v \text{ of } p_1 \Rightarrow e_1 \dots p_n \Rightarrow e_n : \delta_1 t_1} \text{ where } \delta_1 = \vartheta(\dots \vartheta(\delta, t_1, t_2), \dots, t_n)$
FIX	$\frac{\delta m, p_1 : tv_1, \dots, p_n : tv_n \vdash e_1 : \delta t_1 \quad \dots \quad \delta m, p_1 : tv_1, \dots, p_n : tv_n \vdash e_n : \delta t_n}{\delta_1 m, p_1 : \delta_1 tv_1, \dots, p_n : \delta_1 tv_n \vdash e_1 : \delta_1 t_1}$ <p>where tv_1, \dots, tv_n are new type variables not occurring in δm</p> $\frac{\delta_1 m, p_1 : \delta_1 tv_1, \dots, p_n : \delta_1 tv_n \vdash e_1 : \delta_1 t_1 \quad \dots \quad \delta_1 m, p_1 : \delta_1 tv_1, \dots, p_n : \delta_1 tv_n \vdash e_n : \delta_1 t_n}{\delta_1 m \vdash \text{fix } p_1 = e_1, \dots, p_n = e_n \text{ in } e : t} \text{ where } \delta_1 = \vartheta(\dots(\vartheta(\delta, t_1, tv_1), \dots, t_n, tv_n))$
REC	$\frac{m, p_1 : t_1, e_1 : t_1, \dots, p_n : t_n, e_n : t_n \vdash e : t}{m \vdash \text{letrec } p_1 = e_1, \dots, p_n = e_n \text{ in } e : t}$
FTB	$\frac{\delta m \vdash e_1 : \delta t_1 \quad \delta m \vdash e_2 : \delta t_2}{\delta_1 m \vdash \Box e_1 e_2 : \delta_1 t_1} \text{ where } \delta_1 = \vartheta(\delta, t_1, t_2)$

Figure 2: Type Inference Rules

Let the environment m contain:

$0 : int, 1 : int, 2 : int, \dots$
 $= : int \rightarrow int \rightarrow bool$
 $- : int \rightarrow int \rightarrow int$
 $if : bool \rightarrow tv_1 \rightarrow tv_1 \rightarrow tv_1$

Let the substitution function δ be empty. Let $e_1 = \lambda n. if (= n 0) 1 (fac(-n 1))$. Let $FAC = letrec\ fac = e_1\ in\ (fac\ 3)$. Assume that each right-hand side becomes part of the premisses for rules following it. The typing for FAC proceeds as follows:

VAR	$\delta m, n : tv_2$	$\vdash n : tv_2$
APP	$\delta m, n : tv_2$	$\vdash (-n) : int \rightarrow int\ \text{and}\ \delta_1 = (\delta\ \text{and}\ tv_2 = int)$
APP	$\delta_1 m, n : int$	$\vdash (-n\ 1) : int$
APP	$\delta_1 m, n : int$	$\vdash (= n) : int \rightarrow bool$
APP	$\delta_1 m, n : int$	$\vdash (= n\ 0) : bool$
APP	$\delta_1 m, n : int$	$\vdash (if(= n 0)) : tv_1 \rightarrow tv_1 \rightarrow tv_1$
APP	$\delta_1 m, n : int$	$\vdash (if(= n 0) 1) : int \rightarrow int\ \text{and}\ \delta_2 = (\delta_1\ \text{and}\ tv_1 = int)$
VAR	$\delta_2 m, n : int, fac : tv_3$	$\vdash fac : tv_3$
APP	$\delta_2 m, n : int, fac : tv_3$	$\vdash (fac(-n 1)) : int\ \text{and}\ \delta_3 = (\delta_2\ \text{and}\ tv_3 = int \rightarrow tv_4)$
APP	$\delta_3 m, n : int, fac : tv_4 \rightarrow int$	$\vdash (if(= n 0) 1 (fac(-n 1))) : int$
FUN	$\delta_3 m, n : int, fac : tv_4 \rightarrow int$	$\vdash e_1 : int \rightarrow int$
FIX	$\delta_3 m, fac : tv_4 \rightarrow int$	$\vdash e_1 : int \rightarrow int\ \text{and}\ \delta_4 = (\delta_3\ \text{and}\ tv_4 = int)$
APP	$\delta_4 m, fac : int \rightarrow int$	$\vdash (fac\ 3) : int$
REC	$\delta_4 m$	$\vdash FAC : int$

Assume m also contains $cons : tv_1 \rightarrow (list\ tv_1) \rightarrow (list\ tv_1)$ and $nil : (list\ tv_1)$.

The typing for $letrec\ length = \lambda l. case\ l\ of\ nil \Rightarrow 0, (cons\ a\ x) \Rightarrow +1\ (length\ x)$:

APP	δm	$\vdash (+1) : int \rightarrow int$
VAR	$\delta m, a : tv_2$	$\vdash a : tv_2$
VAR	$\delta m, a : tv_2, x : tv_3$	$\vdash x : tv_3$
APP	$\delta m, a : tv_2, x : tv_3$	$\vdash (cons\ a) : (list\ tv_2) \rightarrow (list\ tv_2)\ \text{and}\ \delta_1 = (\delta\ \text{and}\ tv_1 = tv_2)$
APP	$\delta_1 m, a : tv_2, x : tv_3$	$\vdash (cons\ a\ x) : list\ tv_3\ \text{and}\ \delta_2 = (\delta_1\ \text{and}\ tv_2 = tv_3)$
VAR	$\delta_2 m, a : tv_3, x : tv_3, l : tv_4$	$\vdash l : tv_4$
VAR	$\delta_2 m, a : tv_3, x : tv_3,$ $l : tv_4, length : tv_5$	$\vdash length : tv_5$
APP	$\delta_2 m, a : tv_3, x : tv_3,$ $l : tv_4, length : tv_5$	$\vdash (length\ x) : tv_6\ \text{and}\ \delta_3 = (\delta_2\ \text{and}\ tv_5 = tv_3 \rightarrow tv_6)$
APP	$\delta_3 m, a : tv_3, x : tv_3,$ $l : tv_4, length : tv_3 \rightarrow tv_6$	$\vdash (+1\ (length\ x)) : int\ \text{and}\ \delta_4 = (\delta_3\ \text{and}\ tv_6 = int)$
CSE	$\delta_4 m, l : tv_4,$ $length : tv_3 \rightarrow int$	$\vdash (case...) : int\ \text{and}\ \delta_5 = (\delta_4\ \text{and}\ tv_3 = list\ tv_7)$
LAM	$\delta_5 m, l : tv_4,$ $length : list\ tv_7 \rightarrow int$	$\vdash (\lambda l. case...) : list\ tv_7 \rightarrow int\ \text{and}\ \delta_6 = (\delta_5\ \text{and}\ tv_4 = list\ tv_7)$
FIX	$\delta_6 m, length : list\ tv_7 \rightarrow int$	$\vdash (\lambda l. case...) : list\ tv_7 \rightarrow int$

In this case, the last step tells us what the type of the binding of $length$ is.

Figure 3: Examples of Typing

4 Examples of Compiling into IFL

It is worthwhile taking a quick look at how some programs in high level languages look in IFL. High level functional language syntaxes do not vary greatly (which is a good indicator for an intermediate language), and the reader is pointed to the references to find out the workings of a given example.

The first example is from Miranda, and is given in [20]. Figure 4 shows both the Miranda and IFL code for this example.

The second example is in HOPE, from [2]. Figure 5 shows the definition of a tree, and some functions on it. In the reference, these are bundled into a module. The module structure has been ignored in this example (since, as mentioned above), a general module structure has not yet been worked out. In HOPE types must be given explicitly, but the functions shown are easily typechecked in IFL.

```

tree * ::= niltree | node * (tree *) (tree *)

sort = flatten.build

foldr op z = g
  where
    g [] = z
    g (a:x) = op a (g x)

build = foldr insert niltree
  where
    insert b niltree = node b niltree niltree
    insert b (node a s t) = node a (insert b s) t, b <= a
                        = node a s (insert b t), b > a

flatten niltree = []
flatten (node a s t) = flatten s ++ [a] ++ flatten t

```

MIRANDA tree sorting program

```

lettype tree a = sum (niltree, node a (tree a) (tree a))

let sort = comp flatten build

let foldr = λ op z.
  letrec g = λ l. case l of nil => z, (cons a x) => op a (g x)
  in g

let build = letrec insert = λ b t.
  case t of niltree => node b niltree niltree,
  (node a s t) => if b <= a then node a (insert b s) t else
                  if b > a then node a s (insert b t) else fail
  in foldr insert niltree

letrec flatten = λ t. case t of niltree => nil,
  (node a s t) => cons (flatten s, cons (a, flatten t))

```

IFL tree sorting program

Assume for the translation into IFL that the type list *** is defined with the constructors *nil* and *cons*. As well, a composition operator, *comp* and the append operator *append* are defined as:

```

let comp = λ f g x. f (g x)
let append = λ l1 l2. case l1 of nil => l2, (cons a x) => append x (cons a l2)

```

Figure 4: Example of Miranda to IFL Translation

```

data otree == empty ++ tip(num) ++ node(otree#num#otree)

dec insert: num#otree → otree
dec flatten: otree → list num

—insert (n,empty) <= tip(n)
—insert (n,tip(m)) <= n < m then node(tip(n),m,empty)
                        else node (empty,m,tip(n))
—insert (n,node(t1,m,t2)) <= n < m then node(insert (n,t1),m,t2)
                        else node(t1,m,insert(n,t2))

—flatten (empty) <= nil
—flatten (tip(n)) <= [n]
—flatten (node(t1,n,t2)) <= flatten(t1) <> (n::flatten(t2))

```

HOPE tree functions

```

lettype otree = sum (sum (empty,tip num), node otree num otree)

letrec insert = λ(prod(n,tree)).
  case tree of empty => tip n
              (tip m) => if n < m then node (tip n) m empty
                        else node empty m (tip n)
              (node t1 m t2) => if n < m then node (insert (prod(n,t1))) m t2
                        else node t1 m (insert (prod (n,t2)))

letrec flatten = λ tree.
  case tree of empty => nil
              (tip n) => cons n nil
              (node t1 n t2) => append (flatten t1) (cons n (flatten t2))

```

IFL tree functions

Assume for the translation into IFL that the type list $*$ is defined with the constructors `nil` and `cons`. As well, `append` which concatenates two lists can be defined in IFL as:

```

let append = λ l1 l2. case l1 of nil => l2, (cons a x) => append x (cons a l2)

```

Figure 5: Example of HOPE to IFL Translation

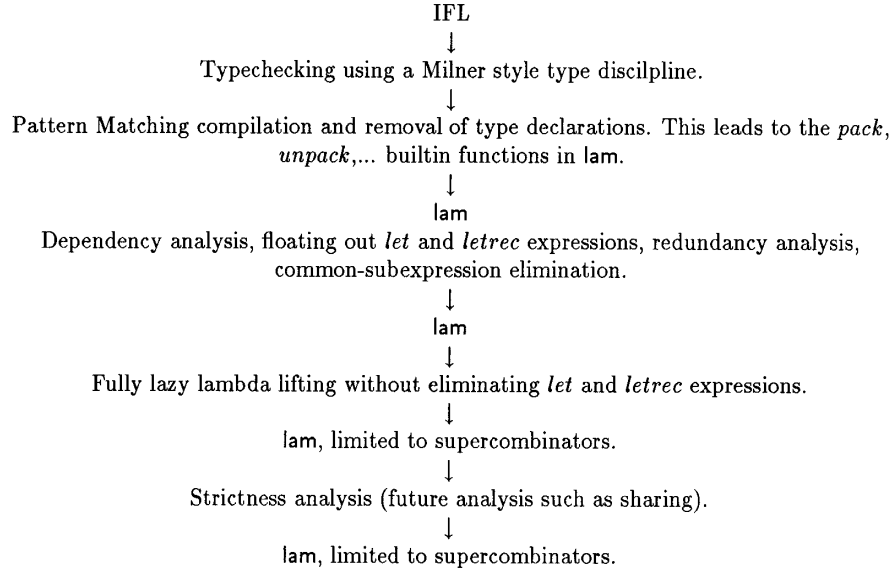
<code>lam</code>	$::=$	<code>bind_b exp_b lam⁺</code>
<code>bind_b</code>	$::=$	<code>let var = exp_b... var = exp_b letrec var = exp_b... var = exp_b</code>
<code>exp_b</code>	$::=$	<code>var const_b exp_bexp_b annot exp_b (exp_b) number string λvar⁺.exp_b let var = exp_b... var = exp_b in exp_b letrec var = exp_b... var = exp_b in exp_b</code>
<code>const_b</code>	$::=$	<code>const pack unpack ...</code>
<code>var</code>	$::=$	<code>alphanumeric*</code>
<code>number</code>	$::=$	<code>digit⁺ -number number . number number e number</code>
<code>string</code>	$::=$	<code>"alphanumeric"</code>

Figure 6: Abstract Syntax for lam

5 An Intermediate Functional Language Compiler

In order to be of optimal use, IFL will need a compiler to transform it into the lowest reasonable representation. As shown in Peyton-Jones, many back end optimizations rely on exploiting `let` and `letrec` structures. Thus, the back end of the IFL compiler has been defined to include these structures inside supercombinator bodies (supercombinators are a restricted class of lambda expressions where lazy evaluation is preserved). The abstract syntax for the back end is given in Figure 6. Note that the interactive commands still exist at this level, although all type declarations and pattern matching have been eliminated. This language is nearly identical with all intermediate functional languages (including those discussed earlier).


The compilation proceeds as follows:



Thus, the output will be fully annotated supercombinator definitions ready for compilation into machine code.

6 Semantic Issues

In order for IFL to be used, its semantics must be fully worked out, including the compiler semantics. Using standard denotational and algebraic semantics it is possible to show that IFL has the same semantics as the lambda calculus (with constants). This is achieved in the following manner:



Goguen et al. [8] have shown that a context-free grammar forms an initial algebra and that semantic functions from the context-free grammar are unique homomorphisms. To show that “meaning” is maintained by the Compile operation we have to show that:

$$\text{rep} \circ M' \circ \text{Compile} = M$$

which, since M and M' are unique homomorphisms, means that you must show that Compile and rep are also homomorphisms.

In order to accomplish the above, the semantics of IFL, compile, and lam must be totally defined. The meaning of an IFL term is well defined only when it is correctly typed, implying that the well-typedness of IFL terms must also be specified. This leads to a large amount of semantics.

7 Conclusion

To date, the following has been accomplished:

- IFL and lam have been defined as shown above. The idea of a high level intermediate language with a type system and a model for the environment have been developed.
- The semantics for IFL have been defined (see appendix). An intuitive and easy semantics for expressions involving abstract data types is forwarded.
- Parts of the semantics for typechecking have been defined.
- The compilation issues mentioned above have been researched, and many of them have been implemented for lam.

Further work includes:

- Completing the semantics of types and typechecking.
- Defining a concrete syntax for IFL.
- Completing an implementation of the IFL to lam translation. Many of the individual functions have been written and only need to be drawn together.
- Showing the translation from a high level language into IFL. (Hopefully the high level language will be Haskell – at this time its syntax is still being worked out, although it is hoped that the abstract syntax will be available soon [17].)

- Showing the translation from lam into machine code. The TIM machine [5][17] is one of the newest machines for functional languages and would be a good target machine. As well, the SECD machine [12][10] is being used for other research at the U of C which would make it an attractive target machine.

In the future I hope to extend this work in several ways:

- Total systems correctness is a recent idea. A complete correctness proof IFL along with some new ideas on how to simplify such a task would be very useful.
- Abstract interpretation is a growing discipline which will have a profound effect on compiler optimizations and other semantic issues.
- Given the inherent parallelism of functional languages, control structures could be added to IFL as annotations and the semantics of parallelism could be worked out. Lots of work has been done in this area but it is not yet totally conquered.
- Sharing analysis, reference counting, path semantics, and other interesting areas using functional languages could be studied using IFL as a base.

8 Appendix: Semantics

The semantics for the lambda calculus are given so that my style of denotational semantics is clear. Then the semantics of lam are defined, since they parallel those of the lambda calculus. IFL is analyzed and a new approach for the semantics of abstract types is discussed. An outline of what is necessary for a type semantics is given, although this work is not yet complete.

8.1 The Lambda Calculus

The syntax of the lambda calculus with constants is given as:

```

expression ::=  variable
               |  constant
               |  expression expression
               |  λ variable. expression
               |  (expression)

```

The domain of meanings must contain all possible values, V , for expressions. Thus, we will need a set of constants, say C , and we will need functions from values to values, $V \rightarrow V$. From Scott's work [19] we know that we can define the domain as:

$$V \cong C + (V \rightarrow V)$$

with the injection function \hookrightarrow_V and the projection function $\hookrightarrow_{V \rightarrow V}$. Using this, the semantics are:

$$\begin{aligned}
\mathcal{L} : \{E \mid E \in \text{expression}\} &\rightarrow V \\
\mathcal{L} \llbracket \text{constant} \rrbracket_\sigma &= K \llbracket \text{constant} \rrbracket \\
\mathcal{L} \llbracket \text{variable} \rrbracket_\sigma &= \sigma \llbracket \text{variable} \rrbracket \\
\mathcal{L} \llbracket \text{exp}_1 \text{ exp}_2 \rrbracket_\sigma &= (\mathcal{L} \llbracket \text{exp}_1 \rrbracket_\sigma \hookrightarrow_{V \rightarrow V}) (\mathcal{L} \llbracket \text{exp}_2 \rrbracket_\sigma) \\
\mathcal{L} \llbracket \lambda \text{ var. exp} \rrbracket_\sigma &= \lambda \text{ val. } \mathcal{L} \llbracket \text{exp} \rrbracket_{(\sigma [\text{val}/\text{var}])} \hookrightarrow_V
\end{aligned}$$

where exp_1 and exp_2 are lambda expressions, var is any variable, and $\sigma [\text{val}/\text{var}]$ means environment σ with val associated with the identifier var . K is a function which maps syntactic constants into C .

8.2 LAM

The following syntactic domains will be used for the semantics of lam:

```

Lam    = { L | L ∈ lam }
Expb   = { E | E ∈ expb }
Bndb  = { B | B ∈ bindb }
Var    = { V | V ∈ var }

```

The semantic domain is the same for both lam and IFL:

$V \cong B_0 + B_1 + \dots + F + P + S + E$	All Values
$B_0 = \{\perp, \text{true}, \text{false}\}$	Booleans
$B_1 = \{\perp, 0, 1, 2, \dots\}$	Numbers
\dots other flat domains of basic values	
$F = V \rightarrow V$	Functions
$P = V * V$	Products
$S = V + V$	Disjoint Unions
$E = \{\perp, \text{error}, \text{untyped}\}$	Errors

The sub-domain E holds error information: **error** implies a compile time error, such as an unbound variable; **untyped** implies that an expression could not be typed. Each of B_i is a flat domain containing constant information.

The semantics for **lam** closely parallel those of the lambda calculus (I will prove here –later– that they are identical). In addition, the semantic function \mathcal{R}_B is defined to model the interactive environment. It produces a finite product of values ($V * (V * (...))$) for each command which evaluates. For binding commands, the environment is updated, but no value appears in the product result. If a binding fails, the old environment is still used.

Define the environment as $Env: Var \rightarrow V$ and represent the environment as σ .

$$\begin{aligned}
\mathcal{R}_B: Lam &\rightarrow Env \rightarrow V \\
\mathcal{R}_B \llbracket super; ine_b \rrbracket \sigma &= (\mathcal{S}_B \llbracket super \rrbracket \sigma, \mathcal{R}_B \llbracket ine_b \rrbracket \sigma) \\
\mathcal{R}_B \llbracket binder; ine_b \rrbracket \sigma &= \text{if } \sigma_1 = \{\text{error}\} \text{ then } \mathcal{R}_B \llbracket ine \rrbracket \sigma \text{ else } \mathcal{R}_B \llbracket ine \rrbracket \sigma_1 \\
&\quad \text{where } \sigma_1 = \mathcal{B}_B \llbracket binder \rrbracket \sigma \\
\\
\mathcal{B}_B: Bnd_b &\rightarrow Env \rightarrow Env \\
\mathcal{B}_B \llbracket let v_1 = e_1 \dots v_n = e_n \rrbracket \sigma &= \sigma [v_1 \dots v_n / \mathcal{S}_B \llbracket s_1 \rrbracket \sigma, \dots \mathcal{S}_B \llbracket s_1 \rrbracket \sigma] \\
\mathcal{B}_B \llbracket letrec v_1 = e_1 \dots v_n = e_n \rrbracket \sigma &= fix (\lambda \sigma_1. \sigma [v_1 \dots v_n / \mathcal{S}_B \llbracket s_1 \rrbracket \sigma_1, \dots \mathcal{S}_B \llbracket s_1 \rrbracket \sigma_1]) \\
\\
\mathcal{S}_B: Lam &\rightarrow Env \rightarrow V \\
\mathcal{S}_B \llbracket var \rrbracket \sigma &= \sigma \text{ var} \\
\mathcal{S}_B \llbracket const \rrbracket \sigma &= \mathcal{K} \llbracket const \rrbracket \sigma \\
\mathcal{S}_B \llbracket number \rrbracket \sigma &= \mathcal{N} \llbracket number \rrbracket \\
\mathcal{S}_B \llbracket string \rrbracket \sigma &= \mathcal{S} \llbracket string \rrbracket \\
\mathcal{S}_B \llbracket s_1 s_2 \rrbracket \sigma &= (\mathcal{S}_B \llbracket s_1 \rrbracket \sigma \hookrightarrow_{V \rightarrow V}) (\mathcal{S}_B \llbracket s_2 \rrbracket \sigma) \\
\mathcal{S}_B \llbracket (s) \rrbracket \sigma &= \mathcal{S}_B \llbracket s \rrbracket \sigma \\
\mathcal{S}_B \llbracket [annot] s \rrbracket \sigma &= \mathcal{A} \llbracket [annot] s \rrbracket \sigma \\
\mathcal{S}_B \llbracket \lambda v_1 \dots v_n. s \rrbracket \sigma &= \lambda s_1 \dots s_n. \mathcal{S}_B \llbracket s \rrbracket (\sigma [v_1 \dots v_n / \mathcal{S}_B \llbracket s_1 \rrbracket \sigma \dots \mathcal{S}_B \llbracket s_n \rrbracket \sigma]) \\
\mathcal{S}_B \llbracket let v_1 = e_1 \dots v_n = e_n in s \rrbracket \sigma &= \mathcal{S}_B \llbracket s \rrbracket (\mathcal{B}_B \llbracket let v_1 = e_1 \dots v_n = e_n \rrbracket \sigma) \\
\mathcal{S}_B \llbracket letrec v_1 = e_1 \dots v_n = e_n in s \rrbracket \sigma &= \mathcal{S}_B \llbracket s \rrbracket (\mathcal{B}_B \llbracket letrec v_1 = e_1 \dots v_n = e_n \rrbracket \sigma)
\end{aligned}$$

\mathcal{K} and \mathcal{N} and \mathcal{S} have their obvious meanings. \mathcal{A} is dependent on the annotation, some of which are discussed later.

8.3 IFL

The semantics of IFL are fairly straightforward except for pattern matching and abstract types. Pattern matching will be explained as it is developed below, but abstract types warrant some discussion before we proceed.

The meaning of an expression is dependent on the meanings of its subexpressions. Since some expressions will contain constructors from type definitions, in order to understand the expression we must know the meaning of the constructor. In the following discussion a constructor is a simply a function. The fact that it is a function whose purpose is to make a new type is not important at this stage since we are not discussing the semantics of types, but of expressions. Later, expressions will be limited to those which are well-typed, and the meaning of a type will then be addressed.

Consider this simplistic abstract type in IFL:

```

abstype list * = sum (nil, cons prod(*, list *)) in
  let hd = λ(cons h t). h,
      tl = λ(cons h t). t;
command;
```

The idea behind an abstract type is that it is only defined for the operations given on it (in this case **list *** is only defined for **hd** and **tl**) and not for any expressions given subsequent to the definition (ie. **command**). For

this to be the case, the constructors (`nil` and `cons`) must be in the environment whenever `hd` or `tl` is called, but at no other time. As well, both `hd` and `tl` must access exactly the same constructor in the environment or they will be working on different types. (ie. we cannot take the easy way out and simply add the constructors to the environment whenever `hd` or `tl` is used since this would result in new constructors for each instance of `hd` and `tl`).

An elegant solution to this problem is to give an indirection to variable access in the environment. This can be done by defining the environment σ to be of type $(Var \rightarrow Int, Int \rightarrow V)$ where Int forms the indirection pointer. Define σ to be (σ_1, σ_2) to represent these two functions, and σvar to be $\sigma_2(\sigma_1 var)$ to represent the value in the environment corresponding to var . In all cases except abstract types, this environment works exactly like the ordinary environment mapping $var \rightarrow V$. (This indirection is often used in denotational semantics to incorporate branches, side effects, jumps, and procedure calls [7]. For abstract types it is slightly more complicated since we also require the use of annotations, as described below).

For an abstract data type we need a mapping in the environment for each constructor. For example, $nil \rightarrow n$ and $cons \rightarrow m$ are the mappings to Int (given that m and n are unique integers) – call these σ_1^a . $n \rightarrow \llbracket nil \rrbracket$ and $m \rightarrow \llbracket cons \rrbracket$ are the mapping from Int to the meanings of the constructors – call these σ_2^a .

The condition that `hd` and `tl` access the same constructors is equivalent to making sure that both access the same unique indirection keys (the integers m and n). This can be accomplished by adding σ_2^a to the environment.

The condition that only `hd` and `tl` can access the constructors means that σ_1^a must be in the environment when `hd` or `tl` are used, but at no other time. Notice that the mapping from $Var \rightarrow Int$ is syntactic in that it can be represented by pairs of values. Since this is the case, σ_1^a can be passed to `hd` and `tl` using an annotation.

This idea is developed by the semantics given below.

The syntactic domains are given below:

Ifl	$= \{ I \mid I \in ifl \}$
Exp	$= \{ E \mid E \in exp \}$
Bnd	$= \{ B \mid B \in bind \}$
Pat	$= \{ P \mid P \in pattern \}$
Tyn	$= \{ T \mid T \in tyname \}$
Typ	$= \{ T \mid T \in type \}$
Con	$= \{ C \mid C \in constructor \}$
Var	$= \{ V \mid V \in var \}$

The semantic domain is the same as the one developed for `lam`.

8.3.1 Expressions

Define a `Binding` in the domain V as:

$$\text{Binding: } (Var \rightarrow (Int + \{\text{error}\}), ((\{\text{error}\} \rightarrow \{\text{error}\}) + (Int \rightarrow V)))$$

A binding maps a variable to a value in the domain using an indirection key (an integer) as discussed above. If the variable is not bound to a key or a key is not bound to a value, the result of the binding is an error.

\mathcal{A} is used to add bindings to the environment. If a new binding is added then any access to the environment will first check the new bindings and then the old ones. \mathcal{A}_O and \mathcal{A}_T are variations of \mathcal{A} which add $Var \rightarrow Int$ or $Int \rightarrow V$ mappings only, not both. \mathcal{N} gives an empty binding where any access will result in an error, since nothing is bound.

$\mathcal{A}: \text{Binding} \rightarrow \text{Binding} \rightarrow \text{Binding}$
 $\mathcal{A}(\sigma_1, \sigma_2)(\sigma_3, \sigma_4) = \lambda \text{var}. (\text{if } (\sigma_1 \text{ var}) = \{\text{error}\} \text{ then } (\sigma_3 \text{ var}) \text{ else } (\sigma_1 \text{ var}),$
 $\quad \text{if } (\sigma_2 (\sigma_1 \text{ var})) = \{\text{error}\} \text{ then } (\sigma_3 (\sigma_4 \text{ var})) \text{ else } (\sigma_2 (\sigma_1 \text{ var})))$

$\mathcal{A}_O: \text{Binding} \rightarrow (\text{Var} \rightarrow \text{Int}) \rightarrow \text{Binding}$
 $\mathcal{A}_O(\sigma_1, \sigma_2) \sigma_3 = \lambda \text{var}. (\text{if } (\sigma_1 \text{ var}) = \{\text{error}\} \text{ then } (\sigma_3 \text{ var}) \text{ else } (\sigma_1 \text{ var}), \sigma_2)$

$\mathcal{A}_T: \text{Binding} \rightarrow (\text{Int} \rightarrow \text{V}) \rightarrow \text{Binding}$
 $\mathcal{A}_T(\sigma_1, \sigma_2) \sigma_4 = \lambda \text{var}. (\sigma_1, \text{if } (\sigma_2 \text{ var}) = \{\text{error}\} \text{ then } (\sigma_4 \text{ var}) \text{ else } (\sigma_2 \text{ var}))$

$\mathcal{N}: \text{Var} \rightarrow \text{V}$
 $\mathcal{N} = \lambda \text{var}. \{\text{error}\}$

\mathcal{R} models the interactive run-time environment. It takes in a countably infinite sequence of IFL terms and outputs a countably infinite product term of values. That is, if a binding is evaluated, no value is returned. The binding either succeeds and the new environment is used from then on, or it fails and the old environment is still used. A “;” is used to separate individual commands inside ifl^+ .

$\mathcal{R}: \text{IfI} \rightarrow \text{V}$
 $\mathcal{R} \llbracket \text{exp}; \text{ifl} \rrbracket \sigma = (\mathcal{E} \llbracket \text{exp} \rrbracket \sigma, \mathcal{R} \llbracket \text{ifl} \rrbracket \sigma)$
 $\mathcal{R} \llbracket \text{bind}; \text{ifl} \rrbracket \sigma = \text{if } \sigma_1 = \{\text{error}\} \text{ then } \mathcal{R} \llbracket \text{ifl} \rrbracket \sigma \text{ else } \mathcal{R} \llbracket \text{ifl} \rrbracket \sigma_1$
 $\quad \text{where } \sigma_1 = \mathcal{B} \llbracket \text{bind} \rrbracket \sigma$

\mathcal{M} is used for pattern matching. Lazy pattern matching is used on single patterns. If the pattern matches, a new environment is returned with each variable in the pattern bound. If the match fails, $\{\text{error}\}$ is returned. An auxiliary function must be defined first:

$\text{SEL } s \text{ i } (s \ a_1, \dots, a_i, \dots, a_n) = a_i$
 $\text{SEL } s \text{ i } (s' \ a_1, \dots, a_n) = \{\text{error}\}$
 $\text{SEL } s \text{ i } x = \{\text{error}\}$
 $\text{SEL } s \text{ i } \perp = \perp$

$\mathcal{M}: \text{Pat} \rightarrow \text{IfI} \rightarrow \text{Binding} \rightarrow \text{Binding}$
 $\mathcal{M} \llbracket p \rrbracket x \{\text{error}\} = \{\text{error}\}$
 $\mathcal{M} \llbracket \text{var} \rrbracket x (\sigma_1, \sigma_2) = (\sigma_1[x/n], \sigma_2[n/\text{var}])$
 $\quad \text{where } n \text{ is a unique integer key}$
 $\mathcal{M} \llbracket \text{const} \rrbracket x \sigma = \text{if } \text{const} = \mathcal{E} \llbracket x \rrbracket \sigma \text{ then } \sigma \text{ else } \{\text{error}\}$
 $\mathcal{M} \llbracket \text{con } p_1 \dots p_n \rrbracket e \sigma = \mathcal{M} \llbracket p_1 \rrbracket (\text{SEL con 1 } e) (\mathcal{M} \llbracket p_2 \rrbracket (\text{SEL con 2 } e)$
 $\quad (\dots \mathcal{M} \llbracket p_n \rrbracket (\text{SEL con } n \ e) \sigma) \dots)$

\mathcal{T} makes bindings for new type constructors.

$\mathcal{T}: \text{Typ} \rightarrow \text{Binding} \rightarrow \text{Binding}$
 $\mathcal{T} \llbracket \text{sum } t_1 \dots t_n \rrbracket \sigma = \mathcal{T} \llbracket t_1 \rrbracket (\mathcal{T} \llbracket t_2 \rrbracket (\dots \mathcal{T} \llbracket t_n \rrbracket \sigma \dots))$
 $\mathcal{T} \llbracket \text{con } t_1 \dots t_n \rrbracket (\sigma_1, \sigma_2) = (\sigma_1[\text{con}/n], \sigma_2[n/\lambda t_1 \dots t_n. (\text{con}, t_1, \dots, t_n)])$
 $\quad \text{where } n \text{ is a unique integer key}$

\mathcal{B} models binding in the environment. The binding for an abstract data type was discussed above and is shown below. As mentioned above, the function from $\text{IfI} \rightarrow \text{Int}$ could be represented in an annotation in the form $[\text{con}_1 ==> n_1 \dots \text{con}_m ==> n_m]$ where each con_i is a constructor and each n_i is unique integer key. For simplicity I have simply included the mapping inside the annotation.

$$\begin{aligned}
B & \text{ Bnd} \rightarrow \text{Binding} \rightarrow \text{Binding} \\
B & \parallel \text{let } p_1 \leftarrow e_1 \quad p_n \leftarrow e_n \parallel o \\
& \quad \lambda f \parallel p_n \parallel e_n \left(\lambda f \parallel p_1 \parallel e_1 \ o \right) \\
B & \parallel \text{letrec } p_1 \leftarrow e_1 \quad p_n \leftarrow e_n \parallel o \\
& \quad \text{fix } (\lambda o_1 \rightarrow \lambda (\lambda f \parallel p_1 \parallel e_1 \ o_1 \left(\lambda (\lambda f \parallel p_1 \parallel e_1 \ o_1 \ o))) \right)) \\
B & \parallel \text{lettype } tn_1 \leftarrow t_1 \quad tn_n \leftarrow t_n \parallel o \\
& \quad \text{fix } (\lambda o_1 \rightarrow \lambda (t_1 \parallel t_1 \parallel o_1 \rightarrow (\lambda (t_n \parallel t_n \parallel o_1 \ o))) \rightarrow) \\
B & \parallel \text{abstype } tn_1 \leftarrow t_1 \quad tn_n \leftarrow t_n \text{ in letbnd } \parallel o \\
& \quad B_{ADT} \parallel \text{letbnd } \parallel o \ (B \parallel \text{lettype } tn_1 \leftarrow t_1 \quad tn_n \leftarrow t_n \parallel \mathcal{N})
\end{aligned}$$

$$\begin{aligned}
B_{ADT} & \text{ Bnd} \rightarrow \text{Binding} \rightarrow \text{Binding} \rightarrow \text{Binding} \\
B_{ADT} & \parallel \text{let}(\text{rec}) \ p_1 \leftarrow e_1 \quad p_n \leftarrow e_n \parallel o \ (o_1, o_2) \\
& \quad B \parallel \text{let}(\text{rec}) \ p_1 \leftarrow [o_1] \ e_1 \quad p_n \leftarrow [o_2] \ e_n \parallel (\lambda t \ o \ o_2)
\end{aligned}$$

C is used for case statements. Successive matches are tried until one is found which does not cause an error.

$$\begin{aligned}
C & \mathbb{V} \rightarrow \mathbb{V} \rightarrow \mathbb{V} \\
C & \{\text{error}\} \text{ f} \rightarrow \text{f} \\
C & \text{ill x} \rightarrow \text{ill}
\end{aligned}$$

K evaluates constants. These are pretty standard and not all of them are given below.

$$\begin{aligned}
K & \text{ Cnst} \rightarrow \mathbb{V} \\
K & \parallel \text{if} \parallel o \rightarrow \text{Abt e if } \mathcal{E} \parallel b \parallel o \rightarrow \text{true then } \mathcal{E} \parallel t \parallel o \text{ else } \mathcal{E} \parallel e \parallel o \\
K & \parallel \lambda \parallel o \rightarrow \text{lambda a b } \mathcal{E} \parallel a \parallel o \rightarrow \mathcal{E} \parallel a \parallel o \\
K & \parallel \text{int} \parallel o \rightarrow \text{integer} \\
K & \parallel \text{fail} \parallel o \rightarrow \{\text{error}\}
\end{aligned}$$

AN is used for annotations. At this time the only annotation described below is for abstract data types. It simply takes the $Var \rightarrow Int$ mapping out of the annotation and adds it to the environment.

$$\begin{aligned}
AN & \text{ Ill} \rightarrow \text{Binding} \rightarrow \mathbb{V} \\
AN & \parallel [o_1] \ \text{exp} \parallel o \rightarrow \mathcal{E} \parallel \text{exp} \parallel (\lambda_0 \ o \ o_1)
\end{aligned}$$

\mathcal{E} models expressions which evaluate to values.

$$\begin{array}{ll}
\mathcal{E}: \text{Exp} \rightarrow \text{Binding} \rightarrow V & \\
\mathcal{E} \llbracket x \rrbracket \{\text{error}\} & = \{\text{error}\} \\
\mathcal{E} \llbracket \text{var} \rrbracket \sigma & = \sigma \llbracket \text{var} \rrbracket \\
\mathcal{E} \llbracket \text{const} \rrbracket \sigma & = K \llbracket \text{constant} \rrbracket \\
\mathcal{E} \llbracket \text{number} \rrbracket \sigma & = N \llbracket \text{number} \rrbracket \\
\mathcal{E} \llbracket \text{string} \rrbracket \sigma & = S \llbracket \text{string} \rrbracket \\
\mathcal{E} \llbracket \text{exp}_1 \text{ exp}_2 \rrbracket \sigma & = (\mathcal{E} \llbracket \text{exp}_1 \rrbracket \sigma \hookrightarrow_{V \rightarrow V}) (\mathcal{E} \llbracket \text{exp}_2 \rrbracket \sigma) \\
\mathcal{E} \llbracket [\text{annot}] \text{exp} \rrbracket \sigma & = \mathcal{ANN} \llbracket \text{annot exp} \rrbracket \sigma \\
\mathcal{E} \llbracket (\text{exp}) \rrbracket \sigma & = \mathcal{E} \llbracket \text{exp} \rrbracket \sigma \\
\mathcal{E} \llbracket \lambda p_1 \dots p_n. \text{exp} \rrbracket \sigma & = \lambda e_1 \dots e_n. \mathcal{E} \llbracket \text{exp} \rrbracket \\
& \quad (B \llbracket \text{let } p_1 = e_1 \dots p_n = e_n \rrbracket \sigma) \hookrightarrow_V \\
\mathcal{E} \llbracket \text{let } p_1 = e_1 \dots p_n = e_n \text{ in exp} \rrbracket \sigma & = \mathcal{E} \llbracket \text{exp} \rrbracket (B \llbracket \text{let } p_1 = e_1 \dots p_n = e_n \rrbracket \sigma) \\
\mathcal{E} \llbracket \text{letrec } p_1 = e_1 \dots p_n = e_n \text{ in exp} \rrbracket \sigma & = \mathcal{E} \llbracket \text{exp} \rrbracket (B \llbracket \text{letrec } p_1 = e_1 \dots p_n = e_n \rrbracket \sigma) \\
\mathcal{E} \llbracket \text{lettype } tn_1 = t_1 \dots tn_n = t_n \text{ in exp} \rrbracket \sigma & = \mathcal{E} \llbracket \text{exp} \rrbracket (B \llbracket \text{lettype } tn_1 = t_1 \dots tn_n = t_n \rrbracket \sigma) \\
\mathcal{E} \llbracket \text{case } v \text{ of } p_1 \Rightarrow e_1 \dots p_n \Rightarrow e_n \rrbracket \sigma & = C (\mathcal{E} \llbracket e_1 \rrbracket (\mathcal{M} \llbracket p_1 \rrbracket (\mathcal{E} \llbracket v \rrbracket \sigma) \sigma), \dots, \\
& \quad C (\mathcal{E} \llbracket e_{n-1} \rrbracket (\mathcal{M} \llbracket p_{n-1} \rrbracket (\mathcal{E} \llbracket v \rrbracket \sigma) \sigma), \\
& \quad (\mathcal{E} \llbracket e_n \rrbracket (\mathcal{M} \llbracket p_n \rrbracket (\mathcal{E} \llbracket v \rrbracket \sigma) \sigma)) \dots) \\
\mathcal{E} \llbracket \Box e_1 e_2 \rrbracket \sigma & = C (\mathcal{E} \llbracket e_1 \rrbracket \sigma, \mathcal{E} \llbracket e_2 \rrbracket \sigma)
\end{array}$$

where S maps strings into the appropriate B_i , N maps numbers into B_1 , and K maps other constants into the other B_i domains.

8.3.2 Typechecking

Work on the semantics of the types for IFL is still underway. The type discipline is a simple Milner style system at this point.

Typechecking semantics simply limit the domain V to those meanings which come from well-typed IFL expressions.

The abstract syntax of a type is given by:

$$\begin{array}{l}
\text{texp} ::= \text{tvar} \mid \text{tcon texp} \dots \text{texp} \\
\text{tvar} ::= \text{"*"} \text{ integer} \\
\text{tcon} ::= \text{sum} \mid \text{prod} \mid \text{fun} \mid \text{alphanumeric}^*
\end{array}$$

A builtin type, such as `int`, is a `tcon` of zero arity.

The syntactic domains for the semantics of typechecking will include `Ifl`, `Exp`, `Bnd`, `Pat`, `Tyn` and `Var` for IFL terms as well as the following domains for type expressions:

$$\begin{array}{l}
\text{Texp} = \{T \mid T \in \text{texp}\} \\
\text{Tvar} = \{V \mid V \in \text{tvar}\} \\
\text{Tcon} = \{C \mid C \in \text{tcon}\}
\end{array}$$

In order to show that a well-typed expression does will not cause any run-time type errors, it will be necessary to show that an syntactic object of a certain type always maps to an semantic object with that same type. That is, an integer should always map into the domain B_1 . The easiest way to do this is to define the domain of type expressions to be the same as the domain for expressions, namely V . In this case, we could map the type "int" to the domain of integers, B_1 .

MacQueen's work with *ideals* [13] allows us to do exactly this. To summarize his work, define a subset I of V to be an *ideal* if and only if:

- 1) $I \neq \{\}$
- 2) $\forall y \in I, x \in V. x \sqsubseteq y \rightarrow x \in I$
- 3) $\forall \langle x \rangle. (\forall n. x_n \in I) \rightarrow \sqcup \langle x \rangle \in I$

\perp is assumed to have every type, thus an ideal cannot be empty. An ideal contains chains beginning with \perp and containing a least fixed point.

Denote the subset of V which contains only ideals as $\mathcal{I}(V)$. Given that I and J are ideals, we can define binary functions on them:

- $I \oplus J$ is the union of the injection of I into $V + V$ and the injection of J into $V + V$.
- $I \otimes J$ is the cross product of I and J , $I \times J$.
- $I \oplus J$ is the function space $\{f \in V \rightarrow V \mid f(I) \subseteq J\}$

Each of these preserve ideality.

Given a mapping from type variables to $\mathcal{I}(V)$, we can define

$$\begin{aligned}
\mathcal{D}: t \in \text{texp} &\rightarrow (v \in \text{tvar} \rightarrow \mathcal{I}(V)) \rightarrow \mathcal{I}(V) \\
\mathcal{D} \llbracket c \rrbracket v &= \mathcal{D}_c \llbracket c \rrbracket \\
\mathcal{D} \llbracket t \rrbracket v &= v \ t \\
\mathcal{D} \llbracket \text{prod } \text{texp}_1 \ \text{texp}_2 \rrbracket v &= \mathcal{D} \llbracket \text{texp}_1 \rrbracket v \otimes \mathcal{D} \llbracket \text{texp}_2 \rrbracket v \\
\mathcal{D} \llbracket \text{sum } \text{texp}_1 \ \text{texp}_2 \rrbracket v &= \mathcal{D} \llbracket \text{texp}_1 \rrbracket v \oplus \mathcal{D} \llbracket \text{texp}_2 \rrbracket v \\
\mathcal{D} \llbracket \text{fun } \text{texp}_1 \ \text{texp}_2 \rrbracket v &= \mathcal{D} \llbracket \text{texp}_1 \rrbracket v \oplus \mathcal{D} \llbracket \text{texp}_2 \rrbracket v \\
\\
\mathcal{D}_c \llbracket \text{bool} \rrbracket &= B_0 \\
\mathcal{D}_c \llbracket \text{int} \rrbracket &= B_1 \\
&\dots
\end{aligned}$$

This ideal model was actually defined in order to handle recursive types, thus, if recursive types are added to the discipline, I may still work with the same model.

The semantics for types are only partially complete, and are thus not included here.

References

- [1] R. Burstall and B. Lampson. A kernel language for abstract types and modules. In G. Khan, D.B. MacQueen, and G. Plotkin, editors, *Semantics of Data Types, LNCS 173*, pages 1–50, 1984.
- [2] R.M. Burstall, D.B. MacQueen, and D.T. Sannella. Hope: An experimental applicative language. Internal Report CSR-6280, University of Edinburgh, May 1980.
- [3] Luca Cardelli and Peter Wegner. On understanding types, data abstraction, and polymorphism. *Computing Surveys*, 17(4):471–520, December 1985.
- [4] Jon Fairbairn. *Design and Implementation of a Simple Typed Language Based on the Lambda Calculus*. PhD thesis, Cambridge University, 1985.
- [5] Jon Fairbairn and Stuart Wray. Tim: A simple, lazy abstract machine to execute supercombinators. In *Functional Languages and Computer Architecture. LNCS 274*, pages 34–45, 1987.
- [6] Anthony J. Field and Peter G. Harrison. *Functional Programming*. Addison-Wesley, 1988.
- [7] Micheal Gordon. *The Denotational Description of Programming Languages*. Springer-Verlag, 1979.
- [8] J. A. Gugen, J. W. Thatcher, E. G. Wagner, and J.B. Wright. Initial algebra semantics and continuous algebra. In *Association for Computing Machinery ??*, pages 68–95, 1977.
- [9] Robert Harper, David MacQueen, and Robin Milner. *Standard ML*. University of Edinburgh, 1987.
- [10] P. Henderson. *Functional Programming – Application and Implementation*. Prentice/Hall International, 1980.
- [11] Paul Hudak. *Compiling Functional Languages*. Tutorial for Functional Programming and Computer Architectures Conference, Portland Oregon, 1987.
- [12] P. Landin. The mechanical evaluation of expressions. *Computer Journal*, 6(4):308–320, 1964.
- [13] D. MacQueen. An ideal model for recursive polymorphic types. In *ACM ??*, pages 165–174, 1983.
- [14] R. Milner. A theory of type polymorphism in programming. Internal Report 9, University of Edinburgh, September 1977.
- [15] Simon Peyton Jones. *The Implementation of Functional Programming Languages*. Prentice-Hall International Series in Computer Science, 1987.
- [16] Simon Peyton Jones. FLIC — a Functional Language Intermediate Code. *ACM SIGPLAN Notices*, 23(8):30–48, August 1988. Also: Internal Note 2048, Department of Computer Science, University College London, Gower St., London WC1E 6BT.
- [17] Simon Peyton Jones. Personal communication. March 1988.
- [18] John C. Reynolds. Three approaches to type structure. In *Springer Verlag LNCS 185*, pages 97 – 138, 1985.
- [19] Dana Scott. Data types as lattices. *Siam J. Comput.*, 5(3):522–587, September 1976.
- [20] D. Turner. Miranda: A non-strict functional language with polymorphic types. In *Functional Programming Languages and Computer Architecture, Nancy France*. Springer-Verlag, 1985.
- [21] Phillip Wadler. How to replace failure by a list of successes. In *Functional Programming Languages and Computer Architecture, Nancy France*. Springer-Verlag, 1985.
- [22] Stuart Wray. *Implementation and Programming Techniques for Functional Languages*. PhD thesis, Cambridge University, 1986.