

The L₃ project

Advanced Compiler Construction
Michel Schinz – 2026-02-19

Project overview

As the semester progresses, you will get:

- parts of an L₃ compiler written in Scala, and
- parts of a virtual machine, written in C.

You will have to:

- do one non-graded, warm-up exercise,
- complete the compiler,
- complete the virtual machine.

The L₃ language

The L₃ language

L₃ is a **Lisp-like** language. Its main characteristics are:

- it is “dynamically typed”,
- it is functional:
 - functions are first-class values, and can be nested,
 - there are few side-effects (exceptions: mutable blocks and I/O),
- it automatically frees memory,
- it is simple but quite powerful.

A taste of L₃

An L₃ function to compute x^y for $x \in \mathbb{Z}, y \in \mathbb{N}$:

```
(defrec pow
  (fun (x y)
    (cond ((= 0 y) 1)
          ((even? y)
           (let ((t (pow x (/ y 2))))
             (* t t)))
          (#t
           (* x (pow x (- y 1)))))))
```

$x^0 = 1$

$x^{2z} = (x^z)^2$

$x^{z+1} = x(x^z)$

Values

L₃ offers four types of atomic values:

1. unit,
 2. booleans,
 3. characters, represented by their Unicode code point,
 4. integers, 31 bits [!] in two's complement.
- and one type of composite value: tagged blocks.

Literal values

"c₁...c_n"

String literal (translated to a block expression, see later).

'c'

Character literal.

... -2 -1 0 1 2 3 ...

Integer literals (also in base 16 with #x prefix, or in base 2 with #b prefix).

#_tag

Block tag (integer literal, automatically computed by compiler)

#t #f

Boolean literals (true and false, respectively).

#u

Unit literal.

Top-level definitions

(def n e)

Top-level non-recursive definition. The expression e is evaluated and its value is bound to name n in the rest of the program. The name n is *not* visible in expression e.

(defrec n f)

Top-level recursive *function* definition. The function expression f is evaluated and its value is bound to name n in the rest of the program. The function can be recursive, i.e. the name n is visible in the function expression f.

Local definitions

(let ((n₁ e₁) ...) b₁ b₂ ...)

Parallel local value definition. The expressions e₁, ... are evaluated in that order, and their values are then bound to names n₁, ... in the body b₁, b₂, ... The value of the whole expression is the value of the last b_i.

(let* ((n₁ e₁) ...) b₁ b₂ ...)

Sequential local value definition. Equivalent to a nested sequence of **let**:

(let ((n₁ e₁)) (let (...) ...))

(letrec ((n₁ f₁) ...) b₁ b₂ ...)

Recursive local function definition. The function expressions f₁, ... are evaluated and bound to names n₁, ... in the body b₁, b₂ ... The functions can be mutually recursive.

Conditional expressions

(if e₁ e₂ e₃)

Two-ways conditional. If e₁ evaluates to a true value (i.e. anything but **#f**), e₂ is evaluated, otherwise e₃ is evaluated. The value of the whole expression is the value of the evaluated branch.

The else branch, e₃, is optional and defaults to **#u** (unit).

(cond (c₁ b_{1,1} b_{1,2} ...) (c₂ b_{2,1} b_{2,2} ...) ...)

N-ways conditional. If c₁ evaluates to a true value, evaluate b_{1,1}, b_{1,2} ...; else, if c₂ evaluates to a true value, evaluate b_{2,1}, b_{2,2} ...; etc. The value of the whole expression is the value of the evaluated branch or **#u** if none of the conditions are true.

Logical expressions

(and e₁ e₂ e₃ ...)

Short-cutting conjunction. If e₁ evaluates to a true value, proceed with the evaluation of e₂, and so on. The value of the whole expression is that of the last evaluated e_i.

(or e₁ e₂ e₃ ...)

Short-cutting disjunction. If e₁ evaluates to a true value, produce that value. Otherwise, proceed with the evaluation of e₂, and so on.

(not e)

Negation. If e evaluates to a true value, produce the value **#f**. Otherwise, produce the value **#t**.

Loops and blocks

(rec n ((n₁ e₁) ...) b₁ b₂ ...)

General loop. Equivalent to:

**(letrec ((n (fun (n₁ ...) b₁ b₂ ...)))
 (n e₁ ...))**

(begin b₁ b₂ ...)

Sequential evaluation. First evaluate expression b₁, discarding its value, then b₂, etc. The value of the whole expression is the value of the last b_i.

Functions and primitives

(fun (n₁ ...) b₁ b₂ ...)

Anonymous function with arguments n₁, ... and body b₁, b₂, ... The return value is the value of the last b_i.

(e e₁ ...)

Function application. Expressions e, e₁, ... are evaluated in order, and then the value of e – which must be a function – is applied to the value of e₁, ...

Note: if e is a simple identifier, a special form of name resolution, based on arity, is used – see later.

(@ p e₁ e₂ ...)

Primitive application. First evaluate expressions e₁, e₂, ... in that order, and then apply primitive p to the value of these expressions.

Arity-based name lookup

A special name lookup rule is used when analysing a function application in which the function is a simple name:

(n e₁ e₂ ... e_k)

In such a case, the name n@k – i.e. the name itself, followed by @, followed by the arity in base 10 – is first looked up, and used instead of n instead if it exists. Otherwise, name analysis proceeds as usual.

This allows a kind of overloading based on arity (although it is *not* overloading per se).

Arity-based name lookup

Arity-based name lookup can for example be used to define several functions to create lists of different lengths:

```
(def list-make@1 (fun (e1) ...))
```

```
(def list-make@2 (fun (e1 e2) ...))
```

and so on for list-make@3, list-make@4, etc.

With these definitions, the following two function applications are both valid:

1. (list-make 1) (invokes list-make@1),

2. (list-make 1 (+ 2 3)) (invokes list-make@2).

However, the following one is *not* valid, unless a definition for the bare name

list-make also appears in scope:

```
(map list-make l)
```

Primitives

L₃ offers the following primitives:

- integer: < <= + - * / % **truncated division/remainder**

- integer: shift-left shift-right and or xor

- polymorphic: = id **identity**

- type tests: block? int? char? bool? unit?

- character: char->int int->char

- I/O: byte-read byte-write

- tagged blocks: block-alloc

block-tag block-length block-get block-set!

Tagged blocks

L₃ offers a single kind of composite values: tagged blocks. They are manipulated with the following primitives:

(@ block-alloc t s)

Allocates an uninitialised block with tag t and length s.

(@ block-tag b)

Returns the tag of block b, as an integer.

(@ block-length b)

Returns the length of block b.

(@ block-get b n)

Returns the *n*th element (0-based) of block b.

(@ block-set! b n v)

Sets the *n*th element (0-based) of block b to v.

Using tagged blocks

Tagged blocks are a low-level data structure. They are not meant to be used directly in programs, but rather as a means to implement more sophisticated data structures like strings, arrays, lists, etc.

The valid tags range from 0 to 255, inclusive. They are assigned automatically by the compiler, based on the (symbolic) tag names appearing in the program.

Valid primitive arguments

Primitives only work correctly when applied to certain arguments, otherwise their behaviour is undefined.

+ - * and or xor : $int \times int \Rightarrow int$

shift-left shift-right : $int \times (int \in \{0, 1, \dots, 31\}) \Rightarrow int$

/ % : $int \times (int \neq 0) \Rightarrow int$

< <= : $int \times int \Rightarrow bool$

= : $\forall \alpha, \beta. \alpha \times \beta \Rightarrow bool$

id : $\forall \alpha. \alpha \Rightarrow \alpha$

int->char : $int \in \{\text{valid Unicode code-points}\} \Rightarrow char$

char->int : $char \Rightarrow int$

Valid primitive arguments

block? int? char? bool? unit? : $\forall \alpha. \alpha \Rightarrow bool$

byte-read : $\Rightarrow int \in \{-1, 0, 1, \dots, 255\}$

byte-write : $int \in \{0, 1, \dots, 255\} \Rightarrow ?$

arbitrary
return value

block-alloc : $(int \in \{0, 1, \dots, 255\}) \times int \Rightarrow block$

block-tag block-length : $block \Rightarrow int$

block-get : $\exists \alpha. block \times int \Rightarrow \alpha$

block-set! : $\forall \alpha. block \times int \times \alpha \Rightarrow ?$


Undefined behaviour

The fact that primitives have undefined behaviour when applied to invalid arguments means that they can do *anything* in such a case.
For example, division by zero can produce an error, crash the program, or produce an arbitrary value like 0.

Grasping the syntax

Like all Lisp-like languages, L₃ "has no syntax", in that its concrete syntax is very close to its abstract syntax.

For example, the L₃ expression on the left is almost a direct transcription of a pre-order traversal of its AST on the right, in which nodes are parenthesised and tagged, while leaves are unadorned.



L₃ EBNF grammar (1)

```
program ::= { def | defrec | expr } expr
def ::= (def ident expr)
defrec ::= (defrec ident fun)
expr ::= fun | let | let* | letrec | rec | begin | if | cond | and | or | not
        | app | prim | ident | num | block-tag | str | chr | bool | unit
exprs ::= expr { expr }
fun ::= (fun ({ ident }) exprs)
let ::= (let ({ (ident expr) }) exprs)
let* ::= (let* ({ (ident expr) }) exprs)
letrec ::= (letrec ({ (ident fun) }) exprs)
rec ::= (rec ident ({ (ident expr) }) exprs)
begin ::= (begin exprs)
```

L₃ EBNF grammar (2)

```
if ::= (if expr expr [ expr ])
cond ::= (cond (expr exprs) {(expr exprs)})
and ::= (and expr expr { expr })
or ::= (or expr expr { expr })
not ::= (not expr)
app ::= (expr { expr })
prim ::= (@ prim-name { expr })
```

L₃ EBNF grammar (3)

```
str ::= "{any character except newline}"
chr ::= 'any character'
bool ::= #t | #f
unit ::= #u
ident ::= identstart { identstart | digit } [ @ digit { digit } ]
identstart ::= a | ... | z | A | ... | Z | | ! | % | & | * | + | -
| . | / | : | < | = | > | ? | ^ | _ | ~
prim-name ::= block-tag | etc.
```

L₃ EBNF grammar (4)

```
num ::= num2 | num10 | num16
num2 ::= #b digit2 { digit2 }
num10 ::= [-] digit10 { digit10 }
num16 ::= #x digit16 { digit16 }
digit2 ::= 0 | 1
digit10 ::= digit2 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
digit16 ::= digit10 | A | B | C | D | E | F | a | b | c | d | e | f
block-tag ::= #_ identstart { identstart | digit }
```

Exercise

Write the L₃ version of the factorial function, defined as:

fact(0) = 1

fact(n) = n · fact(n - 1) [if n > 0]

What does the following (valid) L₃ program compute?

```
((fun (f x) (f x))
 (fun (x) (@+ x 1))
 20)
```

L₃ syntactic sugar

L₃ syntactic sugar

L₃ has a substantial amount of **syntactic sugar**: constructs that can be syntactically translated to other existing constructs. Syntactic sugar does not offer additional expressive power to the programmer, but some syntactical convenience.

For example, L₃ allows **if** expressions without an else branch, which is implicitly taken to be the unit value #u:

$(\text{if } e_1 e_2) \Leftrightarrow (\text{if } e_1 e_2 \text{ \#u})$

Desugaring

Syntactic sugar is typically removed very early in the compilation process – e.g. during parsing – to simplify the language that the compiler has to handle.

This process is known as **desugaring**.

Desugaring can be specified as a function denoted by $\llbracket \cdot \rrbracket$ taking an L₃ term and producing a desugared CL₃ term (CL₃ is *Core L₃*, the desugared version of L₃). To clarify the presentation, L₃ terms appear in orange, CL₃ terms in green, and meta-terms in black.

L₃ desugaring (1)

To simplify the specification of desugaring for whole programs, we assume that all top-level expressions are wrapped sequentially in a single **(program ...)** expression.

```
 $\llbracket (\text{program } (\text{def } n\ e)\ s_1\ s_2\ \dots) \rrbracket =$   
   $(\text{let } ((n\ \llbracket e \rrbracket))\ \llbracket (\text{program } s_1\ s_2\ \dots) \rrbracket)$   
 $\llbracket (\text{program } (\text{defrec } n\ e)\ s_1\ s_2\ \dots) \rrbracket =$   
   $(\text{letrec } ((n\ \llbracket e \rrbracket))\ \llbracket (\text{program } s_1\ s_2\ \dots) \rrbracket)$   
 $\llbracket (\text{program } e\ s_1\ s_2\ \dots) \rrbracket =$   
   $\llbracket (\text{begin } e\ (\text{program } s_1\ s_2\ \dots)) \rrbracket$   
 $\llbracket (\text{program } e) \rrbracket =$   
   $\llbracket e \rrbracket$ 
```

L₃ desugaring (2)

Desugaring sometimes requires the creation of **fresh names**, i.e. names that do not appear anywhere else in the program. Their binding occurrence is underlined in the rules, as illustrated by the one below.

```
 $\llbracket (\text{begin } b_1\ b_2\ b_3\ \dots) \rrbracket =$   
   $(\text{let } ((\underline{t}\ \llbracket b_1 \rrbracket))\ \llbracket (\text{begin } b_2\ b_3\ \dots) \rrbracket)$   
 $\llbracket (\text{begin } b) \rrbracket =$   
   $\llbracket b \rrbracket$ 
```

L₃ desugaring (3)

```

[[let ((n1 e1) ...) b1 b2 ...]] =
  (let ((n1 [[e1]] ...) [(begin b1 b2 ...)]))
[[let* ((n1 e1) (n2 e2) ...) b1 b2 ...]] =
  [[let ((n1 e1)) (let* ((n2 e2) ...) b1 b2 ...)]]
[[let* () b1 b2 ...]] =
  [(begin b1 b2 ...)]
[[letrec ((f1 (fun (n1,1 ...) b1,1 b1,2 ...) ...) b1 b2 ...)] =
  (letrec ((f1 (fun (n1,1 ...) [(begin b1,1 b1,2 ...)])
    ...))
    [(begin b1 b2 ...)])

```

L₃ desugaring (4)

```

[[fun (n1 ...) b1 b2 ...]] =
  (letrec ((f (fun (n1 ...) [(begin b1 b2 ...)]]))
    f)
[[rec n ((n1 e1) ...) b1 b2 ...]] =
  (letrec ((n (fun (n1 ...) [(begin b1 b2 ...)]]))
    (n [[e1]] ...))
[[e e1 ...]] =
  [[e] [[e1]] ...]
[[@ p e1 ...]] =
  (@ p [[e1]] ...)

```

L₃ desugaring (5)

```

[[if e e1]] =
  [[if e e1 #u]]
[[if e e1 e2]] =
  (if [[e] [[e1]] [[e2]])
[[cond (e1 b1,1 b1,2 ...) (e2 b2,1 b2,2 ...) ...]] =
  [[if e1 (begin b1,1 b1,2 ...) (cond (e2 b2,1 b2,2 ...) ...)]]
[[cond ()]] =
  #u

```

L₃ desugaring (6)

```

[[and e1 e2 e3 ...]] =
  [[if e1 (and e2 e3 ...) #f]]
[[and e]] =
  [[e]]
[[or e1 e2 e3 ...]] =
  [[let ((v e1)) (if v v (or e2 e3 ...))]]
[[or e]] =
  [[e]]
[[not e]] =
  [[if e #f #t]]

```

L₃ desugaring (7)

L₃ does not have a string type. It offers string literals, though, which are desugared to blocks of characters.

```
[[ "c1...cn" ] =  
  [(let ((s (@block-alloc #__string n)))  
    (@block-set! s 0 'c1')  
    ...  
    s)]  
[[l]] = if l is a (non-string) literal  
|  
[[n]] = if n is a name  
n
```

strings are
tagged with the tag
named `_string`

L₃ desugaring example

```
[(program (@byte-write (if #t 79 75))  
  (@byte-write (if #f 79 75)))]  
= [(begin (@byte-write (if #t 79 75))  
  (program  
    (@byte-write (if #f 79 75))))]  
= (let ((t [(@byte-write (if #t 79 75))]))  
  [(begin  
    (program  
      (@byte-write (if #f 79 75)))]))]  
= (let ((t (@byte-write (if #t 79 75))))  
  (@byte-write (if #f 79 75)))
```


Exercise

Desugar the following L₃ expression :

```
(rec loop ((i 1))  
  (int-print i)  
  (if (< i 9)  
    (loop (+ i 1)))))
```

The L₃ compiler

L₃ compiler architecture



Note: CL₃, CPS and ASM each designate a *family* of very similar languages, with minor differences between them.

Intermediate languages

The L₃ compiler manipulates a total of four (families of) languages:

1. L₃ is the source language that is parsed, but never exists as a tree – it is desugared to CL₃ immediately,
2. CL₃ – a.k.a. CoreL₃ – is the desugared version of L₃,
3. CPS is the main intermediate language, on which optimizations are performed,
4. ASM is the assembly language of the target (virtual) machine.

The compiler contains interpreters for the last three languages, which is useful to check that a program behaves in the same way as it is undergoes transformation.

These interpreters also serve as semantics for their language.