

Reversibility in Erlang: Imperative Constructs

Pietro Lami¹ Ivan Lanese² Jean-Bernard Stefani¹
Claudio Sacerdoti Coen³ Giovanni Fabbretti¹

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

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

³Univ. of Bologna, 40126 Bologna, Italy

RC22, 05/07/2022

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 undoing 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 necessary 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 send.



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 been 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.

Each Erlang node has a map associating atoms to pids

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

$$\Gamma; \Pi; M \rightarrow \Gamma'; \Pi'; M'$$

stores history information

BACKWARD

$$\Gamma; \Pi; M \leftarrow \Gamma'; \Pi'; M'$$

exploits the stored 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 \perp .

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 \perp are *ghost* tuples, namely alive tuples that have been removed from the map by 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

$$\begin{array}{c}
 \theta, e, S \xrightarrow{\text{register}(\kappa, a, p')} \theta', e', S' \\
 \text{\textit{t fresh}} \quad M_a = \emptyset \quad M_{p'} = \emptyset \quad \text{isAlive}(p', \Pi)
 \end{array}$$

$$\Gamma; \langle p, \text{regS}(\theta, e, S, \{a, p', t, T\}) : h, \theta', e' \{ \kappa \rightarrow \text{true} \}, S' \rangle \mid \Pi; M \rightarrow \\
 \Gamma; \langle p, h, \theta, e, S \rangle \mid \Pi; M \cup \{ \langle a, p', t, T \rangle \}$$

$\overline{\text{REG}}$

$$\text{readop}(t, \Pi) = \emptyset$$

$$\Gamma; \langle p, \text{regS}(\theta, e, S, \{a, p', t, T\}) : h, \theta', e', S' \rangle \mid \Pi; M \cup \{ \langle a, p', t, T \rangle \} \\
 \leftarrow \Gamma; \langle p, h, \theta, e, S \rangle \mid \Pi; M$$

Extended forward and backward semantics

REG

$$\frac{
 \begin{array}{l}
 \theta, e, S \xrightarrow{\text{register}(\kappa, a, p')} \theta', e', S' \\
 t \text{ fresh} \quad M_a = \emptyset \quad M_{p'} = \emptyset \quad \text{isAlive}(p', \Pi)
 \end{array}
 }{
 \Gamma; \langle p, \text{regS}(\theta, e, S, \{\langle a, p', t, T \rangle\}):h, \theta', e' \{ \kappa \rightarrow \text{true} \}, S' \rangle \mid \Pi; M \cup \{\langle a, p', t, T \rangle\}
 }$$

$\overline{\text{REG}}$

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

Table of Contents

- 1 Introduction
- 2 Extending reversible semantics in CauDER
- 3 Case study**
- 4 Future work

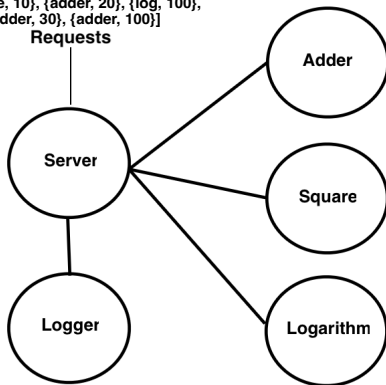
Case study

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

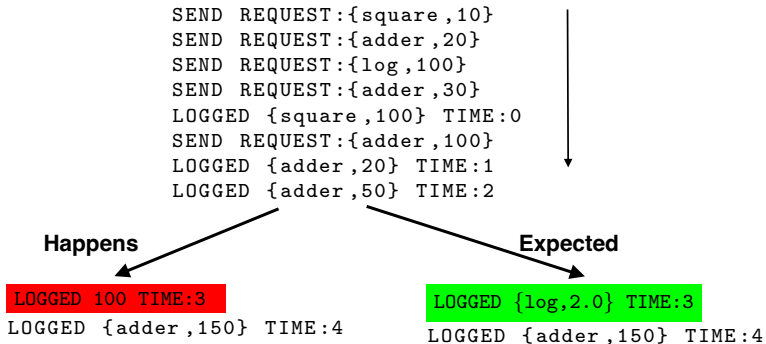
Each service is a process, and they are created only when there is a first request for them.

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

`[(square, 10), {adder, 20}, {log, 100},
{adder, 30}, {adder, 100}]`



Log of the server



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

History of the server in CauDER

In the server history we can see the send of the logarithm request but we do not have the answer from the service as in the other services (e.g. square). So we can do a rollback of message (11)

```
...  
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)  
receive({logged,{{square,100},0}},10)  
send with atom log ({square,100},9)  
receive({reply,{square,100}},6)  
send with atom square (10,4)  
receive({square,10},0)  
...
```

Bad behavior

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.

Detect the error

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 also to send messages to the logger

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

Screenshot

The screenshot displays the CauDER interface with several panels:

- Code:** Shows Erlang code for a process named 'receive'. It includes a `register` call and a `case` statement for handling different atoms.
- Process:** Shows the current process as `Node: nonode@nohost PID: 1 - server:server/0`.
- Actions:** A control panel with tabs for Manual, Automatic, Replay, and Rollback. It includes buttons for 'Roll steps', 'Roll send with', 'Roll receive', 'Roll start', 'Roll spawn', 'Roll variable', 'Roll register', and 'Roll delete'.
- System Info:** Shows 'Nodes' as `[nonode@nohost]` and 'Mailbox' as empty.
- Map Info:** A table showing the state of atoms and their tuples. This section is highlighted with a red border.

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

At the bottom of the interface, it indicates: `Performed 284 (of 1000) forward steps in 24 ms` and `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 P doing a registered operation and another process, Q, doing a (successful) register followed by a delete operation of the same tuple.

Executing P first and the Q 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.

Rollback semantics

$$(U - Satisfy) \frac{S \leftarrow_{p,r,\Psi'} S' \wedge \psi \in \Psi'}{\llbracket S \rrbracket_{\{p,\psi\}:\Psi} \rightsquigarrow \llbracket S' \rrbracket_{\Psi}}$$

$$(U - Act) \frac{S \leftarrow_{p,r,\Psi'} S' \wedge \{p,r\} \notin \Psi'}{\llbracket S \rrbracket_{\{p,\psi\}:\Psi} \rightsquigarrow \llbracket S' \rrbracket_{\{p,\psi\}:\Psi}}$$

$$(Request) \frac{S = \Gamma; \langle p, h, \theta, e, S \rangle \mid \Pi; M \wedge S \not\leftarrow_{p,r,\Psi'} \wedge \{p', \psi'\} = dep(\langle p, h, \theta, e, S \rangle, S)}{\llbracket S \rrbracket_{\{p,\psi\}:\Psi} \rightsquigarrow \llbracket S' \rrbracket_{\{p',\psi'\}:\{p,\psi\}:\Psi}}$$

Dependencies operator

$\text{dep}(\langle -, \text{send}(-, -, -, \{v, \lambda\}) : h, -, -, - \rangle, \Gamma \cup \{(p, p', \{v, \lambda\})\}; -, -)$	$= \{p', \lambda^\downarrow\}$
$\text{dep}(\langle p, \text{sendS}(-, -, -, \{v, \lambda\}, -) : h, -, -, - \rangle, \Gamma \cup \{(p, p', \{v, \lambda\})\}; -, -)$	$= \{p', \lambda^\downarrow\}$
$\text{dep}(\langle -, \text{sendS}(-, -, -, \{a, p, t, \top\}) : h, -, -, - \rangle, -, -; M')$	$= \{p', \text{del}(t)\}$
$\text{dep}(\langle -, \text{spawn}(-, -, -, p') : h, -, -, - \rangle, -, -; \Pi; -, -)$	$= \{p', s\}$
$\text{dep}(\langle -, \text{readS}(-, -, -, M \cup \{a, -, t, \top\}) : h, -, -, - \rangle, -, -; M')$	$= \{p', \text{del}(t)\}$
$\text{dep}(\langle -, \text{readS}(-, -, -, M \cup \{\langle -, p, t, \top \rangle\}) : h, -, -, - \rangle, -, -; M')$	$= \{p', \text{del}(t)\}$
$\text{dep}(\langle -, \text{readF}(-, -, -, a, M) : h, -, -, - \rangle, -, -; M' \cup \{a, -, t, -\})$	$= \{p', \text{regS}(t)\}$
$\text{dep}(\langle -, \text{readM}(-, -, -, M) : h, -, -, - \rangle, -, -; M' \cup \{\langle -, -, t, \top \rangle\})$	$= \{p', \text{regS}(t)\}$
$\text{dep}(\langle -, \text{readM}(-, -, -, M) : h, -, -, - \rangle, -, -; M' \cup \{\langle -, -, t, \perp \rangle\})$	$= \{p', \text{del}(t)\}$
$\text{dep}(\langle -, \text{regS}(-, -, -, \{\langle -, -, t, \top \rangle\}) : h, -, -, - \rangle, -, -; M' \cup \{\langle -, -, t, \perp \rangle\})$	$= \{p', \text{del}(t)\}$
$\text{dep}(\langle -, \text{regS}(-, -, -, \{\langle -, -, t, \top \rangle\}) : h, -, -, - \rangle, -, -; -)$	$= \{p', \text{read}(t)\}$
$\text{dep}(\langle -, \text{del}(-, -, -, \{a, -, t, \top\}, M) : h, -, -, - \rangle, -, -; M' \cup \{a, -, t_a, -\})$	$= \{p', \text{regS}(t_a)\}$
$\text{dep}(\langle -, \text{del}(-, -, -, \{\langle -, p, t, \top \rangle\}, M) : h, -, -, - \rangle, -, -; M' \cup \{\langle -, p, t_p, - \rangle\})$	$= \{p', \text{regS}(t_p)\}$
$\text{dep}(\langle -, \text{del}(-, -, -, \{a, p, t, \top\}, -) : h, -, -, - \rangle, -, -; M')$	$= \{p', \text{readfail}(t)\}$