

# CORE ERLANG 1.0

## language specification

Richard Carlsson      Björn Gustavsson      Erik Johansson  
Thomas Lindgren      Sven-Olof Nyström      Mikael Pettersson  
                                 Robert Virding

November 14, 2000

### Abstract

We describe a core language for the concurrent functional language ERLANG, aptly named “CORE ERLANG”, presenting its grammar and informal static and dynamic semantics relative to ERLANG. We also discuss built-in functions and other open issues, and sketch a syntax tree representation.

## 1 Motivation

CORE ERLANG is an intermediate representation of ERLANG, intended to lie at a level between source code and the intermediate code typically found in compilers.

During its evolution, the syntax of ERLANG has become somewhat complicated, making it difficult to develop programs that operate on the source. Such programs might be new parsers, optimisers that transform source code, and various instrumentations on the source code, for example profilers and debuggers.

CORE ERLANG should meet the following goals:

- CORE ERLANG should be as regular as possible, to facilitate the development of code-walking tools.
- CORE ERLANG should have a clear and simple semantics.
- CORE ERLANG should be straight-forward to translate to every intermediate code used in any ERLANG implementation; similarly, it should be straightforward to translate from ERLANG programs to equivalent CORE ERLANG programs.
- There should exist a well-defined textual representation of CORE ERLANG, with a simple and preferably unambiguous grammar, making it easy to construct tools for printing and reading programs. This representation should be possible to use for communication between tools using different internal representations of CORE ERLANG.

- The textual representation should be easy for humans to read and edit, in case the developer wants to check what the ERLANG source looks like in CORE ERLANG form, inspect – or modify – the results of code transformations, or maybe write some CORE ERLANG code by hand.

These goals force CORE ERLANG to a fairly high level of abstraction. It is not, for example, possible to break down the `receive` construct into operations that operate on the mailbox, since no such (useful) translation would be compatible with all ERLANG implementations.

Section 2 discusses lexical analysis and parsing. Section 3 gives the grammar for the language and section 4 the static semantics. Section 5 describes evaluation of programs and expressions. Section 6 discusses issues that may need further specification. Appendix A contains a quick reference to the language and appendix B lists character escape sequences. Appendix C shows a simple syntax tree representation.

## 2 Lexical analysis and parsing

We discuss the lexical processing of a CORE ERLANG program in terms of operations on a sequence of Unicode [4] characters (of which both ASCII and Latin-1 [3], ISO/IEC 8859-1, are subsets), such that there are no Unicode escapes (written `\uXXXX`, where each `X` is a hexadecimal digit) in the sequence. Note, though, that it is *not* required that tools handling CORE ERLANG source programs use Unicode for input and output; like ERLANG, no part of the written language *per se* requires characters outside the 7-bit ASCII subset. However, in order to support STANDARD ERLANG [2], tools must be able to handle Unicode encodings of CORE ERLANG character, string and atom literals.

We assume that the translation from the sequence of characters into a sequence of tokens, suitable for parsing according to the grammar of the following section, is straightforward, being very similar to that of ERLANG. In CORE ERLANG, atom literals are always single-quoted, to avoid any possible confusion with keywords. Comments on any source code line in CORE ERLANG, like in ERLANG, begin with the leftmost percent character `%`, (`\u0025`) on that line that does not occur as part of an atom, string or character literal, and continue up to (but not including) the first line terminator following that character. Comments are ignored by the tokenisation, in effect only recognising the line terminator.

A line terminator is defined as the longest sequence of input characters consisting of exactly one ASCII CR (`\u000d`), one ASCII LF (`\u000a`), or one CR followed by one LF. Line terminators are generally treated as whitespace, except in atom and string literals where line terminators are not allowed. CORE ERLANG does not distinguish between periods (`.` characters) that are followed by whitespace (called *FullStop* tokens in the ERLANG Reference Manual [1]) and those that are not (*i. e.*, ordinary separator periods).

The tokenisation should concatenate all adjacent string literals (these may be separated by any number of whitespace characters, line terminators and comments). Thus, the text `"Hey" "Ho"` denotes the same string literal as `"HeyHo"`. This allows strings to be split over several lines.

## 3 Grammar

This section describes the basic grammar of CORE ERLANG programs. An effort has been made to reduce the language as far as it is possible – and practical – while maintaining readability and preserving most of the lexical conventions of ERLANG. For instance, it can be noted that the syntactic distinction between variables and function names is not really necessary, but makes the connection between exported function names and calls to locally bound functions more obvious to the eye than if “plain” ERLANG-style variables were used for all bindings.

### 3.1 Notation

Literals are described using ordinary regular expressions, where ? stands for “zero or one occurrence of”, + for “one or more repetitions of”, and ellipsis (...) indicates repetition over a range of characters; no other symbols are used except parentheses and the standard | for alternative choices and \* for zero or more repetitions.

For some widely used grammar rules we use abbreviations, such as *i* for *Integer* and *v* for *AnnotatedVariable*. The abbreviations are given within parentheses by the corresponding rules.

To further keep the presentation compact, we use ellipsis notation with indices (“ $x_1 \dots x_n$ ”) instead of giving explicit recursive rules for (possibly empty) sequences.

### 3.2 Lexical definitions

```
sign ::= + | -
digit ::= 0 | 1 | ... | 9
uppercase ::= A | ... | Z | \u00c0 | ... | \u00d6 | \u00d8 | ... | \u00de
lowercase ::= a | ... | z | \u00df | ... | \u00f6 | \u00f8 | ... | \u00ff
inputchar ::= any character except CR and LF
control ::= \u0000 | ... | \u001f
space ::= \u0020
namechar ::= uppercase | lowercase | digit | @ | _
escape ::= \ (octal | (^ ctrlchar) | escapechar)
octaldigit ::= 0 | 1 | ... | 7
octal ::= octaldigit (octaldigit octaldigit?)?
ctrlchar ::= \u0040 | ... | \u005f
escapechar ::= b | d | e | f | n | r | s | t | v | " | ' | \
```

An escape sequence `\ octal` denotes the character whose Unicode value is given by the octal numeral. An escape sequence `\ ^ ctrlchar` denotes the character whose Unicode value is 64 less than that of the *ctrlchar*. The meanings of escape sequences `\ escapechar` are defined in appendix B.

### 3.3 Terminals

*Integer (i):*

*sign?* *digit*<sup>+</sup>

*Float:*

*sign?* *digit*<sup>+</sup> . *digit*<sup>+</sup> ((**E** | **e**) *sign?* *digit*<sup>+</sup>)?

*Atom (a):*

' ((*inputchar* except *control* and \ and ') | *escape*)\* '

*Char:*

\$ ((*inputchar* except *control* and *space* and \) | *escape*)

*String:*

" ((*inputchar* except *control* and \ and ") | *escape*)\* "

*VariableName:*

(*uppercase* | (- *namechar*)) *namechar*\*

Note that a single underscore character '\_' is not a valid *VariableName*.

### 3.4 Non-terminals

*AnnotatedModule:*

*Module*  
( *Module* -| [ *c*<sub>1</sub> , ... , *c*<sub>*n*</sub> ] ) (*n* ≥ 0)

*Module:*

module *a* *ModuleHeader* *ModuleBody* end

*ModuleHeader:*

*Exports* *Attributes*

*Exports:*

[ *FunctionName*<sub>1</sub> , ... , *FunctionName*<sub>*n*</sub> ] (*n* ≥ 0)

*FunctionName (a/i):*

*a* / *i*

where *a* is called the *identifier*, and *i* the *arity*.

*Attributes:*

attributes [ *ModuleAttribute*<sub>1</sub> , ... , *ModuleAttribute*<sub>*n*</sub> ] (*n* ≥ 0)

*ModuleAttribute:*

$a = c$

where  $a$  is called the *key*, and  $c$  the *value* of the attribute.

*ModuleBody:*

$FunctionDefinition_1 \cdots FunctionDefinition_n$  ( $n \geq 0$ )

*FunctionDefinition:*

$AnnotatedFunctionName = AnnotatedFun$

*AnnotatedFunctionName:*

$FunctionName$

$( FunctionName -| [ c_1 , \dots , c_n ] )$  ( $n \geq 0$ )

*AnnotatedFun:*

$Fun$

$( Fun -| [ c_1 , \dots , c_n ] )$  ( $n \geq 0$ )

*Constant (c):*

$AtomicLiteral$

$\{ c_1 , \dots , c_n \}$  ( $n \geq 0$ )

$[ c_1 , \dots , c_n ]$  ( $n \geq 1$ )

$[ c_1 , \dots , c_{n-1} | c_n ]$  ( $n \geq 2$ )

*AtomicLiteral:*

$Integer$

$Float$

$Atom$

$Nil$

$Char$

$String$

*Nil:*

$[ ]$

*AnnotatedVariable (v):*

$VariableName$

$( VariableName -| [ c_1 , \dots , c_n ] )$  ( $n \geq 0$ )

*AnnotatedPattern* (*p*):

$v^1$

*Pattern*

( *Pattern* -| [  $c_1, \dots, c_n$  ] ) ( $n \geq 0$ )

*Pattern*:

*AtomicLiteral*

{  $p_1, \dots, p_n$  } ( $n \geq 0$ )

[  $p_1, \dots, p_n$  ] ( $n \geq 1$ )

[  $p_1, \dots, p_{n-1} \mid p_n$  ] ( $n \geq 2$ )

$v = p$

where the last form  $v = p$  is called an *alias pattern*.

*Expression* (*e*):

AnnotatedValueList

AnnotatedSingleExpression

*AnnotatedValueList*:

*ValueList*

( *ValueList* -| [  $c_1, \dots, c_n$  ] ) ( $n \geq 0$ )

*ValueList*:

< *AnnotatedSingleExpression*<sub>1</sub>, ..., *AnnotatedSingleExpression*<sub>n</sub> > ( $n \geq 0$ )

*AnnotatedSingleExpression*:

*SingleExpression*

( *SingleExpression* -| [  $c_1, \dots, c_n$  ] ) ( $n \geq 0$ )

*SingleExpression*:

*AtomicLiteral*

*VariableName*

*FunctionName*

*Tuple*

*List*

*Let*

---

<sup>1</sup>The separation of variables from other patterns is necessary to keep the grammar LALR(1)

*Case*

*Fun*

*Letrec*

*Application*

*InterModuleCall*

*PrimOpCall*

*Try*

*Receive*

*Protected*

*Sequencing*

*Catch*

*Tuple:*

$\{ e_1, \dots, e_n \}$  ( $n \geq 0$ )

Note that this includes the 0-tuple  $\{ \}$  and 1-tuples  $\{x\}$ .

*List:*

$[ e_1, \dots, e_n ]$  ( $n \geq 1$ )

$[ e_1, \dots, e_{n-1} \mid e_n ]$  ( $n \geq 2$ )

*Let:*

`let Variables = e1 in e2`

*Variables:*

$v$

$\langle v_1, \dots, v_n \rangle$  ( $n \geq 0$ )

*Case:*

`case e of AnnotatedClause1 ... AnnotatedClausen end` ( $n \geq 1$ )

*AnnotatedClause:*

*Clause*

$( \textit{Clause} \mid [ c_1, \dots, c_n ] )$  ( $n \geq 0$ )

*Clause:*

*Patterns Guard -> e*

*Patterns:*

`p`  
`< p1 , ... , pn >` ( $n \geq 0$ )

*Guard:*

`when e`

*Fun:*

`fun ( v1 , ... , vn ) -> e` ( $n \geq 0$ )

Note that there is no `end` keyword terminating the expression.

*Letrec:*

`letrec FunctionDefinition1 ... FunctionDefinitionn in e` ( $n \geq 0$ )

*Application:*

`apply e0 ( e1 , ... , en )` ( $n \geq 0$ )

*InterModuleCall:*

`call e'1 : e'2 ( e1 , ... , en )` ( $n \geq 0$ )

*PrimOpCall:*

`primop a ( e1 , ... , en )` ( $n \geq 0$ )

*Try:*

`try e1 catch (v1, v2) -> e2`

*Receive:*

`receive AnnotatedClause1 ... AnnotatedClausen Timeout end` ( $n \geq 0$ )

*Timeout:*

`after e1 -> e2`

where e<sub>1</sub> is called the *expiry expression* and e<sub>2</sub> the *expiry body*.

*Protected:*

`protected e`

*Sequencing:*

`do e1 e2`

*Catch:*

`catch e`



## 4 Static semantics

### 4.1 Annotations

An annotation ( $\bullet$  -| [ $c_1, \dots, c_n$ ]) associates a list of constant literals  $c_1, \dots, c_n$  with the enclosed phrase  $\bullet$ . Annotations are always optional; leaving out an annotation is equivalent to specifying an empty annotation list. The interpretation of annotations on program phrases is implementation-dependent.

### 4.2 Module definitions

The general form of a module definition is:

```
module a [FunctionName1, ..., FunctionNamen]  
  attributes [ModuleAttribute1, ..., ModuleAttributen']  
  FunctionDefinition1 ... FunctionDefinitionn'  
end
```

(cf. p. 4), where the atom  $a$  is the name of the module.

For each *FunctionName*  $a/i$  listed in the *Exports* declaration, it is a compile-time error if the function name  $a/i$  does not occur on the left-hand side of a *FunctionDefinition* in the corresponding *ModuleBody*.

For each *ModuleAttribute*  $a = c$  listed in the *Attributes* declaration, it is a compile-time error if there exists more than one *ModuleAttribute* with the same key  $a$  in the list. The interpretation of module attributes is implementation-dependent.

Each *FunctionDefinition* in the *ModuleBody* associates a *FunctionName*  $a_k/i_k$  with a *Fun*  $f_k$ . It is a compile-time error if the number of parameters of the right-hand side  $f_k$  does not equal the left-hand side arity  $i_k$ . The scope of each such function definition is the whole of the corresponding *Module*; see evaluation of *InterModuleCall* expressions (p. 14) for details. It is a compile-time error if the same function name  $a/i$  occurs on the left-hand side of two function definitions  $D_j, D_k, j \neq k$ , in a *ModuleBody*  $D_1 \dots D_n$ . (Cf. *Letrec* expressions, p. 10.) A function name thus defined in the module body is said to be *exported* by the module if and only if it is also in the *Exports* declaration of the module.

### 4.3 Atomic literals

A *String* (p. 4) is defined as shorthand for the corresponding list of *Char* literals (cf. *List*, p. 7). *E.g.*, "Hi!" denotes the list [ $\$H, \$i, \$!$ ]. Also recall that the tokenisation process will concatenate all adjacent *String* literals that are separated only by whitespace and/or comments.

### 4.4 Lists

For lists in general, the following equivalences are defined:

$$[x_1, \dots, x_n] ::= [x_1, \dots, x_n \mid []]$$

for  $n \geq 1$ , and

$$[x_1, \dots, x_{n-1} \mid x_n] ::= [x_1, \dots, x_{n-2} \mid [x_{n-1} \mid x_n]]$$

for  $n \geq 3$ . Thus, every list (p. 7) can be equivalently written on a unique normal form using only the *list constructor* primitive<sup>2</sup> `[ • | • ]`, and the constant literal `[]` (*Nil*). This also applies to lists in constants (p. 5) and patterns (p. 6).

## 4.5 Expressions

- For a *VariableName* or *FunctionName* expression, it is a compile-time error if the occurrence is not within the scope of a corresponding binding. A *VariableName* can only be bound in a *Let*, a *Clause*, a *Fun*, or a *Try*, while a *FunctionName* can only be bound in a *FunctionDefinition* of a *Module* or a *Letrec*.
- In *Fun* (p. 8) expressions `fun (v1, ..., vn) -> e`, and in *Let* (p. 7) expressions `let <v1, ..., vn> = e1 in e2`, no variable name may occur more than once in  $v_1, \dots, v_n$ . Likewise, in *Try* (p. 8) expressions `try e1 catch (v1, v2) -> e2`,  $v_1$  and  $v_2$  must be distinct variables. A *Let* expression `let <v> = e1 in e2` may equivalently be written `let v = e1 in e2`.
- In a *Letrec* (p. 8) expression `letrec D1 ... Dn in e`, it is a compile-time error if the same *FunctionName*  $a/i$  occurs on the left-hand side of two function definitions  $D_j, D_k, j \neq k$ , in  $D_1 \dots D_n$ .
- In a *Case* (p. 7) expression `case e of Clause1 ... Clausen end`, it is a compile-time error if not all clauses of the expression have the same number of patterns (cf. section 4.6).
- In a *Receive* (p. 8) expression, on the general form:

`receive Clause1 ... Clausen after e1 -> e2 end`

it is a compile-time error if some clause of the expression does not have exactly one pattern (cf. section 4.6).

- A *Protected* (p. 8) expression `protected e` may only occur as part of a clause guard (cf. p. 16).

## 4.6 Clauses and patterns

A *Clause* (p. 7) has the general form `<p1, ..., pn> when e1 -> e2`, where  $e_1$  is known as the *guard* and  $e_2$  as the *body* of the clause.  $e_2$  is any expression, whereas  $e_1$  is a restricted expression that must be valid as a CORE ERLANG clause guard (see section 5.6 for details). If  $n$  is 1, the clause can equivalently be written `p1 when e1 -> e2`.

Each  $p_i$  is a *Pattern* (p. 6) consisting of variables, atomic literals, tuple and list constructors, and alias patterns. *No variable name may occur more than once in the patterns  $p_1, \dots, p_n$  of a clause.* Pattern matching is described in section 5.4.

---

<sup>2</sup>Usually called *cons*.

## 5 Dynamic semantics

CORE ERLANG is a higher-order functional language operating on the same data types as ERLANG. As in ERLANG, functions are identified by the pair of the identifier *and* the arity. However, while in ERLANG a function call evaluates to a single value, in CORE ERLANG the result of evaluating an expression is an *ordered sequence*, written  $\langle x_1, \dots, x_n \rangle$ , of zero, one or more values  $x_i$ . A sequence is not in itself a value; thus it is not possible to create sequences of sequences. For simplicity we denote any single-value sequence  $\langle x \rangle$  by  $x$  where no confusion can ensue. If an expression  $e$  always evaluates to a sequence of values  $\langle x_1, \dots, x_n \rangle$ , then we define the *degree* of  $e$  to be the length  $n$  of this sequence.

An *environment*  $\rho$  is a mapping from names to ERLANG values; *e.g.*,  $\rho = [v \mapsto \text{'true'}]$  maps the single variable name  $v$  to the atom  $\text{'true'}$ . We write  $\rho_1\rho_2$  to denote the extension of  $\rho_1$  by the elements of  $\rho_2$ , such that if  $v \mapsto x$  is in  $\rho_1$  and  $v \mapsto y$  is in  $\rho_2$ , then only the latter is in  $\rho_1\rho_2$ . To simplify the presentation, in the context of environments all names are assumed to be without annotations.

### 5.1 Programs and processes

A CORE ERLANG *program* consists of an unordered set of definitions of distinctly named modules (*cf. Module*, p. 4). Execution of a program is performed by evaluating an initial expression `call  $a_1:a_2(x_1, \dots, x_n)$` , where  $a_1$  and  $a_2$  are atoms and  $x_1, \dots, x_n$  are any values (*cf. InterModuleCall*, p. 14), in an empty environment. The program execution ends when the evaluation of the initial call is completed, either normally, yielding a final result (the interpretation of which is implementation-dependent), or abruptly, by causing an exception to be raised that is not caught by a *Try* expression (*cf. p. 15*) in the program.

Each particular instance of a program execution is associated with some specific *process*. We define a process to be an object with a unique *identity* and a mutable *state*. The state of a process is assumed to contain a *mailbox* object, but otherwise its details are implementation-dependent. A mailbox is an ordered sequence of values, such that its contents may be inspected, a value may be appended to the sequence, and any value (at any position) may be removed from the sequence; no other operations are allowed. The state of a process, including the mailbox, may be mutated at any point during its lifetime, as a side effect of program execution or by other causes; this is also implementation-dependent.

The set of module definitions constituting the program is mutable, and at any time, module definitions may be added, removed or replaced,<sup>3</sup> maintaining the invariant that each module definition in the set is distinctly named. A definition  $m$  with name  $a$  in the set at any time, is generally referred to as the *latest version of  $a$*  at that time.

---

<sup>3</sup>Variously known as “dynamic code replacement”, “run-time code replacement”, and “hot code loading”.

## 5.2 Evaluation of expressions

Argument evaluation in CORE ERLANG is *strict*, *i. e.*, all arguments to an operator are completely evaluated before the evaluation of the operator begins; furthermore, the evaluation order of arguments is *always undefined*, unless otherwise stated (notably in *Let* expressions, *Case* expressions, *Receive* expressions and *Try* expressions). The degree of any expression used as argument to another is unless otherwise stated expected to be 1 (one); if the degree of an expression does not match its use, the behaviour is undefined.

Every expression is assumed to be evaluated in a given environment  $\rho$ , mapping all free variables and function names in the expression to ERLANG values.

Expression evaluation can either terminate normally, yielding a sequence of values, or *abruptly*, by raising an exception. An *exception* is an object having two components called the *tag* and the *value*; both components must be ERLANG values. Except for *Try* expressions (see p. 15), *Protected* expressions (p. 16) and *Catch* expressions (p. 17), if the evaluation of an immediate subexpression  $e'$  of some expression  $e$  terminates abruptly with exception  $x$ , then evaluation of  $e$  also terminates abruptly with exception  $x$ .

*ValueList*:

$$\langle s_1, \dots, s_n \rangle$$

where each  $s_i$  is a *SingleExpression*, which must have degree 1.

This evaluates to the sequence  $\langle x_1, \dots, x_n \rangle$  where for  $i \in [1, n]$ ,  $s_i$  evaluates to  $x_i$ . The degree of the *ValueList* expression is thus  $n$ .

*AtomicLiteral*:

This evaluates to the ERLANG value denoted by the literal. *Nil* (p. 5) denotes the empty list, which is a unique constant whose type is distinct from all other constants; it is thus not *e. g.* an atom. *Char* literals (p. 4) may be interpreted as denoting integer values representing character codes, but implementations may instead support a distinct character type.

*VariableName*:

This evaluates to the value to which the *VariableName*  $v$  is bound in the environment  $\rho$ , that is, the value of  $\rho(v)$ .

*FunctionName*:

$$a/i$$

This evaluates to the *closure* to which the *FunctionName*  $a/i$  is bound in the environment  $\rho$ , that is, the value of  $\rho(a/i)$ . See also *Application* (p. 14) and *Fun* (p. 13).

*Tuple*:

$$\{e_1, \dots, e_n\}$$

This evaluates to the ERLANG  $n$ -tuple  $\{x_1, \dots, x_n\}$ , where for  $i \in [1, n]$ ,  $e_i$  evaluates to  $x_i$ . Note that a 1-tuple  $\{x\}$  is distinct from the value  $x$ , and that the 0-tuple  $\{\}$  is a unique value.

*List:*

`[e1 | e2]`

This evaluates to the ERLANG list constructor `[x1 | x2]`, where for  $i \in [1, 2]$ ,  $e_i$  evaluates to  $x_i$ . See section 4.4 for details on list notation.

*Let:*

`let <v1, ..., vn> = e1 in e2`

$e_1$  is evaluated in the environment  $\rho$ , yielding a sequence  $\langle x_1, \dots, x_n \rangle$  of values.  $e_2$  is then evaluated in the environment  $\rho[v_1 \mapsto x_1, \dots, v_n \mapsto x_n]$ .  $e_1$  must be completely evaluated before evaluation of  $e_2$  begins, unless interleaving their evaluation yields no observable difference. The result is that of  $e_2$  if evaluation of both expressions completes normally.

Note that if for all  $i \in [1, n]$ ,  $v_i$  is not used in  $e_2$ , the expression is effectively a sequencing operator (*cf.* *Sequencing*, p. 17), evaluating  $e_1$  before  $e_2$  but discarding its value.

If  $e_1$  does not have degree  $n$ , the behaviour is undefined.

*Case:*

`case e of P1 when g1 -> b1 ... Pn when gn -> bn end`

where each  $P_i$ ,  $i \in [1, n]$ , is a sequence  $\langle p_{i1}, \dots, p_{ik} \rangle$  of patterns, for some fixed  $k$  (*cf.* p. 10).

The switch expression  $e$  is first evaluated in the environment  $\rho$ . If this succeeds, yielding a sequence  $\langle x_1, \dots, x_k \rangle$  of values, that sequence is then tried against the clauses of the *Case* in environment  $\rho$  as described in section 5.5.

If clause selection succeeds with selected clause  $j$  and mapping  $\rho'$  as result, the body  $b_j$  is evaluated in the environment  $\rho\rho'$ , and the result of that evaluation is the result of the *Case* expression.

If no clause can be selected, or if  $e$  does not have degree  $k$ , the behaviour is undefined.

*Fun:*

`fun (v1, ..., vn) -> e`

This evaluates to the *closure*<sup>4</sup> defined by abstracting the expression  $e$  with respect to the parameters  $v_1, \dots, v_n$  in the environment  $\rho$ ; see also *Application* (p. 14).

*Letrec:*

`letrec a1/i1 = f1 ... an/in = fn in e`

---

<sup>4</sup>A closure is defined as the pair consisting of: a) the program code of the function, and b) the environment in which it should be evaluated.

where for  $k \in [1, n]$ , each  $a_k/i_k$  is a *FunctionName* and each  $f_k$  a *Fun*.

The result of evaluating the *Letrec* in environment  $\rho$  is the result of evaluating expression  $e$  in the environment  $\rho'$ , which is the smallest environment such that:

- for each  $x$  in the domain of  $\rho$ , *except*  $x \in \{a_1/i_1, \dots, a_n/i_n\}$ ,  $\rho'(x)$  is equal to  $\rho(x)$
- for each  $a_k/i_k \in \{a_1/i_1, \dots, a_n/i_n\}$ ,  $\rho'(a_k/i_k)$  is equal to the result of evaluating the corresponding *Fun* expression  $f_k$  in the environment  $\rho'$  itself.

(Note that this definition of  $\rho'$  is circular; however, also note that only *Fun* expressions can be bound by a *Letrec*.)

*Application:*

`apply`  $e_0(e_1, \dots, e_n)$

where  $e_0$  evaluates to a closure  $f$  (*cf.* *Fun*, p. 13).

All of  $e_0, e_1, \dots, e_n$  are evaluated in the environment  $\rho$ . Assume that  $e_1, \dots, e_n$  evaluate to values  $x_1, \dots, x_n$ , respectively, and that  $f$  is the result of evaluating an expression `fun`  $(v_1, \dots, v_k) \rightarrow e'$  in an environment  $\rho'$ . Evaluation of the application is then performed by evaluating  $e'$  in the environment  $\rho'[v_1 \mapsto x_1, \dots, v_n \mapsto x_n]$ , if  $n = k$ .

If  $e_0$  does not evaluate to a closure, or if the number  $n$  of arguments in the application is not equal to the arity  $k$  of  $f$ , the behaviour is implementation-dependent.

If the code defining the function of the closure is no longer available<sup>5</sup> at the time of evaluation of the application, the behaviour is implementation-dependent.

*InterModuleCall:*

`call`  $e'_1:e'_2(e_1, \dots, e_n)$

where  $e'_1$  and  $e'_2$  evaluate to atoms  $a_1$  and  $a_2$ , respectively.

All of  $e'_1, e'_2$  and  $e_1, \dots, e_n$  are evaluated in the environment  $\rho$ . Let  $m$  be the latest version of the module named by  $a_1$  at the time of evaluation of the *InterModuleCall* expression. If the *ModuleBody*  $D_1 \dots D_k$  of  $m$  (*cf.* p. 5) contains a *FunctionDefinition* defining the name  $a_2/n$ , and  $a_2/n$  is also in the *Exports* declaration of  $m$ , then let the closure  $f$  be the result of evaluating the expression `letrec`  $D_1 \dots D_k$  `in`  $a_2/n$ , in the empty environment.<sup>6</sup> The *InterModuleCall* expression is then equivalent to an *Application* `apply`  $f(e_1, \dots, e_n)$ .

<sup>5</sup>An implementation could use a garbage collection scheme to safely remove unused code. Another strategy, used by current ERLANG implementations, is to force the removal of code which has been superseded twice by a newer version. This so-called *purging* of code might however be unsafe, unless extra runtime checks are done.

<sup>6</sup>The domain of the environment of such a closure is simply the function names defined by the module, and it is therefore not necessary to represent the closure explicitly.

If  $a_2/n$  is not defined and exported by  $m$ , the behaviour of the inter-module call expression is implementation-dependent.

If  $e'_1$  and  $e'_2$  do not both evaluate to atoms, the behaviour is implementation-dependent.

*PrimOpCall*:

```
primop a(e1, ..., en)
```

$e_1, \dots, e_n$  are evaluated in the environment  $\rho$  to values  $x_1, \dots, x_n$ , respectively. The primitive operation to be performed is identified by the name  $a$  and the number  $n$  of arguments (its arity).

Evaluation of a *PrimOpCall* is always implementation-dependent and may depend on the values  $x_1, \dots, x_n$ , the state of the associated process (*e.g.* the mailbox), or the external state (*i.e.*, the world). The evaluation may have side effects, and may complete abruptly by raising an exception (*cf.* *Try*, below).

*Try*:

```
try e1 catch (v1, v2) -> e2
```

$e_1$  is evaluated in the environment  $\rho$ , and if that evaluation completes normally, its result is also the result of the *Try* expression; otherwise, if evaluation of  $e_1$  completes abruptly with exception  $x$ ,  $e_2$  is evaluated in the environment  $\rho[v_1 \mapsto t, v_2 \mapsto u]$ , where  $t$  is the *tag* and  $u$  the *value* of  $x$  (*cf.* p. 12), and the result of that evaluation becomes the result of the *Try* expression.

CORE ERLANG defines no specific way of raising exceptions, but given a primitive operation named *e.g.* `raise`, of arity 2, with the effect of always terminating abruptly by raising an exception whose corresponding tag and value are the actual parameters to the call, we could define:

```
erlang:exit(R) ::= primop 'raise'('EXIT', R)
erlang:throw(R) ::= primop 'raise'('THROW', R)
```

for the ERLANG built-in standard functions `exit/1` and `throw/1`.

*Receive*:

```
receive <p1> when g1 -> b1 ... <pn> when gn -> bn
after e1 -> e2 end
```

Evaluation of a *Receive* is divided into stages, as follows:

1. First, the expiry expression  $e_1$  is evaluated to a value  $t$  in the environment  $\rho$ .  $t$  must be either a nonnegative integer or the atom `'infinity'`, otherwise the behaviour is implementation-dependent.
2. Next, each value in the mailbox (of the associated process), in first-to-last order, is tried one at a time against the clauses `<p1> when g1 -> b1 ... <pn> when gn -> bn` in the environment  $\rho$ , as described in section 5.5, until one of the following occurs:

- If for some value  $M_k$  at position  $k$  in the mailbox and some  $i \in [1, n]$ , clause selection succeeds yielding a selected clause  $i$  and a mapping  $\rho'$ , then the element at position  $k$  is first deleted from the mailbox, and expression  $b_i$  is evaluated in environment  $\rho\rho'$  to yield the value of the *Receive*.
  - If there are no remaining values to be tried in the mailbox, then either if  $t$  is the integer 0 (zero), or  $t$  is a positive integer and  $t$  or more milliseconds have passed since the transition from stage 1 to stage 2 was made, the expiry body  $e_2$  is evaluated in environment  $\rho$  to yield the value of the *Receive*; otherwise stage 3 is entered.
3. The evaluation of the *Receive* is at this point suspended, and may be resumed when either or both of the following has occurred:
- One or more values have been appended to the mailbox.
  - $t$  or more milliseconds have passed since the transition from stage 1 to stage 2 was made, when  $t$  is a positive integer.

The evaluation then again enters stage 2, where this time only those values in the mailbox (if any) should be tried that have not been tried since the latest transition from stage 1 to stage 2 was made. (Note that any subsequent *Receive* will thus start over from the first value in the mailbox, and not continue where any previous *Receive* finished.)

*A Receive may never be evaluated as part of the evaluation of a clause guard of another Receive.* The removal of a message from the mailbox is a *side effect*, and this is not allowed in a guard. Even more importantly, two *Receive* expressions being evaluated in a nested fashion using the same mailbox could want to select and remove the same message, and it is not obvious how such conflicts could be resolved in a consistent way useful to the programmer. Another, lesser complication would be that the evaluation would have to be able to track nested timeouts to any depth.

Because the timeout limit  $t$  (when  $t$  is a positive integer) is *soft*, *i. e.*, a lower bound only, an implementation is free to allow any number of values to be appended to the mailbox while evaluation is suspended in stage 3, even after the timeout limit has expired. However, implementations should in general attempt to detect timeouts as soon as possible.

It can be noted that it is quite possible for an implementation to signal timeouts by simply appending a unique value, associated with a particular active *Receive*, to the corresponding mailbox, causing the second wake-up condition of stage 3 to be subsumed by the first. However, unselected timeout messages will then need to be garbage collected from the mailboxes in order to prevent cases of unbounded growth.

*Protected:*

`protected e`

A *Protected* expression may only occur as part of a clause guard.  $e$  is evaluated in the environment  $\rho$ , and if that evaluation completes normally, its



result is also the result of the *Protected* expression; otherwise, if evaluation of  $e$  completes abruptly, the *Protected* expression evaluates to the constant ‘false’.

Outside of clause guards, this behaviour can be simulated by the expression

```
try e catch (v1, v2) -> 'false'
```

where  $v_1$  and  $v_2$  are distinct variables (*cf. Try*, p. 15).

### 5.3 Standard syntactic sugar

This section describes CORE ERLANG expressions that are defined in terms of the primitives that have been described above, but which are nevertheless included in the language for convenience (usually referred to as “syntactic sugar”).

*Sequencing:*

```
do e1 e2
```

This is equivalent to `let <v1, ..., vn> = e1 in e2`, where  $n$  is the degree of  $e_1$ , and the variables  $v_1, \dots, v_n$  do not occur free in  $e_2$ . Thus,  $e_1$  is evaluated before  $e_2$ , but its result is not used (*cf. Let*, p. 13).

*Catch:*

```
catch e
```

This is equivalent to

```
try e
catch (v1, v2) ->
  case v1 of
    'THROW' when 'true' -> v2
    v3 when 'true' -> {v1, v2}
  end
```

where  $v_1$  and  $v_2$  are distinct variables (*cf. Try*, p. 15), and  $v_3$  is not the same variable as  $v_2$ . This encodes the behaviour of ERLANG `catch` expressions.

### 5.4 Pattern matching

Pattern matching recursively matches the structures of a sequence of values  $x_1, \dots, x_n$  against a corresponding sequence of patterns (*cf. p. 6*)  $p_1, \dots, p_n$ , either succeeding, yielding as result a mapping from the variables in  $p_1, \dots, p_n$  to subterms of  $x_1, \dots, x_n$ , or otherwise failing. No variable name may occur more than once in the sequence of patterns.

- A sequence of patterns  $p_1, \dots, p_n$  matches a sequence of values  $x_1, \dots, x_n$ , yielding the mapping  $\rho_1 \cdots \rho_n$ , if and only if for all  $i \in [1, n]$ ,  $p_i$  matches  $x_i$  yielding the mapping  $\rho_i$ .

- An *AtomicLiteral* pattern  $p$  matches a value  $x$ , yielding the empty mapping  $[\ ]$ , if and only if  $p$  denotes  $x$ .
- A *VariableName* pattern  $p$  *always* matches a value  $x$ , yielding the mapping  $[p \mapsto x]$ .
- A tuple pattern  $\{p_1, \dots, p_n\}$  matches a value  $\{x_1, \dots, x_n\}$ , yielding the mapping  $\rho_1 \cdots \rho_n$ , if and only if for all  $i \in [1, n]$ ,  $p_i$  matches  $x_i$  yielding the mapping  $\rho_i$ .
- A list constructor pattern  $[p_1 \mid p_2]$  matches a value  $[x_1 \mid x_2]$ , yielding the mapping  $\rho_1\rho_2$ , if and only if for  $i \in [1, 2]$ ,  $p_i$  matches  $x_i$  yielding the mapping  $\rho_i$ .
- An alias pattern  $v = p$  matches a value  $x$ , yielding the mapping  $\rho[v \mapsto x]$ , if and only if  $p$  matches  $x$  yielding the mapping  $\rho$ .

## 5.5 Clause selection

First, recall that in CORE ERLANG, a variable may occur at most once in the patterns of a single clause, and note that pattern variables are always binding occurrences; variables cannot be repeated or imported in patterns, as they can in ERLANG.

Given a sequence  $x_1, \dots, x_k$  of switch values and an environment  $\rho$ , a sequence of clauses

$$P_1 \text{ when } g_1 \text{ -> } b_1 \cdots P_n \text{ when } g_n \text{ -> } b_n$$

where each  $P_i$  is a sequence  $\langle p_{i1}, \dots, p_{ik} \rangle$  of patterns, is tried in left-to-right order as follows:

If the pattern sequence  $p_{i1}, \dots, p_{ik}$  is matched successfully against  $x_1, \dots, x_k$ , yielding a mapping  $\rho' = [v_1 \mapsto x'_1, \dots, v_m \mapsto x'_m]$ , where  $v_1, \dots, v_m$  are exactly the variables occurring in  $p_{i1}, \dots, p_{ik}$ , each bound in  $\rho'$  to some subterm  $x'_j$  of  $x_1, \dots, x_k$  as the result of the pattern matching (*cf.* section. 5.4 for details), then the expression `protected`  $g_i$  is evaluated in the environment  $\rho\rho'$ . If the result is `'true'`, clause selection succeeds, yielding the selected clause  $i$  and mapping  $\rho'$  as result. If the result is `'false'`, the next clause in order is tried; if no clause remains, clause selection fails.

## 5.6 Clause guards

A CORE ERLANG clause guard must not have observable side effects and should evaluate in bounded (preferably constant or linear) time. Because all guards are evaluated within a `protected` expression, if the evaluation of a guard does not complete normally, the raised exception is implicitly caught and discarded, and the value `'false'` is used for the result, thus failing the clause quietly.<sup>7</sup> If a clause guard evaluates to a value other than `'true'` or `'false'`, the behaviour is undefined.

At present, CORE ERLANG clause guards and all their subexpressions are restricted to *Protected* together with the following subset of expressions:

<sup>7</sup>It is the explicit intention that a compiler might utilise a more efficient error handling mechanism in the restricted case of guard evaluation.

- *AtomicLiteral*
- *VariableName* and *FunctionName*
- *Tuple*
- *List*
- *Let* (and thus also *Sequencing*)
- *PrimOpCall*
- *InterModuleCall*

where for any *PrimOpCall* `primop a(e1, ..., en)`, the primitive operation *a/n* must not have observable side effects, and for any *InterModuleCall* `call e1:e2(e1, ..., en)`, *e<sub>1</sub>* and *e<sub>2</sub>* must be atom literals such that the function named *e<sub>2</sub>/n* in module *e<sub>1</sub>* is trusted to exist and not have observable side effects.<sup>8</sup> The set of trusted functions and primitive operations is implementation-dependent; in implementations of ERLANG, it typically includes those so called built-in functions (BIFs) that are classified as “guard BIFs”, and type tests; see the ERLANG Reference Manual [1] for details.

## 6 Open issues

This section discusses known issues that may warrant further specification in future versions of this document.

### 6.1 Source code portability

Because several details of the semantics of CORE ERLANG have been defined as implementation-dependent, it is possible for an implementation to expect a particular behaviour for each of those details. (Typical examples of expected behaviour could be that an exception on a particular form is raised, or that an attempt is made to load missing code.) Therefore, CORE ERLANG code generated by one implementation (*e.g.*, by translation from ERLANG source code) might not be suitable as input to another implementation that makes different assumptions.

At present, there is no canonical translation from ERLANG to CORE ERLANG, which preserves the semantics of the ERLANG program while making as few assumptions as possible about implementation-dependent behaviour in CORE ERLANG. Furthermore, the naming conventions for the ERLANG operators and type tests (*cf.* section 6.3) need to be standardised in order to create a canonical translation.

---

<sup>8</sup>Not all such “remote” functions must have actual implementations in existing modules, but may instead be aliases for built-in operations known to the compiler.

## 6.2 Syntax for binary objects

The latest addition to the ERLANG language, as of this writing, is a compact syntax for composition and decomposition of binary-type objects. This can easily be expanded into primitive operations during translation from ERLANG to CORE ERLANG, but it is not yet clear whether it is actually preferable that the expansion is performed at that level, or whether a similar syntax should be added to the CORE ERLANG language in order to facilitate optimisations on such expressions.

## 6.3 Built-in functions

The ERLANG language specifies a large number of so-called built-in functions (BIFs), together with a set of unary and binary operators, and boolean type test operations that are recognized only in clause guards. Most BIFs, but not all, currently belong to the `erlang` module. Some BIFs may be used in clause guard expressions. Some BIFs are recognised by the compiler as if implicitly declared as imported, thus not needing to be qualified by their module names.

BIFs are predefined functions supplied with the implementation, but do not have to be implemented in any particular way – they can be inline-expanded by the compiler, implemented in another language such as C, or be regular ERLANG functions. The only requirement on a BIF is that it “must not be redefined during the lifetime of a node” [1], which makes it possible for an implementation to “use all information in the description of the BIF to make execution efficient”.

All BIFs have a “home module”, making it possible to dynamically call also those BIFs that are not implemented as regular ERLANG functions by their module and function names. The ERLANG operators and type tests, however, do not at present have corresponding documented home module and function names, but this will be added in future releases.

The CORE ERLANG representation of an explicit call to an ERLANG BIF that is not a regular ERLANG function can therefore be either of `call a1:a2(...)`, where  $a_1$  and  $a_2$  are atom literals, or `primop a(...)`, where  $a$  is an atom literal. In the former case, it is then assumed that the compiler will recognise the call as an alias for a built-in operation and eventually generate appropriate code for making the call, possibly by first rewriting it as a `primop` call. The names and semantics of `primop` operations are however always implementation-dependent, and it can be expected that programs operating on CORE ERLANG code will be more portable if the form `call a1:a2(...)` is used and retained for as long as possible in the compilation process.

In order to extend the portability of programs that operate on CORE ERLANG code, it will be necessary to parameterise information about built-in functions. Because the ERLANG language keeps evolving, and because different ERLANG implementations may not have the exact same sets of predefined functions, it is generally not a good idea to hard-code assumptions about BIFs. Instead, such information should in as much as possible be moved to separate modules, so that when porting a CORE ERLANG analysis or transformation from one ERLANG implementation to another, only these modules need rewriting. It is then possible that a set of standard modules for BIF information could be agreed on, which could be assumed to be supplied by every ERLANG implementation.

## References

- [1] J. Barklund, R. Virding, *ERLANG 4.7.3 Reference Manual*, draft version 0.7. June 1999.
- [2] J. Barklund, R. Virding, *Specification of the STANDARD ERLANG programming language*, draft version 0.7. June 1999.
- [3] ISO/IEC, *Information processing – 8-bit single-byte coded graphic character sets*, 1987. Reference number ISO 8879:1987.
- [4] The Unicode Consortium, *The Unicode Standard, Version 2.0*. Addison-Wesley, Reading, Mass., 1996.

## A Quick reference

This section gives an informal overview of the elements of the language.

### A.1 Comments

Example:

```
% This is a comment; it ends just before the line break.
```

### A.2 Constant literals

Examples:

**Integers:** 8, +17, 299792458, -4711

**Floating-point numbers:** 0.0, 2.7182818, -3.14, +1.2E-6,  
-1.23e12, 1.0e+9

**Atoms:** 'foo', 'Bar', 'foo bar', '', '%#\010@\n! ',  
'\_hello\_world'

**Character literals:** \$A, \$\$, \$\n, \$\s, \$\\", \$\12, \$\101, \$^A

**Strings:** "Hello, world!", "Two\nlines", "",  
"Ring\^G" "My\7" "Bell\007!"

### A.3 Variables

Examples:

```
X, Bar, Value_2, One2Three, Stay@home, _hello_world
```

### A.4 Keywords

after	apply	attributes	call	case
catch	do	end	fun	in
let	letrec	module	of	primop
protected	receive	try	when	

### A.5 Separators

(	)	{	}	[	]	<	>
,	:		/	=	->	-	

### A.6 Annotations

(• -| [ *const*<sub>1</sub> , ... , *const*<sub>*n*</sub> ] )

## A.7 Programs and expressions

```

module ::= module Atom [fnamei1, ..., fnameik]
           attributes [Atom1 = const1, ..., Atomm = constm]
           fname1 = fun1 ... fnamen = funn end
fname ::= Atom / Integer
const ::= lit | [const1 | const2] | {const1, ..., constn}
lit ::= Integer | Float | Atom
        | Char | String | []
fun ::= fun (var1, ..., varn) -> vals
var ::= VariableName
vals ::= expr | <expr1, ..., exprn>
expr ::= var | fname | lit | fun
        | [vals1 | vals2] | {vals1, ..., valsn}
        | let vars = vals1 in vals2
        | case vals of clause1 ... clausen end
        | letrec fname1 = fun1 ... fnamen = funn in vals
        | apply vals0 (vals1, ..., valsn)
        | call vals'1:vals'2 (vals1, ..., valsn)
        | primop Atom (vals1, ..., valsn)
        | try vals1 catch (var1, var2) -> vals2
        | receive clause1 ... clausen after vals1 -> vals2 end
        | protected vals
        | do vals1 vals2
        | catch vals
vars ::= var | <var1, ..., varn>
clause ::= pats when vals1 -> vals2
pats ::= pat | <pat1, ..., patn>
pat ::= var | lit | [pat1 | pat2] | {pat1, ..., patn}
        | var = pat

```

## B Escape sequences

This table shows the Unicode character values for the escape sequences defined by CORE ERLANG; they are the same as in ERLANG.

<code>\ b</code>	<code>\u0008</code>	(backspace, BS)
<code>\ d</code>	<code>\u007f</code>	(delete, DEL)
<code>\ e</code>	<code>\u001b</code>	(escape, ESC)
<code>\ f</code>	<code>\u000c</code>	(form feed, FF)
<code>\ n</code>	<code>\u000a</code>	(linefeed, LF)
<code>\ r</code>	<code>\u000d</code>	(carriage return, CR)
<code>\ s</code>	<code>\u0020</code>	(space, SPA)
<code>\ t</code>	<code>\u0009</code>	(horizontal tab, HT)
<code>\ v</code>	<code>\u000b</code>	(vertical tab, VT)
<code>\ "</code>	<code>\u0022</code>	(double quote, ")
<code>\ '</code>	<code>\u0027</code>	(apostrophe/single quote, ')
<code>\ \</code>	<code>\u005c</code>	(backslash, \)



## C Syntax tree representation

The following schema describes a representation of CORE ERLANG syntax trees, suitable for most general purposes. For brevity, define  $x^i$  to mean an ordered sequence  $(x_1, \dots, x_i)$ , for  $i \geq 0$ , and  $x^*$  to mean any sequence in  $\bigcup_{i=0}^{\infty} \{x^i\}$ .

```
module ::= module (atom a) fname* attr* def *
fname  ::= fname a i
attr   ::= (atom a, const)
const  ::= lit const*
lit    ::= int i | float f
        | atom a | char c
        | nil | cons | tuple
def    ::= (fname, fun)
fun    ::= fun v* w
v      ::= var s
w      ::= e | values e*
e      ::= v | fname | lit w* | fun
        | let v* w w
        | case w clause*
        | letrec def* w
        | apply w w*
        | call w w w*
        | primop a w*
        | try w v* w
        | receive clause* w w
        | protected w
        | do w w
        | catch w
clause ::= clause pat* w w
pat    ::= v | lit pat* | alias v pat
```

where  $a$  stands for an atom,  $i$  for an integer,  $f$  for a floating-point number,  $c$  for a character and  $s$  for a string.

Implicitly, each constructor above should also have an additional field for an associated list of annotations (`atom a`, `const`).

## Index

- " , *see* strings
- \$ , *see* character literals
- % , *see* comments
- ' , *see* atom literals
- . , *see* period characters
- [ ] , *see* nil
  
- abrupt termination, 12
- alias pattern, 6
- annotations, 9, 22
- application, *see* functions, application
- argument evaluation
  - order, *see* evaluation order
  - strictness, 12
- ASCII, 2
- atom literals, 4
- attributes, *see* module attributes
  
- BIFs, 20
- binary syntax, 20
- built-in operations, *see* primitive operation calls
  
- calling functions
  - application of functional value, 8, 14
  - inter-module calls, 8, 14
- case expressions, 7, 10, 13
- catch expressions, 8, 17
- catching errors, *see* error handling
- character literals, 4, 12
- clause guards, 8, 10
  - evaluation, 18
  - failing, 18
  - restrictions on, 18
  - result of evaluation, 18
- clauses, 7, 10, 18
  - annotations on, 7
  - body, 10
  - evaluation, 13
  - guard, *see* clause guards
- closure, 12–14
  - application of, *see* functions, application
  - unavailable code, *see* functions, calling purged code
  
- code replacement, 11, 14
- comments, 2, 22
- compile-time error
  - arity mismatch in function definition, 9
  - exported name not defined, 9
  - multiply defined function, 9, 10
  - multiply defined module attribute, 9
  - unbound function name, 10
  - unbound variable, 10
  - wrong number of patterns in `case`, 10
  - wrong number of patterns in `receive`, 10
- cons, *see* list constructor
- constant literals, 5, 22
  
- data types, 11
- degree of expression, 11
- do expressions, 8, 17
  
- empty list, *see* nil
- environment, 11
- ERLANG
  - BIFs, 20
  - `receive`, 2
  - tokenisation, 2
- error handling
  - general, *see* try expressions
  - in guards, *see* protected expressions
- escape sequences, 3, 24
- evaluation order
  - in case expressions, 13
  - in let expressions, 13
  - in protected expressions, 17
  - in receive expressions, 15
  - in try expressions, 15
  - of arguments, 12
  - of clauses, 18
- exceptions, 12
  - catching, *see* error handling
  - raising, 15
- execution, *see* programs
- expiry expression, 8
- expressions, 6, 10, 23

- evaluation
  - Application*, 14
  - AtomicLiteral*, 12
  - Case*, 13
  - Catch*, 17
  - Fun*, 13
  - FunctionName*, 12
  - InterModuleCall*, 14
  - Let*, 13
  - Letrec*, 13
  - List*, 13
  - PrimOpCall*, 15
  - Protected*, 16
  - Receive*, 15
  - Sequencing*, 17
  - Try*, 15
  - Tuple*, 12
  - ValueList*, 12
  - VariableName*, 12
- guards, 10
- result of evaluation, 11
- syntax, *see* grammar

floating-point literals, 4

full stop tokens, 2

functions
 

- abstraction, *see* fun expressions
- application, 8, 14
- calling purged code, 14
- defining
  - in *Letrec* expressions, 8, 13
  - in module definitions, 4
- exported, 9
- function names, 3, 4, 11, 12
- inter-module call, 8, 14

fun expressions, 8, 10, 13, 14

grammar, 3
 

- AnnotatedClause*, 7
- AnnotatedFun*, 5
- AnnotatedFunctionName*, 5
- AnnotatedModule*, 4
- AnnotatedPattern*, 6
- AnnotatedSingleExpression*, 6
- AnnotatedValueList*, 6
- AnnotatedVariable*, 5
- Application*, 8
- Atom*, 4
- AtomicLiteral*, 5
- Attributes*, 4
- Case*, 7
- Catch*, 8
- Char*, 4
- Clause*, 7
- Constant*, 5
- Exports*, 4
- Expression*, 6
- Float*, 4
- Fun*, 8
- FunctionDefinition*, 5
- FunctionName*, 4
- Guard*, 8
- Integer*, 4
- InterModuleCall*, 8
- Let*, 7
- List*, 7
- Module*, 4
- ModuleAttribute*, 5
- ModuleBody*, 5
- ModuleHeader*, 4
- Nil*, 5
- Pattern*, 6
- Patterns*, 8
- PrimOpCall*, 8
- Protected*, 8
- Receive*, 8
- Sequencing*, 8
- SingleExpression*, 6
- String*, 4
- Timeout*, 8
- Try*, 8
- Tuple*, 7
- ValueList*, 6
- VariableName*, 4
- Variables*, 7

guards, *see* clause guards

hot-code loading, *see* code replacement

implementation-dependent behaviour
 

- application of non-closure, 14
- arity mismatch, 14
- calling purged code, 14
- improper value of expiry expression, 15
- meaning of module attributes, 9
- meaning of program annotations, 9

- meaning of program result, 11
  - module or function name in `call`
    - not an atom, 15
  - process state, 11
  - target of inter-module call not found, 15
- `infinity`, 15
- input encoding, 2
- integer literals, 4
- inter-module calls, *see* functions
- ISO/IEC 8859-1, *see* Latin-1
- keywords, 22
- lambda abstraction, *see* `fun` expressions
- latest version, *see* module
- Latin-1, 2
- `let` expressions, 7, 10, 13
- `letrec` expressions, 8, 10, 13
- line terminators, 2
- list constructor, 10
- lists, 7, 9, 13
  - in constant literals, 5
  - in patterns, 6
- literals, 5, 9
- mailbox, *see* processes, 11
- module
  - attributes, 5
  - definition, 4
  - garbage collecting, 14
  - latest version, 11
  - purging, 14
- `nil`, 5, 10, 12
- open issues, 19
- operations, *see* primitive operation calls
- pattern matching, 17–18
- patterns, 6, 8, 10
- period characters, 2
- portability of source code, 19
- primitive operation calls, 8, 15
- processes, 11
  - identity, 11
  - mailbox, 11
  - mutating the state of, 11
  - state, 11
- programs, 11, 23
  - execution, 11
- `protected` expressions, 8, 10, 16
- quick reference, 22
- `receive` expressions, 8, 10, 15
- scope of function definitions, 9
- separators, 2, 22
- sequencing of expressions
  - explicit, 8, 17
  - implicit, *see* `let` expressions
- STANDARD ERLANG, 2
- strictness, *see* argument evaluation
- strings, 4
- switching, *see* `case` expressions
- syntactic sugar, 17
- syntax
  - grammar, 3
  - representation, 25
- timeout, *see* expiry expression
- tokenisation, 2
- `try` expressions, 8, 10, 15
- tuples, 7, 12
  - in constant literals, 5
  - in patterns, 6
- undefined behaviour
  - degree mismatch, 12, 13
  - evaluation order, 12
  - no matching `case` clause, 13
  - non-boolean clause guard, 18
- Unicode, 2
- value sequences, 6, 11
- variables, 4, 5, 22
  - binding by pattern matching, 18
  - explicit binding, *see* `let` expressions
  - importing, 18
  - in patterns, 10, 18
- versions of code, 14
- whitespace, 2