

*H/*Direct: A Binary Foreign Language Interface for Haskell

Sigbjorn Finne

Daan Leijen

Erik Meijer

Simon Peyton Jones

April 16, 1998

Abstract

*H/*Direct is a foreign-language interface for the purely functional language Haskell. Rather than rely on host-language type signatures, *H/*Direct compiles Interface Definition Language (IDL) to Haskell stub code that marshals data across the interface. This approach allows Haskell to call both C and COM, and allows a Haskell component to be wrapped in a C or COM interface. IDL is a complex language and language mappings for IDL are usually described informally. In contrast, we provide a relatively formal and precise definition of the mapping between Haskell and IDL.

This paper has been submitted to the International Conference on Functional Programming 1998 (ICFP'98).

1 Introduction

A foreign-language interface provides a way for programs written in one language to call, or be called by, programs written in another. Programming languages that do not supply a foreign-language interface die a slow, lingering death — good languages die more slowly than bad ones, but they all die in the end.

In this paper we describe a new foreign-language for the functional programming language Haskell. In contrast to earlier foreign-language interfaces for Haskell, such as Green Card [5], we describe a design based on a standard Interface Definition Language (IDL). We discuss the reasons for this decision in Section 2.

Our interface provides direct access to libraries written in C (or any other language using C's calling convention), and makes it possible to write Haskell procedures that can be called from C. The same tool also makes it allows us to call COM components directly from Haskell [9], or to seal up Haskell programs as a COM component. (COM is Microsoft's component object model; it offers a language-independent interface standard between software components. The interfaces of these components are written in IDL.)

*H/*Direct generates Haskell stub code from IDL interface descriptions. It is carefully designed to be independent of the particular Haskell implementation. To maintain this independence, *H/*Direct requires the implementation to support a primitive foreign-language interface mechanism, expressed using a (non-standard) Haskell `foreign` declaration;

*H/*Direct provides the means to leverage that primitive facility into the full glory of IDL.

Because they cater for a variety of languages, foreign-language interfaces tend to become rich, complex, incomplete, and described only by example. The main contribution of this paper is to provide (part of) a formal description of the interface. This precision encompasses not only the programmer's-eye view of the interface, but also its implementation. The bulk of the paper is taken up with this description.

2 Background

The basic way in which almost any foreign-language interface works is this. The signature of each foreign-language procedure is expressed in some formal notation. From this signature, stub code is generated that *marshals* the parameters “across the border” between the two languages, calls the procedure using the foreign language's calling convention, and then un-marshals the results back across the border. Dealing with the different calling conventions of the two languages is usually the easy bit. The complications come in the parameter marshalling, which transforms data values built by one language into a form that is comprehensible to the other.

A major design decision is the choice of notation in which to describe the signatures of the procedures that are to be called across the interface. There are three main possibilities:

- *Use the host language (Haskell, in our case).* That is, write a Haskell type signature for the foreign function, and generate the stub code from it. Green Card uses this approach [5], as does J/Direct [8] (Microsoft's foreign-language interface for Java).
- *Use the foreign language (say C).* In this case the stub code must be generated from the C prototype for the procedure. SWIG [1] uses this approach.
- *Use a separate Interface Definition Language (IDL),* designed specifically for the purpose.

We discuss the first two possibilities in Section 2.1 and the third in Section 2.2.

2.1 Using the host or foreign language

At first sight the first two options look much more convenient than the third, because the caller is written in one language and the callee in the other, so the interface is conveniently expressed for at least one of them. Here, for example, is how J/Direct allows Java to make foreign-language calls:

```
class ShowMsgBox {
    public static void main(String args[])
    {
        MessageBox(0,"Hello!","Java Messagebox",0);
    }

    /** @dll.import("USER32") */
    private static native
    int MessageBox( int hwndOwner, String text
                  , String title, int fuStyle
                  );
}
```

The `dll.import` directive tells the compiler that the `Java MessageBox` method will link to the native Windows `USER32.DLL`. The parameter marshaling (for example of the strings) is generated based on the Java type signature for `MessageBox`.

The fatal flaw is that *it is invariably impossible, in general, to generate adequate stub code based solely on the type signature of a procedure in one language or the other*. There are three kinds of difficulties.

1. First, some practically-important languages, notably C, have a type system that is too weak to express the necessary distinctions. For example:
 - The stub code generator must know the mode of each parameter — `in`, `in out`, or `out` — because each mode demands different marshaling code.
 - Some pointers have a significant `NULL` value while others do not. Some pointers point to values that can (and sometimes should) be copied across the border, while others refer to mutable locations whose contents must not be copied.
 - There may be important inter-relationships between the parameters. For example, one parameter might point to an array of values, while another gives the number of elements in the array. The marshaling code needs to know about such dependencies.
2. On the other hand, it may not even be enough to give the signature in a language with an expressive type system, such as Haskell. The trouble is that the type signature still says too little about the foreign procedures type signature. For example, is the result of a Haskell procedure returned as the result of the foreign procedure, or via an `out`-parameter of that procedure? In the case of J/Direct, when a record is passed as an argument, Java's type signature is not enough to specify the layout of the record because Java does not specify the layout of the fields of an object and the garbage collector can move the object around in memory.

3. The signature of a foreign procedure may say too little about allocation responsibilities. For example, if the caller passes a data structure to the callee (such as a string), can the latter assume that the structure will still be available after the call? Does the caller or callee allocate space to hold the results?

In an earlier paper we described Green Card, whose basic approach was to use Haskell as the language in which to give the type signatures for foreign procedures [5]. To deal with the issues described above we provided ways of augmenting the Haskell type signature to allow the programmer to “customise” the stub code that would be generated. However, Green Card grew larger and larger — and we realised that what began as a modest design was turning into a full-scale language.

2.2 Using an IDL

Of course, we are not the first to encounter these difficulties. The standard solution is to use a separate Interface Definition Language (IDL) to describe the signatures of procedures that are to be called across the border. IDLs are rich and complicated, for precisely the reasons described above, but they are at least somewhat standardised and come with useful tools. We focus on the IDL used to describe COM interfaces [10], which is closely based on DCE IDL[7]. Another popular IDL dialect is the one defined by OMG as part of the CORBA specification[11], and we intend to provide support for this using the translation from OMG to DCE IDL defined by [13, 12].

Like COM, but unlike CORBA¹, we take the view that the IDL for a foreign procedure defines a *language-independent, binary interface to the foreign procedure* — a sort of *lingua franca*. The interface thus defined is supposed to be complete: it covers calling convention, data format, and allocation rules. It may be necessary to generate stub code on both sides of the border, to marshal parameters into the IDL-mandated format, and then on into the format demanded by the foreign procedure. But these two chunks of marshaling code can be generated separately, each by a tool specialised to its host language. By design, however, IDL's binary conventions are more or less identical to C's, so marshaling on the C side is hardly ever necessary.

Here, for example, is the IDL describing the interface to a function `foo`:

```
int foo( [out] long* l
        , [string, in] char* s
        , [in, out] double* d
        );
```

The parts in square brackets are called *attributes*. In this case they describe the mode of each parameter, but there are a rich set of further attributes that give further (and often essential) information about the type of the parameters. For example, the `string` attribute tells that the parameter `s` points to a null-terminated array of characters, rather than pointing to a single character.

¹CORBA does not define a binary interface. Rather, each ORB vendor provides a *language binding* for a number of supported languages. This language binding essentially provides the marshaling required to an ORB-specific common calling convention. If you want to use a language that the ORB vendor does not support, you are out of luck.




Figure 1: The big picture

2.3 Overview

The “big picture” is given by Figure 1. The interface between Haskell and the foreign language is specified in IDL. This IDL specification is read by *H/Direct*, which then produces Haskell and C² source files files containing Haskell and C *stub code*.

H/Direct can generate stub code that allows Haskell to call C, or C to call Haskell. It can also generate stub code that allows Haskell to create and invoke COM components, and that allows COM components to be written in Haskell. Much of the work in all four cases concerns the marshaling of data between C and Haskell, and that is what we concentrate in this paper.

Since *H/Direct* generates Haskell source code, how does it express the actual foreign-language call (or entry for the inverse case)? We have extended Haskell with a *foreign* declaration that asks the Haskell implementation to generate code for a foreign-language call (or entry) [2]. The *foreign* declaration deals with the most primitive layer of marshaling, which is necessarily implementation dependent; *H/Direct* generates all the implementation-independent marshaling.

To make all this concrete, suppose we have the following IDL interface specification:

```

typedef struct { int x,y; } Point;

void Move( [in,out,ref] Point* p );

```

If asked to generate stub code to enable Haskell to call function *Move*, *H/Direct* will generate the following (Haskell) code:

```

data Point = Point { x,y::Int }
marshalPoint :: Point -> IO (Ptr Point)
marshalPoint = ...

unmarshalPoint :: Ptr Point -> IO Point
unmarshalPoint = ...

move :: Point -> IO Point
move p =
  
```

²For the sake of definiteness we concentrate on C as the foreign language in this paper.

```

do{ a <- marshalPoint p
  ; primMove a
  ; r <- unmarshalPoint a
  ; hdFree
  ; return r
}

foreign import stdcall "Move"
  primMove :: Ptr Point -> IO ()
  
```

This code illustrates the following features:

- For each IDL declaration, *H/Direct* generates one or more Haskell declarations.
- From the IDL procedure declaration *Move*, *H/Direct* generates a Haskell function *move* whose signature is intended to be “what the user would expect”. In particular, the Haskell type signature is expressed using “*high-level*” types; that is, Haskell equivalents of the IDL types. For example, the signature for *move* uses the Haskell record type *Point*. The translation for a procedure declaration is discussed in Section 3.
- The body of the procedure marshals the parameters into their “*low-level*” types, before calling the “*low-level*” Haskell function *primMove*. The latter is defined using a *foreign* declaration; the Haskell implementation generates code for the call to the C procedure *Move*. Section 4 specifies the high-level and low-level type corresponding to each IDL type.
- A “*low-level*” type is still a perfectly first-class Haskell type, but it has the property that it can trivially be marshalled across the border. There is fixed set of primitive “*low-level*” types, including *Int*, *Float*, *Char* and so on. *Addr* is a low-level type that holds a raw machine address. The type constructor *Ptr* is just a synonym for *Addr*:

```

type Ptr a = Addr
addPtr :: Ptr a -> Int -> Ptr b
  
```

The type argument to *Ptr* is used simply to allow *H/Direct* to document its output somewhat, by giving the “*high-level*” type that was marshalled into that *Addr*. Section 5 describes how high-level types are marshalled to and from their low-level equivalents.

- From an IDL *typedef* declaration, *H/Direct* generates a corresponding Haskell type declaration together with some marshalling functions. In general, a marshalling function transforms a “*high-level*” Haskell value (in this case *Point*) into a “*low-level*” Haskell value (in this case *Ptr Point*). These marshalling functions are in the *IO* monad because, as we shall see, they often work often imperatively by allocating some memory and explicitly filling it in, so as to construct a memory layout that matches the interface specification. The translations for *typedef* declarations are discussed in Section 6.
- The function *hdFree* :: *IO* () simply releases all the memory allocated by the marshalling functions.

So much for our example. The difficulty is that IDL is a complex language, so it is not always straightforward to guess

the Haskell type that will correspond to a particular IDL type, nor to generate correct marshalling code. (The former is important to the programmer, the latter only to *H/Direct* itself.) Our goal in this paper is to give a systematic translation of IDL to Haskell stub code.

To simplify translation we assume that the IDL source is brought into a standard form, that is, we factor the translation into a translation of full IDL to a core subset and a translation from core IDL to Haskell. In particular, we assume that: `out` parameters always have an explicit “*”, the pointer default is manifested in all pointer types, and all enumeration have value fields. (The details are unimportant.)

IDL is a large language, and space precludes giving a complete translation here. We do not even give a syntax for IDL, relying on the left-hand sides of the translation rules to specify the syntax we treat. However, the framework we give here is sufficient to treat the whole language, and our implementation does so.

3 Procedure declarations

The translation function $\mathcal{D}[\cdot]$ maps an IDL declaration into one or more Haskell declarations. We begin with IDL procedure declarations. To start with, we concentrate on allowing Haskell to call C; we discuss other variants in Section 7. Here is the translation rule for procedure declarations:

```

 $\mathcal{D}[t\_res\ f([in]\ t\_in, [out]\ t\_out, [in,out]\ t\_inout)]$ 
 $\mapsto$ 
 $\mathcal{T}[f] :: \mathcal{T}[t\_in] \rightarrow \mathcal{T}[t\_inout]$ 
 $\quad \rightarrow \text{IO}(\mathcal{T}[t\_out], \mathcal{T}[t\_inout], \mathcal{T}[t\_res])$ 
 $\mathcal{N}[f] = \lambda m \rightarrow \lambda n \rightarrow$ 
 $\quad \text{do } \{ a \leftarrow \mathcal{M}[t\_in] m$ 
 $\quad ; b \leftarrow \mathcal{O}[t\_out]$ 
 $\quad ; c \leftarrow \mathcal{M}[t\_inout] n$ 
 $\quad ; r \leftarrow \text{primN}[f] a b c$ 
 $\quad ; x \leftarrow \mathcal{U}[t\_out] b$ 
 $\quad ; y \leftarrow \mathcal{U}[t\_inout] c$ 
 $\quad ; z \leftarrow \mathcal{U}[t\_res] r$ 
 $\quad ; \text{hdFree}$ 
 $\quad ; \text{return } (x, y, z)$ 
 $\quad \}$ 
 $\text{foreign import stdcall primN}[f]$ 
 $\quad :: \mathcal{B}[t\_in] \rightarrow \mathcal{B}[t\_out] \rightarrow \mathcal{B}[t\_inout]$ 
 $\quad \rightarrow \text{IO} \mathcal{B}[t\_res]$ 

```

Despite our claim of formality, the fully formal version of this rule has an inconvenient number of subscripts. Instead, we illustrate by giving one parameter of each mode (`[in]`, `[out]`, and `[in, out]`); more complex cases are handled exactly analogously. The translation produces a Haskell function that takes one argument for each IDL `[in]` or `[in, out]` parameter, and returns one result of each IDL `[out]` or `[in, out]` parameter, plus one result for the IDL result (if any). In general, foreign functions can perform side effects, so the result type is in the `IO` monad. We are considering adding a (non- standard) attribute `[pure]`, that declares the procedure to have no side effects; in this case, the Haskell procedure can simply return a tuple rather than an `IO` type.

The generic translation for procedure declaration uses several auxiliary translation schemes:

$t : b$	<i>basic type</i>
$ n$	<i>type names</i>
$ \{attr\} +]t^*$	<i>pointer type</i>
$attr : \text{unique} \text{ref} \text{ptr}$	
$ \text{string} \text{size_is}(e)$	

Figure 2: IDL type syntax

- The translation scheme $\mathcal{T}[t]$ gives the “high-level” Haskell type corresponding to the IDL type t .
- The translation scheme $\mathcal{N}[n]$ does the name mangling required to translate IDL identifiers to valid Haskell identifiers. For instance to account for the fact that Haskell function names must begin with a lower-case letter.
- The translation scheme $\mathcal{B}[t]$ gives the “low-level” Haskell type corresponding to the IDL type t .
- The translation scheme $\mathcal{M}[t] :: \mathcal{T}[t] \rightarrow \text{IO} \mathcal{B}[t]$ generates Haskell code that marshals a value of IDL type t from its high-level type $\mathcal{T}[t]$ to its low-level form $\mathcal{B}[t]$. This is used to marshal all the in-parameters of the procedure (`[in]` and `[in,out]`).
- The translation scheme $\mathcal{U}[t] :: \mathcal{B}[t] \rightarrow \text{IO} \mathcal{T}[t]$ generates Haskell code that un-marshals a value of IDL type t . This is used to un-marshal all the out-parameters of the procedure, and its result (if any). $\mathcal{M}[\cdot]$ and $\mathcal{U}[\cdot]$ are mutual inverses (upto memory allocation).
- In addition, for `[out]` parameters the caller is required to allocate a location to hold the result. $\mathcal{O}[t^*] :: \text{IO} \text{Ptr} \mathcal{B}[t]$ is Haskell code that allocates enough space to contain a value of IDL type t .

We will define these functions in detail in Section 4, but first we deal with type declarations.

4 Mapping for types

Next, we turn our attention to the translations $\mathcal{T}[\cdot]$ and $\mathcal{B}[\cdot]$ that translate IDL types to Haskell types, which are given in Figure 3.

Translating base types, which have direct Haskell analogues, is easy. The high-level and low-level type translations coincide, except that the high-level representation of IDL’s 8-bit characters is Haskell’s 16 bit `Char` type. To give more precise mapping we have extended Haskell with new base types: `Word8`, `Word16`, and so on. Similarly, IDL type names are translated to the (Haskell-mangled) name of the corresponding Haskell type.

Matters start to get murkier when we meet pointers. Since a pointer is always passed to an from C as a machine address, the low-level translation of all pointer types is simply a raw machine address:

$$\mathcal{B}[t^*] \mapsto \text{Ptr } \mathcal{T}[t]$$

$B[\text{short}]$	\mapsto	Int32
$B[\text{unsigned short}]$	\mapsto	Word32
$B[\text{float}]$	\mapsto	Float
$B[\text{double}]$	\mapsto	Double
$B[\text{char}]$	\mapsto	Word8
$B[\text{wchar}]$	\mapsto	Char
$B[\text{boolean}]$	\mapsto	Bool
$B[\text{void}]$	\mapsto	()
$B[[\text{attr}]\,t^*]$	\mapsto	Ptr $\mathcal{T}[t]$

$\mathcal{T}[\text{char}]$	\mapsto	Char
$\mathcal{T}[b]$	\mapsto	$B[b]$
$\mathcal{T}[n]$	\mapsto	$N[n]$
$\mathcal{T}[[\text{ref}]\,t^*]$	\mapsto	$\mathcal{T}[t]$
$\mathcal{T}[[\text{unique}]\,t^*]$	\mapsto	Maybe $\mathcal{T}[t]$
$\mathcal{T}[[\text{ptr}]\,t^*]$	\mapsto	Ptr $\mathcal{T}[t]$
$\mathcal{T}[[\text{string}]\,\text{char}^*]$	\mapsto	String
$\mathcal{T}[[\text{size_is}(v)]\,t^*]$	\mapsto	$[\mathcal{T}[t]]$

Figure 3: Type translations

(Recall that $\text{Ptr } t$ is just an abbreviation for Addr , but the Ptr form is somewhat more informative.)

In contrast, the high-level translation of pointers depends on what type of pointer is concerned. IDL has no fewer than five kinds of pointer, distinguished by their attributes! We treat them one at a time (refer in each case to Figure 3):

- A value of IDL type $[\text{ref}]\,t^*$ is the unique pointer, or indirection, to a value of type t . Since pointers are implicit in Haskell, the corresponding high-level Haskell type is just $\mathcal{T}[t]$.
- The IDL type $[\text{unique}]\,t^*$ is exactly the same as $[\text{ref}]\,t^*$, except that the pointer can be `NULL`. The natural way to represent this possibility in Haskell is using the `Maybe` type. The latter is a standard Haskell type defined like this:

```
data Maybe a = Nothing | Just a
```

- An IDL value of type $[\text{ptr}]\,t^*$ is the address of a value that might be shared, and might contain cycles. It is far from clear how such a thing should be marshalled, so we adopt a simple convention:

```
 $\mathcal{T}[[\text{ptr}]\,t^*] \mapsto \text{Ptr } \mathcal{T}[t]$ 
```

That is, $[\text{ptr}]$ values are not moved across the border at all. Instead they are represented by a value of type $\text{Ptr } \mathcal{T}[t]$, a raw machine address.

This is often useful. For a start, some libraries implement an abstract data type, in which the client is expected to manipulate only pointers to the values. Similarly, COM interface pointers should be treated simply as addresses. Finally, some operating system procedures (notably those concerned with windows) return such huge structures that a client might want to marshal them back selectively.

- A value of type $[\text{string}]\,\text{char}^*$ is the address of a null-terminated sequence of characters. (Contrast $[\text{ref}]\,\text{char}^*$, which is the address of a single character.) The corresponding Haskell type is, of course, `String`. The `[\text{string}]` attribute applies to the following array types `char`, `byte`, `unsigned short`, `unsigned long`, `structs` with `byte` (only!) fields and, in Microsoft-only IDL, `wchar`.

- Sometimes a procedure takes a parameter that is a pointer to an array of values, where another parameter of the procedure gives the size of the array. For example:

```
void DrawPolygon
  ( [in, size_is(nPoints)] Point* points
  , [in] int nPoints
  );
```

The `[\text{size_is}(nPoints)]` attribute tells that the second parameter, `nPoints`, gives the size of the array. (This is quite like the `[\text{string}]` case, except that the size of the array is given separately, whereas strings have a sentinel at the end.) We translate arrays to Haskell lists.

While each of these variants has a reasonable rationale, we have found the plethora of IDL pointer types to be a rich source of confusion. The translations in Figure 3 look innocuous enough, but we have found them extremely helpful in clarifying and formalising just exactly what the translation of an IDL type should be.

Even if the translation are not quite “right” (whatever that means), we now have a language in which to discuss variants. For example, it may eventually turn out that the IDL `[\text{ptr}]` attribute is conventionally used for subtly different purposes than the ones we suggest above. If so, the translations can readily be changed, and the changes explained to programmers in a precise way.

5 Marshalling

In the translation of the IDL type signature for a procedure (Section 3), we invoked marshalling functions $\mathcal{M}[\cdot]$ and $\mathcal{U}[\cdot]$ for each of the types involved. Now that we have defined the high and low-level translations of each type, the marshalling code is relatively easy to define. In this section we define these marshalling functions. Lack of space precludes us from giving complete details so we will concentrate mostly on marshalling basic types.

Marshalling a structured value consists, as we shall see, of two steps: allocate some memory in the *parameter-marshalling area* to hold the value, and then actually marshal the Haskell value into that memory. The translations are much more elegant if we define auxiliary schemes, $\mathcal{W}[\cdot]$ and $\mathcal{R}[\cdot]$, that perform this “by-reference” marshalling. We also need a number of functions to manipulate the parameter-marshalling area. More precisely:

$\mathcal{W}[t] :: \text{Ptr } \mathcal{T}[t] \rightarrow \mathcal{T}[t] \rightarrow \text{IO } ()$ marshals its second argument into the memory location(s) pointed to by its first argument; the latter is a raw machine address.

$\mathcal{R}[t] :: \text{Ptr } \mathcal{T}[t] \rightarrow \text{IO } \mathcal{T}[t]$ unmarshals a value of IDL type t out of memory location(s) pointed to by its argument. $\mathcal{W}[\cdot]$ and $\mathcal{R}[\cdot]$ are mutually inverse (upto memory allocation).

$\mathcal{S}[t] :: \text{Int}$ is the number of bytes occupied by an IDL value of type t . The function $\mathcal{O}[\cdot]$, mentioned in Section 3, is defined thus:

$$\mathcal{O}[[attr] t*] \mapsto \text{hdAlloc } \mathcal{S}[t]$$

$\text{hdAlloc} :: \text{Int} \rightarrow \text{IO } (\text{Ptr } a)$ allocates the specified number of bytes in the parameter-marshalling area, returning a pointer to the allocated area.

$\text{hdWrite } b :: \text{Ptr } \mathcal{T}[b] \rightarrow \mathcal{T}[b] \rightarrow \text{IO } ()$, where b is a basic type, marshals a value of IDL type b into the specified memory location(s).

$\text{hdRead } b :: \text{Ptr } \mathcal{T}[b] \rightarrow \text{IO } \mathcal{T}[b]$, where b is a basic type, unmarshals a value of IDL type t .

$\text{hdFree} :: \text{IO } ()$ frees the whole parameter-marshalling area.

With these definitions in mind, Figure 4 gives the marshalling schemes. We omit the schemes for `[size_is]` because it is tiresomely complicated. Apart from that, the translations are easy to read:

- For basic types there is no marshalling to do, except that we must convert between the 16-bit Haskell `Char` and 8-bit IDL `char` types.
- Marshalling a `typedef'd` type can be done by invoking its marshalling function.
- Marshalling a `[ref]` pointer is done by allocating some memory with `hdAlloc`, and then marshalling the value into it with $\mathcal{W}[\cdot]$. Unmarshalling is similar, except that there is no allocation step; we just invoke $\mathcal{R}[\cdot]$.
- Dealing with `[unique]` pointers is similar, except that we have to take account of the possibility of a `NULL` value.

Again, it is very helpful to have a precise language in which to discuss these translations. Though they look simple, we can attest that it is very easy to get confused by pointers to pointers to things, and we have far greater confidence in our implementation as a result of writing the translations formally.

6 Type declarations

On top of the primitive base types, IDL supports the definition of a number of constructed types. For example

```
typedef int trip[3];
typedef struct TagPoint { int x,y; } Point;
typedef enum { Red=0, Blue=1, Green=2 } RGB;
typedef union _floats switch (int ftype) {
    case 0: float f;
    case 1: double d;
} Floats;
```

$\mathcal{M}[t] :: \mathcal{T}[t] \rightarrow \text{IO } \mathcal{B}[t]$	
$\mathcal{M}[\text{char}]$	$\mapsto \text{marshallChar}$
$\mathcal{M}[b]$	$\mapsto \text{return}$
$\mathcal{M}[n]$	$\mapsto \text{marshalln}$
$\mathcal{M}[[\text{ref}] t*]$	$\mapsto \backslash x \rightarrow$ $\text{do}\{ \text{px} \leftarrow \text{hdAlloc } \mathcal{S}[t]$ $; \mathcal{W}[t] \text{ px } x \}$
$\mathcal{M}[[\text{unique}] t*]$	$\mapsto \backslash x \rightarrow$ $\text{case } x \text{ of}$ $\text{Nothing} \rightarrow \text{return nullPtr}$ $\text{Just } y \rightarrow \mathcal{M}[[\text{ref}] t*] y$
$\mathcal{M}[[\text{ptr}] t*]$	$\mapsto \text{return}$
$\mathcal{M}[[\text{string}] t*]$	$\mapsto \text{marshallString}$
$\mathcal{W}[t] :: \text{Ptr } \mathcal{T}[t] \rightarrow \mathcal{T}[t] \rightarrow \text{IO } ()$	
$\mathcal{W}[b]$	$\mapsto \text{hdWrite } b$
$\mathcal{W}[[attr] t*]$	$\mapsto \backslash p \ x \rightarrow$ $\text{do}\{ a \leftarrow \mathcal{M}[[attr] t*] \ x$ $; \text{hdWriteAddr } p \ a \}$
$\mathcal{U}[t] :: \mathcal{B}[t] \rightarrow \text{IO } \mathcal{T}[t]$	
$\mathcal{U}[\text{char}]$	$\mapsto \text{unmarshallChar}$
$\mathcal{U}[b]$	$\mapsto \text{return}$
$\mathcal{U}[n]$	$\mapsto \text{unmarshalln}$
$\mathcal{U}[[\text{ref}] t*]$	$\mapsto \mathcal{R}[t]$
$\mathcal{U}[[\text{unique}] t*]$	$\mapsto \backslash p \rightarrow$ $\text{if } p == \text{nullPtr} \text{ then}$ return Nothing else $\text{do}\{ x \leftarrow \mathcal{R}[t] \ p$ $; \text{return (Just } x \text{)} \}$
$\mathcal{U}[[\text{ptr}] t*]$	$\mapsto \text{return}$
$\mathcal{U}[[\text{string}] t*]$	$\mapsto \text{unmarshallString}$
$\mathcal{R}[t] :: \text{Ptr } \mathcal{T}[t] \rightarrow \text{IO } \mathcal{T}[t]$	
$\mathcal{R}[b]$	$\mapsto \text{hdRead } b$
$\mathcal{R}[[attr] t*]$	$\mapsto \backslash p \rightarrow$ $\text{do}\{ a \leftarrow \text{hdReadAddr } p$ $; \mathcal{U}[[attr] t*] a \}$

Figure 4: The marshalling schemes

```

t : t[e]                                array type
| enum {
|   tag1 = v1, ..., tagn = vn}      enumeration
| struct tag {
|   f1 : t1; ...; fn : tn;}          record type
| union tag1
|   switch ( b tag2 ) {
|     case v1:t1 f1; ... case vn:tn fn; }  union type

```

Figure 5: IDL constructed type syntax

which declares array, record, enumeration and union (or sum) types, respectively. Figure 5 shows the syntax of IDL's constructed types.

The translation provides rules for converting between IDL constructed types into corresponding Haskell representations. To ease the task of defining this type mapping, we assume that each constructed type appears as part of an IDL type declaration. In general, a type declaration has the following form:

```
typedef t name;
```

declaring *name* to be a synonym for the type *t*, which is either a base type or one of the above constructed types. A type declaration for an IDL type *t* gives rise to the definition of the following Haskell declarations:

- A Haskell type declaration for the Haskell type $\mathcal{N}[\text{name}]$, such that $\mathcal{T}[\text{name}] = \mathcal{N}[\text{name}]$.
- `marshallN[name] :: T[name] -> IO B[t]` which implements the $\mathcal{M}[\]$ scheme for converting from the Haskell representation $\mathcal{T}[\text{t}]$ to the IDL type *t*.
- `unmarshallN[name] :: B[t] -> IO T[name]` which implements the dual $\mathcal{U}[\]$ scheme for unmarshalling.
- `marshallN[name]At :: Ptr B[t] -> T[name]` -> `IO ()` for performing by-reference marshalling of the constructed type.
- `unmarshallN[name]At :: Ptr B[t] -> IO T[name]` which implements the $\mathcal{R}[\]$ scheme for unmarshalling a constructed type by-reference.
- `sizeofN[name] :: Int`, a constant holding the size of the external representation of the type (in 8-bit bytes.)

The general rules for converting type declarations into Haskell types is presented in Figure 6. Here is what they generate when applied:

- In the case of a type declaration for a base type, this merely defines a type synonym. For example

```
typedef int year;
```

is translated into the type synonym

```
type Year = Int
```

plus marshalling functions for `Year`.

- For a record type such as `Point`:

```
typedef struct TagPoint {int x,y;} Point;
```

generates a single constructor Haskell data type:

```
data Point = TagPoint { x:: Int, y::Int }
```

In addition to this, the $\mathcal{D}[\]$ scheme generates a collection of marshalling functions, including `marshallPoint`:

```

marshallPoint :: Point -> IO (Ptr Point)
marshallPoint (Point x y) =
  do{ ptr <- hdAlloc sizeofPoint
      ; let ptr1 = addPtr ptr 0
      ; marshallIntAt ptr1 x
      ; let ptr2 = addPtr ptr1 sizeofInt
      ; marshallIntAt ptr2 y
      ; return ptr
    }

```

It marshals a `Point` by allocating enough memory to hold the external representation of the point. The size of the record type is computed as follows:

```

sizeofPoint :: Int32
sizeofPoint = structSize [sizeofInt,sizeofInt]

```

where `structSize` is a (platform specific) function that computes the size of a `struct` given the field sizes.³

`Point`'s two fields are marshalled into the external representation of `Point` by calling the by-reference marshaller for the basic type `Int`, supplying a pointer that has been appropriately offset.

- For the union type example given at the start of Section 6, the following Haskell type is generated:

```
data Floats = F Float | D Double
```

together with actions for marshalling between the algebraic type and a union (omitting the type signatures for the by-reference marshallers):

```

marshallFloats :: Floats -> IO (Ptr Floats)
unmarshallFloats :: Ptr Floats -> IO Floats

```

The external representation of a union is normally a `struct` containing the discriminant and enough room to accommodate the largest member of the union. In the case of `Floats`, the external representation must be large enough to contain and `int` and a `double`.

- Enumerations have a direct Haskell equivalent as algebraic data types with nullary constructors. For example, the `RGB` declaration:

³Similarly, a function that returns the offsets at which to marshal each field into is also provided. Due to lack of space, `marshallPoint` makes the simplifying assumption that structures contain no internal padding.

```


$$\begin{aligned}
\mathcal{D}[\text{typedef } t \text{ name;}] &\rightarrow \text{type } \mathcal{N}[\text{name}] = \mathcal{T}[t] \\
&\quad \text{marshall} \mathcal{N}[\text{name}] = \text{marshall} \mathcal{T}[t] \\
&\quad \text{marshall} \mathcal{N}[\text{name}] \text{At} = \text{marshall} \mathcal{T}[t] \text{At} \\
&\quad \text{unmarshall} \mathcal{N}[\text{name}] = \text{unmarshall} \mathcal{T}[t] \\
&\quad \text{unmarshall} \mathcal{N}[\text{name}] \text{At} = \text{unmarshall} \mathcal{T}[t] \text{At} \\
&\quad \text{sizeof} \mathcal{N}[\text{name}] = \mathcal{S}[t]
\end{aligned}$$


$$\begin{aligned}
\mathcal{D}[\text{typedef } t \text{ name}[dim];] &\rightarrow \text{type } \mathcal{N}[\text{name}] = [ \mathcal{T}[t] ] \\
&\quad \text{marshall} \mathcal{N}[\text{name}] = \text{marshall} \text{Array } dim \text{ marshall} \mathcal{T}[t] \text{At} \\
&\quad \text{marshall} \mathcal{N}[\text{name}] \text{At} = \text{marshall} \text{ArrayAt } dim \text{ marshall} \mathcal{T}[t] \text{At} \\
&\quad \text{unmarshall} \mathcal{N}[\text{name}] = \text{unmarshall} \text{Array } dim \text{ unmarshall} \mathcal{T}[t] \text{At} \\
&\quad \text{unmarshall} \mathcal{N}[\text{name}] \text{At} = \text{unmarshall} \text{ArrayAt } dim \text{ unmarshall} \mathcal{T}[t] \text{At} \\
&\quad \text{sizeof} \mathcal{N}[\text{name}] = dim * \mathcal{S}[t]
\end{aligned}$$


$$\begin{aligned}
\mathcal{D}[\text{typedef struct tag}\{ \dots; t \text{ field; } \dots \} \text{ name;}] &\rightarrow \\
&\quad \text{data } \mathcal{N}[\text{name}] = \mathcal{N}[\text{tag}]\{ \dots, \mathcal{N}[\text{field}_i] :: \mathcal{T}[t_i], \dots \} \\
&\quad \text{marshall} \mathcal{N}[\text{name}] \text{ rec} = \text{do} \\
&\quad \quad \text{ptr} \leftarrow \text{hdAlloc } \mathcal{S}[\text{name}] \\
&\quad \quad \text{marshall} \mathcal{N}[\text{name}] \text{At } \text{ptr } \text{rec} \\
&\quad \quad \text{return } \text{ptr}
\end{aligned}$$


$$\begin{aligned}
&\quad \text{marshall} \mathcal{N}[\text{name}] \text{At } \text{ptr } (\mathcal{N}[\text{tag}]\{ \dots, \mathcal{N}[\text{field}_i], \dots \}) = \text{do} \\
&\quad \quad \text{let } \text{ptr}_1 = \text{addPtr } \text{ptr } 0 \\
&\quad \quad \dots \\
&\quad \quad \text{let } \text{ptr}_i = \text{addPtr } \text{ptr}_{i-1} \mathcal{S}[t_{i-1}] \\
&\quad \quad \mathcal{W}[t_i] \text{ ptr}_i \text{ field}_i \\
&\quad \quad \dots \\
&\quad \quad \text{return } ()
\end{aligned}$$


$$\begin{aligned}
\text{unmarshall} \mathcal{N}[\text{name}] &= \text{unmarshall} \mathcal{N}[\text{name}] \text{At} \\
\text{unmarshall} \mathcal{N}[\text{name}] \text{At } \text{ptr} &= \text{do} \\
&\quad \text{let } \text{ptr}_1 = \text{addPtr } \text{ptr } 0 \\
&\quad \quad \dots \\
&\quad \quad \text{let } \text{ptr}_i = \text{addPtr } \text{ptr}_{i-1} \mathcal{S}[t_{i-1}] \\
&\quad \quad \mathcal{N}[\text{field}_i] \leftarrow \mathcal{R}[t_i] \text{ ptr}_i \\
&\quad \quad \dots \\
&\quad \quad \text{return } (\mathcal{N}[\text{tag}] \dots \mathcal{N}[\text{field}_i] \dots) \\
&\quad \text{sizeof} \mathcal{N}[\text{name}] = \text{structSize } [\dots, \mathcal{S}[\text{field}_i], \dots]
\end{aligned}$$


$$\begin{aligned}
\mathcal{D}[\text{typedef enum } \{ \dots, alt = value, \dots \} \text{ name;}] &\rightarrow \\
&\quad \text{data } \mathcal{N}[\text{name}] = \dots \mid \mathcal{N}[\text{alt}] \mid \dots \\
&\quad \text{marshall} \mathcal{N}[\text{name}] \text{ x} = \\
&\quad \quad \text{case } \text{x} \text{ of } \{ \dots; \mathcal{N}[\text{alt}] \rightarrow \mathcal{N}[\text{value}]; \dots \} \\
&\quad \quad \text{unmarshall} \mathcal{N}[\text{name}] \text{ x} = \\
&\quad \quad \text{case } \text{x} \text{ of } \{ \dots; \mathcal{N}[\text{value}] \rightarrow \text{return } \mathcal{N}[\text{alt}]; \dots \} \\
&\quad \quad \text{unmarshall} \mathcal{N}[\text{name}] \text{At } \text{ptr} = \text{do} \\
&\quad \quad \quad \text{v} \leftarrow \text{hdReadInt } \text{ptr} \\
&\quad \quad \quad \text{unmarshall} \mathcal{N}[\text{name}] \text{ v} \\
&\quad \quad \quad \text{sizeof} \mathcal{N}[\text{name}] = \text{sizeofint}
\end{aligned}$$


```

Figure 6: Translating declarations

```

typedef enum {red=0,green=1,blue=2} RGB;

is translated into the Haskell type

data RGB = Red | Green | Blue

with concrete representation  $B[\text{RGB}] = \text{Int32}$ 

The marshalling actions simply map between the nullary constructors and Int32:

marshallRGB :: RGB -> IO Int32
marshallRGB nm =
  return (case { Red -> 0; Green -> 1; Blue ->2} )

unmarshallRGB :: Int32 -> IO RGB
unmarshallRGB v =
  case v of
    0 -> return Red
    1 -> return Green
    2 -> return Blue
    _ -> fail (userError ...)

```

7 The inverse mapping

Once marshalling and un-marshalling functions are defined for each data type, it is not hard to reverse the mapping and build code that allows C to call Haskell. The translation for a `typedef` remains unchanged, but the translation for an IDL procedure declaration is reversed. Since the procedure is being implemented in Haskell, its `[in]`-parameters are un-marshalled, the Haskell procedure is called, its results are marshalled, and returned to the caller. (We omit the details, but the translation rule can be expressed just as we did in Section 3. For example, the `Move` IDL declaration of that Section would be compiled to the following Haskell code:

```

foreign export stdcall "Move"
  primMove :: Ptr Point -> IO ()

primMove a =
  do { p <- unmarshallPoint a
      ; q <- move p
      ; marshallPointAt a q
      ; return ()
    }

move :: Point -> IO Point
move = error "Not yet implemented"

```

The `foreign export` declaration asks the Haskell compiler to make `Move` externally callable with a `stdcall` interface. `primMove` does the marshalling, before calling `move`, which should be provided by the programmer.

We are also interested in allowing Haskell programs to create and invoke COM objects, and in allowing a Haskell program to be sealed up inside a COM object. This too is a straightforward extension. There are a couple of wrinkles, however:

- COM methods conventionally return a value of type `HRESULT`, which is used to signal exceptional conditions. *H/Direct* “knows” about `HRESULT` and reflects its exceptional values into exceptions in Haskell’s `IO` monad.

- COM methods are invoked indirectly, through a vector table. To support this the Haskell `foreign` declaration has to be extended to allow indirect calls. For example, the Haskell-to-COM side looks like this:

```

foreign import stdcall
  dynamic primFoo :: Addr -> ...

```

The keyword `dynamic` replaces the static name of the foreign function, and the address of the function is instead passed as the first argument to `primFoo`. The `foreign export` case is similar.

- Lastly, there are several design choices concerning what the programmer has to write to implement a COM object. Does she write a collection of functions that take the object state as their first argument? Or does she write a single function that returns a record of all the methods of the object?

8 Status and conclusions

H/Direct is now our fourth attempt at a foreign-language interface for Haskell. The first was `ccall`, a limited and low-level extension roughly equivalent to `foreign import` [3]. The second was Green Card, which gradually turned into a domain-specific language [5]. The third was a pre-cursor to *H/Direct*, Red Card, which was specifically aimed at interfacing Haskell to COM objects, [4, 6]. *H/Direct* embodies the lessons we have learned: strive for implementation-independence; avoid inventing new languages; the customer is always right.

We do not claim great originality for these observations. What is new in this paper is a much more precise description of the mapping between Haskell and IDL than is usually given. This precision has exposed details of the mapping that would otherwise quite likely have been mis-implemented. Indeed, the specification of how pointers are translated exposed a bug in our current implementation of *H/Direct*. It also allows us automatically to support nested structures and other relatively complicated types, without great difficulty. These aspects often go un-implemented in other foreign-language interfaces.

We are well advanced on an implementation of *H/Direct*. We can parse and type-check the whole of Microsoft IDL, and can generate stubs that allow Haskell to call C and COM. We have not yet implemented the reverse mapping, but we expect to do so in the next few months.

Acknowledgements

We thank Conal Elliott for playing the vital role of Friendly Customer; much of our motivation derives from his desired applications. We thank EPSRC and Microsoft for their support, both of equipment and manpower. Erik Meijer would like to thank the PacSoft group at the Oregon Graduate Institute for their hospitality during the final phases of writing this paper.

References

- [1] D. Beazley. SWIG and automated C/C++ scripting extensions. *Dr. Dobb's Journal*, Februar 1998.
- [2] Sigbjorn Finne et al. A primitive foreign function interface for Haskell. In preparation, preliminary specification available from <http://www.dcs.gla.ac.uk/~sof/primitive.ps.gz>, March 1998.
- [3] Simon L. Peyton Jones and Philip Wadler. Imperative functional programming. In *POPL 20*, pages 71–84, 1993.
- [4] Simon Peyton Jones, Erik Meijer, and Daan Leijen. Scripting COM components from Haskell. In *Proceedings of ICSR5*, 1998.
- [5] Simon Peyton Jones, Thomas Nordin, and Alastair Reid. Green Card: a foreign-language interface for Haskell. In *Proc. Haskell Workshop*, 1997.
- [6] Daan Leijen. Red card: Interfacing Haskell with COM. Master's thesis, University of Amsterdam, 1998.
- [7] X/Open Company Ltd. *X/Open Preliminary Specification X/Open DCE: Remote Procedure Call*, 1993.
- [8] Microsoft, <http://www.microsoft.com/java/>. *SDK for Java*, 1998.
- [9] Microsoft Press. *Developing for Microsoft Agent*, 1998.
- [10] Dale Rogerson. *Inside COM*. Microsoft Press, 1997.
- [11] Jon Siegel. *CORBA Fundamentals and Programming*. John Wiley & Sons, 1996.
- [12] A Vogel, B Gray, and K Duddy. Understanding any IDL, lesson one: DCE and CORBA. In *Proceedings of SDNE'96*, 1996.
- [13] A Vogel and B Grey. Translating DCE IDL in OMG IDL and vice versa. Technical Report 22, CRC for Distributed Systems Technology, Brisbane, 1995.