



University of  
St Andrews

# A Concurrency System for IDRIS & ERLANG

Archibald Samuel Elliott

School of Computer Science

University of St Andrews

St Andrews, Scotland

ashe@st-andrews.ac.uk



## Abstract

Concurrent programming is notoriously difficult, due to needing to reason not only about the sequential progress of any algorithms, but also about how information moves between concurrent agents.

What if programmers were able to reason about their concurrent programs and statically verify both sequential and concurrent guarantees about those programs' behaviour? That would likely reduce the number of bugs and defects in concurrent systems.

I propose a system combining dependent types, in the form of the IDRIS programming language, and the Actor model, in the form of ERLANG and its runtime system, to create a system which allows me to do exactly this. By expressing these concurrent programs and properties in IDRIS, and by being able to compile IDRIS programs to ERLANG, I can produce statically verified concurrent programs.

Given that these programs are generated for the ERLANG runtime system, I also produce verified, flexible IDRIS APIs for ERLANG/OTP behaviours.

What's more, with this description of the IDRIS code generation interface, it is now much easier to write code generators for the IDRIS compiler, which will bring dependently-typed programs to even more platforms and situations. I, for one, welcome our new dependently-typed supervisors.

## 1. IDRIS to ERLANG Compiler

My compiler uses the existing IDRIS code generation system. There are three possible IDRIS intermediate representations that I could generate code from: a high-level IR with lambdas and laziness, a defunctionalised IR with only fully-applied functions, and an applicative-normal form IR.

In my case, we used the defunctionalised IR to generate ERLANG source code from, according to the following translation,  $\mathcal{E} \llbracket \langle \text{expression} \rangle \rrbracket$ .

$\mathcal{V} \llbracket \langle \text{variable} \rangle \rrbracket$  turns IR variables into valid ERLANG variable names;  $\mathcal{C} \llbracket \langle \text{name} \rangle ; \langle \text{expression} \rangle^* \rrbracket$  creates constructors;  $\mathcal{N} \llbracket \langle \text{name} \rangle \rrbracket$  turns IR names into valid ERLANG Atoms;  $\mathcal{L} \llbracket \langle \text{constant} \rangle \rrbracket$  turns IR constants into ERLANG expressions; and  $\mathcal{O} \llbracket \langle \text{operation} \rangle ; \langle \text{expression} \rangle^* \rrbracket$  turns primitive operators into their ERLANG equivalent.

```

 $\mathcal{E} \llbracket \langle \text{variable} \rangle \rrbracket \Rightarrow \mathcal{V} \llbracket \langle \text{variable} \rangle \rrbracket$ 
 $\mathcal{E} \llbracket \langle \text{name} \rangle (\langle \text{expression} \rangle^*) \rrbracket \Rightarrow \mathcal{C} \llbracket \langle \text{name} \rangle ; \langle \text{expression} \rangle^* \rrbracket$  when  $\langle \text{name} \rangle$  is a Constructor
 $\Rightarrow \mathcal{N} \llbracket \langle \text{name} \rangle \rrbracket (\mathcal{E} \llbracket \langle \text{expression} \rangle_0 \rrbracket, \dots, \mathcal{E} \llbracket \langle \text{expression} \rangle_{n-1} \rrbracket)$  otherwise
 $\mathcal{E} \llbracket \text{let } \langle \text{name} \rangle := \langle \text{expression} \rangle_0 \text{ in } \langle \text{expression} \rangle_1 \rrbracket \Rightarrow \mathcal{V} \llbracket \text{global } \langle \text{name} \rangle = \text{begin } \mathcal{E} \llbracket \langle \text{expression} \rangle_0 \rrbracket \text{ end, } \mathcal{E} \llbracket \langle \text{expression} \rangle_1 \rrbracket \rrbracket$ 
 $\mathcal{E} \llbracket \text{update } \langle \text{name} \rangle := \langle \text{expression} \rangle \rrbracket \Rightarrow \mathcal{E} \llbracket \langle \text{expression} \rangle \rrbracket$ 
 $\mathcal{E} \llbracket \langle \text{expression} \rangle_n \rrbracket \Rightarrow \text{element } (n+2, \mathcal{E} \llbracket \langle \text{expression} \rangle \rrbracket)$ 
 $\mathcal{E} \llbracket \text{new } \langle \text{name} \rangle (\langle \text{expression} \rangle^*) \rrbracket \Rightarrow \mathcal{C} \llbracket \langle \text{name} \rangle ; \langle \text{expression} \rangle^* \rrbracket$ 
 $\mathcal{E} \llbracket \text{case } \langle \text{expression} \rangle \text{ of } \langle \text{alternative} \rangle^* \text{ end} \rrbracket \Rightarrow \text{case } \mathcal{E} \llbracket \langle \text{expression} \rangle \rrbracket \text{ of}$ 
 $\quad \mathcal{A} \llbracket \langle \text{alternative} \rangle_0 \rrbracket ;$ 
 $\quad \dots;$ 
 $\quad \mathcal{A} \llbracket \langle \text{alternative} \rangle_{n-1} \rrbracket$ 
 $\quad \text{end}$ 
 $\mathcal{E} \llbracket \text{chkcase } \langle \text{expression} \rangle \text{ of } \langle \text{alternative} \rangle^* \text{ end} \rrbracket \Rightarrow \mathcal{E} \llbracket \text{case } \langle \text{expression} \rangle \text{ of } \langle \text{alternative} \rangle^* \text{ end} \rrbracket$ 
 $\mathcal{E} \llbracket \langle \text{constant} \rangle \rrbracket \Rightarrow \mathcal{L} \llbracket \langle \text{constant} \rangle \rrbracket$ 
 $\mathcal{E} \llbracket \text{foreign } \langle \text{fdesc} \rangle_0 (\langle \text{fdesc} \rangle_1 (\langle \text{farg} \rangle^*) \rrbracket) \rrbracket \Rightarrow \langle \text{fdesc} \rangle_1 (\mathcal{E} \llbracket \langle \text{expression} \rangle_2 \rrbracket, \dots, \mathcal{E} \llbracket \langle \text{expression} \rangle_{n+1} \rrbracket)$ 
 $\mathcal{E} \llbracket \text{operator } \langle \text{operator} \rangle (\langle \text{expression} \rangle^*) \rrbracket \Rightarrow \mathcal{O} \llbracket \langle \text{operator} \rangle ; \langle \text{expression} \rangle^* \rrbracket$ 
 $\mathcal{E} \llbracket \text{nothing} \rrbracket \Rightarrow \text{'undefined'}$ 
 $\mathcal{E} \llbracket \text{error } \langle \text{string} \rangle \rrbracket \Rightarrow \text{erlang:error } (\langle \text{string} \rangle)$ 
 $\mathcal{A} \llbracket \text{match } \langle \text{name} \rangle_0 (\langle \text{name} \rangle^* \rightarrow \langle \text{expression} \rangle) \rrbracket \Rightarrow \mathcal{C} \llbracket \langle \text{name} \rangle_0 ; \text{global } \langle \text{name} \rangle_1 \dots \text{global } \langle \text{name} \rangle_n \rrbracket \rightarrow \mathcal{E} \llbracket \langle \text{expression} \rangle \rrbracket$ 
 $\mathcal{A} \llbracket \langle \text{constant} \rangle \rightarrow \langle \text{expression} \rangle \rrbracket \Rightarrow \mathcal{L} \llbracket \langle \text{constant} \rangle \rrbracket \rightarrow \mathcal{E} \llbracket \langle \text{expression} \rangle \rrbracket$ 
 $\mathcal{A} \llbracket \text{default} \rightarrow \langle \text{expression} \rangle \rrbracket \Rightarrow \_ \rightarrow \mathcal{E} \llbracket \langle \text{expression} \rangle \rrbracket$ 

```

## Foreign Calls

In IDRIS, the type of `IO a` is in fact parameterised by the types that the foreign function interface uses. So for instance, `IO a` is the type of computations that use C-like types in their foreign functions.

ERLANG doesn't have a type system like C's, so instead we defined our new computation type, `EIO a` which is parameterised by the types that ERLANG uses, so we can make foreign calls to ERLANG code.

## 2. Concurrency in IDRIS

Of course, our IDRIS to ERLANG compiler is uninteresting unless we can use it to write concurrent programs.

The system I settled on implements the Actor model — ERLANG's underlying concurrency model — as a type `Process l a`, where `l` is the type of messages that process can receive (something I like to refer to as the language of a process), and `m` is the type of what the computation creates. The type of `m` is in fact a `Type*` so that we can use IDRIS' Uniqueness Types to provide guarantees about protocol progress.

We can then embed important details about how spawning, sending and receiving work into the types of the functions that do so, for instance:

```

1 -- The Process Type, from above
  data Process : Type → Type* → Type

-- A reference to a process, as used for messaging.
5 -- This embeds the process' language in its type.
  data ProcRef : Type → Type

-- Spawn gives us a reference to the spawned process,
-- indexed by the correct language.
10 spawn : Process l' a → Process l (ProcRef l')

-- Receive returns us a value with a type that matches
-- the process' language.
receive : Process l l

15 -- Send can only send messages of the correct type
send : ProcRef l' → l' → Process l ()

```

For instance, this program will only run concurrently (it uses some basic constructs on top of `Process` to provide RPC functionality):

```

1 module Main

import ErlPrelude
import Erlang.Process
5 import Erlang.RPC

-- This is the handler in the RPC process.
-- 'i' is the current count,
-- 'Nothing' will read the current count,
-- 'Just x' will add 'x' to the count and return the new count,
-- The returned pair contains the reply and the new count
counter_rpc : Int → Maybe Int → Process (Maybe Int) (Int, Int)
counter_rpc i Nothing = return (i, i)
counter_rpc i (Just x) = return (i+x, i+x)

15 test_proc : Process Int ()
test_proc = do counter ← spawn_rpc counter_rpc 0
              x ← rpc counter (Just 3)
              lift ◦ putStrLn $ "x (3) = " ++ (cast x)
              y ← rpc counter (Nothing)
              lift ◦ putStrLn $ "y (3) = " ++ (cast x)
              return ()

20 main : EIO ()
25 main = run test_proc

```

## 3. ERLANG Concurrency Patterns

ERLANG/OTP contains several other higher-order actor patterns which most programmers structure their applications around. For instance, there are behaviours for generalised servers (`gen_server`), generalised finite state machines (`gen_fsm`), and generalised event handlers (`gen_event`). We can put types on these patterns, especially adding languages for these processes. In this way, we use dependent typ-

ing to allow for flexible but strong types about various parts of the languages.

For instance, with a generalised server, there are two ways of communicating with it. Firstly, there is a synchronous call, which will receive a reply. Then there is also an asynchronous cast, which will not receive a reply.

We can encode this inside a more advanced language type, like so:

```

1 -- A Gen Server language
-- cl is the call type
-- (cl → Type) computes the reply type from the call value
-- the third type is the cast type
5 data GSL : (cl:Type) → (cl → Type) → Type → Type

-- Type of computations that can interact with a Gen Server
GSP : Type → Type

10 -- A Gen Server reference includes the language
data GSRef : (GSL _ _ _) → Type

-- Making a synchronous call to a Gen Server
call : {l:GSL cl cr ct} → GSRef l → (m:cl) → GSP (cr m)

15 -- Making an asynchronous cast to a Gen Server
cast : {l:GSL _ _ ct} → GSRef l → ct → GSP Unit

```

## Conclusion

I have shown that IDRIS is an entirely adequate programming language for concurrent programming. With the new ERLANG code generation system that I have built, we can now write and run safe, flexible actor-based programs which conform to statically proven guarantees.

I have built an IDRIS to ERLANG compiler, which supports concurrent IDRIS code. Along the way I have documented exactly how to produce a compiler and how my compiler works.

I have also devised a way to model the Actor model in IDRIS' dependent types in a lightweight way such that we can run the resultant concurrent programs on ERLANG, including doing inter-process RPC.

I looked into a heavily dependently-typed model of the three main ERLANG/OTP behaviours, which I finished, but are not executable.

I wanted to look into modelling other concurrency calculi in IDRIS, but was unable to do so, and likewise I was unable to devise a Hoare-like logic for ERLANG and its runtime system. I hope that the work in my dissertation paves the way for future research into concurrent programming with IDRIS and other dependently-typed programming languages.