Etos: an Erlang to Scheme compiler

Marc Feeley and Martin Larose
Université de Montréal
C.P. 6128 succursale centre-ville
Montréal H3C 3J7, Canada
{feeley,larosem}@iro.umontreal.ca

August 18, 1997

Abstract

The programming languages Erlang and Scheme have many common features, yet the performance of the current implementations of Erlang appears to be below that of good implementations of Scheme. This disparity has prompted us to investigate the translation of Erlang to Scheme. In this paper we describe the design and implementation of the **Etos** Erlang to Scheme compiler and compare its performance to other systems. On most benchmark programs, Etos outperforms all currently available implementations of Erlang.

1 Introduction

Erlang [3] and Scheme [13, 5] have some obvious differences (e.g. infix vs. prefix syntax, pattern matching vs. access functions, `catch/throw` vs. `call/cc`, concurrency) but also a large number of similarities (e.g. use of functional style, dynamic typing, automatic memory management, data types). Erlang has been mostly developed internally at Ericsson and as a result there is a limited choice of compilers. As the implementers of these compilers freely admit [1], "Performance has always been a major problem". On the other hand there are many implementations of Scheme available [15] and the good compilers appear to generate faster code than the Erlang compilers available from Ericsson (for example Hartel *et al.* [11] has shown that the "pseudoknot" benchmark compiled with Ericsson's BEAM/C 6.0.4 is about 5 times slower than when compiled with the Gambit-C 2.3 Scheme compiler).

Because of the strong similarity between Erlang and Scheme and the availability of several good Scheme compilers, we have begun the implementation of an Erlang to Scheme compiler called "Etos". This paper explains the major design issues of such a compiler, how these are solved in Etos 1.4, and the performance of the compiler compared to other Erlang compilers.

2 Portability vs Efficiency

Early on we decided that portability of the compiler was important in order to maximize its usefulness and allow experiments across platforms. Etos is written in standard Scheme [5] and the generated programs conform fairly closely to the standard (the discrepancies are explained later).

It is clear however that better performance can be achieved if non-standard features of the target Scheme implementation are exploited (in particular the existence of fast operations on fixed precision integers, i.e. fixnums). To allow for this, the generated code contains calls to Scheme macros whose definition depends on the target Scheme implementation. The appropriate macro definition file is supplied when the Scheme program is compiled. This avoids the need to recompile the Erlang program from scratch when the target Scheme implementation is changed. For example, the Erlang addition operator is translated to a Scheme call to the `erl-add` macro. The macro call `(erl-add x y)` may simply expand to a call to a generic addition procedure which adds `x` and `y`, or if fixnum arithmetic is available, expand to an inline expression which performs a fixnum addition if `x` and `y` are fixnums and otherwise calls the generic addition

procedure.

Using a macro file also allows to move some of the code generation details out of the compiler and into the macro file, making it easy to experiment and tune the compiler. For example the representation of Erlang data types can easily be changed by modifying the macro definitions.

3 Direct Translation

We also wanted the translation to be direct so that Erlang features would map into the most natural Scheme equivalent. This has several benefits:

- Erlang and Scheme source code can be mixed more easily in an application if the calling convention and data representation are similar. Special features of Scheme (such as first-class continuations and assignment) and language extensions (such as a C-interface and special libraries) can then be accessed easily.
- The generated code can be read and debugged by humans.
- A comparison of compiler technology between Erlang and Scheme compilers will be fairer because the Scheme compiler will process a program with roughly the same structure as the Erlang compiler.

When a direct translation is not possible, we tried to generate Scheme code with a structure that we felt would be compiled efficiently by most Scheme compilers. Nevertheless there is often a run time overhead in the generated Scheme code that makes it slower than if the application had been written originally in Scheme. For example, Erlang's "`<`" operator is generic (it works on numbers as well as lists and other data types) but in most application programs it is only used to compare numbers. The code generated by Etos can't use Scheme's "`<`" primitive directly because it works on numbers only.

4 Data Types

The most important Erlang data types have a direct equivalent in standard Scheme:

Erlang	Scheme
integer	exact integer
float	inexact real
atom	symbol
list	list
tuple	vector
function	procedure

4.1 Numbers

Scheme numbers are organized into a class hierarchy: integer $\subseteq$ rational $\subseteq$ real $\subseteq$ complex. Independently of their class, numbers have an "exactness". For instance `2.0` denotes the inexact number 2 and `1/2` denotes the exact number 0.5. Scheme exact integers correspond to Erlang integers. In both Scheme and Erlang, integers can be of limited range (24 bits minimum required by Erlang) but typically they are implemented as bignums which have unlimited precision. Scheme inexact reals correspond to Erlang floats.

An unfortunate consequence of this representation is that testing for an Erlang integer or float translates into two tests in standard Scheme (i.e. `(and (integer? x) (exact? x))` tests if `x` is an exact integer).

Scheme's rational and complex numerical types are not needed as they do not exist in Erlang.

4.2 Atoms

Scheme symbols can be used to represent Erlang atoms. Both can contain arbitrary characters and symbols can be compared for equality efficiently with the `eq?` predicate (which is simply a pointer comparison in many implementations of Scheme). The Scheme procedures `string->symbol` and `symbol->string` are equivalent to the Erlang built-in functions `list_to_atom` and `atom_to_list` except that the former deals with strings (which in Scheme is a data type separate from lists).

One complication is that Scheme is a case-insensitive language. Variables and symbols in the source of Scheme programs are stripped of their case. A simple solution for variables is to prefix uppercase letters with an escape character (i.e. `^`), so that the Erlang variable `ListOfFloats` becomes `^list^of^floats` in Scheme.

The only way to force a particular case for symbols in Scheme is to use the procedure `string->symbol`. Constants containing atoms (e.g. the constant list `[one,two]`) are created at run time using `string->symbol`. This is done by storing the objects created into global variables once in the initialization phase of the Scheme program and references to these globals replace references to the constants. Constants not containing atoms get converted to Scheme constants.

Alternative representations for atoms which were rejected are:

- Strings: no special treatment for uppercase letters is needed but the equality test is much more expensive.
- Symbols with escape character for uppercase letters: requires an unnatural and inefficient translation of `list_to_atom` and `atom_to_list`.

4.3 Lists

Lists are handled similarly in Scheme and Erlang. In Scheme, lists are made up of the empty list (i.e. `'()`) and pairs created with the binary `cons` primitive or the variable arity `list` primitive. The primitives `car` and `cdr` extract the head and tail of a list.

4.4 Tuples

Scheme vectors are the obvious counterpart of tuples. Vectors are constructed either with the variable arity `vector` primitive (Erlang's `{...}`), the `list->vector` primitive (Erlang's `list_to_tuple`), or the `make-vector` primitive (which creates a vector of length computed at run time).

A minor incompatibility is that tuples are indexed from 1 (with the `element` builtin function) and Scheme vectors are indexed from 0 (with the `vector-ref` primitive).

A more serious problem is that lists and vectors are the only compound data structures in standard Scheme. Since the Erlang data types port, pid, reference, and binary don't have a direct counterpart in Scheme, they must be implemented using lists or vectors. We have used vectors to implement these data types (as well as tuples and functions) because their content can be accessed in constant time. The first element of the vector is a symbol which indicates the type and the data associated with the type is in the remaining elements. Thus the tuple `{1,2,3}` is represented by the Scheme vector `#(tuple 1 2 3)`. Note that with this representation, tuple indexing does not require a run time decrement of the index to access an element. However, an Erlang type test translates to two Scheme tests (for example `(and (vector? x) (eq? (vector-ref x 0) 'tuple))` tests if `x` is a tuple).

A more space efficient representation which is based on Scheme's ability to test object identity with `eq?` is to use no tag for tuples and a special tag for non-tuples:

```
(define pid-tag (vector 'pid))

(define make-pid
  (lambda (...)
    (vector pid-tag ...)))

(define pid?
  (lambda (x)
    (and (vector? x)
         (> (vector-length x) 0)
         (eq? (vector-ref x 0) pid-tag))))
```

This representation was not used because type testing (which is a frequent operation in pattern-matching) is more expensive in this representation. One more test is required for non-tuples (as shown above) and many more tests for tuples:

```
(define tuple?
  (lambda (x)
    (and
     (vector? x)
     (or (= (vector-length x) 0)
         (let ((tag (vector-ref x 0)))
           (not
            (or (eq? tag function-tag)
                (eq? tag port-tag)
                (eq? tag pid-tag)
                (eq? tag reference-tag)
                (eq? tag binary-tag)))))))))
```

4.5 Functions

Scheme procedures are the obvious counterpart of Erlang functions. Erlang functions are of fixed arity so the variable arity mechanism of Scheme is not necessary. Both Erlang and Scheme can create and call functional objects (known as "closures" in the Scheme community).

Unfortunately, this direct representation does not support error detection. Erlang's general function calling mechanism needs to ensure that the function that is being called is of the appropriate arity, and signal an error if it isn't. Because there is no standard way in Scheme to extract the arity of a procedure or to trap the application of a procedure to the wrong number of arguments, functional objects are represented as a tagged vector which contains the function's arity and the corresponding Scheme closure.

Note that toplevel functions of a module contain the arity information in their name and no arity test is needed when they are called. For example the function `bar` of arity 2 in module `foo` will be translated to a Scheme lambda-expression of arity 2 bound to the global variable `foo:bar/2` (which is a valid variable name in Scheme). When the compiler encounters a call such as `foo:bar(1,2)`, it will translate it to a Scheme call to `foo:bar/2` which is guaranteed to be bound to a procedure of arity 2.

4.6 Other Types

The other Erlang data types (port, binary, record) are only partially implemented in Etos 1.4 but this is mostly because of a lack of time. They can be represented with tagged Scheme vectors as shown above.

Note that Erlang records are just syntactic sugar for tuples so no special representation is required for them.

5 Front End

To ensure compatibility with existing Erlang compilers, Etos' parser specification was derived from the one for the JAM interpreter and processed by our own Scheme parser generator [6, 4]. The original parser constructs a parse tree built of tuples. Because Etos needs to attach semantic information on the nodes of the parse tree, a conversion phase was added to extend the tree nodes with additional fields. This conversion also computes the bound variables at each node and performs constant propagation and constant folding. Constant propagation and folding are mainly needed to avoid allocation of structures which are constant, such as in:

```
f(X) -> Y = {1,2}, [X,Y,3,4].
```

which gets compiled as though it were:

```
f(X) -> [X|[{1,2},3,4]].
```

The list `[{1,2},3,4]` is represented internally as the Scheme constant list `'(#(tuple 1 2) 3 4)`.

Following this, the free variables before and after each node are computed. This is done as a separate pass because the bound variable analysis requires a left-to-right traversal of the parse tree, whereas the free variable analysis requires a right-to-left traversal. The free variables are needed to efficiently translate `case`, `if`, and `receive` expressions, which is explained in the next section.

6 Binding and Pattern Matching

6.1 Binding in Erlang

Erlang's approach for binding variables is a relic of its Prolog heritage. Binding is an integral part of pattern matching. Once it is bound by a pattern matching operation, a variable can be referenced in the rest of a function clause (unless it has become an "unsafe" variable, see below). For example, in

```
f({A,B}) -> [X,X,X] = A, B+X.
```

the function `f` will pattern match its sole argument with a two-tuple. In the process, the variables `A` and `B` get bound to the first and second element respectively. After this, `A` is referenced and pattern matched with a list containing three times the same element. Note that the first occurrence of `X` binds `X` to the first element of the list and the remaining occurrences reference the variable.

6.2 Binding in Scheme

In Scheme the basic binding construct is the lambda-expression and binding occurs when a procedure is called, as in:

```
((lambda (x) (* x x))
 3)
```

Here the variable `x` is bound to `3` when the closure returned by evaluating the lambda-expression is called with `3`. Scheme also has the binding constructs `let`, `let*` and `letrec` but these are simply syntactic sugar

Erlang syntactic categories:
- `<const>`: constant
- `<ubvar>`: unbound variable
- `<bvar>`: bound variable
- `<expr1>`, `<expr2>`: arbitrary expressions
- `<pat1>`, `<pat2>`: arbitrary patterns
- `<fn>`: function name

Expression translation:

$E($ `<const>` $, k) =$ `(`k $C($`<const>`$))$

$E($ `<bvar>` $, k) =$ `(`k $N($`<bvar>`$))$

$E($ `<pat1>=<expr1>` $, k) = E($ `<expr1>` $,$ `(lambda (v1) (`$P($ `<pat1>` $,$ `(`k `v1) , (erl-exit-badmatch) ) v1)) )`

$E($ `<expr1>,<expr2>` $, k) = E($ `<expr1>` $,$ `(lambda (v1)` $E($ `<expr2>` $, k))$ `)`

$E($ `<expr1>+<expr2>` $, k) = E($ `<expr1>` $,$ `(lambda (v1)` $E($ `<expr2>` $,$ `(lambda (v2) (`k `(erl-add v1 v2))) ))` $)$

$E($ `<fn>(<expr1>)` $, k) = E($ `<expr1>` $,$ `(lambda (v1) (`k `(`$N($`<fn>`$)$`/1 v1))) )`

Pattern-matching translation:

$P($ `<ubvar>` $, s, f) =$ `(lambda (`$N($`<ubvar>`$))` $s$`)`

$P($ `<bvar>` $, s, f) =$ `(lambda (v1) (if (erl-eq-object? v1` $N($`<bvar>`$))` $s$ $f$`))`

$P($ `[]` $, s, f) =$ `(lambda (v1) (if (erl-nil? v1)` s f`))`

$P($ `[<pat1>|<pat2>]` $, s, f) =$ `(lambda (v1)`

```
                              (if (erl-cons? v1)
                                (P( <pat1>, (P( <pat2>, s, f) (erl-tl v1)), f) (erl-hd v1))
                                f))
```

Auxiliary functions:
- $C(const)$: translate an Erlang constant to Scheme
- $N(name)$: translate an Erlang variable or function name to Scheme

Note:
- vn stands for a freshly created variable which will not conflict with other variables.

Figure 1: Simplified translation algorithm for a subset of Erlang.

---

for lambda-expressions and calls. For example the previous expression is equivalent to:

```
(let ((x 3))
  (* x x))
```

6.3 Translation of Binding and Pattern Matching

To translate an Erlang binding operation to Scheme it is necessary to nest the evaluation of the "rest of the function clause" inside the binding construct. This can be achieved by performing a partial CPS conversion, as shown in Figure 1.

The translation function E takes two parameters: the Erlang expression to translate and a Scheme lambda-expression denoting the continuation which consumes the result of the Erlang expression. E returns the equivalent Scheme expression. E makes use of the function P to translate pattern-matching. P's arguments are: the pattern to match and the success and failure Scheme expressions. P returns a one argument Scheme lambda-expression which pattern matches its argument to the pattern, and returns the value of the success expression if there is a match and returns the value of the failure expression otherwise.

When an Erlang function is translated, E is called on each function clause to translate the right hand side with the initial continuation `(lambda (x) x)` (i.e. the identity function). Note that the continuation k and all lambda-expressions generated in the translation are always inserted in the function position of a call. This implies that in the resulting Scheme code all the lambda-expressions generated can be expressed with the `let` binding construct (except for those gen-

erated in the translation of functional objects, which is not shown). To correctly implement tail-calls, an additional translation rule is used to eliminate applications of the identity function, i.e.

`((lambda (x) x)` Y`)` $\rightarrow Y$

The translation algorithm is not a traditional CPS conversion because function calls remain in direct style (i.e. translated Erlang functions do not take an additional continuation argument). This partial CPS conversion is only used to translate Erlang binding to Scheme binding. An interesting property of function E is that it embeds k in the scope of all Scheme bindings generated, so that these bindings can be accessed by k. Similarly, P always embeds s (the success expression) in the scope of all Scheme bindings generated. For example, consider the Erlang expression:

```
[X|Y] = foo:f(A), X+bar:g(Y)
```

This is translated to the following Scheme expression (if we assume that `A` is a bound variable):

```
(let ((v7 ^a))
  (let ((v5 (foo:f/1 v7)))
    (let ((v6 v5))
      (if (erl-cons? v6)
        (let ((^x (erl-hd v6)))
          (let ((^y (erl-tl v6)))
            (let ((v1 v5))
              (let ((v2 ^x))
                (let ((v4 ^y))
                  (let ((v3 (bar:g/1 v4)))
                    (erl-add v2 v3)))))))
        (erl-exit-badmatch)))))
```

Note that there are many useless bindings in this code. In the actual implementation, the translator keeps track of constants, bound variables and singly referenced expressions and propagates them to avoid useless bindings. With this improvement the Scheme code generated is:

```
(let ((v5 (foo:f/1 ^a)))
  (if (erl-cons? v5)
    (erl-add (erl-hd v5)
             (bar:g/1 (erl-tl v5)))
    (erl-exit-badmatch))))))
```

This is close to what we would expect a Scheme programmer to write.

6.4 Translation of Conditionals

Conditional expressions (i.e. `case`, `if`, and `receive`) must be handled carefully to avoid code duplication. Consider the following Erlang expression:

```
case X of
  1 -> Y = X*2;
  Z -> Y = X+1
end,
X*Y
```

The `case` expression will select one of the two bindings of `Y` based on the value of `X`. After the `case`, `Y` is a bound variable that can be referenced freely. On the other hand `Z` is not accessible after the `case` because it does not receive a value in all clauses of the case (it is an "unsafe" variable after the `case`).

The `case` construct could be implemented by adding to the translation function E a rule like Figure 2a. Note that the continuation k is inserted once in the generated code for each clause of the `case`. This leads to code duplication which is a problem if the `case` is not the last expression in the function body and the `case` has more than one clause. If the function body is a sequence of n binary `case` expressions, some of the code will be duplicated 2^n times.

This code explosion can be avoided by factoring the continuation so that it appears only once in the generated code. A translation rule like Figure 2b would almost work. The reason it is incorrect is that k is no longer nested in the scope of the binding constructs generated for the `case` clauses, so the bindings they introduce are not visible in k.

A correct implementation has to transfer these bindings to k. This can be done by a partial lambda-lifting of k as shown in Figure 2c. The arguments of the lambda-lifted k (i.e. `vk`) are the result of the `case` (i.e. `vr`) and the set of bound variables that are added by the clauses of the `case` and referenced in k (i.e. AV). Each clause of the case simply propagates these bindings to `vk`. AV can be computed easily from the free variables (it is the difference between the set of free variables after the `case` and the set of free variables after the selector expression). The lambda-lifting is partial because `vk` may still have free variables after the transformation.

This lambda-lifting could be avoided by using assignment. Dummy bindings to the variables AV would be introduced just before the first pattern

```
E( case <expr0> of       ,k) = E( <expr0> , (lambda (v0)
     <pat1> -> <expr1>;                      (P( <pat1> ,
     <pat2> -> <expr2>                           E( <expr1> ,k),   ;;; duplication of k
   end                                           (P( <pat2> ,
                                                     E( <expr2> ,k),   ;;; duplication of k
                                                     (erl-exit-case-clause) )
                                                 v0) )
                                               v0)))
```

a) Inefficient translation of the `case` construct.

```
E( case <expr0> of       ,k) = E( <expr0> , (lambda (v0)
     <pat1> -> <expr1>;                      (let ((vk k))   ;;; k not in right scope
     <pat2> -> <expr2>                         (P( <pat1> ,
   end                                             E( <expr1> , vk ),
                                                   (P( <pat2> ,
                                                       E( <expr2> , vk ),
                                                       (erl-exit-case-clause) )
                                                    v0) )
                                                 v0))))
```

b) Incorrect translation of the `case` construct.

```
E( case <expr0> of       ,k) = E( <expr0> , (lambda (v0)
     <pat1> -> <expr1>;                      (let ((vk (lambda (vr AV...) (k vr))))
     <pat2> -> <expr2>                         (P( <pat1> ,
   end                                             E( <expr1>, (lambda (vr) (vk vr AV...)) ),
                                                   (P( <pat2> ,
                                                       E( <expr2>, (lambda (vr) (vk vr AV...)) ),
                                                       (erl-exit-case-clause) )
                                                    v0) )
                                                 v0))))
```

Where $AV...$ is the set of bound variables that are added by the clauses of the `case` and referenced in k.

c) Correct translation of the `case` construct.

Figure 2: Translation of the `case` construct.

matching operation. Assignment would be used to set the value of these variables in the clauses of the `case`. This solution was rejected because many Scheme systems treat assignment less efficiently than binding. This is because of assignment conversion (traditionally performed to implement `call/cc` correctly) and generational GC.

In the actual implementation of the conditional constructs, the patterns are analyzed to detect common tests and factor them out so that they are only executed once. For example the translation of the following `case` expression will only contain one test that `X` is a pair:

```
case X of
  [1|Y] -> ...;
  [2|Z] -> ...
end
```

7 Errors and `catch/throw`

The traditional way of performing non-local exits in Scheme is to use first-class continuations. A `catch` is translated to a call to Scheme's `call/cc` procedure which captures the current continuation. This "escape" continuation is stored in the process descriptor after saving the current escape continuation for when the `catch` returns. A `throw` simply calls the current

escape continuation with its argument. When control resumes at a `catch` (either because of a normal return or a `throw`), the saved escape continuation is restored in the process descriptor.

8 Concurrency

First-class continuations are also used to implement concurrency. The state of a process is maintained in a process descriptor. Suspending a process is done by calling `call/cc` to capture its current continuation and storing this continuation in the process descriptor. By simply calling a suspended process' continuation, the process will resume execution.

Three queues of processes are maintained by the runtime system: the ready queue (processes that are runnable), the waiting queue (processes that are waiting for a message to arrive in their mailbox), and the timeout queue (processes which are waiting for a message with timeout). The timeout queue is a priority queue, ordered on the time of timeout, so that timeouts can be processed efficiently.

There is no standard way in Scheme to deal with time and timer interrupts. To simulate preemptive scheduling the runtime system keeps track of the function calls and causes a context switch every so many calls. When using the Gambit-C Scheme system, which has primitives to install timer interrupt handlers, a context switch occurs at the end of the time slice, which is currently set to 100 msecs.

9 Limitations

Due to its experimental and preliminary nature, Etos 1.4 does not implement Erlang fully. Most notably, these features of Erlang are not implemented:

1. Macros, records, ports, and binaries.
2. Process registry and dictionary.
3. Dynamic code loading.
4. Several built-in functions and libraries.
5. Distribution (all Erlang processes must be running in a single user process).

10 Performance

10.1 Benchmark Programs

To measure the performance of our compiler we have used mostly benchmark programs from other Erlang compilers. We have added two benchmarks (`ring` and `stable`) to measure the performance of messaging and processes. Unfortunately, we were not able to use the "Estone" benchmark [16] because it uses ports which are not implemented in Etos.

- `barnes` (iterated 10 times): Simulates gravitational force between 1000 bodies.
- `fib` (iterated 50 times): Computes 30th fibonacci number with a recursive function.
- `huff` (iterated 5000 times): Compresses and uncompresses a 38 byte string with the Huffman encoder.
- `length` (iterated 100000 times): Tail recursive function that returns the length of a 2000 element list.
- `nrev` (iterated 20000 times): Naive reverse of a 100 element list.
- `pseudoknot` (iterated 3 times): Floating-point intensive application taken from molecular biology [11].
- `qsort` (iterated 50000 times): Sorts 50 integers using the Quicksort algorithm.
- `ring` (iterated 100 times): Creates a ring of 10 processes which pass around a token 100000 times.
- `smith` (iterated 30 times): Matches a DNA sequence of length 32 to 100 other sequences of length 32. Uses the Smith-Waterman algorithm.
- `stable` (iterated 5000 times): Solves the stable marriage problem concurrently with 10 men and 10 women. Creates 20 processes which send messages in fairly random patterns.
- `tak` (iterated 1000 times): Recursive integer arithmetic Takeuchi function. Calculates `tak(18,12,6)`.

10.2 Erlang Compilers

Etos was coupled with the Gambit-C Scheme compiler version 2.7a [8]. We will first briefly describe the

Gambit-C compiler.

The Gambit programming system combines an interpreter and a compiler fully compliant to R^{4}RS and IEEE specifications. The Gambit-C compiler translates Scheme programs to portable C code which can run on a wide variety of platforms. Gambit-C also supports some extensions to the Scheme standard such as an interface to C which allows Scheme code to call C routines and vice versa.

The Gambit-C compiler performs many optimizations, including automatic inlining of user procedures, allocation coalescing, and unboxing of temporary floating point results. The compiler also emits instructions in the generated code to check for stack overflows and external events such as user or timer interrupts. The time between each check is guaranteed to be bound.

Gambit-C includes a memory management system based on a stop and copy garbage collector which grows and shrinks the heap as the demands of the programs change. The user can force a minimum and/or maximum heap size with a command line argument. Scheme objects are encoded in a machine word (usually 32 bits), where the lower two bits are the primary type tag. All heap allocated objects are prefixed with a header which gives the length and secondary type information of the object. Characters and strings are represented using the Unicode character set (i.e. 16 bit characters).

The implementation of continuations uses a lazy copying strategy. Continuation frames are allocated in a small area called the "stack cache". This area is managed like a stack (i.e. LIFO allocation) except when the `call/cc` procedure is called. All frames in the stack cache upon entry to `call/cc` can no longer be deallocated. When control returns to such a frame, it is copied to the top of the stack cache. Finally, when the stack cache overflows (because of repeated calls to `call/cc` or because of a deep recursion), the garbage collector is called to move all reachable frames from the stack cache to the heap.

We have compared Etos version 1.4 [7] with three implementations of Erlang compilers:

- Hipe version 0.27 [14], an extension of the JAM bytecode compiler that selectively compiles bytecodes to native code;
- BEAM/C version 4.5.2 [12], compiles Erlang code to C using a register machine as intermediate;
- JAM version 4.4.1 [2], a bytecode compiler for a stack machine.

Program	Etos (secs)	Time relative to Etos		
		Hipe	BEAM	JAM
`fib`	31.50	1.15	1.98	8.33
`huff`	9.74	1.48	5.01	24.81
`length`	11.56	2.07	3.44	34.48
`smith`	10.79	2.17	3.37	13.06
`tak`	13.26	1.12	4.37	11.09
`barnes`	9.18	2.08	–	4.07
`pseudoknot`	16.75	2.37	–	3.18
`nrev`	22.10	.84	1.83	10.98
`qsort`	14.97	.96	3.88	15.38
`ring`	129.68	.30	.31	1.92
`stable`	21.27	1.16	.64	2.43

Figure 3: Execution time of benchmarks

10.3 Execution Time

The measurements were made on a Sun UltraSparc 143 MHz with 122 Mb of memory. Each benchmark program was run 5 times and the average was taken after removing the best and worse times.

The Scheme code generated by Etos is compiled with Gambit-C 2.7a and the resulting C code is then compiled with `gcc` 2.7.2 using the option `-O1`. The executable binary sets a fixed 10 Mb heap.

The results are given in Figure 3. They show that Etos outperforms the other Erlang compilers on most benchmarks. If we subdivide the benchmarks according to the language features they stress, we can explain the results further:

- `fib`, `huff`, `length`, `smith` and `tak`, which are integer intensive programs, take advantage of the efficient treatment of fixnum arithmetic in Gambit-C and from the inlining of functions. Etos is up to two times faster than Hipe, 5 times faster than BEAM/C, and 35 times faster than JAM.
- On the floating point number benchmarks, `barnes` and `pseudoknot`, Etos is also faster than the other Erlang implemenations. In this case

Etos is a little over two times faster than Hipe. These programs crashed when compiled with BEAM/C.

- List processing is represented by `nrev` and `qsort`. On these programs Hipe is a little faster than Etos (4% to 16%), which is still roughly two to four times faster than BEAM/C. Etos' poor performance is only partly attributable to its implementation of lists:

  1. Gambit-C represents lists using 3 word long pairs as opposed to 2 words on the other systems. Allocation is longer and the GC has more data to copy.
  2. Gambit-C guarantees that interrupts are checked at bound intervals [9] which is not the case for the other systems. For example, the code generated by Gambit-C for the function `app` (the most time consuming function of the `nrev` benchmark) tests interrupts twice as often as Hipe (i.e. on function entry and return).
  3. The technique used by Gambit-C to implement proper tail-recursion in C imposes an overhead on function returns as well as calls between modules. For `nrev` the overhead is high because most of the time is spent in a tight non-tail recursive function. Independent experiments [10] have shown that this kind of program can be sped up by a factor of two to four when native code is generated.

- Finally `ring` and `stable` manipulate processes. Here we see a divergence in the results. Hipe is roughly three times faster than Etos on `ring`. Etos performs slightly better than Hipe on `stable` but is not as fast as BEAM/C. We suspect that our particular way of using `call/cc` to implement processes is the main reason for Etos' poor performance:

  1. When a process' mailbox is empty, a `receive` must call the runtime library which then calls `call/cc` to suspend the process. These intermodule calls are rather expensive in Gambit-C. It would be better to inline the `receive` and `call/cc`.
  2. Scheme's interface to `call/cc` (which receives a closure and must allocate a closure) adds considerable overhead to the underlying `call/cc` mechanism.

11 Future Work

The Etos compiler is still in its infancy. The features in Section 9 need to be added and obviously other changes will be needed for the upcoming Erlang 5.0 specification. There are also some interesting avenues we want to explore.

An interesting extension to Etos is to add library functions to access Gambit-C's C-interface from Erlang code. Interfacing Erlang, Scheme and C code will then be easy.

The Gambit-C side of the compilation can also be improved. In certain cases the Scheme code generated by Etos could be compiled better by Gambit-C (its optimizations were tuned to the style of code Scheme programmers tend to write). It is worth considering new optimizations and extensions specifically designed for Etos's output. For example, a more efficient interface to `call/cc` could be designed. Moreover we think the performance of Etos will improve by a factor of two on average when we start using a native code back-end for Gambit. We are also working on a hard real-time garbage collector and a generational collector to improve the response time.

12 Conclusions

The preliminary version of Etos shows promising results. It performs very well on integer and floating point arithmetic, beating all other currently available implementations of Erlang. Its performance on list processing and process management is not as good but we think this can be improved in a number of ways.

These results are not all that surprising: Scheme and Erlang offer very similar features (data types, functional style, dynamic typing) and their differences (pattern matching, escape methods, concurrency) can be eliminated by a fairly straightforward compilation process. Scheme appears to be well suited as a target for an Erlang compiler.

Acknowledgements

This work was supported in part by grants from the Natural Sciences and Engineering Research Council of Canada and the Fonds pour la formation de chercheurs et l'aide à la recherche.

References

[1] J. L. Armstrong. The development of erlang. In *Proceedings of the International Conference on Functional Programming*, pages 196–203, Amsterdam, June 1997.

[2] J. L. Armstrong, B. O. Däcker, S. R. Virding, and M. C. Williams. Implementing a functional language for highly parallel real-time applications. In *Proceedings of Software Engineering for Telecommunication Switching Systems*, Florence, April 1992.

[3] J. L. Armstrong, S. R. Virding, C. Wikström, and M. C. Williams. *Concurrent Programming in Erlang*. Prentice Hall, second edition edition, 1996.

[4] D. Boucher. Lalr-scm. Available at `ftp.iro.umontreal.ca` in `pub/parallele/boucherd`.

[5] W. Clinger and J. Rees [editors]. Revised^4 Report on the Algorithmic Language Scheme. *Lisp Pointers*, 4(3):1–55, July-September 1991.

[6] D. Dubé. SILex, user manual. Available at `ftp.iro.umontreal.ca` in `pub/parallele`.

[7] M. Feeley. Etos version 1.4. Compiler available at `ftp.iro.umontreal.ca` in `pub/parallele/etos/etos-1.4`.

[8] M. Feeley. Gambit-C version 2.7a, user manual. Compiler available at `ftp.iro.umontreal.ca` in `pub/parallele/gambit/gambit-2.7`.

[9] M. Feeley. Polling efficiently on stock hardware. In *Proceedings of the Functional Programming and Computer Architecture*, pages 179–187, Copenhagen, June 1993.

[10] M. Feeley, J. Miller, G. Rozas, and J. Wilson. Compiling Higher-Order Languages into Fully Tail-Recursive Portable C. Technical Report 1078, Département d'informatique et de recherche opérationelle, Université de Montréal, 1997.

[11] P. H. Hartel, M. Feeley, M. Alt, L. Augustsson, P. Baumann, M. Beemster, E. Chailloux, C. H. Flood, W. Grieskamp, J. H. G. Van Groningen, K. Hammond, B. Hausman, M. Y. Ivory, R. E. Jones, J. Kamperman, P. Lee, X. Leroy, R. D. Lins, S. Loosemore, N. Röjemo, M. Serrano, J.-P. Talpin, J. Thackray, S. Thomas, P. Walters, P. Weis, and P. Wentworth. Benchmarking implementations of functional languages with "Pseudoknot", a float-intensive benchmark. *Journal of Functional Programming*, 6(4):621–655, 1996.

[12] B. Hausman. Turbo Erlang: approaching the speed of C. In Evan Tick and Giancarlo Succi, editors, *Implementations of Logic Programming Systems*, pages 119–135. Kluwer Academic Publishers, 1994.

[13] IEEE Standard for the Scheme Programming Language. IEEE Standard 1178-1990, IEEE, New York, 1991.

[14] E. Johansson, C. Jonsson, T. Lindgren, J. Bevemyr, and H. Millroth. A pragmatic approach to compilation of Erlang. UPMAIL Technical Report 136, Uppsala University, Sweden, July 1997.

[15] The Internet Scheme Repository. `http://www.cs.indiana.edu/scheme-repository`.

[16] C. Wikstrom. Estone, an erlang benchmark. Available at `http://www.ericsson.se:800/cslab/~klacke/estone/`.