

# Causal-Consistent Debugging of Distributed Erlang Programs

Giovanni Fabbretti <sup>1</sup>, Ivan Lanese <sup>2</sup>, Jean-Bernard Stefani <sup>1</sup>

<sup>1</sup>SPADES Team, INRIA

<sup>2</sup>University of Bologna/FOCUS Team, INRIA

Reversible Computation, 07-07-2021

# Introduction

CauDER is a reversible causal-consistent debugger for the Erlang programming language.

Distinctive features of CauDER:

- Reversibility of systems composed by several processes
- Causal-consistent rollback

Our goal: extend CauDER with support for distribution

# Motivations

Concurrent and distributed systems are everywhere and both are well known for their intrinsic difficulties.

Hence we need effective tools while writing code.

# The Erlang language

Erlang, developed in 1986 by Ericsson, is a concurrent, distributed, functional programming language, based on message passing.

It is probably the most popular programming language that implements the **actor model**.

Erlang owes its success to three aspects: the support of concurrency and distribution, the facilities to do error-handling and the OTP libraries.

# Reversibility

**Causal Consistency:** before undoing an action we must ensure that all of its consequences, if any, have been undone.

Each process embeds a memory that contains information about past states of its computation.

## Our contribution

CauDEr, in its first version, addressed the functional and concurrent fragment of the Erlang language.

This work extends CauDEr and the theory behind it with the support for distributed programs.

# Table of Contents

- 1 Introduction
- 2 Background**
- 3 Distributed CauDER
- 4 Future work

# A deeper look inside CauDER

CauDER implements three semantics:

- a forward semantics that defines Erlang's behavior and stores information in the history
- a backward semantics that ensures that we undo only actions whose consequences have been already undone
- a rollback semantics which automatically undoes all the consequences of an action



# Semantics structure

## Definition (System)

A system is defined as  $\Gamma; \Pi$  where

- $\Gamma$  is the global mailbox
- $\Pi$  is a pool of processes

## Definition

A process is defined as:  $\langle p, h, \theta, e, q \rangle$  where

- $p$  is the process id
- $h$  is the process history
- $\theta$  is the process environment
- $e$  is the expression to evaluate
- $q$  is the local mailbox

# Causal dependencies

We say that there is a dependency between two actions in two cases:

- they cannot be executed in the opposite order
- by executing them in the opposite order the result would change

(A simplification of) The dependencies that rose in the functional and concurrent fragment of the language are that a receive of a message depends on its send.

# Table of Contents

- 1 Introduction
- 2 Background
- 3 Distributed CauDER**
- 4 Future work

# The distributed primitives

The three primitives for distributed computing that we support are the following:

- `slave:start/{1,2}` starts a slave node
- `erlang:node/0` returns the name of the local node
- `erlang:nodes/0` returns the names of the connected nodes

With the introduction of nodes we also updated the `spawn`, which now takes an extra parameter that represents the node where it must be performed.

# Causal dependencies

Informally, one could summarise the causal dependencies of the distributed primitives as follow:

- 1 every action of process  $p$  depends on the (successful) spawn of  $p$ ;
- 2 a (successful) spawn on node  $nid$  depends on the start of  $nid$ ;
- 3 a (successful) start of node  $nid$  depends on previous failed spawns on the same node, if any (if we swap the order, the spawn will succeed);
- 4 a failed start of node  $nid$  depends on its (successful) start;
- 5 a nodes reading a set  $\Omega$  depends on the start of all nids in  $\Omega$ , if any.

## Semantics structure

### Definition

A system is defined as  $\Gamma; \Pi; \Omega$  where

- $\Omega$  is the set of connected nodes
- the other elements are as before

### Definition

A process is defined as:  $\langle nid, p, h, \theta, e, q \rangle$  where

- $nid$  is the id of the node where the process is running
- the other elements are as before

## Extended forward and backward semantics

$$\begin{array}{c}
 (StartS) \quad \frac{\theta, e \xrightarrow{\text{start}(\kappa, nid')} \theta', e' \quad nid' \notin \Omega}{\Gamma; \langle nid, p, h, \theta, e \rangle \mid \Pi; \Omega \xrightarrow{p, \text{start}(nid'), \{s, st_{nid'}\}} \Gamma; \langle nid, p, \text{start}(\theta, e, \text{succ}, nid') : h, \theta', e' \{ \kappa \mapsto nid' \} \rangle \mid \Pi; \{nid'\} \cup \Omega} \\
 \\
 (\overline{StartS}) \quad \frac{\Gamma; \langle nid, p, \text{start}(\theta, e, \text{succ}, nid') : h, \theta', e' \rangle \mid \Pi; \Omega \cup \{nid'\}}{\text{if } \text{spawns}(nid', \Pi) = [] \wedge \text{reads}(nid', \Pi) = [] \wedge \text{failed\_starts}(nid', \Pi) = [] \text{ then } \Gamma; \langle nid, p, h, \theta, e \rangle \mid \Pi; \Omega}
 \end{array}$$

## Extended forward and backward semantics

(StartS)

$$\theta, e \xrightarrow{\text{start}(\kappa, \text{nid}')} \theta', e' \quad \text{nid}' \notin \Omega$$

$$\frac{}{\Gamma; \langle \text{nid}, p, h, \theta, e \rangle \mid \Pi; \Omega \xrightarrow{p, \text{start}(\text{nid}'), \{s, st_{\text{nid}'}\}} \Gamma; \langle \text{nid}, p, \text{start}(\theta, e, \text{succ}, \text{nid}') : h, \theta', e' \{ \kappa \mapsto \text{nid}' \} \rangle \mid \Pi; \{ \text{nid}' \} \cup \Omega}$$

(StartS̄)

$$\frac{\Gamma; \langle \text{nid}, p, \text{start}(\theta, e, \text{succ}, \text{nid}') : h, \theta', e' \rangle \mid \Pi; \Omega \cup \{ \text{nid}' \}}{\begin{array}{l} \leftarrow_{p, \text{start}(\text{nid}'), \{s, st_{\text{nid}'}\}} \Gamma; \langle \text{nid}, p, h, \theta, e \rangle \mid \Pi; \Omega \\ \text{if } \text{spawns}(\text{nid}', \Pi) = [] \wedge \text{reads}(\text{nid}', \Pi) = [] \wedge \text{failed\_starts}(\text{nid}', \Pi) = [] \end{array}}$$



## Extended forward and backward semantics

$$(\text{StartS}) \quad \frac{\theta, e \xrightarrow{\text{start}(\kappa, \text{id}')} \theta', e' \quad \text{id}' \notin \Omega}{\Gamma; \langle \text{id}, p, h, \theta, e \rangle \mid \Pi; \Omega \xrightarrow{p, \text{start}(\text{id}'), \{s, st_{\text{id}'}} \Gamma; \langle \text{id}, p, \text{start}(\theta, e, \text{succ}, \text{id}') : h, \theta', e' \{ \kappa \mapsto \text{id}' \} \rangle \mid \Pi; \{ \text{id}' \} \cup \Omega}$$

$$(\overline{\text{StartS}}) \quad \begin{array}{l} \Gamma; \langle \text{id}, p, \text{start}(\theta, e, \text{succ}, \text{id}') : h, \theta', e' \rangle \mid \Pi; \Omega \cup \{ \text{id}' \} \\ \leftarrow_{p, \text{start}(\text{id}'), \{s, st_{\text{id}'}} \Gamma; \langle \text{id}, p, h, \theta, e \rangle \mid \Pi; \Omega \\ \text{if } \text{spawns}(\text{id}', \Pi) = [] \wedge \text{reads}(\text{id}', \Pi) = [] \wedge \text{failed\_starts}(\text{id}', \Pi) = [] \end{array}$$

## Extended forward and backward semantics

$$\text{(StartS)} \quad \frac{\theta, e \xrightarrow{\text{start}(\kappa, \text{id}')} \theta', e' \quad \text{id}' \notin \Omega}{\Gamma; \langle \text{id}, p, h, \theta, e \rangle \mid \Pi; \Omega \xrightarrow{p, \text{start}(\text{id}'), \{s, \text{st}_{\text{id}'}\}} \Gamma; \langle \text{id}, p, \text{start}(\theta, e, \text{succ}, \text{id}') : h, \theta', e' \{ \kappa \mapsto \text{id}' \} \rangle \mid \Pi; \{ \text{id}' \} \cup \Omega}$$

$$\overline{\text{(StartS)}} \quad \frac{\Gamma; \langle \text{id}, p, \text{start}(\theta, e, \text{succ}, \text{id}') : h, \theta', e' \rangle \mid \Pi; \Omega \cup \{ \text{id}' \}}{\begin{array}{l} \leftarrow_{p, \text{start}(\text{id}'), \{s, \text{st}_{\text{id}'}\}} \Gamma; \langle \text{id}, p, h, \theta, e \rangle \mid \Pi; \Omega \\ \text{if } \text{spawns}(\text{id}', \Pi) = [] \wedge \text{reads}(\text{id}', \Pi) = [] \wedge \text{failed\_starts}(\text{id}', \Pi) = [] \end{array}}$$

# Extended forward and backward semantics

$$\begin{array}{c}
 \text{(StartS)} \\
 \hline
 \theta, e \xrightarrow{\text{start}(\kappa, \text{id}')} \theta', e' \quad \text{id}' \notin \Omega \\
 \Gamma; \langle \text{id}, p, h, \theta, e \rangle \mid \Pi; \Omega \rightarrow_{p, \text{start}(\text{id}'), \{s, \text{st}_{\text{id}'}\}} \\
 \Gamma; \langle \text{id}, p, \text{start}(\theta, e, \text{succ}, \text{id}') : h, \theta', e' \{ \kappa \mapsto \text{id}' \} \rangle \mid \Pi; \{ \text{id}' \} \cup \Omega
 \end{array}$$

$$\begin{array}{c}
 \overline{\text{(StartS)}} \\
 \Gamma; \langle \text{id}, p, \text{start}(\theta, e, \text{succ}, \text{id}') : h, \theta', e' \rangle \mid \Pi; \Omega \cup \{ \text{id}' \} \\
 \leftarrow_{p, \text{start}(\text{id}'), \{s, \text{st}_{\text{id}'}\}} \Gamma; \langle \text{id}, p, h, \theta, e \rangle \mid \Pi; \Omega \\
 \text{if } \boxed{\text{spawns}(\text{id}', \Pi) = [] \wedge \text{reads}(\text{id}', \Pi) = [] \wedge \text{failed\_starts}(\text{id}', \Pi) = []}
 \end{array}$$

## Extended forward and backward semantics

$$(StartS) \frac{\theta, e \xrightarrow{\text{start}(\kappa, nid')} \theta', e' \quad nid' \notin \Omega}{\Gamma; \langle nid, p, h, \theta, e \rangle \mid \Pi; \Omega \xrightarrow{p, \text{start}(nid'), \{s, st_{nid'}\}} \Gamma; \langle nid, p, \text{start}(\theta, e, \text{succ}, nid') : h, \theta', e' \{ \kappa \mapsto nid' \} \rangle \mid \Pi; \{nid'\} \cup \Omega}$$

$$\overline{(StartS)} \quad \Gamma; \langle nid, p, \text{start}(\theta, e, \text{succ}, nid') : h, \theta', e' \rangle \mid \Pi; \Omega \cup \{nid'\}$$

if  $\neg_{p, \text{start}(nid'), \{s, st_{nid'}\}} \Gamma; \langle nid, p, h, \theta, e \rangle \mid \Pi; \Omega$

if  $\text{spawns}(nid', \Pi) = [] \wedge \text{reads}(nid', \Pi) = [] \wedge \text{failed\_starts}(nid', \Pi) = []$

# Rollback Semantics

The rollback semantics allows us to reach a past state of the computation of the system, such past state is specified as an action performed by a process.

Some of the considered requests are:

- $\{p, \lambda^{\downarrow}\}$ : the receive of the message uniquely identified by  $\lambda$ ;
- $\{p, st_{nid}\}$ : the successful start of node  $nid$ ;
- $\{p, sp_{p'}\}$ : the spawn of process  $p'$ .

A system in rollback mode is denoted as  $\llbracket \mathcal{S} \rrbracket_{\Psi}$

# Rollback Semantics

$$\frac{\mathcal{S} \leftarrow_{p,r,\Psi'} \mathcal{S}' \wedge \psi \in \Psi'}{\llbracket \mathcal{S} \rrbracket_{\{p,\psi\}:\Psi} \rightsquigarrow \llbracket \mathcal{S}' \rrbracket_{\Psi}}$$

$$\frac{\mathcal{S} \leftarrow_{p,r,\Psi'} \mathcal{S}' \wedge \psi \notin \Psi'}{\llbracket \mathcal{S} \rrbracket_{\{p,\psi\}:\Psi} \rightsquigarrow \llbracket \mathcal{S}' \rrbracket_{\{p,\psi\}:\Psi}}$$

$$\frac{\mathcal{S} = \Gamma; \langle nid, p, h, \theta, e \rangle \mid \Pi; \Omega \wedge \mathcal{S} \not\leftarrow_{p,r,\Psi'} \wedge \{p', \psi'\} = \text{bwd\_dep}(\langle nid, p, h, \theta, e \rangle, \mathcal{S})}{\llbracket \mathcal{S} \rrbracket_{\{p,\psi\}:\Psi} \rightsquigarrow \llbracket \mathcal{S}' \rrbracket_{\{p',\psi'\}:\{p,\psi\}:\Psi}}$$

## Rollback semantics: dependencies operator

The dependencies operator does pattern matching on the history item and given the system computes the request to undo the consequences.

$$\text{bwd\_dep}(\langle -, -, \text{nodes}(-, -, \Omega') : h, -, - \rangle, -, \Pi; \{nid'\} \cup \Omega) = \{\text{parent}(nid', \Pi), st_{nid'}\}$$

where  $nid' \notin \Omega'$

# Table of Contents

- 1 Introduction
- 2 Background
- 3 Distributed CauDER
- 4 Future work**



## Future work

Still many features of Erlang that are to be covered (monitors, links, errors, etc.) to have a debugger that can be used on real programs.

Another line of research worth being investigated is the study of automatic ways to derive a reversible semantics starting from a non-reversible one.

The end

*Thank you for the attention!*