

# Controlling Reversibility in $HO\pi$

Ivan Lanese <sup>1</sup>   **Claudio Antares Mezzina** <sup>2</sup>   Alan Schmitt<sup>2</sup>  
Jean-Bernard Stefani <sup>2</sup>

Focus Team, University of Bologna/INRIA, Italy

INRIA Grenoble-Rhône-Alpes, France

September 8, 2011

Concur 2011

# Roadmap

- 1 Reversibility
- 2 A small reversible language
- 3 Controlling Rollback

What if we could **undo** every action ?

Language support for Recovery Oriented Computing (ROC)

# Undo manifestations

- Application undo
- System undo (ROC)
- Logs and checkpoints
- Transaction rollback
- Distributed rollback-recovery

## Enabling defeasible partial agreements

- Systematic rollback-recovery
- Different transaction models

# The case for reversibility

## Enabling causality tracking

- Debugging
- Diagnosis
- Simulation
- Fault isolation

# Roadmap

- 1 Reversibility
- 2 A small reversible language**
- 3 Controlling Rollback

## Syntax

$P, Q ::= \mathbf{0}$	<i>null process</i>
$X$	<i>variable</i>
$\nu a. P$	<i>new name</i>
$(P \mid Q)$	<i>parallel composition</i>
$a\langle P \rangle$	<i>message</i>
$(a(X) \triangleright P)$	<i>trigger</i>

$$a \in \mathcal{N}$$



Computation is done by process passing

$$a\langle P \rangle \mid a(X) \triangleright X \rightarrow P$$

After communication no information about **previous** state of the configuration.

Computation is done by process passing

$$a\langle P \rangle \mid a(X) \triangleright X \rightarrow P$$

After communication no information about **previous** state of the configuration.

# A reversible $HO_{\pi}:\rho_{\pi}$

To make  $HO_{\pi}$  reversible:

- Log each action (message receipt)
- Uniquely identify action participants

## Syntax

$$P, Q ::= \mathbf{0} \mid X \mid \nu a. P \mid (P \mid Q) \mid a\langle P \rangle \mid (a(X) \triangleright P)$$

$M, N ::=$	<i>configurations</i>
$\mathbf{0}$	<i>null configuration</i>
$\mid \nu u. M$	<i>restriction</i>
$\mid (M \mid N)$	<i>parallel</i>
$\mid \kappa : P$	<i>thread</i>
$\mid [m; k]$	<i>memory</i>

$$a \in \mathcal{N}, k \in \mathcal{K}, u \in \mathcal{N} \cup \mathcal{K}$$

Intuitions:

- $\kappa : P$  thread of computation (process)  $P$  uniquely identified by tag  $\kappa$
- $[m; k]$  log of occurrence  $k$  of action  $m$  (message receipt)

## Reduction rules

FORWARD: 
$$\frac{m = (\kappa_1 : a\langle P \rangle) \mid (\kappa_2 : a(X) \triangleright Q)}{(\kappa_1 : a\langle P \rangle) \mid (\kappa_2 : a(X) \triangleright Q) \rightarrow \nu k. (k : Q\{P/X\}) \mid [m; k]}$$

BACKWARD: 
$$(k : P) \mid [m; k] \rightsquigarrow m$$

## Syntactical equivalence

$$(E.TAGP) \quad k : \prod_{i=1}^n \tau_i \equiv \nu \tilde{h}. \prod_{i=1}^n (\langle h_i, \tilde{h} \rangle \cdot k : \tau_i) \quad \tilde{h} = \{h_1, \dots, h_n\}$$

$$\tau = a\langle P \rangle \quad \vee \quad \tau = a(X) \triangleright P$$

Rule E.TAGP builds unique identifiers for parallel branches

## Theorem (Full reversal)

$M \rightarrow^* N$  if and only if  $N \rightarrow^* M$

$$\rightarrow = \Rightarrow \cup \rightsquigarrow$$



# Roadmap

- 1 Reversibility
- 2 A small reversible language
- 3 Controlling Rollback**

In  $\rho\pi$  back tracking is **uncontrolled**

- When do we get back?
  - normal execution should go **forward**
  - enable backward execution with a specific primitive
- How far do we get back?
  - each action is logged and **uniquely identified**
  - $[M; k]$  as **checkpoint**: roll  $k$

In  $\rho\pi$  back tracking is **uncontrolled**

- When do we get back?
  - normal execution should go **forward**
  - enable backward execution with a specific primitive
- How far do we get back?
  - each action is logged and **uniquely identified**
  - $[M; k]$  as **checkpoint**: roll  $k$

In  $\rho\pi$  back tracking is **uncontrolled**

- When do we get back?
  - normal execution should go **forward**
  - enable backward execution with a specific primitive
- How far do we get back?
  - each action is logged and **uniquely identified**
  - $[M; k]$  as **checkpoint**: **roll  $k$**

# How the roll primitive should work?

Naive implementation of roll  $k$ :

- 1 **collect** all the processes caused by  $k$
- 2 **delete** them
- 3 **restore** the content of memory  $[M; k]$

## Reduction rules

$$\text{(COM)} \quad \frac{m = (\kappa_1 : a\langle P \rangle) \mid (\kappa_2 : a(X) \triangleright_{\gamma} Q)}{(\kappa_1 : a\langle P \rangle) \mid (\kappa_2 : a(X) \triangleright_{\gamma} Q) \rightarrow \nu k. (k : Q\{^{k,P}/_{\gamma,X}\}) \mid [m; k]}$$

$$\text{(NAIVE)} \quad \frac{N \triangleright k \quad \text{complete}(N \mid [m; k] \mid (\kappa : \text{roll } k))}{N \mid [m; k] \mid (\kappa : \text{roll } k) \rightsquigarrow m \mid N \not\downarrow k}$$

# How the roll primitive should work?

Naive implementation of roll  $k$ :

- 1 **collect** all the processes caused by  $k$
- 2 **delete** them
- 3 **restore** the content of memory  $[M; k]$

## Reduction rules

$$\text{(COM)} \quad \frac{m = (\kappa_1 : a\langle P \rangle) \mid (\kappa_2 : a(X) \triangleright_{\gamma} Q)}{(\kappa_1 : a\langle P \rangle) \mid (\kappa_2 : a(X) \triangleright_{\gamma} Q) \rightarrow \nu k. (k : Q^{\{k, P / \gamma, X\}}) \mid [m; k]}$$

$$\text{(NAIVE)} \quad \frac{N \blacktriangleright k \quad \text{complete}(N \mid [m; k] \mid (\kappa : \text{roll } k))}{N \mid [m; k] \mid (\kappa : \text{roll } k) \rightsquigarrow m \mid N \not\downarrow k}$$


# Controlling Rollback

$k_1 : a\langle 0 \rangle$     $(k_2 : a(X) \triangleright_{\gamma} b\langle roll \ \gamma \rangle)$     $(k_3 : b(X) \triangleright c\langle 0 \rangle | X)$



# Controlling Rollback


$k_1 : a\langle 0 \rangle$     $(k_2 : a(X) \triangleright_\gamma b\langle \text{roll } \gamma \rangle)$     $(k_3 : b(X) \triangleright c\langle 0 \rangle | X)$



$[k_1 : M \mid k_2 : N; \mathbf{k}] \mid k : b\langle \text{roll } k \rangle$




# Controlling Rollback



# Controlling Rollback

$k_1 : a\langle 0 \rangle$     $(k_2 : a(X) \triangleright_\gamma b\langle \text{roll } \gamma \rangle)$     $(k_3 : b(X) \triangleright c\langle 0 \rangle | X)$




# Controlling Rollback


$k_1 : a\langle 0 \rangle$      $(k_2 : a(X) \triangleright_{\gamma} b\langle roll \ \gamma \rangle)$      $(k_3 : b(X) \triangleright c\langle 0 \rangle | X)$



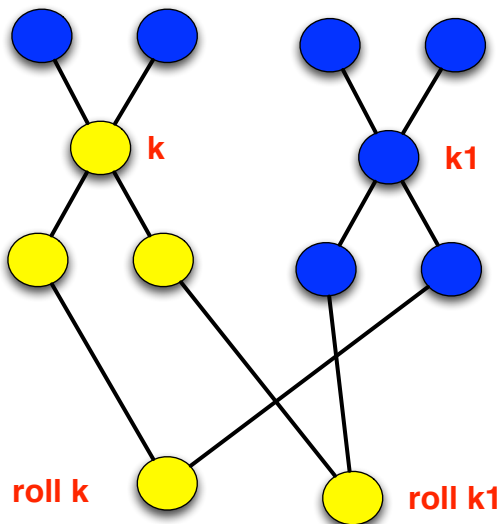
# Concurrent rolls




# Concurrent rolls



# Concurrent rolls



# Concurrent rolls



## Problem with concurrent **rolls**

- naive semantics is not unsound with respect to  $\rho\pi$
- naive semantics is **insufficiently permissive** with respect to  $\rho\pi$



# Why a better semantics?

Abstract semantics should

- be as liberal as possible
- not unduly restrict implementations

An implementation of naive:


- will reach same backward states
- hard to characterize
- artificial, do not correspond to any policy

## Reduction rules: HL semantics


(START)  $(\kappa_1 : \text{roll } k) \mid [m; k] \rightsquigarrow (\kappa_1 : \text{roll } k) \mid [m; k]^\bullet$

(ROLL) 
$$\frac{N \blacktriangleright k \quad \text{complete}(N \mid [m; k])}{N \mid [m; k]^\bullet \rightsquigarrow m \mid N \not\downarrow k}$$


# Concurrent rolls




# Concurrent rolls




# Concurrent rolls





# Concurrent rolls



# Concurrent rolls





## Soundness of backward reductions

if  $M \rightsquigarrow_T^* M'$  then  $M' \twoheadrightarrow^* M$ , with  $M$  and  $M'$  unmarked

## Completeness of backward reductions


if  $\phi(M) \rightsquigarrow_T^* N$  then  $M \rightsquigarrow_{HL}^* M'$  with  $M' = \phi(N)$ .


- $\phi(.) : HL \rightarrow \rho\pi$
- $T$  is a set of keys
- $\rightsquigarrow_T$  is a  $\rho\pi$  reduction restricted to keys caused by  $T$

- **sound** and **complete** with respect to  $\rho\pi$
- good as specification semantics
- relies on global checks and atomic steps


far away from a real implementation


- local checks
- use of asynchronous notifications
- rollback in two phases:
  - 1 top-down visit to gather all the causally dependent processes
  - 2 bottom-up visit to undo communications (one by one)







# LL example













## Correspondence HL vs LL

$$M_{HL} \approx_{LL}^c M$$

## Full Abstraction

$$j(M)_{LL} \approx_{LL}^c j(N) \text{ iff } M_{HL} \approx_{HL}^c N$$

- $M, N$  consistent  $HL$  configuration
- $j(\cdot)$  injection from  $HL$  to  $LL$

We presented a fine grained rollback primitive for  $HO\pi$

- built on top of  $\rho\pi$
- sound and complete with respect to  $\rho\pi$

We presented a distributed semantics

- closer to an implementation
- **fully abstract** with respect to the high level semantics

- encode  $\text{roll}\pi$  into  $HO\pi$  (✓)
- refine and relate the LL semantics to checkpoint and rollback schemes
  - localities, failures ...
- mix **dynamic compensations** and controlled rollback

Thank you!

Questions?

## LL Reduction Rules 1/2

(START)  $(\kappa_1 : \text{roll } k) \mid [m; k] \rightsquigarrow_{LL} (\kappa_1 : \text{roll } k) \mid [m; k]^\bullet \mid \text{rl } k$

(SPAN)  $\text{rl } \kappa_1 \mid [\kappa_1 : P \mid M; k]^\circ \rightsquigarrow_{LL} [[\kappa_1 : P] \mid M; k]^\circ \mid \text{rl } k$

(BRANCH) 
$$\frac{\langle h_i, \tilde{h} \rangle \cdot k \text{ occurs in } M}{\text{rl } k \mid M \rightsquigarrow_{LL} \prod_{h_i \in \tilde{h}} \text{rl } \langle h_i, \tilde{h} \rangle \cdot k \mid M}$$

## LL Redution Rules 2/2

(UP)  $\text{rl } \kappa_1 \mid (\kappa_1 : P) \rightsquigarrow_{LL} [\kappa_1 : P]$       (STOP)  $[m; k]^\circ \mid [k : P] \rightsquigarrow_{LL} m$

(GB)  $\nu k. \prod_{i \in I} \text{rl } \kappa_i \equiv 0 \quad \kappa_i = k \vee \kappa_i = \langle h, \tilde{u} \rangle \cdot k$

(TAGPFR)  $[k : \prod_{i=1}^n \tau_i] \equiv_{LL} \nu \tilde{h}. \prod_{i=1}^n [\langle h_i, \tilde{h} \rangle \cdot k : \tau_i]$   $\tilde{h} = \{h_1, \dots, h_n\}$   $n \geq 2$