

Causal-Consistent Replay Debugging for Message Passing Programs

Ivan Lanese, Adrián Palacios & Germán Vidal

Università Bologna & [Universitat Politècnica de València](#)

(paper presented at FORTE 2019)

Roadmap

Goal: debugging technique for concurrent programs

- 1 A simple (eager) functional language with message-passing concurrency (subset of Erlang)
- 2 Logging semantics: records the order in which messages are delivered to each process
- 3 Reversible semantics: allows us to explore back and forth the recorded execution in a causal-consistent way (i.e., an action cannot be undone until all the actions that depend on it have already been undone)
- 4 Controlled (replay/rollback) semantics: where the user can specify the actions to replay/undo → CauDEr

Roadmap

The Language

Syntax (sequential)
Syntax (concurrent)
Core Erlang
Semantics

Logging Semantics

Reversible Semantics

Uncontrolled
Controlled

Reversible Debugging

Conclusions

Goal: debugging technique for concurrent programs

- 1 A simple (eager) functional language with message-passing concurrency (subset of Erlang)
- 2 Logging semantics: records the order in which messages are delivered to each process
- 3 Reversible semantics: allows us to explore back and forth the recorded execution in a causal-consistent way (i.e., an action cannot be undone until all the actions that depend on it have already been undone)
- 4 Controlled (replay/rollback) semantics: where the user can specify the actions to replay/undo → CauDEr

Roadmap

The Language

Syntax (sequential)
Syntax (concurrent)
Core Erlang
Semantics

Logging Semantics

Reversible Semantics

Uncontrolled
Controlled

Reversible Debugging

Conclusions

Goal: debugging technique for concurrent programs

- 1 A simple (eager) functional language with message-passing concurrency (subset of Erlang)
- 2 **Logging semantics:** records the order in which messages are delivered to each process
- 3 **Reversible semantics:** allows us to explore back and forth the recorded execution in a causal-consistent way (i.e., an action cannot be undone until all the actions that depend on it have already been undone)
- 4 **Controlled (replay/rollback) semantics:** where the user can specify the actions to replay/undo → CauDEr

Roadmap

The Language

Syntax (sequential)
Syntax (concurrent)
Core Erlang
Semantics

Logging Semantics

Reversible Semantics

Uncontrolled
Controlled

Reversible Debugging

Conclusions

Goal: debugging technique for concurrent programs

- 1 A simple (eager) functional language with message-passing concurrency (subset of Erlang)
- 2 **Logging semantics:** records the order in which messages are delivered to each process
- 3 **Reversible semantics:** allows us to explore back and forth the recorded execution in a **causal-consistent** way (i.e., an action cannot be undone until all the actions that depend on it have already been undone)
- 4 **Controlled (replay/rollback) semantics:** where the user can specify the actions to replay/undo → **CauDEr**

Roadmap

The Language

Syntax (sequential)
Syntax (concurrent)
Core Erlang
Semantics

Logging Semantics

Reversible Semantics

Uncontrolled
Controlled

Reversible Debugging

Conclusions

Goal: debugging technique for concurrent programs

- 1 A simple (eager) functional language with message-passing concurrency (subset of Erlang)
- 2 **Logging semantics:** records the order in which messages are delivered to each process
- 3 **Reversible semantics:** allows us to explore back and forth the recorded execution in a **causal-consistent** way (i.e., an action cannot be undone until all the actions that depend on it have already been undone)
- 4 **Controlled (replay/rollback) semantics:** where the user can specify the actions to replay/undo → **CauDEr**

The language

We consider a simple functional and concurrent programming language similar to Erlang

- No shared memory, only **message passing** (asynchronous communication)
- Each process has a **local queue** (mailbox)
- A **system** is a collection of processes

Sequential Erlang in 5 examples

append/2

```
append([H|T], L) -> [H|append(T, L)];  
append([], L) -> L.
```

Variables start with an uppercase letter

Function names and atoms (i.e., constants) start with a lowercase letter

Alternative definition:

append/2

```
append(A, B) -> case A of  
    [H|T] -> [H|append(T, B)];  
    [] -> B  
end.
```


Sequential Erlang in 5 examples

append/2

```
append([H|T], L) -> [H|append(T, L)];  
append([], L) -> L.
```

Variables start with an uppercase letter

Function names and atoms (i.e., constants) start with a lowercase letter

Alternative definition:

append/2

```
append(A, B) -> case A of  
    [H|T] -> [H|append(T, L)];  
    [] -> L  
end.
```

Sequential Erlang in 5 examples

Roadmap

The
Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging
Semantics

Reversible
Semantics

Uncontrolled

Controlled

Reversible
Debugging

Conclusions

toint/1

```
toint({s,N}) -> int(N) + 1;  
toint(zero) -> 0.
```

E.g., `toint({s,{s,{s,zero}}})` evaluates to 3

No user-defined algebraic data types (so we cannot write `s(s(s(zero)))`)

Main data types: numbers, atoms, lists, and tuples

Sequential Erlang in 5 examples

Roadmap

The
Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging
Semantics

Reversible
Semantics

Uncontrolled

Controlled

Reversible
Debugging

Conclusions

factorial/1

```
factorial(N) when N > 0 -> N * factorial(N - 1);  
factorial(1)             -> 0.
```

Besides pattern matching, we can have **guards**

Only built-in functions are allowed in guards

Sequential Erlang in 5 examples

minmax/1

```
minmax(L) -> Min = lists:min(L),  
             Max = lists:max(L),  
             {Min,Max}.
```

Sequence e_1, \dots, e_n evaluates all expressions, returns the evaluation of e_n

Equation $pat = exp$ evaluates exp and perform pattern matching with pat

Sequential Erlang in 5 examples

Roadmap

The
Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging
Semantics

Reversible
Semantics

Uncontrolled

Controlled

Reversible
Debugging

Conclusions

inclist/1

```
inclist(L) -> lists : map(fun(x) -> x + 1 end, L).
```

Higher-order functions

Anonymous functions

No partial applications

Concurrency features

- `spawn/1` and `spawn/3`: creates a new process as a side-effect and returns the pid of the new process
- `self/0`: returns the pid of the current process
- `pid ! val`: sends `val` to process `pid` as a side-effect and returns `val`
- `receive ... end`: waits for a message that matches some pattern (otherwise, blocks execution) and returns the expression in the selected branch

Concurrent Erlang in 1 example

```
main()    -> S = spawn(fun() -> server([]) end),  
          client(S).
```

```
client(S) -> S ! {self(), {add,paper}},  
            S ! {self(), {add,pencil}},  
            S ! {self(), take},  
            receive  
              X -> X  
            end.
```

```
server(L) -> receive  
             {_, {add, Item}} -> server([Item|L]);  
             {C, take} -> C ! hd(L), server(tl(L))  
           end.
```

From Erlang to Core Erlang

Roadmap

The Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging Semantics

Reversible Semantics

Uncontrolled

Controlled

Reversible Debugging

Conclusions

Core Erlang is an **intermediate representation** used during the compilation of Erlang programs

It is a convenient representation for defining analyses and other tools

Not as readable as Erlang...

From Erlang to Core Erlang

erlang

```
a(42)  ->  ok;  
a(N)   ->  M = N + 1, a(M).
```

core erlang

```
'a'/1 = fun(_@c0) ->  
  case _@c0 of  
    < 42 > when 'true' -> 'ok'  
    < _@c2 > when 'true' -> let < _@c3 > = call 'erlang':'+'(N, 1)  
                          in apply 'a'/1 (_@c3)  
  end
```

Essentially: one clause per function, case for pattern matching, let for sequences, apply for function applications, call for built-in calls, etc

From Erlang to Core Erlang

erlang

```
a(42)  ->  ok;  
a(N)   ->  M = N + 1, a(M).
```

core erlang

```
'a'/1 = fun(_@c0) ->  
  case _@c0 of  
    < 42 > when 'true' -> 'ok'  
    < _@c2 > when 'true' -> let < _@c3 > = call 'erlang':'+'(N, 1)  
                          in apply 'a'/1 (_@c3)  
  end
```

Essentially: one clause per function, case for pattern matching, let for sequences, apply for function applications, call for built-in calls, etc

Core Erlang syntax

We consider a subset of Core Erlang with this syntax:

```
Module ::= module Atom = fun1, ..., funn
  fun ::= fname = fun (X1, ..., Xn) → expr
  fname ::= Atom/Integer
  lit ::= Atom | Integer | Float | []
  expr ::= Var | lit | fname | [expr1|expr2] | {expr1, ..., exprn}
          | call expr (expr1, ..., exprn) | apply expr (expr1, ..., exprn)
          | case expr of clause1; ...; clausem end
          | let Var = expr1 in expr2 | receive clause1; ...; clausen end
          | spawn(expr, [expr1, ..., exprn]) | expr1 ! expr2 | self()
  clause ::= pat when expr1 → expr2
  pat ::= Var | lit | [pat1|pat2] | {pat1, ..., patn}
```

Core Erlang syntax

We consider a subset of Core Erlang with this syntax:

```
Module ::= module Atom = fun1, ..., funn
      fun ::= fname = fun (X1, ..., Xn) → expr
      fname ::= Atom/Integer
      lit ::= Atom | Integer | Float | []
      expr ::= Var | lit | fname | [expr1|expr2] | {expr1, ..., exprn}
            | call expr (expr1, ..., exprn) | apply expr (expr1, ..., exprn)
            | case expr of clause1; ... ; clausem end
            | let Var = expr1 in expr2 | receive clause1; ... ; clausen end
            | spawn(expr, [expr1, ..., exprn]) | expr1 ! expr2 | self()
      clause ::= pat when expr1 → expr2
      pat ::= Var | lit | [pat1|pat2] | {pat1, ..., patn}
```

Some preliminary definitions

Definition (process)

//no local queue!

A process is a triple $\langle p, \theta, e \rangle$ where

- p is the pid of the process
- θ is an environment
- e is the expression to be reduced

Definition (system)

A system is denoted by $\Gamma; \Pi$, where

- Γ models the network & local queues (global mailbox)
- Π is a pool of processes

Γ is a multiset of triples ($sender_pid, target_pid, message$)

We often use $\Gamma; \langle p, \theta, e \rangle$ & Π to denote an arbitrary system

Some preliminary definitions

Definition (process)

//no local queue!

A process is a triple $\langle p, \theta, e \rangle$ where

- p is the pid of the process
- θ is an environment
- e is the expression to be reduced

Definition (system)

A system is denoted by $\Gamma; \Pi$, where

- Γ models the network & local queues (**global mailbox**)
- Π is a pool of processes

Γ is a multiset of triples (*sender_pid, target_pid, message*)

We often use $\Gamma; \langle p, \theta, e \rangle$ & Π to denote an arbitrary system

Core Erlang Semantics

Two-level (reduction) semantics:

- Semantics of **expressions** (sequential & concurrent)
- Semantics of **systems**

For concurrent actions, we face the following problems:

- 1 we don't know the result of the actions (**fresh variables**)
- 2 we must perform side effects (**labels**)

Labels

- At expression level, transitions for concurrent actions are labelled with enough information
- At system level, labels are used to perform the associated actions

Core Erlang Semantics

Two-level (reduction) semantics:

- Semantics of **expressions** (sequential & concurrent)
- Semantics of **systems**

For concurrent actions, we face the following problems:

- 1 we don't know the result of the actions (**fresh variables**)
- 2 we must perform side effects (**labels**)

Labels

- At expression level, transitions for **concurrent actions are labelled** with enough information
- At system level, **labels are used** to perform the associated actions

Expression semantics: sequential expressions

Roadmap

The Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging Semantics

Reversible Semantics

Uncontrolled

Controlled

Reversible Debugging

Conclusions

$$\begin{array}{c}
 \text{(Var)} \frac{}{\theta, X \xrightarrow{\tau} \theta, \theta(X)} \quad \text{(Tuple)} \frac{\theta, e_i \xrightarrow{\ell} \theta', e'_i}{\theta, \{\overline{v_{1,i-1}}, e_i, \overline{e_{i+1,n}}\} \xrightarrow{\ell} \theta', \{\overline{v_{1,i-1}}, e'_i, \overline{e_{i+1,n}}\}} \\
 \text{(List1)} \frac{\theta, e_1 \xrightarrow{\ell} \theta', e'_1}{\theta, [e_1 | e_2] \xrightarrow{\ell} \theta', [e'_1 | e_2]} \quad \text{(List2)} \frac{\theta, e_2 \xrightarrow{\ell} \theta', e'_2}{\theta, [v_1 | e_2] \xrightarrow{\ell} \theta', [v_1 | e'_2]} \\
 \text{(Let1)} \frac{\theta, e_1 \xrightarrow{\ell} \theta', e'_1}{\theta, \text{let } X = e_1 \text{ in } e_2 \xrightarrow{\ell} \theta', \text{let } X = e'_1 \text{ in } e_2} \quad \text{(Let2)} \frac{}{\theta, \text{let } X = v \text{ in } e \xrightarrow{\tau} \theta[X \mapsto v], e} \\
 \text{(Case1)} \frac{\theta, e \xrightarrow{\ell} \theta', e'}{\theta, \text{case } e \text{ of } c_1; \dots; c_n \text{ end} \xrightarrow{\ell} \theta', \text{case } e' \text{ of } c_1; \dots; c_n \text{ end}} \quad \text{(Case2)} \frac{\text{match}(v, c_1, \dots, c_n) = \langle \theta_i, e_i \rangle}{\theta, \text{case } v \text{ of } c_1; \dots; c_n \text{ end} \xrightarrow{\tau} \theta \theta_i, e_i} \\
 \text{(Apply1)} \frac{\theta, e_i \xrightarrow{\ell} \theta', e'_i \quad i \in \{1, \dots, n\}}{\theta, \text{apply } a/n (\overline{v_{1,i-1}}, e_i, \overline{e_{i+1,n}}) \xrightarrow{\ell} \theta', \text{apply } a/n (\overline{v_{1,i-1}}, e'_i, \overline{e_{i+1,n}})} \\
 \text{(Apply2)} \frac{\mu(a/n) = \text{fun } (X_1, \dots, X_n) \rightarrow e}{\theta, \text{apply } a/n (v_1, \dots, v_n) \xrightarrow{\tau} \{X_1 \mapsto v_1, \dots, X_n \mapsto v_n\}, e}
 \end{array}$$

Sending a message

Roadmap

The Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging Semantics

Uncontrolled

Controlled

Reversible Semantics

Uncontrolled

Controlled

Reversible Debugging

Conclusions

(expression semantics)

$$(Send1) \frac{\theta, e_1 \xrightarrow{\ell} \theta', e'_1}{\theta, e_1 ! e_2 \xrightarrow{\ell} \theta', e'_1 ! e_2} \quad \frac{\theta, e_2 \xrightarrow{\ell} \theta', e'_2}{\theta, v_1 ! e_2 \xrightarrow{\ell} \theta', v_1 ! e'_2}$$

$$(Send2) \frac{}{\theta, v_1 ! v_2 \xrightarrow{\text{send}(v_1, v_2)} \theta, v_2}$$

(system semantics)

$$(Send) \frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e'}{\Gamma; \langle p, \theta, e \rangle \& \Pi \hookrightarrow \Gamma \cup \{(p, p', v)\}; \langle p, \theta', e' \rangle \& \Pi}$$

Sending a message

(expression semantics)

$$(\text{Send1}) \frac{\theta, e_1 \xrightarrow{\ell} \theta', e'_1}{\theta, e_1 ! e_2 \xrightarrow{\ell} \theta', e'_1 ! e_2} \quad \frac{\theta, e_2 \xrightarrow{\ell} \theta', e'_2}{\theta, v_1 ! e_2 \xrightarrow{\ell} \theta', v_1 ! e'_2}$$

$$(\text{Send2}) \frac{}{\theta, v_1 ! v_2 \xrightarrow{\text{send}(v_1, v_2)} \theta, v_2}$$

(system semantics)

$$(\text{Send}) \frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e'}{\Gamma; \langle p, \theta, e \rangle \& \Pi \hookrightarrow \Gamma \cup \{(p, p', v)\}; \langle p, \theta', e' \rangle \& \Pi}$$

Spawning a process

Roadmap

The Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging

Semantics

Reversible
Semantics

Uncontrolled

Controlled

Reversible
Debugging

Conclusions

(expression semantics)

$$(\text{Spawn}) \frac{}{\theta, \text{spawn}(a/n, [v_1, \dots, v_n]) \xrightarrow{\text{spawn}(\kappa, a/n, [\overline{v}_n])} \theta, \kappa}$$

(system semantics)

$$(\text{Spawn}) \frac{\theta, e \xrightarrow{\text{spawn}(\kappa, a/n, [\overline{v}_n])} \theta', e' \quad p' \text{ is a fresh pid}}{\Gamma; \langle p, \theta, e \rangle \& \Pi \hookrightarrow \Gamma; \langle p, \theta', e' \{ \kappa \mapsto p' \} \rangle \& \langle p', \theta', \text{apply } a/n (\overline{v}_n) \rangle \& \Pi}$$

Spawning a process

(expression semantics)

$$(\text{Spawn}) \frac{}{\theta, \text{spawn}(a/n, [v_1, \dots, v_n]) \xrightarrow{\text{spawn}(\kappa, a/n, [\overline{v}_n])} \theta, \kappa}$$

(system semantics)

$$(\text{Spawn}) \frac{\theta, e \xrightarrow{\text{spawn}(\kappa, a/n, [\overline{v}_n])} \theta', e' \quad p' \text{ is a fresh pid}}{\Gamma; \langle p, \theta, e \rangle \& \Pi \hookrightarrow \Gamma; \langle p, \theta', e' \{ \kappa \mapsto p' \} \rangle \& \langle p', \theta', \text{apply } a/n (\overline{v}_n) \rangle \& \Pi}$$

Receiving a message

Roadmap

The Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging

Semantics

Reversible

Semantics

Uncontrolled

Controlled

Reversible

Debugging

Conclusions

(expression semantics)

$$(Receive) \frac{}{\theta, \text{receive } cl_1; \dots; cl_n \text{ end} \xrightarrow{\text{rec}(\kappa, \overline{cl_n})} \theta, \kappa}$$

(system semantics)

$$(Receive) \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl_n})} \theta', e' \quad \text{matchrec}(\theta, \overline{cl_n}, v) = (\theta_i, e_i)}{\Gamma \cup \{(p', p, v)\}; \langle p, \theta, e \rangle \& \Pi \hookrightarrow \Gamma; \langle p, \theta' \theta_i, e' \{ \kappa \mapsto e_i \} \rangle \& \Pi}$$

Receiving a message

Roadmap

The Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging

Semantics

Reversible Semantics

Uncontrolled

Controlled

Reversible Debugging

Conclusions

(expression semantics)

$$(Receive) \frac{}{\theta, \text{receive } cl_1; \dots; cl_n \text{ end} \xrightarrow{\text{rec}(\kappa, \overline{cl_n})} \theta, \kappa}$$

(system semantics)

$$(Receive) \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl_n})} \theta', e' \quad \text{matchrec}(\theta, \overline{cl_n}, v) = (\theta_i, e_i)}{\Gamma \cup \{(p', p, v)\}; \langle p, \theta, e \rangle \& II \hookrightarrow \Gamma; \langle p, \theta' \theta_i, e' \{ \kappa \mapsto e_i \} \rangle \& II}$$

Logging semantics

In concurrent languages, replaying a particular computation might be difficult (even impossible) given the **nondeterminism** of the language

We tag messages with unique identifiers

$$v \mapsto \{v, \ell\}, \text{ where } \ell \text{ is fresh}$$

A log $\mathcal{L}(d)$ of a derivation d is a sequence of items $\text{spawn}(p)$, $\text{send}(\ell)$ or $\text{rec}(\ell)$ for each process in d

(logs are local to each process)

Logging semantics

In concurrent languages, replaying a particular computation might be difficult (even impossible) given the **nondeterminism** of the language

We tag messages with unique identifiers

$$v \mapsto \{v, \ell\}, \quad \text{where } \ell \text{ is fresh}$$

A **log** $\mathcal{L}(d)$ of a derivation d is a sequence of items **spawn**(p), **send**(ℓ) or **rec**(ℓ) for each process in d

(logs are local to each process)

(Seq)

$$\frac{\theta, e \xrightarrow{\tau} \theta', e'}{\Gamma; \langle p, \theta, e \rangle \mid \Pi \hookrightarrow_{p, \text{seq}} \Gamma; \langle p, \theta', e' \rangle \mid \Pi}$$

(Send)

$$\frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e' \text{ and } \ell \text{ is a fresh symbol}}{\Gamma; \langle p, \theta, e \rangle \mid \Pi \hookrightarrow_{p, \text{send}(\ell)} \Gamma \cup \{(p, p', \{v, \ell\})\}; \langle p, \theta', e' \rangle \mid \Pi}$$

(Receive)

$$\frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl}_n)} \theta', e' \text{ and } \text{matchrec}(\theta, \overline{cl}_n, v) = (\theta_i, e_i)}{\Gamma \cup \{(p', p, \{v, \ell\})\}; \langle p, \theta, e \rangle \mid \Pi \hookrightarrow_{p, \text{rec}(\ell)} \Gamma; \langle p, \theta' \theta_i, e' \{ \kappa \mapsto e_i \} \rangle \mid \Pi}$$

(Spawn)

$$\frac{\theta, e \xrightarrow{\text{spawn}(\kappa, a/n, [\overline{v}_n])} \theta', e' \text{ and } p' \text{ is a fresh pid}}{\Gamma; \langle p, \theta, e \rangle \mid \Pi \hookrightarrow_{p, \text{spawn}(p')} \Gamma; \langle p, \theta', e' \{ \kappa \mapsto p' \} \rangle \mid \langle p', id, \text{apply } a/n (\overline{v}_n) \rangle \mid \Pi}$$

(Self)

$$\frac{\theta, e \xrightarrow{\text{self}(\kappa)} \theta', e'}{\Gamma; \langle p, \theta, e \rangle \mid \Pi \hookrightarrow_{p, \text{self}} \Gamma; \langle p, \theta', e' \{ \kappa \mapsto p \} \rangle \mid \Pi}$$

(implemented by a program instrumentation)

Causally equivalent derivations

$t_1 = (s_1 \xrightarrow{p_1, r_1} s'_1)$ happened before $t_2 = (s_2 \xrightarrow{p_2, r_2} s'_2)$, in symbols $t_1 \rightsquigarrow t_2$, if one of the following conditions holds:

- $p_1 = p_2$ and t_1 comes before t_2 ;
- $r_1 = \text{spawn}(p)$ and $p_2 = p$;
- $r_1 = \text{send}(\ell)$ and $r_2 = \text{rec}(\ell)$.

t_1 and t_2 are **independent** if $t_1 \not\rightsquigarrow t_2$ and $t_2 \not\rightsquigarrow t_1$

d_1 and d_2 are **causally equivalent** ($d_1 \approx d_2$) if d_1 can be obtained from d_2 by switching consecutive independent transitions

Given (coinitial) derivations d_1 and d_2 , $\mathcal{L}(d_1) = \mathcal{L}(d_2)$ iff $d_1 \approx d_2$

Causally equivalent derivations

$t_1 = (s_1 \xrightarrow{p_1, r_1} s'_1)$ happened before $t_2 = (s_2 \xrightarrow{p_2, r_2} s'_2)$, in symbols $t_1 \rightsquigarrow t_2$, if one of the following conditions holds:

- $p_1 = p_2$ and t_1 comes before t_2 ;
- $r_1 = \text{spawn}(p)$ and $p_2 = p$;
- $r_1 = \text{send}(\ell)$ and $r_2 = \text{rec}(\ell)$.

t_1 and t_2 are **independent** if $t_1 \not\rightsquigarrow t_2$ and $t_2 \not\rightsquigarrow t_1$

d_1 and d_2 are **causally equivalent** ($d_1 \approx d_2$) if d_1 can be obtained from d_2 by switching consecutive independent transitions

Given (coinitial) derivations d_1 and d_2 , $\mathcal{L}(d_1) = \mathcal{L}(d_2)$ iff $d_1 \approx d_2$

Reversible Semantics

Processes have the form $\langle p, \omega, h, \theta, e \rangle$
with ω a *log* and h a *history*

A history h is a sequence of terms headed by constructors **seq**, **send**, **rec**, **spawn**, and **self**, and whose arguments are the information required to (deterministically) undo the step

Uncontrolled forward semantics

(Send)

$$\frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e'}{\Gamma; \langle p, \text{send}(\ell) : \omega, h, \theta, e \rangle \mid \Pi \rightarrow_{p, \text{send}(\ell), \{s, \ell \uparrow\}} \Gamma \cup \{(p, p', \{v, \ell\})\}; \langle p, \omega, \text{send}(\theta, e, p', \{v, \ell\}) : h, \theta', e' \rangle \mid \Pi}$$

(Receive)

$$\frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl}_n)} \theta', e' \text{ and } \text{matchrec}(\theta, \overline{cl}_n, v) = (\theta_i, e_i)}{\Gamma \cup \{(p', p, \{v, \ell\})\} \langle p, \text{rec}(\ell) : \omega, h, \theta, e \rangle \mid \Pi \rightarrow_{p, \text{rec}(\ell), \{s, \ell \downarrow\}} \Gamma; \langle p, \omega, \text{rec}(\theta, e, p', \{v, \ell\}) : h, \theta' \theta_i, e' \{ \kappa \mapsto e_i \} \rangle \mid \Pi}$$

(Spawn)

$$\frac{\theta, e \xrightarrow{\text{spawn}(\kappa, a/n, \overline{v}_n)} \theta', e' \text{ and } \omega' = \text{trace}(d, p')}{\Gamma; \langle p, \text{spawn}(p') : \omega, h, \theta, e \rangle \mid \Pi \rightarrow_{p, \text{spawn}(p'), \{s, \text{sp}_{p'}\}} \Gamma; \langle p, \omega, \text{spawn}(\theta, e, p') : h, \theta', e' \{ \kappa \mapsto p' \} \rangle \mid \langle p', \omega', [], id, \text{apply } a/n (\overline{v}_n) \rangle \mid \Pi}$$

Uncontrolled forward semantics

Roadmap

The Language

- Syntax (sequential)
- Syntax (concurrent)
- Core Erlang
- Semantics

Logging Semantics

Reversible Semantics

- Uncontrolled
- Controlled

Reversible Debugging

Conclusions

(Receive)

$$\frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{cl}_n)} \theta', e' \text{ and } \text{matchrec}(\theta, \overline{cl}_n, v) = (\theta_i, e_i)}{\Gamma \cup \{(p', p, \{v, \ell\})\} \langle p, \text{rec}(\ell) : \omega, h, \theta, e \rangle \mid \Pi \rightarrow_{p, \text{rec}(\ell), \{s, \ell \downarrow\}} \Gamma; \langle p, \omega, \text{rec}(\theta, e, p', \{v, \ell\}) : h, \theta' \theta_i, e' \{ \kappa \mapsto e_i \} \rangle \mid \Pi}$$

Uncontrolled backward semantics

Roadmap

The Language

- Syntax (sequential)
- Syntax (concurrent)
- Core Erlang
- Semantics

Logging Semantics

Reversible Semantics

- Uncontrolled
- Controlled

Reversible Debugging

Conclusions

$$\overline{(Receive)} \quad \frac{\Gamma; \langle p, \omega, \text{rec}(\theta, e, p', \{v, \ell\}) : h, \theta', e' \rangle \mid \Pi}{\vdash_{p, \text{rec}(\ell), \{s, \ell \downarrow\} \cup \mathcal{V}} \Gamma \cup \{(p', p, \{v, \ell\})\}; \langle p, \text{rec}(\ell) : \omega, h, \theta, e \rangle \mid \Pi}$$

where $\mathcal{V} = \text{Dom}(\theta') \setminus \text{Dom}(\theta)$

Some results...

Coinitial derivations are cofinal
iff they are causally equivalent

Misbehaviors are preserved
by all causally equivalent derivations

Controlled replay/rollback semantics

Roadmap

The Language

Syntax (sequential)
Syntax (concurrent)
Core Erlang
Semantics

Logging Semantics

Reversible Semantics

Uncontrolled
Controlled

Reversible Debugging

Conclusions

We allow the user to start a replay/rollback until a particular action is performed, e.g.,

- $\{p, s\}$: **one step** backward/forward of process p
- $\{p, \ell^\uparrow\}$: a backward/forward derivation of process p up to the **sending of the message** tagged with ℓ
- $\{p, \ell^\downarrow\}$: a backward/forward derivation of process p up to the **reception of the message** tagged with ℓ
- $\{p, sp_{p'}\}$: a backward/forward derivation of process p up to the **spawning of the process** with pid p'
- $\{p, X\}$: a backward derivation of process p up to the **introduction of variable X**
- ...

Controlled semantics takes a **stack of requests** (initially one)

It is defined as a **layer on top of the uncontrolled semantics**:

- If a process can perform a step satisfying the request on top of the stack → do it and remove the request
- If a process can perform a step but it doesn't satisfy the request → update the system but keep the request
- If a step on the process is not possible → track dependencies and add new requests on top of the stack

Controlled semantics takes a **stack of requests** (initially one)

It is defined as a **layer on top of the uncontrolled semantics**:

- If a process can perform a step satisfying the request on top of the stack → do it and remove the request
- If a process can perform a step but it doesn't satisfy the request → update the system but keep the request
- If a step on the process is not possible → track dependencies and add new requests on top of the stack

Controlled semantics takes a **stack of requests** (initially one)

It is defined as a **layer on top of the uncontrolled semantics**:

- If a process can perform a step satisfying the request on top of the stack → do it and remove the request
- If a process can perform a step but it doesn't satisfy the request → update the system but keep the request
- If a step on the process is not possible → track dependencies and add new requests on top of the stack

Controlled semantics takes a **stack of requests** (initially one)

It is defined as a **layer on top of the uncontrolled semantics**:

- If a process can perform a step satisfying the request on top of the stack → do it and remove the request
- If a process can perform a step but it doesn't satisfy the request → update the system but keep the request
- If a step on the process is not possible → track dependencies and add new requests on top of the stack

Reversible debugging

Two components: code instrumentation (logging)
+ causal-consistent reversible debugger (CauDEr)

<https://github.com/mistupv/tracer>

<https://github.com/mistupv/cauder/tree/replay>

Roadmap

The Language

Syntax (sequential)

Syntax (concurrent)

Core Erlang

Semantics

Logging Semantics



Reversible Semantics

Uncontrolled

Controlled

Reversible Debugging

Conclusions



A Note on the implementation

Current prototypes show **good potential**, but more implementation effort is still required:

- move from Core Erlang to Erlang (or add pretty printing)
- graphical representation of traces
- consider more Erlang features: links, monitors, built-in's, input/output, behaviours, etc

Conclusions & future work

Promising approach for (causal-consistent) reversible debugging of message passing concurrent programs

Most ideas are applicable to other concurrent languages

Some ideas for **future work**:

- deal with (partially) unknown modules, trusted components, etc
- combine it with program slicing / automatic bug location
- keep improving the implementation

Conclusions & future work

Promising approach for (causal-consistent) reversible debugging of message passing concurrent programs

Most ideas are applicable to other concurrent languages

Some ideas for **future work**:

- deal with (partially) unknown modules, trusted components, etc
- combine it with program slicing / automatic bug location
- keep improving the implementation

Thanks for your attention!

Questions?