

Foreign Interface for PLT Scheme

Eli Barzilay
Northeastern University

Dmitry Orlovsky
Northeastern University

Abstract

Even a programmer devoted to Scheme may prefer using foreign libraries in certain situation. Connecting the two worlds involves glue code, usually using C, which requires significant programming efforts and system expertise. In this paper we describe a PLT Scheme extension for interacting with foreign code, designed around a simple philosophy: *stay in the fun world*, even if it is no longer a safe sand box. Our system relieves the programmer from low-level technicalities while keeping the benefits of Scheme as a better programming environment compared to C.

1 Introduction

Scheme has proved itself as a useful and fun language, good for both general-purpose and domain-specific usages. However, schemers cannot assume a closed system; other languages will always exist, leading to a need for interfacing with functionality that is accessible through foreign libraries. Such libraries come in many different flavors, but the popular ‘least common denominator’ has been, and still is, plain C libraries¹. Our goal is to create a mechanism within Scheme for smooth interfacing with such foreign libraries.

1.1 Background

A foreign interface is a piece of *glue code*, intended to make it possible to use functionality written in one language (often C) available to programs written in another (usually high-level) language. Such glue code involves low-level details that users of high-level languages usually take for granted. For example:

- marshaling objects to and from foreign code,
- managing memory and other resources,
- dealing with different calling conventions, implicit function arguments, etc.

¹Different languages can be used to create foreign libraries, “C” is only used as a generic label.

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

Fifth Workshop on Scheme and Functional Programming. September 22, 2004, Snowbird, Utah, USA. Copyright 2004 Eli Barzilay.

Foreign function interfaces are subsystems that create such glue code, simplifying an otherwise tedious and error-prone task.

1.2 Foreign Interfaces

There are lots of existing foreign function interfaces; Urban’s FFI survey [17], although an incomplete project, provides a good discussion of such systems and relevant issues. Generally speaking, such interfaces can be classified as either static code generators or dynamic foreign interfaces. In principle, the two are quite similar:

- A *static* foreign interface is created and compiled statically, *before* running the program that intends to use it;
- A *dynamic* interface is created at run-time, *while* the application is running.

In practice, the differences are more dramatic:

- A static interface is usually implemented using a C compiler. The advantage of this approach is that it is easy interface foreign code, as most of it is intended to be linked in using a compiler (for example, C header files are used to describe an interface), and since most languages are implemented in C, they provide convenient facilities for calling C functions. Disadvantages of the static approach include being restricted to the pre-compiled interface, requiring either a compiler or a platform dependent binary distribution for such code.
- A dynamic interface is generated at run-time, leading to the obvious advantage of requiring no C compiler or binary distributions. This has a significant effect on dynamic languages like Scheme, where single running REPL can be used to connect to different libraries, supporting exploratory programming in a natural way. The disadvantage of this approach is that it requires more (platform-dependent) low-level work such as stack management and creating stubs (glue functions), while not getting the usual support from a C compiler.

The issues that need addressing are essentially the same ones described in Section 1.1, only the approach differs. The technical issues involved in an interface implementation make static interface generators more popular. It should be noted that it is common to call these systems “foreign *function* interfaces” — in the following text we prefer “foreign interfaces” as these interfaces deal with accessing foreign objects as well as foreign calls.

In both the static and the dynamic cases, it is desirable to have some description of the foreign entities, usually functions, in a way that can help automate the process of generating the glue layer. In this context a “function” can be viewed differently depending on your point of view: from the low-level side, a function is simply a pointer and a description of how it is called; from the high-level Scheme

side, it is an object that is expected to have the usual function semantics. *Interface description languages* (IDLs) have a major role in foreign interface systems — these are languages that express arbitrary function behaviors for both of these viewpoints:

- On the C side, there is the type definition of the function, and possibly additional information such as input/output pointers, object ownership, etc.
- In addition to this, there are details that are related to the Scheme side. For example, automatic memory management issues, value marshaling, dealing with aggregates (vectors and structs), and creating new object types.
- On the Scheme side, the result is a plain procedure, like any other Scheme procedure object.

Ideally, the IDL that is used to describe the interface is rich enough to express both views while providing enough information to completely automate the interface generation.

1.3 Implementing a Dynamic Interface

The low-level mechanics of foreign function calls are usually very demanding: managing functions at the binary level is inherently platform dependent, and can even require assembly code or other compiler-specific hacks. Statically, these problems are not too difficult: simply generate C glue code, and let the C compiler do its usual work. Doing this efficiently in a dynamic fashion is difficult, since it is usually not desirable to drag a complete C compiler into your run-time. Dealing with the dynamic aspects of foreign functions is greatly simplified using a library that handles the low-level details: we use `libffi` [11], a library that supports foreign function call-outs and call-backs.

- A call-out is a normal function call. In a dynamic setting, we create a “call-interface” object which specifies (binary) input/output types; this object can be used with an arbitrary function pointer and an array of input values to perform a call-out to the function and retrieve its result. Doing this requires manipulating the stack and knowing how a function is called, these are details that `libffi` deals with.
- A call-back is trickier. Our Scheme implementation has several fixed C-level functions which can implement arbitrary Scheme evaluation. A callback is, however, a simple function pointer — no additional information is available. Modern systems (e.g. Gnome) that use callbacks allow user to register a function pointer together with an arbitrary data pointer, but there is no standard way for this. A proper solution is one that allows creating general “C closures” — combining a function and a data pointer into a single new function pointer. Again, this is technically challenging, as it requires generating stub functions at run-time, which, when applied to some arguments, call the packaged function with the packaged data pointer and the arguments. Again, `libffi` provides the required magic.

`libffi` is maintained and distributed as part of the GCC project, but its goal is to provide a portable library. We use it for all platforms that PLT Scheme targets, including Windows (using a slightly adapted version that works with Microsoft’s compiler, courtesy of the Thomas Heller [13]).

1.4 Outline

In Section 2 we state the goal of our work, emphasizing our main design principle. Section 3 describes our implementation, both the C part of the code and the complementing Scheme module.

Section 4 demonstrates how our system copes with some of the common and uncommon situations that interface programmers deal with. We conclude with a related work comparison, and outline future plans.

2 Goal: Use Foreign Libraries, Avoid C

Our design follows a simple principle: keep C-level functionality to a minimum. The core of a system for interfacing foreign libraries must itself be written in C, but we try to make such functionality available to Scheme as soon as possible, putting more responsibility on the Scheme level. When dealing with the many details of the interface, mainly type declarations and data marshaling, there is a natural tendency to make a system that is rich in features. We avoid dealing with such complexities in C when possible, providing just enough of an interface that makes it possible to do it in Scheme instead. The combination of a dynamic interface and a minimalistic C-level implementation that should be complemented by Scheme code are the main features that make our approach unique.

Switching more responsibility to Scheme comes with benefits that are familiar to Scheme programmers, but there is an additional advantage that is important in this particular case: the important issue is generating glue code that bridges the gap between foreign libraries and the high-level language. In the static case this involves either complex yet limited C preprocessor acrobatics (e.g., SWIG [1] goes as far as implementing its own C parser). On the other hand, Scheme already comes with a superior syntax system, and PLT Scheme makes this even better with additional language features (syntax objects, module system, etc). This syntax system is much easier for implementing sophisticated glue code with, especially considering our target crowd which undoubtedly feels more at home in the Scheme world.

For example, consider the issue of primitive foreign types that are handled by an interface. Once we can move C integers from Scheme to C and back, we might consider extending the system to deal with C enumerations. This raises a few questions regarding the interface design — how should this C definition:

```
typedef enum { foo1, foo2, foo3 } foo;
```

be available for Scheme code?

- Should we provide three integer bindings? If so, how do we deal with name clashes?
- Otherwise, should we use a mapping from strings to integers? Maybe use symbols? What about enumerated values that are or-able bit patterns? How should such a map be implemented: as a linked list? A vector of constant names? A hash table?

Answers to these questions determine the nature of the C implementation; once it is written, trying alternatives lead to significant maintenance costs. Our design keeps such complications away from the C level, pushing them up to the Scheme side where there are better ways to deal with them. For example, the C level part of our interface does not commit to a specific implementation for enumerations — it simply exposes C integers. Different strategies are then implemented in the Scheme part, resulting in easier code maintenance. In addition, some Scheme aspects are less accessible from C, making a Scheme solution even more attractive. for example, implementing enumerations as bindings that use the module system to avoid global name-space pollution, or implementing them as syntax objects (removing run-time lookup costs) are both much harder to implement in C than in Scheme.

Another important factor in the complexity of the C implementa-

and 'c-read' argument, and its value is bound to 'buf'.

- The third argument has a value expression so the value that is passed on to the foreign function is always the length of the second (string) argument.

'c-read' is therefore a Scheme procedure that expects two arguments and returns an integer, by arranging for properly calling the foreign 'read'.

In some rare cases, an interface needs to have better control of the wrapper's argument list — which is the purpose of the optional '<args> ::' prefix: it specifies the arguments to the resulting wrapper function. For example, if 'read' were to expect the buffer size first, we would use this _fun type:

```
(_fun (fd buf) ::
  (fd : _int)
  (_int = (string-length buf))
  (buf : _string)
  -> _int)
```

Note that identifiers are important here, as they connect the foreign inputs with the wrapper's inputs. The '<args>' part can also be used to specify normal Scheme argument lists, including optional arguments.

A second '->' marker denotes a result expression different than the one that the foreign function returned. This expression can use any bound values and arguments, as well as the foreign result value (if given an identifier). For example, the 'modf' interface given above is better written with _fun as:

```
(define c-modf
  (get-ffi-obj "modf" "libc.so.6"
    (_fun _double (p : _pointer = (malloc _double))
      -> (r : _double)
      -> (values (ptr-ref p _double) r))))
```

The fact that we can insert any Scheme expression for the return value makes it easy to change such definitions so they use alternative ways for assembling the return values, for example, changing 'values' to 'cons' in the above. If this was implemented in C, such changes would require more work.

The similarity between the _fun syntax and 'let*' is not incidental: _fun assembles a wrapper function that contains a single 'let*' expression, which evaluates the various expressions, binding the results to specified identifiers. For example, the usage of _fun in the last example expands to:

```
(_cprocedure (list _double _pointer) _double
  (lambda (ffi)
    (lambda (tmp15)
      (let* ((p (malloc _double))
             (r (ffi tmp15 p)))
        (values (ptr-ref p _double) r))))))
```

This satisfies the efficiency requirement: only one extra function call is wrapped around the foreign call.

3.2.7 Additional IDL Features: Custom Function Types

The _fun facility handles some common cases where we need to bridge a gap between the foreign function and Scheme code that uses it, but there are additional cases that are not addressed. For example, the 'modf' interface code above represents such a common situation — output pointers that are used by foreign code to return multiple values. We therefore extend the _fun syntax further, by making it interact with special 'custom function types' that

Custom Type	Usage
_ptr	input, output, or input/output pointers
_box	similar to an input/output _ptr, but modifies the Scheme box contents (PLT Scheme has a mutable box type. Note that we don't need to associate Scheme boxes with 'shadow' pointers: either copy values, or use a _pointer instead of a box)
_list, _vector	marshal lists and vectors as C pointers
_bytes	uses Scheme byte-strings (raw, non-Unicode strings)
_?	a special non-type intended for saving intermediate interface results

Table 4. Simple types defined by the Scheme module

are themselves syntaxes — such types can install pieces of code that are used before and after the foreign call, possibly modifying the corresponding value. In the case of output pointers we want to allocate some memory before the foreign call and dereference it afterward, a task that is achieved by the _ptr custom type. _ptr is a syntax with usages that has the following form:

```
(_ptr <mode> <type-expr>)
<mode> ::= i | o | io
```

The <mode> specifies an input, output, or input/output pointer. In the 'modf' case, we use an output pointer:

```
(define c-modf
  (get-ffi-obj "modf" "libc.so.6"
    (_fun _double (p : (_ptr o _double))
      -> (r : _double) -> (values p r))))
```

The code that is generated by this _fun syntax is similar to the previous code,

```
(lambda (tmp15)
  (let* ((p (malloc _double))
         (r (ffi tmp15 p))
         (p (ptr-ref p _double)))
    (values p r)))
```

but notice that we don't need to explicitly allocate a double or dereference the pointer.

The custom function types that are provided by the 'foreign' are listed in Table 4. Further details on these types can be found in our user manual.

As mentioned above, Custom types are implemented as syntaxes. _fun tries to expand each type expression it encounters, and if an expansion is identified as a custom type, then it has certain forms that contain the relevant pieces of code. A custom type expansion is a '(<key:> <val> ...)' sequence where all of the <key:>s are from a short list of known keys. Each key interacts with generated wrapper functions in a different way, which affects how its corresponding argument is treated:

type: specifies the foreign type to be used (#f can be used to make this not participate in the foreign call).

expr: specifies an expression to be used for arguments of this type, removing it from wrapper arguments.

bind: specifies a name that is bound to the original argument if it is required later (e.g., _box needs to refer to the original box).

1st-arg: specifies a name that can be used to refer to the first argument of the foreign call (good for common cases where the first argument has a special meaning, e.g., for method calls).

prev-arg: similar to 1st-arg, but refers to the previous argument.

pre: a pre-foreign code chunk that is used to change the argument’s value.

post: a similar post-foreign code chunk.

The following is the implementation of the `_ptr` custom type from the ‘foreign’ module. It is provided to roughly demonstrate how this is done; again, complete details are given in the user manual.

```
(define-syntax _ptr
  (syntax-rules (i o io)
    [(_ i t)
     ;; input: malloc a pointer, set its value from the argument
     (type: _pointer
      pre: (x => (let ([p (malloc t)]) (ptr-set! p t x) p)))]
    [(_ o t)
     ;; output: malloc a pointer on entry, dereference on exit
     (type: _pointer
      pre: (malloc t)
      post: (x => (ptr-ref x t)))]
    [(_ io t)
     ;; input/output: like output, but set its contents on entry
     (type: _pointer
      pre: (x => (let ([p (malloc t)]) (ptr-set! p t x) p))
      post: (x => (ptr-ref x t)))]))
```

All of the special custom types provided by ‘foreign’ are defined this way.

To conclude: our `_fun` satisfies all requirements mentioned above for a good IDL: it is easy to read and write, it can express all wrapper interactions that other IDLs can express and more, it is efficient, and extensible by the ability to add new custom types that handle new kinds of processing. As expected from a syntax transformer that performs some substantial work, it carries some conceptual overhead, but we believe that overall it is better than the C processing alternatives since Scheme is superior in its syntactical abstraction capabilities.

4 Usage Examples

With the implementation of our system, we provide a few (mostly Linux) library interfaces. This was used to test the implementation, motivating the overall design. We now describe a few examples of using our system, all based on these interface implementations.

Syntactic Abstractions

C provides some (limited) degree of syntactic abstraction, whereas Scheme truly shines in this area. When a complete library interface is desired (rather than pulling out a few useful functions), repetition is common. Writing interfaces in Scheme makes such problems almost non-existent — for example, our ImageMagick interface uses a simple macro:

```
(define-syntax defmagick
  (syntax-rules (:)
    [(_ id : x ...)
     (define id
      (get-ffi-obj 'id libwand (_fun x ...)))]))
```

to make interface definitions easier.

Defining new syntaxes can help in other, less common situations. For example, KSM [4] has a `clang:sym` form that exposes a foreign library variable as a Scheme binding. Using PLT Scheme macros, we can achieve this functionality in Scheme using a macro that defines the C ‘variable’ as a macro¹⁰:

¹⁰From the MzScheme [9, Section 12.1] manual: The ‘syntax-idrules’ form has the same syntax as ‘syntaxrules’, except that each pattern is used in its entirety (instead of starting with a key-

```
(define-syntax defcvar
  (syntax-rules ()
    [(_ var lib type)
     (define-syntax var
      (syntax-id-rules (set!)
        [(set! var1 val1)
         (set-ffi-obj! 'var lib type val1)]
        [(var . xs)
         ((get-ffi-obj 'var lib type) . xs)]
        [var (get-ffi-obj 'var lib type)])))))
```

and verify that it is working properly:

```
> (defcvar z "x.so" _int)
> z
0
> (set! z 123)
> z
123
> ((get-ffi-obj "getz" "x.so" (_fun -> _int)))
123
```

where the C code that is compiled into “x.so” is:

```
int z = 0;
int getz() { return z; }
```

Using Types

C types in our system are somewhat lighter than expected: there is only a loose correlation between these types and Scheme object types. A type in our context can simply mean a different way of marshaling Scheme values to/from C, for example, the `_file` type from Section 3.2.1 is simply a different way to marshal MzScheme path objects which are normally used with `_path`. No C-level support is needed for such cases: there are no new binary tags involved, and no new object representations at the implementation level, meaning that it is extremely cheap to create such type descriptors. A common usage of types is therefore as a simple mechanism to add hook on the translation process.

For example, the ImageMagick library specifies a ‘MagickWand’ type, which is always being manipulated as a ‘MagickWand*’ pointer. There are functions that return a pointer to a newly created ‘MagickWand’ object, and these objects must be destroyed with the ‘DestroyMagickWand’ function. To do this automatically, we define a `_MagickWand` type using `_pointer` and providing a new translation when going from C to Scheme, one that uses ‘registerfinalizer’ to make the GC use ‘DestroyMagickWand’ when reclaiming the pointer object¹¹:

```
(define _MagickWand
  (make-c-type _pointer
   #f ; Scheme->C translation is the same as _cpointer
   (lambda (ptr)
     (if ptr
      (begin (register-finalizer ptr destructor) ptr)
      (error '_MagickWand "got a NULL pointer"))))
```

We can make this even better with a new `cpointer` type which uses an appropriate tag to identify these pointers and make sure that we don’t confuse pointers to internal ImageMagick objects of different types. The following definition uses ‘definecpointertype’ (see Section 3.2.2) to create a type that tags all pointers when they are

word placeholder that is ignored).

¹¹This assumes that there is no way to get a second pointer object that refers to the same ‘MagickWand’ object, so care should be taken with functions that can create such aliases.

moved from the foreign side to Scheme, and check the tag when sending a Scheme pointer object out to foreign code.

```
(define-cpointer-type _MagickWand #f #f12
  (lambda (ptr)
    (if ptr
      (begin (register-finalizer ptr destructor) ptr)
      (error '_MagickWand "got a NULL pointer"))))
```

A different example of using a new type comes from our TCL interface: the `Tcl_Eval` function returns a status integer, indicating a possible error. In our implementation, we define `evaltcl` as:

```
(define eval-tcl
  (get-ffi-obj "Tcl_Eval" libtcl
    (_fun (interp : _interp = (current-interp))
      (expr : _string)
      -> _tclret)))
```

using the following `_tclret` definition:

```
(define _tclret
  (make-ctype (_enum '(ok error return ...))
    (lambda (x) (error "tclret: only for returning"))
    (lambda (x)
      (when (eq? x 'error)
        (error 'tcl (get-string-result
          (current-interp))))
      x)))
```

which effectively translates a TCL error into a Scheme exception.

Note that the TCL interface uses a Scheme parameter ‘`current-interp`’ as the value of the first argument to ‘`TCL_Eval`’. We can make this implicit by defining a new custom type syntax, using the ‘`expr:`’ keyword:

```
(define-syntax _cur-interp
  (syntax-id-rules ()
    [_ (type: _interp expr: (current-interp))]))
(define eval-tcl
  (get-ffi-obj "Tcl_Eval" libtcl
    (_fun _cur-interp (expr : _string) -> _tclret)))
```

Using Custom Types

Custom types are intended to be used in situations where simple independent processing of each argument is insufficient. For example, many functions in the ImageMagick interface return a ‘`status`’ integer that indicates if there was an error. If an error has occurred, the main object involved in the function invocation should be used to retrieve the error message and severity. One way to deal with this situation is to save the object in a place accessible right after the foreign call, like a parameter. This is essentially what the TCL interface does, where `_tclret` uses a parameter to get the error message. The ImageMagick interface is different — instead of a single implicit context parameter, it fits more an object-oriented style, where each method call happens in its object’s context.

As a result, a good interface must be able to provide a relation between different arguments, namely the result value (to be checked for an error) and the first argument (providing the current object context). This is done using the `1st-arg:` keyword of a custom type which specifies an identifier that will be bound to the first argument:

```
(define-syntax _status
  (syntax-id-rules (_status)
    [_status
      (type: _bool
        1st-arg: 1st
        post: (r => (unless r
          (raise-wand-exception 1st))))]))
```

Memory Management

Usually, there are important aspects of the library interface that are not fully specified. Memory management issues often fall under this category. For example, a naive interface might behave in a surprising way:

```
> (define crypt
  (get-ffi-obj "crypt" "libcrypt"
    (_fun _string _string -> _string)))
> (define a (crypt "foo1" "23"))
> a
"23.kLNfMwUW0Q"
> (define b (crypt "foo4" "56"))
> b
"568.5HohJYC0g"
> a                               ; a is modified!
"568.5HohJYC0g"
> (string-set! a 0 #\X) ; verify that a and b
> (list a b)                ; are the same string
("X68.5HohJYC0g" "X68.5HohJYC0g")
> (eq? a b)                 ; ...but not quite the same
#f
```

Using a simple SWIG interface, made using the C prototype declaration for ‘`crypt`’:

```
extern char *crypt(const char *key, const char *salt);
```

suffers from this problem too. The reason for this strange behavior is that both our interface implementation and SWIG’s generated code use MzScheme’s ‘`make_string_without_copying`’ function, which simply wraps an existing C string in a Scheme string object. The standard Unix `crypt` function returns a pointer to its own static string, making the above interaction create two Scheme string objects that point to this static buffer — but the Scheme objects are still different. This can be dangerous as it breaks an implementation assumption, so some solution is required. Changing the implementation to use ‘`scheme_make_string`’ would not be acceptable in the general case since it leads to an expensive overhead. In addition, there are other foreign functions (e.g., `getcwd`) that can allocate a return string, and blindly copying it will cause a memory leak (the allocated string is not in GC-controlled memory).

Using our system simplifies such a solution since we don’t have to break out of Scheme, we can simply use a new type¹³:

```
(define _string/copy
  (make-ctype _string #f
    (lambda (x) (string-append x #""))))
```

We can solve numerous problems in a similar way, for example, using semaphores to avoid problems with the single `crypt` buffer, or creating a new `_string/free` that copies a string and freeing the previous GC-invisible one.

¹²Use `_cpointer` as a base type, no extra translation when going to from Scheme to C, and register the destructor on the way back.

¹³Note that this is not relevant now, since our system is part of the Unicode-enabled MzScheme, so Scheme strings are stored in Unicode format, meaning that they are always copied.

5 Related Work

The first and foremost advantage that our foreign interface has over existing implementations, is the fact that it is truly dynamic. This means that functionality that traditionally is available only via C code is available to Scheme programmers, which makes for a compiler- and architecture-independent system. Furthermore, the dynamic aspect of the system allows for playing with foreign extensions dynamically, modifying and debugging the interface at run-time¹⁴. Exploratory programming is therefore possible, hence the overall development cycle becomes much lighter.

A second advantage comes from the fact that we use Scheme. Using a language with robust syntactical abstractions makes it possible to provide an IDL-like interface for interface programmers, with features that can go beyond capabilities of conventional IDLs [18, 16]. Having syntactic abstractions in the language makes it possible for users to extend their own code using new constructs, including ones that are unique to a single library, in contrast to fixed IDLs that are either fixed, or used through a primitive facility like the C preprocessor.

Dynamic interfaces are not as common as static interfaces. Existing dynamic systems, for example the Allegro CL foreign function interface [10] and Python’s `ctype` module, do not provide the low-level C-substitute features that we do. Urban’s FFI survey [17], although a little out-dated, provides an excellent overview on existing systems and implementation issues. It is interesting to note an SML interface system [2] as another, somewhat similar system to ours. Similar to our design, the main idea is *data-level interoperability* [8] — making raw C data available to the high-level language, but our system differs in a few important aspects:

- Our design is built around the idea of enabling arbitrary C-like unsafe code — whereas Blume’s system uses SML’s type system to enhance interaction with foreign code.
- Our system goes one step further in giving users more power. “If you can do it in C, then we will let you do it in Scheme” rather than “Some C-level operations are useful enough that we let you use them”.
- Blume’s system is limited to SML’s syntactic framework, where we use Scheme’s capabilities for creating IDL-like syntax.

We focus our comparison on static interface generators such as GreenCard [14], G-Wrap [3], and SWIG [1]. There are Scheme systems that fall under this category too by providing support for combining Scheme and C code, for example, Gambit-C¹⁵ [5] and KSM [4]. Most notably, SRFI 50 [15] attempts to standardize this approach, possibly making it possible for different Scheme implementations to share C code. These systems make it possible to write Scheme code that is converted to C code, so it is easy to write such ‘Scheme’ code that calls C functions as if they were plain function calls. Some of these systems lack a code generation component that is derived by an IDL or some equivalent, but they can all be seen as static code generators.

We now focus on SWIG as a popular system that can be used for multiple high-level languages. A simple translation using SWIG requires the user to compile (through the SWIG parser) a C header file with a SWIG interface file, resulting in C code that is then, yet again, compiled using a C compiler, to produce a C module that is finally imported into Scheme. In contrast to the static approach,

func.	Glue Type	CPU	Real	GC
crypt	SWIG	38%	4%	-34%
	Handwritten C glue	53%	49%	0%
sqadd	SWIG	55%	57%	0%
	Handwritten C glue	60%	61%	0%

Table 5. Comparison of overhead time

our ‘foreign’ library makes it possible for a Scheme developer to quickly open up a C library, pull out a few procedure objects and start an interactive development session.

It could be argued that a simpler, more user friendly system comes at a price of expensive overhead, leading to an inherent sacrifice of performance. Testing out two simple benchmarks, we found that the interface overhead of our system is just slightly slower as a compiled interface that was generated by SWIG, which itself has an almost identical overhead to hand-written glue code.

Our results are summarized in Table 5. Two functions were used for this analysis — the first is the `crypt` function taken from the standard Unix `libcrypt`: consuming two strings and producing an encrypted string result. The second is a simple C function, `sqadd`, that performs an addition of two integer squares. We measured a million executions of `crypt` and 30 million executions of `sqadd`, performing each test for 16 rounds beginning with a fresh MzScheme process, discarding the 6 extreme timings and averaging the other 10. The percentages are computed as: $\frac{Time_{PLT} - Time_{RawC}}{Time_{SWIG} - Time_{RawC}} - 1$ where $Time_{PLT}$ is the averaged running time of our interface, $Time_{SWIG}$ is the average running time of SWIG, and $Time_{RawC}$ is that of an implementation of comparable repetition loops in C. The same computation was used to compare our system against handwritten C glue code.

As Table 5 shows, our system is about 1.5 times slower than SWIG, and, in most cases the handwritten glue code. The biggest performance hit is in the simple arithmetic function, where the actual foreign code does much less than the interface code. Situations like this should rarely occur since the usual case of using a foreign library is when it can do some substantial work that is otherwise hard to achieve in Scheme.

While issues of timing and performance are important, aspects such as implementation complexity and ease of use must also be considered. Comparing our system to SWIG and interfaces that use an IDL, it becomes clear that our implementation is better in at least one aspect. One advantage that our system provides over the static approaches is the ability to specify additional functionality using new user-defined types that involve arbitrary translation code. The main point here is that such translations are written in the high-level language itself rather than dealing with the intricacies of the C implementation.

In addition, regardless of interface design and syntactical complexity, our implementation is better because the interfacing mechanism itself is in a high order language: making it possible to include arbitrary Scheme code as part of the foreign call specification. This is further enhanced by the fact that we use Scheme since it is possible to create new syntactic abstractions to deal with new requirements. Either with SWIG interface files or with an IDL, the interface developer is still confined by C and C-like code with its known shortcomings when it comes to dealing with complex problems.

¹⁴As long as no fatal errors occur.

¹⁵Some parts of this were ported to PLT’s MzC compiler.

6 Future Work

C++ Libraries Currently, there is support only for plain C libraries. Depending on implementation details, it can be feasible to interface C++ libraries. This might involve plenty of details regarding object layout, inheritance, virtual function tables, name mangling, etc. Hopefully, these issues can be addressed in Scheme so we might not need any further enhancements to the C part of our implementation.

Parsing C One of the main disadvantage of our system is that it is not using C, so we cannot use C include files as rough interface specifications. We plan to investigate a simple C header-file parser that will parse files into s-expressions, which can be used to automate some aspects of interface generation (A working parser prototype exists). Such a parser does not need to be fast and efficient, since parsing can be done at syntax expansion time, eliminating any run-time speed costs. In addition, note that as usual with other interface generators, this will almost never mean that an interface can be fully automated, as header files do not provide enough information — this situation might improve if we target some IDL language instead (most use similar syntax).

Memory Management Issues Currently, our system works well with both versions of PLT Scheme: the one that uses the Bohm conservative garbage collector and the one that uses a precise moving collector. However, there are still issues that interface writers need to be aware of. In time, we will gain more experience writing interfaces, which will motivate further functionality that will make this easier — our goal is, of course, making GC-related issues as transparent as possible for interface writers.

One aspect of this, is dealing with struct objects that might contain GC-able pointers. We have a plan to deal with this, effectively making it possible to specify in Scheme a map of pointer offsets that the garbage collector should be aware of, making it treat new Scheme-defined structs properly.

Additional Scheme Support There are some areas in which additional Scheme support is needed. For example, an array of structs is hard to deal with — there is no way to get to one such struct and modify it, since accessing it will create a copy. We believe that it is possible to write Scheme code that will make this possible, by not pulling out a struct copy, but rather provide forms that will use nested reference indexes, where some are vector indexes and some are struct field names. If we can make this composable, it would be possible to deal with them in an easy way — without resorting to pointer aliasing¹⁶.

An additional area where additional support is needed, is when dealing with foreign functions that block. MzScheme contains a few hooks that are intended to be used when it is embedded as a library, these hooks can be used for calling blocking foreign functions as well.

Using Contracts PLT Scheme has support for procedure contracts [7] which could be used to enhance the robustness of library interfaces. Specifically, we want to treat contract violations in modules that use the ‘foreign’ module as more severe, as these are equivalents of C bugs, which might result in a crash. A module would also need some way of declaring it as a proper interface, meaning that code that uses it should not be blamed for crashes. Alternatively, code that is not intended as an interface (i.e., code that provides functionality for interface modules) should propagate the property of contract violation severity.

¹⁶The precise garbage collector makes it impossible to get a pointer to the internal part of an allocated block

Assembly Code Generation Working our way to native just-in-time compilation, we plan on adding machine-code generation ability to PLT Scheme. We will interface this functionality via the ‘foreign’ module. Furthermore, some of the interface aspects can be implemented in assembly when runtime is important.

7 References

- [1] D. M. Beazley. SWIG: An easy to use tool for integrating scripting languages with C and C++. In *Proceedings of the 4th USENIX Tcl/Tk Workshop*, pages 129–139, July 1996.
- [2] Matthias Blume. No-longer-foreign: Teaching an ML compiler to speak C “natively”. In *BABEL’01: First workshop on multi-language infrastructure and interoperability*, September 2001.
- [3] Rob Browning. G-Wrap home page. <http://www.nongnu.org/g-wrap/>.
- [4] Hangil Chang. KSM-Scheme home page. <http://square.umin.ac.jp/hchang/ksm/>.
- [5] Marc Feeley. Gambit Scheme system. <http://www.iro.umontreal.ca/gambit/>.
- [6] Marc Feeley. SRFI 4: Homogeneous numeric vector datatypes. <http://srfi.schemers.org/srfi-4/>.
- [7] Robert Bruce Findler and Matthias Felleisen. Contracts for higher-order functions. In *ACM SIGPLAN International Conference on Functional Programming*, 2002.
- [8] Kathleen Fisher, Riccardo Pucella, and John Reppy. Data-level interoperability. Bell Labs Technical Memorandum, April 2000.
- [9] Matthew Flatt. *PLT MzScheme: Language Manual*. PLT, August 2004. Version 208.
- [10] Franz Lisp. Foreign function interface. <http://www.franz.com/support/documentation/6.1/doc/foreign-functions.htm>.
- [11] Anthony Green. The libffi home page. <http://sources.redhat.com/libffi/>.
- [12] Samuel P. Harbison. *Modula-3*. Prentice-Hall, 1992.
- [13] Thomas Heller. The ctypes module. <http://python.net/crew/theller/ctypes/>.
- [14] Simon Peyton Jones, Thomas Nordin, and Alastair Reid. GreenCard: a foreign-language interface for Haskell. In J. Launchbury, editor, *2nd Haskell Workshop*, 1997.
- [15] Richard Kelsey and Michael Sperber. SRFI 50: Mixing scheme and c. <http://srfi.schemers.org/srfi-50/>.
- [16] The Open Group. *CAE Specification, DCE 1.1: Remote Procedure Call*, chapter 4. The Open Group, October 1997.
- [17] Reini Urban. Design issues for foreign function interfaces. <http://xarch.tu-graz.ac.at/autocad/lisp/ffis.html>, Last updated at 2004.
- [18] A. Vogel, B. Gray, and K. Duddy. Understanding any IDL — lesson one: DCE and CORBA. In *Proceedings of the Third International Workshop on Services in Distributed and Networked Environments (SDNE’96)*, 1996.

Acknowledgments

We would like to thank Matthew Flatt: this work would not be possible without his help, especially with GC-related issues. The comments and suggestions made by the reviewers have been extremely helpful, Mike Sperber was particularly helpful in the process of revising this text.