



# Erlang: An Overview

## Part 1 – Sequential Erlang

Thanks to Richard Carlsson for the original version of many slides in this part



# Erlang buzzwords

- Functional (strict)
- Single-assignment
- Dynamically typed
- Concurrent
- Distributed
- Message passing
- Soft real-time
- Fault tolerant
- Shared-nothing
- Automatic memory management (GC)
- Virtual Machine (BEAM)
- Native code (HiPE)
- Dynamic code loading
- Hot-swapping code
- Multiprocessor support
- OTP (Open Telecom Platform) libraries
- Open source

# Background

- Developed by Ericsson, Sweden
  - Experiments 1982-1986 with existing languages
    - Higher productivity, fewer errors
    - Suitable for writing (large) telecom applications
    - Must handle concurrency and error recovery
  - No good match - decided to make their own
    - 1986-1987: First experiments with own language
    - Erlang (after Danish mathematician A. K. Erlang)
    - 1988-1989: Internal use
    - 1990-1998: Erlang sold as a product by Ericsson
  - Open Source (MPL-based license) since 1998
    - Development still done by Ericsson



# Erlang at Uppsala University

- High Performance Erlang (HiPE) research group
  - Native code compiler
    - back-ends: SPARC, x86, x86\_64, PowerPC, PowerPC-64, ARM
  - Program analysis and optimization
  - Runtime system improvements
  - Language development and extensions
  - Programming and static analysis tools
- Most results from the HiPE project have been included in the official Erlang distribution



# Hello, World!

```
%% File: hello.erl
-module(hello) .
-export([run/0]) .

-spec run() -> 'ok' .
run() -> io:format("Hello, World!\n") .
```

- '%' starts a comment
- '.' ends each declaration
  - module name, export list, function spec, function declaration
- Every function must be in a module
  - One module per source file
  - Source file name is module name + ".erl"
- ':' used for calling functions in other modules



# Running Erlang

```
$ erl
Erlang (BEAM) emulator version 5.10.3

Eshell V5.10.3 (abort with ^G)
1> 6*7.
42
2> halt().
$
```

- The Erlang VM emulator is called 'erl'
- The interactive shell lets you write any Erlang expressions and run them (must end with '.'')
- The "1>", "2>", etc. is the shell input prompt
- The "halt()" function call exits the emulator



# Compiling a module

```
$ erl
Erlang (BEAM) emulator version 5.10.3

Eshell V5.10.3 (abort with ^G)
1> c(hello) .
{ok,hello}
2>
```

- The “`c(Module)`” built-in shell function compiles a module and loads it into the system
  - If you change something and do “`c(Module)`” again, the new version of the module will replace the old
- There is also a standalone compiler called “`erlc`”
  - Running “`erlc hello.erl`” creates “`hello.beam`”
  - Can be used in a normal Makefile



# Running a program

```
Eshell V5.10.3 (abort with ^G)
1> c(hello) .
{ok,hello}
2> hello:run() .
Hello, World!
ok
3>
```

- Compile all your modules
- Call the exported function that you want to run, using “`module:function(...)`.”
- The final value is always printed in the shell
  - “ok” is the return value from `io:format(...)`





# A recursive function

```
-module(factorial) .  
-export([fact/1]) .  
  
-spec fact(non_neg_integer()) -> pos_integer() .  
fact(N) when N > 0 ->  
    N * fact(N-1) ;  
fact(0) ->  
    1 .
```

- Variables start with upper-case characters!
- ';' separates function clauses; last clause ends with '.'
- Variables are local to the function clause
- Pattern matching and 'when' guards to select clauses
- Run-time error if no clause matches (e.g.,  $N < 0$ )
- Run-time error if N is not an integer



# Tail recursion with accumulator

```
-module(factorial) .  
-export([fact/1]) .  
  
-spec fact(non_neg_integer()) -> pos_integer() .  
fact(N) -> fact(N, 1) .  
  
fact(N, Fact) when N > 0 ->  
    fact(N-1, Fact*N) ;  
fact(0, Fact) ->  
    Fact .
```

- The *arity* is part of the function name: `fact/1`≠`fact/2`
- Non-exported functions are local to the module
- Function definitions cannot be nested (as in C)
- Last call optimization is performed: the stack does not grow if the result is the value of another function call

# Recursion over lists

```
-module(list) .  
-export([last/1]) .  
  
-spec last([T,...]) -> T.  
last([Element]) -> Element;  
last([_|Rest]) -> last(Rest) .
```

- Pattern matching selects components of the data
- “\_” is a “don't care”-pattern (not a variable)
- “[Head|Tail]” is the syntax for a single list cell
- “[]” is the empty list (often called “nil”)
- “[X, Y, Z]” is a list with exactly three elements
- “[X, Y, Z|Tail]” has three or more elements



# List recursion with accumulator

```
-module(list) .
-export([reverse/1]) .

-spec reverse([T]) -> [T] .
reverse(List) -> reverse(List, []).

reverse([Head|Tail], Acc) ->
    reverse(Tail, [Head|Acc]);
reverse([], Acc) ->
    Acc.
```

- The same syntax is used to *construct lists*
- Strings are simply lists of Unicode characters
  - "Hello" = [\$H, \$e, \$l, \$l, \$o] = [72,101,108,108,111]
  - "" = []
- All list functions can be used on strings

```
12345  
-9876  
16#ffff  
2#010101  
$A  
0.0  
3.1415926  
6.023e+23
```

- Arbitrary-size integers (but usually just one word)
- #-notation for base-N integers (max base = 36)
- \$-notation for character codes (ISO-8859-1)
- Normal floating-point numbers (standard syntax)
  - cannot start with just a '.', as in e.g. C



# Atoms

```
true           % Boolean
false          % Boolean
ok             % used as "void" value
hello_world
doNotUseCamelCaseInAtoms
'This is also an atom'
'foo@bar.baz'
```

- Must start with lower-case character or be quoted
- Single-quotes are used to create arbitrary atoms
- Similar to hashed strings
  - Use only one word of data (just like a small integer)
  - Constant-time equality test (e.g., in pattern matching)
  - At run-time: `atom_to_list(Atom)`, `list_to_atom(List)`

```
{ }  
{42}  
{1,2,3,4}  
{movie, "Yojimbo", 1961, "Kurosawa"}  
{foo, {bar, X},  
      {baz, Y},  
      [1,2,3,4,5]}
```

- Tuples are the main data constructor in Erlang
- A tuple whose 1<sup>st</sup> element is an atom is called a *tagged tuple* - this is used like constructors in ML
  - Just a convention – but almost all code uses this
- The elements of a tuple can be any values
- At run-time: `tuple_to_list(Tup)`, `list_to_tuple(List)`



# Other data types

- Functions
  - Anonymous and other
- Byte and bit strings
  - Sequences of bits
  - `<<0,1,2,...,255>>`
- Process identifiers
  - Usually called 'Pids'
- References
  - Unique “cookies”
  - `R = make_ref()`
- No separate Booleans
  - atoms `true/false`
- Erlang values in general are often called “terms”
- All terms are ordered and can be compared with `<`, `>`, `==`, `:=:`, etc.





# Type tests and conversions

```
is_integer(X)
is_float(X)
is_number(X)
is_atom(X)
is_tuple(X)
is_pid(X)
is_reference(X)
is_function(X)
is_list(X)    % [] or [_|_]
```

```
atom_to_list(A)
list_to_tuple(L)
binary_to_list(B)
```

```
term_to_binary(X)
binary_to_term(B)
```

- Note that `is_list` only looks at the first cell of the list, not the rest
- A list cell whose tail is not another list cell or an empty list is called an “improper list”.
  - Avoid creating them!
- Some conversion functions are just for debugging: avoid!
  - `pid_to_list(Pid)`



# Built-in functions (BIFs)

```
length(List)
tuple_size(Tuple)
element(N, Tuple)
setelement(N, Tuple, Val)
```

```
abs(N)
round(N)
trunc(N)
```

```
throw(Term)
halt()
```

```
time()
date()
now()
```

```
self()
spawn(Function)
exit(Term)
```

- Implemented in C
- All the type tests and conversions are BIFs
- Most BIFs (not all) are in the module “erlang”
- Many common BIFs are auto-imported (recognized without writing “erlang: . . .”)
- Operators (+, -, \*, /, ...) are also really BIFs



# Standard libraries

## Application Libraries

- erts
  - erlang
- kernel
  - code
  - file, filelib
  - inet
  - os
- stdlib
  - lists
  - dict, ordict
  - sets, ordsets, gb\_sets
  - gb\_trees
  - ets, dets

- Written in Erlang
- “Applications” are groups of modules
  - Libraries
  - Application programs
    - Servers/daemons
    - Tools
    - GUI system: wx

# Expressions

```
%% the usual operators
```

```
(X + Y) / -Z * 10 - 1
```

```
%% boolean
```

```
X and not Y or (Z xor W)
```

```
(X andalso Y) orelse Z
```

```
%% bitwise operators
```

```
((X bor Y) band 15) bsl 2
```

```
%% comparisons
```

```
X /= Y           % not !=
```

```
X =< Y           % not <=
```

```
%% list operators
```

```
List1 ++ List2
```

- Boolean and/or/xor are *strict* (always evaluate both arguments)
- Use `andalso/orelse` for short-circuit evaluation
- “`=: =`” for equality, not “`=`”
- We can always use parentheses when not absolutely certain about the precedence

# Fun expressions

```
F1 = fun () -> 42 end
42 = F1 ()

F2 = fun (X) -> X + 1 end
42 = F2 (41)

F3 = fun (X, Y) ->
      {X, Y, F1}
    end

F4 = fun ({foo, X}, Y) ->
        X + Y;
      ({bar, X}, Y) ->
        X - Y;
      (_, Y) ->
        Y
    end

F5 = fun f/3

F6 = fun mod:f/3
```

- Anonymous functions (lambda expressions)
  - Usually called “funs”
- Can have several arguments and clauses
- All variables in the patterns are *new*
  - *All variable bindings in the fun are local*
  - *Variables bound in the environment can be used in the fun-body*



# Pattern matching with '='

```
Tuple = {foo, 42, "hello"},  
{X, Y, Z} = Tuple,
```

```
List = [5, 5, 5, 4, 3, 2, 1],  
[A, A | Rest] = List,
```

```
Struct = {foo, [5,6,7,8], {17, 42}},  
{foo, [A|Tail], {N, Y}} = Struct
```

- Successful matching binds the variables
  - But only if they are not already bound to a value!
  - A new variable can also be repeated in a pattern
  - Previously bound variables can be used in patterns
- Match failure causes runtime error (badmatch)

# Case switches

```
case List of
  [X|Xs] when X >= 0 ->
    X + f(Xs);
  [_X|Xs] ->
    f(Xs);
  [] ->
    0;
  _ ->
    throw(error)
end
```

```
%% boolean switch:
case Bool of
  true -> ... ;
  false -> ...
end
```

- Any number of clauses
- Patterns and guards, just as in functions
- ';' separates clauses
- Use “\_” as catch-all
- Variables may also begin with underscore
  - Signals “I don't intend to use the value of this variable”
  - Compiler won't warn if this variable is not used
- OBS: Variables may be already bound in patterns!



# If switches and guard details

```
if
  0 =< X, X < 256 ->
    X + f(Xs);
true ->
  f(Xs)
end
```

The above construct is better written as

```
case 0 =< X and X < 256 of
  true ->
    X + f(Xs);
false ->
  f(Xs)
end
```

- Like a case switch without the patterns and the “when” keyword
- Need to use “true” as catch-all guard (Ugly!)
- Guards are special
  - Comma-separated list
  - Only specific built-in functions (and all operators)
  - No side effects





# List comprehensions

```
%% map
[f(X) || X <- List]

%% filter
[X || X <- Xs, X > 0]
```

```
Eshell v5.10.3 (abort ...^G)
1> L = [1,2,3].
[1,2,3]
2> [X+1 || X <- L].
[2,3,4]
3> [2*X || X <- L, X < 3].
[2,4]
4> [X+Y || X <- L, Y <- L].
[2,3,4,3,4,5,4,5,6]
```

- Left of the “||” is an *expression template*
- “Pattern <- List” is a *generator*
  - Elements are picked from the list in order
- The other expressions are *Boolean filters*
- If there are multiple generators, you get all combinations of values



# List comprehensions: examples

```
%% quicksort of a list
qsort([]) -> [];
qsort([P|Xs]) ->
    qsort([X || X <- Xs, X =< P])
    ++ [P]    % pivot element
    ++ qsort([X || X <- Xs, P < X]).
```

```
%% generate all permutations of a list
perms([]) -> [[]];
perms(L) ->
    [[X|T] || X <- L, T <- perms(L -- [X])].
```

- Using comprehensions we get very compact code  
...which sometimes can take some effort to understand  
Try writing the same code without comprehensions



# Bit strings and comprehensions

- Bit string pattern matching:

```
case <<8:4, 42:6>> of
  <<A:7/integer, B/bits>> -> {A,B}
end
```

```
case <<8:4, 42:6>> of
  <<A:3/integer, B:A/bits, C/bits>> -> {A,B,C}
end
```

- Bit string comprehensions:

```
<< <<x:2>> || <<x:3>> <= Bits, X < 4 >>
```

- Of course, one can also write:

```
[ <<x:2>> || <<x:3>> <= Bits, X < 4 ]
```

# Catching exceptions

```
try
  lookup(X)
catch
  not_found ->
    use_default(X);
  exit:Term ->
    handle_exit(Term)
end

%% with 'of' and 'after'
try lookup(X, File) of
  Y when Y > 0 -> f(Y);
  Y -> g(Y)
catch
  ...
after
  close_file(File)
end
```

- Three classes of exceptions
  - throw: user-defined
  - error: runtime errors
  - exit: end process
  - Only catch throw exceptions, normally (implicit if left out)
- Re-thrown if no catch-clause matches
- “after” part is always run (side effects only)



# Old-style exception handling

```
Val = (catch lookup(X)),  
  
case Val of  
  not_found ->  
    %% probably thrown  
    use_default(X);  
  {'EXIT', Term} ->  
    handle_exit(Term);  
  _ ->  
    Val  
end
```

- “catch Expr”
  - Value of “Expr” if no exception
  - Value X of “throw(X)” for a throw-exception
  - “{'EXIT', Term}” for other exceptions
- Hard to tell what happened (not safe)
- Mixes up errors/exits
- In lots of old code

# Record syntax

```
-record(foo,  
      {a = 0 :: integer(),  
       b      :: integer()}).  
{foo, 0, 1} = #foo{b = 1}  
  
R = #foo{}  
{foo, 0, undefined} = R  
  
{foo, 0, 2} = R#foo{b=2}  
  
{foo, 2, 1} = R#foo{b=1, a=2}  
  
0 = R#foo.a  
undefined = R#foo.b  
  
f(#foo{b = undefined}) -> 1;  
f(#foo{a = A, b = B})  
  when B > 0 -> A + B;  
f(#foo{}) -> 0.
```

- Records are just a syntax for working with tagged tuples
- You don't have to remember element order and tuple size
- Good for internal work within a module
- Not so good in public interfaces (users must have same definition!)

# Preprocessor

```
-include("defs.hrl").  
  
-ifndef(PI).  
-define(PI, 3.1415926).  
-endif.  
  
area(R) -> ?PI * (R*R).  
  
-define(foo(X), {foo,X+1}).  
  
{foo,42} = ?foo(41)  
  
%% pre-defined macros  
?MODULE  
?LINE
```

- C-style token-level preprocessor
  - Runs after tokenizing, but before parsing
- Record definitions often put in header files, to be included
- Use macros mainly for constants
- Use functions instead of macros if you can (compiler can inline)

# Maps

- Compound terms with a variable number of key-value associations (introduced in Erlang/OTP 17)

```
Eshell V6.2.1 (abort ...^G)
1> M1 = #{name=>"kostis", age=>42, children=>[]}.
#{age => 42,children => [],name => "kostis"}
2> maps:get(age, M1) .
42
3> M2 = maps:update(age, 43, M1) .
#{age => 43,children => [],name => "kostis"}
4> M2#{age := 44, children := ["elina"]}.
#{age => 44,children => ["elina"],name => "kostis"}
5> maps:keys(M2) .
[age,children,name]
6> maps:values(M2) .
[43,[],"kostis"]
7> #{age := Age, children := []} = M1, Age.
42
```



# Dialyzer: A defect detection tool

- A static analyzer that identifies discrepancies in Erlang code bases
  - code points where something is wrong
    - often a bug
    - or in any case something that needs fixing
- Fully automatic
- Extremely easy to use
- Fast and scalable
- Sound for defect detection
  - “Dialyzer is never wrong”



# Dialyzer

- Part of the Erlang/OTP distribution since 2007
- Detects
  - Definite type errors
  - API violations
  - Unreachable and dead code
  - Opacity violations
  - Concurrency errors
    - Data races (`-Wrace_conditions`)
- Experimental extensions with
  - Stronger type inference: type dependencies
  - Detection of message passing errors & deadlocks





# How to use Dialyzer

- First build a PLT (needs to be done once)

```
> dialyzer --build_plt --apps erts kernel stdlib
```

- Once this finishes, analyze your application

```
> cd my_app  
> erlc +debug_info -o ebin src/*.erl  
> dialyzer ebin
```

- If there are unknown functions, you may need to add more Erlang/OTP applications to the PLT

```
> dialyzer --add_to_plt --apps mnesia inets
```