

A general approach to derive uncontrolled reversible semantics

DCore meeting Paris, 2-3 March 2020

Ivan Lanese and Doriana Medić

February 27, 2020

- In this work we are interested in obtaining a general approach for transforming a given forward semantics into the reversible one.
- In particular, our aim is:
 - to obtain a general approach applicable on different models. We use as the case studies Core Erlang and Higher-Order π -calculus (the result holds for π -calculus and CCS, as well);
 - to show that the general framework enjoys the properties as Loop Lemma, Square Lemma and Causal Consistency. In that way in the newly obtained reversible semantics, mentioned properties shall hold by construction;
 - to extend the reversible semantics of Core Erlang to support constructs for error handling mechanism based on links

The forward model needs to satisfy the following requirements.

- The syntax of a system is given in two levels. The basic level is represented by the entities P, Q , which could be processes, messages, etc.
- The high-level syntax of the system is defined as:

$$N ::= P \mid (N_1 \mid N_2) \mid \text{opn}(N_1, \dots, N_n)$$

where

- $N_1 \mid N_2$ symbolises a parallel composition of systems N_1 and N_2 ;
- $\text{opn}(N_1, \dots, N_n)$ represents the n -ary operations applied on systems N_1, \dots, N_n ;
- We write the system N as a term $T[P_1, \dots, P_n]$ to highlight the fact that system is based on the entities P_1, \dots, P_n .

- Structural congruence rules defined on the forward system are in the following shape:

$$\text{(Struct)} \quad T[P_1, \dots, P_n] \equiv T'[P_1, \dots, P_n]$$

- The semantics of the system is given in terms of a reduction semantics with the following rule schema.

$$\text{(ACT)} \quad P_1 \mid \dots \mid P_n \rightarrow Q_1 \mid \dots \mid Q_m$$

$$\text{(OPN)} \quad \frac{N_i \rightarrow N'_i}{\text{opn}(N_0, \dots, N_i, \dots, N_n) \rightarrow \text{opn}(N_0, \dots, N'_i, \dots, N_n)}$$

$$\text{(PAR)} \quad \frac{N_1 \rightarrow N'_1}{N_1 \mid N_2 \rightarrow N'_1 \mid N_2}$$

$$\text{(EQUIV)} \quad \frac{N \equiv N' \quad N \rightarrow N_1 \quad N_1 \equiv N'_1}{N' \rightarrow N'_1}$$

- The reversible system R is defined as:

$$R ::= k : P \mid (R_1 \mid R_2) \mid \text{opn}(R_1, \dots, R_n) \mid H$$

$$H ::= [R_h; K]$$

where

- k_1, \dots, k_n are unique keys identifying the entities;
 - H is the memory called history, recording the past state of the system R_h and the set of the fresh keys generated with the forward execution, denoted with K .
- As in the forward model, we shall write the system R as a term $T[k_1 : P_1, \dots, k_n : P_n]$ to highlight the fact that system is based on the identified entities $k_1 : P_1, \dots, k_n : P_n$.
- Structural congruence rules defined on the system are in the following shape:

$$\text{(Struct)} \quad T[k_1 : P_1, \dots, k_n : P_n] \equiv T'[k_1 : P_1, \dots, k_n : P_n]$$

The forward rules of the uncontrolled reversible semantics are:

$$(F\text{-ACT}) \frac{P_1 \mid \dots \mid P_n \rightarrow Q_1 \mid \dots \mid Q_m \quad j_1, \dots, j_m \text{ are fresh keys}}{k_1 : P_1 \mid \dots \mid k_n : P_n \rightarrow j_1 : Q_1 \mid \dots \mid j_m : Q_m \mid [k_1 : P_1 \mid \dots \mid k_n : P_n; j_1, \dots, j_m]}$$

$$(F\text{-OPN}) \frac{R_i \rightarrow R'_i \mid H \quad (\text{key}(R'_i) \setminus \text{key}(R_i)) \cap (\text{key}(R_0, \dots, R_{i-1}, R_{i+1}, \dots, R_n) = \emptyset)}{\text{opn}(R_0, \dots, R_i, \dots, R_n) \rightarrow \text{opn}(R_0, \dots, R'_i, \dots, R_n) \mid H}$$

$$(F\text{-PAR}) \frac{R_1 \rightarrow R'_1 \quad (\text{key}(R'_1) \setminus \text{key}(R_1)) \cap \text{key}(R_2) = \emptyset}{R_1 \mid R_2 \rightarrow R'_1 \mid R_2}$$

$$(F\text{-EQUIV}) \frac{R \equiv R' \quad R \rightarrow R_1 \quad R_1 \equiv R'_1}{R' \rightarrow R'_1}$$

- function $\text{key}(\cdot)$ compute the set of keys;
- in the rule (F-ACT) the history is $H = [k_1 : P_1 \mid \dots \mid k_n : P_n; j_1, \dots, j_m]$.

The backward rules of the uncontrolled reversible semantics are:

$$(B-ACT) \frac{H = [k_1 : P_1 \mid \dots \mid k_n : P_n; j_1, \dots, j_m]}{j_1 : Q_1 \mid \dots \mid j_m : Q_m \mid H \rightsquigarrow k_1 : P_1 \mid \dots \mid k_n : P_n}$$

$$(B-OPN) \frac{R'_i \mid H \rightsquigarrow R_i}{\text{opn}(R_0, \dots, R'_i, \dots, R_n) \mid H \rightsquigarrow \text{opn}(R_0, \dots, R_i, \dots, R_n)}$$

$$(B-PAR) \frac{R'_1 \rightsquigarrow R_1}{R'_1 \mid R_2 \rightsquigarrow R_1 \mid R_2}$$

$$(B-EQUIV) \frac{R \equiv R' \quad R \rightsquigarrow R_1 \quad R_1 \equiv R'_1}{R' \rightsquigarrow R'_1}$$

Definition

Given a history $H = [k_1 : P_1 \mid \dots \mid k_n : P_n; K]$, the set of all keys belonging to the history H , written as $\text{key}(H)$, is defined with

$$\text{key}(H) = \{k_1\} \cup \dots \cup \{k_n\} \cup K.$$

Transitions are labeled with H and we have that transition t of a system R is $t : R \xrightarrow{H} R'$, where H is the memory created or consumed by transition.

Definition (Concurrent transitions)

Two cointial transitions $R \xrightarrow{H'} R'$ and $R \xrightarrow{H''} R''$ are *concurrent* if $\text{key}(H') \cap \text{key}(H'') = \emptyset$. Transitions which are not concurrent are in *conflict*.

Lemma (Loop Lemma)

For every reachable system R and forward transition $t : R \xrightarrow{H} R'$, there exists the backward transition $t^\bullet : R' \overset{H}{\rightsquigarrow} R$ and vice versa.

Lemma (Square Lemma)

If $t_1 : R \xrightarrow{H'} R'$ and $t_2 : R \xrightarrow{H''} R''$ are two cointial concurrent transitions, there exist two cofinal transitions $t_2/t_1 : R' \xrightarrow{H''} R'''$ and $t_1/t_2 : R'' \xrightarrow{H'} R'''$.

Theorem (Causal Consistency)

Given two cointial derivations d_1 and d_2 , $d_1 \sim d_2$ if and only if d_1 and d_2 are cofinal.

- Danos, V., Krivine, J.: Reversible communicating systems. In: CONCUR. LNCS, vol. 3170, pp. 292–307. Springer (2004)

Example:

- **Higher-Order π -calculus**

$$a\langle P_1 \mid P_2 \rangle \mid a(X) \triangleright X \rightarrow P_1 \mid P_2$$

- **$\rho\pi$ calculus**

$$\begin{aligned}
 (k_1 : a\langle P_1 \mid P_2 \rangle) \mid (k_2 : a(X) \triangleright X) &\rightarrow \nu k. k : (P_1 \mid P_2) \mid [M; k] \\
 &\equiv \nu k, h_1, h_2. (\langle h_1, \tilde{h} \rangle \cdot k : P_1 \mid \langle h_2, \tilde{h} \rangle \cdot k : P_2)
 \end{aligned}$$

where $M = (k_1 : a\langle P_1 \mid P_2 \rangle) \mid (k_2 : a(X) \triangleright X)$ and $\tilde{h} = \{h_1, h_2\}$

- **Our reversible semantics for HO π**

$$k_1 : a\langle P_1 \mid P_2 \rangle \mid k_2 : a(X) \triangleright X \rightarrow j_1 : P_1 \mid j_2 : P_2 \mid [R_h; \{j_1, j_2\}]$$

where $R_h = k_1 : a\langle P_1 \mid P_2 \rangle \mid k_2 : a(X) \triangleright X$

The forward Erlang system is defined as a pool of processes and floating messages with the following grammar:

$$N := \langle p, \theta, e \rangle \mid (p, p', v) \mid (N_1 \mid N_2)$$

where

- $\langle p, \theta, e \rangle$ represents a process with its pid p , environment θ and an expression e to be evaluated;
- every process has a unique pid;
- (p, p', v) represents a floating message with content v sent from the process with pid p to the one with pid p' .
Floating messages are the messages in the system after they are sent and before they are received.

$$\text{(SEQ)} \frac{\theta, e \xrightarrow{\tau} \theta', e'}{\langle p, \theta, e \rangle \hookrightarrow \langle p, \theta', e' \rangle}$$

$$\text{(REC)} \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{c\bar{n}})} \theta', e' \quad \text{and} \quad \text{matchrec}(\theta, \overline{c\bar{l}_n}, \nu) = (\theta_i, e_i)}{\langle p', p, \{\nu, k''\} \rangle \mid \langle p, \theta, e \rangle \hookrightarrow \langle p, \theta'_i, e'_i \{ \kappa \mapsto e_i \} \rangle}$$

$$\text{(SEND)} \frac{\theta, e \xrightarrow{\text{send}(p', \nu)} \theta', e' \quad k'' \text{ is a fresh symbol}}{\langle p, \theta, e \rangle \hookrightarrow \langle p, \theta', e' \rangle \mid \langle p, p', \{\nu, k''\} \rangle}$$

$$\text{(SPAWN)} \frac{\theta, e \xrightarrow{\text{spawn}(\kappa, f/n, [\overline{v_n}]}) \theta', e' \quad p' \text{ is a fresh pid}}{\langle p, \theta, e \rangle \hookrightarrow \langle p, \theta', e' \{ \kappa \mapsto p' \} \rangle \mid \langle p', id, \text{apply } f/n (\overline{v_n}) \rangle}$$

$$\text{(SELF)} \frac{\theta, e \xrightarrow{\text{self}(\kappa)} \theta', e'}{\langle p, \theta, e \rangle \hookrightarrow \langle p, \theta', e' \{ \kappa \mapsto p \} \rangle}$$

$$\text{(PAR)} \frac{N \hookrightarrow N' \quad \text{pid}(N') \cap \text{pid}(N_1) = \emptyset}{N \mid N_1 \hookrightarrow N' \mid N_1}$$

- Ivan Lanese, Adrián Palacios and Germán Vidal (2019): Causal-consistent replay debugging for message passing programs. In: Technical report, DSIC, Universitat Politècnica de Valencia

The reversible Erlang system is defined as:

$$R := k : \langle p, \theta, e \rangle \mid k : (p, p', v) \mid (R \mid H) \mid (R_1 \mid R_2)$$

where

- k is the unique key identifying the entity;
- H is a history

$$(F\text{-SEQ}) \frac{\theta, e \xrightarrow{\tau} \theta', e' \quad k' \text{ is a fresh key}}{k : \langle p, \theta, e \rangle \rightarrow k' : \langle p, \theta', e' \rangle \mid [k : \langle p, \theta, e \rangle; k']}$$

$$(F\text{-SEND}) \frac{\theta, e \xrightarrow{\text{send}(p', v)} \theta', e' \quad k', k'' \text{ are fresh keys}}{k : \langle p, \theta, e \rangle \rightarrow k' : \langle p, \theta', e' \rangle \mid k'' : (p, p', v) \mid [k : \langle p, \theta, e \rangle; k', k']}$$

$$(F\text{-REC}) \frac{\theta, e \xrightarrow{\text{rec}(\kappa, \overline{c}_n)} \theta', e' \quad \text{and} \quad \text{matchrec}(\theta, \overline{c}_n, v) = (\theta_i, e_i) \quad k' \text{ is a fresh key}}{k'' : (p', p, v) \mid k : \langle p, \theta, e \rangle \rightarrow k' : \langle p, \theta'_i, e'_i \{ \kappa \mapsto e_i \} \rangle \mid [k'' : (p', p, v) \mid k : \langle p, \theta, e \rangle; k']}$$

$$(F\text{-SPAWN}) \frac{\theta, e \xrightarrow{\text{spawn}(\kappa, f/n, [\overline{v}_n])} \theta', e' \quad p' \text{ is a fresh pid, } k', k'' \text{ are fresh keys}}{k : \langle p, \theta, e \rangle \rightarrow k' : \langle p, \theta', e' \{ \kappa \mapsto p' \} \rangle \mid k'' : \langle p', id, \text{apply } f/n (\overline{v}_n) \rangle \mid [k : \langle p, \theta, e \rangle; k', k']}$$

$$(F\text{-SELF}) \frac{\theta, e \xrightarrow{\text{self}(\kappa)} \theta', e' \quad k' \text{ is a fresh key}}{k : \langle p, \theta, e \rangle \rightarrow k' : \langle p, \theta', e' \{ \kappa \mapsto p \} \rangle \mid [k : \langle p, \theta, e \rangle; k']}$$

$$(F\text{-PAR}) \frac{R \rightarrow R' \mid H \quad \text{pid}(R') \cap \text{pid}(R_1) = \emptyset \quad \text{and} \quad (\text{key}(R') \setminus \text{key}(R)) \cap \text{key}(R_1) = \emptyset}{R \mid R_1 \rightarrow R' \mid R_1 \mid H}$$

- we show that if two transitions are concurrent in the uncontrolled reversible logging semantics [1], they are concurrent in our reversible semantics (i.e. we prove the correspondence between two conflict relations):

Theorem

Given two different coinital transitions t_1 and t_2 . They are in conflict in the uncontrolled reversible logging semantics [1] if and only if they are in conflict in our uncontrolled reversible semantics.


- we proved that two reversible semantics are strong backward and forward barbed equivalent.

Theorem


The reversible logging semantics of [1] and our reversible semantics are strong backward and forward barbed equivalent.

[1] Ivan Lanese, Adrián Palacios and Germán Vidal (2019): Causal-consistent replay debugging for message passing programs. In: Technical report, DSIC, Universitat Politecnica de Valencia

- What is a link?
 - Link can be seen as a bidirectional path between two processes along which error signals travel.
- Why links?
 - Links are one of the main factors standing behind Erlang's remote error handling mechanisms.
 - Remote error handling is an important feature of Erlang. In contrast to Local error handling, the systems have a higher degree of separation of concerns.



Local Error Handling



Remote Error Handling

Differences with the simple Erlang model (given few slides before):

- the syntax of Erlang is enriched with the function `spawn_link()`;
- the process needs to keep track of all the processes linked with it i.e. process becomes defined as $\langle p, \theta, e, l \rangle$, where l is a set of links;
- notion of *terminated process*.
The process $\langle p, \theta, e, l \rangle$ terminated normally if $e = v$ or it terminated abnormally if $e = r$.
- the forward system is defined as:

$$E := \langle p, \theta, e, l \rangle \mid (p, p', v) \mid (E_1 \mid E_2)$$

- the reversible Link system for Erlang is defined as:

$$R := k : \langle p, \theta, e, l \rangle \mid k : (p, p', v) \mid (R \mid H) \mid (R_1 \mid R_2)$$

- High-level semantics (system rules):

$$(F\text{-ERRP}) \frac{e \in \{v, r\} \quad \wedge \quad p_1 \in I \quad \wedge \quad k', k'_1, k_2 \text{ are fresh keys}}{k : \langle p, \theta, e, l \rangle \mid k_1 : \langle p_1, \theta_1, e_1, l_1 \rangle \rightarrow k' : \langle p, \theta, e, l \setminus p_1 \rangle \mid k'_1 : \langle p_1, \theta_1, e_1, l_1 \setminus p \rangle \mid k_2 : \langle p, p_1, e \rangle \mid [k : \langle p, \theta, e, l \rangle \mid k_1 : \langle p_1, \theta_1, e_1, l_1 \rangle; k', k'_1, k_2]}$$

$$(F\text{-SPLINK}) \frac{\theta, e \xrightarrow{\text{spawn_link}(\kappa, f/n, [\overline{v_n}]})} \theta', e' \quad p' \text{ is a fresh pid} \quad k', k'' \text{ are fresh keys}}{k : \langle p, \theta, e, l \rangle \rightarrow k' : \langle p, \theta', e' \{ \kappa \mapsto p' \}, l \cup \{ \kappa \mapsto p' \} \rangle \mid k'' : \langle p', id, \text{apply } f/n (\overline{v_n}), \emptyset \cup p \rangle \mid [k : \langle p, \theta, e, l \rangle; k', k'']}$$

- The semantics of evaluation of sequential expressions [1]:

$$(CALL3) \frac{\text{eval}(op, v_1, \dots, v_n) = r}{\theta, \text{call } op(v_1, \dots, v_n) \xrightarrow{\tau} \theta, r}$$

-In this way we capture the error in built in functions.

- The semantics concerning the evaluation of the concurrent expressions of [1]:

-the new rules for the function `spawn_link()` are very similar to the rules for function `spawn()`.

[1] Ivan Lanese, Adrián Palacios and Germán Vidal (2019): Causal-consistent replay debugging for message passing programs. In: Technical report, DSIC, Universitat Politecnica de Valencia

Current work:

- We are working on defining the Reversible Link Semantics for Erlang;
- We are finalising proofs;

Future work:

- Investigating if our approach (with some modifications) could be applied to some other models as Klaim;
- Adding the link semantics into CauDEr.