Reversibility in Erlang: Imperative Constructs

Pietro Lami[1] Ivan Lanese[2] Jean-Bernard Stefani[1]
Claudio Sacerdoti Coen[3] Giovanni Fabbretti[1]

[1]Univ. Grenoble Alpes, INRIA, CNRS, Grenoble INP, LIG, 38000 Grenoble, France

[2]Focus Team, Univ. of Bologna, INRIA, 40126 Bologna, Italy

[3]Univ. of Bologna, 40126 Bologna, Italy

DCore, 27/05/2022
Accepted to RC22

Table of Contents

1 Introduction

2 Extending reversible semantics in CauDEr

3 Case study

4 Future work

Reversibility

In a sequential system reversibility can be obtained by recursively undo the last action. This definition is not suitable in concurrent systems, since the last action may not be well-defined.

Causal-consistent reversibility

Any action can be undone provided that all its effects (if any) have been undone before

The idea is that each process saves all the information to restore past states.

Causal dependencies

There is a dependency between two actions when:

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

Example (Dependency)

Receiving a message depends on its sent.

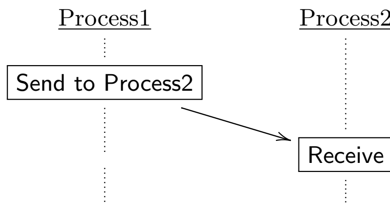

CauDEr

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

Erlang is a functional, concurrent and distributed programming language based on the actor paradigm.

CauDEr implements:

- a forward semantics
- a backward semantics
- a rollback semantics

Our Contribution

We extend CauDEr and its underlying theory by adding the support for some imperative primitives that have not considered before.

From the technical point of view, we show that the dependencies introduced by shared memory are not trivial.

Map

In our extension, atoms and pids are central:

- an atom is a literal constant;
- Pid is an abbreviation for process identifier

In Erlang, a pid can be associated to an atom.

Both atoms and pids in the map must be **unique**.

Imperative primitives

We extend CauDEr with support for the following built-in Erlang functions (BIFs):
given an atom a and a pid p

- `register`(a, p): inserts the pair $\langle a, p\rangle$ in the map;
- `unregister`(a): removes the (unique) pair $\langle a, p\rangle$ from the map;
- `whereis`(a): returns the associated pid p;
- `registered`$()$: returns the list of all the atoms in the map.

Table of Contents

1 Introduction
2 Extending reversible semantics in CauDEr
3 Case study
4 Future work

System definition

We define a reduction semantics with two types of reductions:

FORWARD
$\Gamma; \Pi; M \rightharpoonup \Gamma'; \Pi'; M'$
stores the history information

BACKWARD
$\Gamma; \Pi; M \leftharpoondown \Gamma'; \Pi'; M'$
exploits the history information

Definition (System)

A system is a tuple $\Gamma; \Pi; M$.

- Γ is the global mailbox,
- Π is the pool of processes,
- M is the map.

Process definition

Definition (Process)

A process is a tuple $\langle p, h, \theta, e, S \rangle \in \Pi$, where:

- p is the process pid,
- h is the process history,
- θ is the process environment,
- e is the expression under evaluation,
- S is a stack of process environments.

Extended Map

M is a quadruple $\langle a, p, t, s\rangle$ where a and p are the atom-pid registered, t is a unique identifier for the tuple and s can be either $\top$ or $\bot$.

Unique identifiers t are used to distinguish identical tuples existing at different times.

Tuples whose last field is $\top$ match the ones in the standard semantics, we call them *alive* tuples.

Those with $\bot$ are *ghost* tuples, namely alive tuples that have been removed from the map in a past forward action

Causal dependencies

The causal dependencies of the imperative primitives are:

1. between write actions that involve the same atom or pid;
2. between a write action and a read action that involve the same atom or pid;

Example

1. register(a, p) and unregister(a) are causally dependent
2. register(a, p) and whereis(a) are causally dependent

Extended forward and backward semantics

REG

$$\frac{\theta, e, S \xrightarrow{\text{register}(\kappa,a,p')} \theta', e', S' \qquad t \text{ fresh} \qquad \mathsf{M}_a = \emptyset \qquad \mathsf{M}_{p'} = \emptyset \qquad \text{isAlive}(p', \Pi)}{\Gamma; \langle p, h, \theta, e, S\rangle \mid \Pi; \mathsf{M} \rightharpoonup \Gamma; \langle p, \text{regS}(\theta, e, S, \{\langle a, p', t, \top\rangle\}){:}h, \theta', e'\{\kappa \to \text{true}\}, S'\rangle \mid \Pi; \mathsf{M} \cup \{\langle a, p', t, \top\rangle\}}$$

$\overline{\text{REG}}$

$$\frac{\text{readop}(t, \Pi) = \emptyset}{\Gamma; \langle p, \text{regS}(\theta, e, S, \{\langle a, p', t, \top\rangle\}){:}h, \theta', e', S'\rangle \mid \Pi; \mathsf{M} \cup \{\langle a, p', t, \top\rangle\} \leftharpoondown \Gamma; \langle p, h, \theta, e, S\rangle \mid \Pi; \mathsf{M}}$$

Extended forward and backward semantics

REG

$$\frac{\theta, e, S \xrightarrow{\mathsf{register}(\kappa, a, p')} \theta', e', S' \quad t \text{ fresh} \quad \mathsf{M}_a = \emptyset \quad \mathsf{M}_{p'} = \emptyset \quad \mathsf{isAlive}(p', \Pi)}{\Gamma; \langle p, h, \theta, e, S\rangle \mid \Pi; \mathsf{M} \rightharpoonup \Gamma; \langle p, \mathsf{regS}(\theta, e, S, \{\langle a, p', t, \top\rangle\}){:}h, \theta', e'\{\kappa \to \mathsf{true}\}, S'\rangle \mid \Pi; \mathsf{M} \cup \{\langle a, p', t, \top\rangle\}}$$

$\overline{\text{REG}}$

$$\frac{\mathsf{readop}(t, \Pi) = \emptyset}{\Gamma; \langle p, \mathsf{regS}(\theta, e, S, \{\langle a, p', t, \top\rangle\}){:}h, \theta', e', S'\rangle \mid \Pi; \boxed{\mathsf{M} \cup \{\langle a, p', t, \top\rangle\}} \leftharpoondown \Gamma; \langle p, h, \theta, e, S\rangle \mid \Pi; \mathsf{M}}$$

Table of Contents

1 Introduction

2 Extending reversible semantics in CauDEr

3 Case study

4 Future work

Case study

We consider a simple server dispatching requests to various mathematical services, and logging the results of the evaluation on a logger.

Each service is a process, and they are created only when there is a first request for them.

The logger keeps track of the values it receives, and answers each request with the sequential number of the element in the log

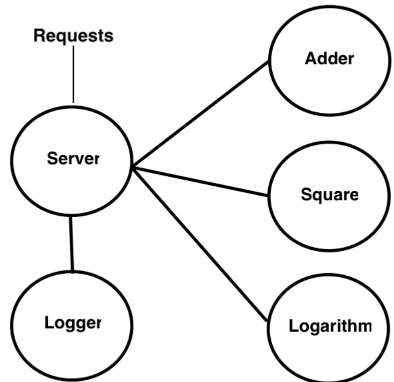

Expected result:

```
SEND REQUEST:{square,10}
SEND REQUEST:{adder,20}
SEND REQUEST:{log,100}
SEND REQUEST:{adder,30}
LOGGED {square,100} TIME:0
SEND REQUEST:{adder,100}
LOGGED {adder,20} TIME:1
LOGGED {adder,50} TIME:2
LOGGED {log,2.0} TIME:3
LOGGED {adder,150} TIME:4
```

Result:

```
SEND REQUEST:{square,10}
SEND REQUEST:{adder,20}
SEND REQUEST:{log,100}
LOGGED 100 TIME:0 <---
LOGGED {square,100} TIME:1
SEND REQUEST:{adder,30}
SEND REQUEST:{adder,100}
LOGGED {adder,20} TIME:2
LOGGED {adder,50} TIME:3
LOGGED {adder,150} TIME:4
```

We see that the there is an error in the behavior of the program.

In the server history we can see the send of the logarithm request but we do not have the answer from the logger as in the other services (e.g. adder).

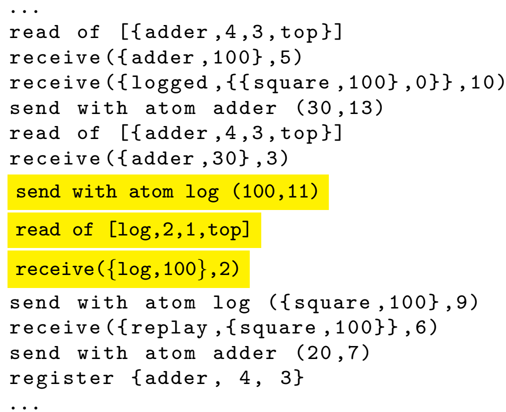
```
...
read of [{adder,4,3,top}]
receive({adder,100},5)
receive({logged,{{square,100},0}},10)
send with atom adder (30,13)
read of [{adder,4,3,top}]
receive({adder,30},3)
send with atom log (100,11)
read of [log,2,1,top]
receive({log,100},2)
send with atom log ({square,100},9)
receive({replay,{square,100}},6)
send with atom adder (20,7)
register {adder, 4, 3}
...
```

So we can do a rollback of message (11)

After the rollback we arrive at:

```
                ...
case whereis(Atom) of
        undefined ->
                register(Atom, spawn(?MODULE, Atom, [])),
                send(Atom,Val);
        _ ->
                send(Atom,Val)
end
```

We are in the branch where the atom is already registered. But this is the first time that we ask for the logarithm service.

So we can rollback the registration of atom log and we arrive at:

```
register(server, spawn(?MODULE, server, [])),
register(log, spawn(?MODULE, logger, [0, []])),
```

We understand that we use the same atom to identify two different processes.
Then the behavior is incorrect because we use the atom log for the logarithm service but sending messages to the logger.

$$\text{log (logarithm)} \neq \text{log (logger)}$$

Screenshot

Code

```
14     {replay, Ris} ->
15       io:format("RESULT:~p\n", [Ris]),
16       log ! Ris;
17     {Atom, Val} ->
18       io:format("SEND REQUEST:~p\n", [{Atom, Val}]),
19       case whereis(Atom) of
20         undefined ->
21           register(Atom,spawn(?MODULE, Atom,[])),
22           Atom ! Val;
23         _ ->
24           Atom ! Val
```

receive

Process Info

Bindings

Name	Value

Stack

```
server:server/0
server:server/0
server:server/0
server:server/0
server:server/0
server:server/0
server:server/0
server:server/0
```

Log

History

```
receive({logged,1},12)
send with atom log ({adder,20},13)
receive({replay,{adder,20}},8)
send with atom log (100,11)
read of [{log,2,1}]
receive({log,100},2)
receive({logged,0},10)
send with atom log ({square,100},9)
receive({replay,{square,100}},6)
send with atom adder (20,7)
```

Map Info

Node Map

Atom	PID	Tuple ID
server	1	0
log	2	1
square	3	2
adder	4	3

Actions

Process

Node: nonode@nohost PID: 1 - server:server/0

Manual | Automatic | Replay | Rollback

Steps: 1 — Roll steps
Uid: — Roll send
Atom-Uid {log,11} — Roll send with
Uid: 0 — Roll receive
Node: — Roll start
Pid: 1 — Roll spawn
Name: — Roll variable
Element: {adder,4,3} — Roll register
Element: — Roll delete

System Info

Nodes | Mailbox

[nonode@nohost]

Trace | Roll Log

```
receive(22)
receive(20)
receive(18)
receive(16)
receive(5)
receive(14)
receive(3)
receive(12)
```

Performed 284 (of 1000) forward steps in 24 ms — Ln 11, Col 1 — Alive 4, Dead 1

Table of Contents

1. Introduction
2. Extending reversible semantics in CauDEr
3. Case study
4. Future work

Future work

- in some cases sequences of transitions would commute, but their composing transitions do not;
- currently this creates dependencies since commuting is possible only between single transitions

Example

A `registered()` operation does not commute with `register(a,_)` followed by `unregister(a)`.

Thank you for the attention

Imperative primitives

We extend CauDEr with support for the following built-in Erlang functions (BIFs):

- `register`: given an atom a and a pid p, it inserts the pair $\langle a, p \rangle$ in the map and returns the atom **true**; if it is not possible an exception is raised;
- `unregister`: given an atom a, it removes the (unique) pair $\langle a, p \rangle$ from the map and returns **true** if the atom a is found, raises an exception otherwise;
- `whereis`: given an atom, it returns the associated pid if it exists, the atom **undefined** otherwise;
- `registered`: returns a list (possibly empty) of all the atoms in the map.

Example (Register followed by delete)

Consider a process do a `registered` operation and another process performs a (successful) `register` followed by a delete operation of a same tuple.

Executing the first process and then second one or vice versa would lead to the same state.

We want to distinguish these two computations, since undoing the `registered` should not always be possible.

Ghost tuples are our solution to this problem.

Server

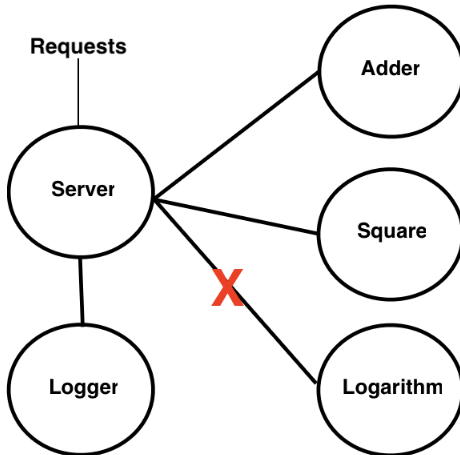

Server Expected

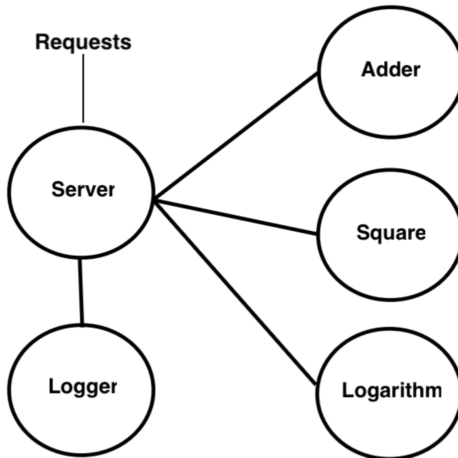

Rollback semantics

$$(U - Satisfy)\frac{\mathcal{S} \leftharpoondown_{p,r,\Psi'} \mathcal{S}' \ \wedge \ \psi \in \Psi'}{[\![\mathcal{S}]\!]_{\{p,\psi\}:\Psi} \rightsquigarrow [\![\mathcal{S}']\!]_{\Psi}}$$

$$(U - Act)\ \frac{\mathcal{S} \leftharpoondown_{p,r,\Psi'} \mathcal{S}' \ \wedge \ \{p,r\} \notin \Psi'}{[\![\mathcal{S}]\!]_{\{p,\psi\}:\Psi} \rightsquigarrow [\![\mathcal{S}']\!]_{\{p,\psi\}:\Psi}}$$

$$(Request)\ \frac{\begin{array}{c}\mathcal{S} = \Gamma; \langle p, h, \theta, e, S\rangle \mid \Pi; \mathrm{M} \ \wedge \ \mathcal{S} \not\leftharpoondown_{p,r,\Psi'} \ \ \wedge \\ \{p',\psi'\} = dep(\langle p, h, \theta, e, S\rangle, \mathcal{S})\end{array}}{[\![\mathcal{S}]\!]_{\{p,\psi\}:\Psi} \rightsquigarrow [\![\mathcal{S}']\!]_{\{p',\psi'\}:\{p,\psi\}:\Psi}}$$

Dependencies operator

$$
\begin{array}{ll}
\mathsf{dep}(<_, \mathsf{send}(_,_,_,\{v,\lambda\}){:}h,_,_,_>, \Gamma \cup \{(p,p',\{v,\lambda\})\};_;_) & = \{p', \lambda^{\Downarrow}\} \\
\mathsf{dep}(<p, \mathsf{sendS}(_,_,_,\{v,\lambda\},_){:}h,_,_,_>, \Gamma \cup \{(p,p',\{v,\lambda\})\};_;_) & = \{p', \lambda^{\Downarrow}\} \\
\mathsf{dep}(<_, \mathsf{sendS}(_,_,_,_,\{\langle a,p,t,\top\rangle\}){:}h,_,_,_>,_;_;\mathsf{M}') & = \{p', \mathsf{del}(t)\} \\
\mathsf{dep}(<_, \mathsf{spawn}(_,_,_,p'){:}h,_,_,_>,_;\Pi;_;_) & = \{p', s\} \\
\mathsf{dep}(<_, \mathsf{readS}(_,_,_,\mathsf{M} \cup \{\langle a,_,t,\top\rangle\}){:}h,_,_,_>,_;_;\mathsf{M}') & = \{p', \mathsf{del}(t)\} \\
\mathsf{dep}(<_, \mathsf{readS}(_,_,_,\mathsf{M} \cup \{\langle _,p,t,\top\rangle\}){:}h,_,_,_>,_;_;\mathsf{M}') & = \{p', \mathsf{del}(t)\} \\
\mathsf{dep}(<_, \mathsf{readF}(_,_,_,a,\mathsf{M}){:}h,_,_,_>,_;_;\mathsf{M}' \cup \{\langle a,_,t,_\rangle\}) & = \{p', \mathsf{regS}(t)\} \\
\mathsf{dep}(<_, \mathsf{readM}(_,_,_,\mathsf{M}){:}h,_,_,_>,_;_;\mathsf{M}' \cup \{\langle _,_,t,\top\rangle\}) & = \{p', \mathsf{regS}(t)\} \\
\mathsf{dep}(<_, \mathsf{readM}(_,_,_,\mathsf{M}){:}h,_,_,_>,_;_;\mathsf{M}' \cup \{\langle _,_,t,\bot\rangle\}) & = \{p', \mathsf{del}(t)\} \\
\mathsf{dep}(<_, \mathsf{regS}(_,_,_,\{\langle _,_,t,\top\rangle\}){:}h,_,_,_>,_;_;\mathsf{M}') & = \{p', \mathsf{del}(t)\} \\
\mathsf{dep}(<_, \mathsf{regS}(_,_,_,\{\langle _,_,t,\top\rangle\}){:}h,_,_,_>,_;_;_) & = \{p', \mathsf{read}(t)\} \\
\mathsf{dep}(<_, \mathsf{del}(_,_,_,\{\langle a,_,t,\top\rangle\},\mathsf{M}){:}h,_,_,_>,_;_;\mathsf{M}' \cup \{\langle a,_,t_a,_\rangle\}) & = \{p', \mathsf{regS}(t_a)\} \\
\mathsf{dep}(<_, \mathsf{del}(_,_,_,\{\langle _,p,t,\top\rangle\},\mathsf{M}){:}h,_,_,_>,_;_;\mathsf{M}' \cup \{\langle _,p,t_p,_\rangle\}) & = \{p', \mathsf{regS}(t_p)\} \\
\mathsf{dep}(<_, \mathsf{del}(_,_,_,\{\langle a,p,t,\top\rangle\},_){:}h,_,_,_>,_;_;\mathsf{M}') & = \{p', \mathsf{readfail}(t)\}
\end{array}
$$