

Concurrent Flexible Reversibility

Ivan Lanese¹ Michael Lienhardt² **Claudio Antares Mezzina**³
Alan Schmitt⁴ Jean-Bernard Stefani⁴

¹Focus Team, University of Bologna/INRIA, Italy

²PPS Laboratory Paris

³**Fondazione Bruno Kessler Trento, Italy**

⁴INRIA France

March 21, 2013

ESOP 2013

Roadmap

- 1 Reversibility
- 2 A flexible reversible language: $\text{croll}\pi$
- 3 Results

Language support for Recovery Oriented Computing (ROC)

- Application undo (e.g. [Ctrl+Z](#))
- System undo (e.g. [Windows Undo System Restore](#))
- Logs and checkpoints
- Transaction rollback
- Distributed rollback-recovery

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

Claim

We want to use reversibility as a unifying framework to build dependable systems

Reversibility


The possibility of executing a computation both in the standard, forward direction, and in the backward direction, going back to a past state

In a **sequential** setting, reversibility is simply undoing (recursively) the last action

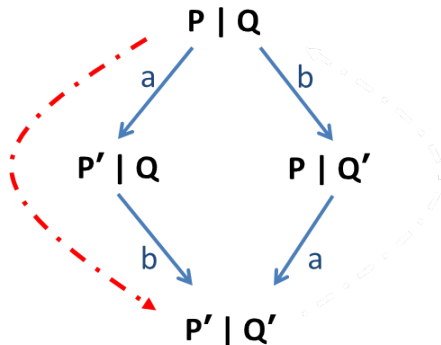
In a **concurrent** setting, one cannot simply undo the **last action**

- Not clear what is the last action
- Independent threads are reversed independently
- Causal dependencies should be respected
 - First undo the consequences and then the causes

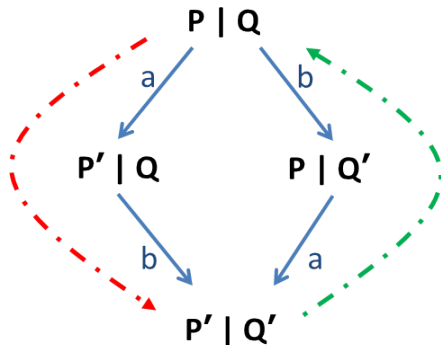
Concurrent Causal Reversibility



Concurrent Causal Reversibility



Concurrent Causal Reversibility



Previous works on reversibility proposed:


① **uncontrolled** reversibility

- show how to go back and forth
- no hint on when to go forward or backward
- more useful to understand basics of concurrent reversibility than as programming languages


② **controlled** reversibility

- backward computations are enabled in case of an error (when)
- the computation should undo the events that caused the error (how far to go)
- the computation that led to an error can be re-executed (may diverge)


Controlled Reversibility: example




Controlled Reversibility: example



Controlled Reversibility: example



Controlled Reversibility: example



Problems

- How do I avoid repeating the same erroneous computations?
- How do I specify what to do after a backward computation?

Solution (our idea)

Mixing controlled reversibility and compensations


Problems

- How do I avoid repeating the same erroneous computations?
- How do I specify what to do after a backward computation?


Solution (our idea)

Mixing controlled reversibility and compensations


Flexible Reversibility: example



Flexible Reversibility: example



Flexible Reversibility: example



Roadmap

1 Reversibility

2 A flexible reversible language: $\text{croll}\pi$

3 Results

Syntax

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

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

- Message passing communications
- Higher Order communications (sent values are processes)

To introduce reversibility in $HO\pi$:

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

To control reversibility:

- Use of a specific primitive roll γ where γ is an action identifier

To specify compensations:

- Use messages with alternatives

Syntax: Processes

P, Q	$a\langle P \rangle \div C$	<i>message</i>
	$(a(X) \triangleright_{\gamma} P)$	<i>trigger</i>
	roll γ	<i>roll</i>
	roll k	<i>active roll</i>
	X	<i>variable</i>
	$(P \mid Q)$	<i>parallel composition</i>
	$\nu a. P$	<i>new name</i>
	0	<i>null process</i>
	$C ::= 0$	<i>empty alternative</i>
	$a\langle P \rangle \div 0$	<i>message alternative</i>

Syntax: Configurations

$M, N ::=$		<i>configurations</i>
$k : P$		<i>thread</i>
$[\mu; k]$		<i>memory</i>
$k \prec (k_1, k_2)$		<i>connector</i>
$\nu u. M$		<i>restriction</i>
$(M \mid N)$		<i>parallel</i>
$\mathbf{0}$		<i>null configuration</i>

$\mu ::= (k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q)$

$u \in \mathcal{I} \quad a \in \mathcal{N} \quad k \in \mathcal{K}$

- $k : P$ thread of computation (process) P uniquely identified by tag k
- $[(k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_\gamma Q); k]$ action identified by k
- **roll** k reverts all the effects of a memory (message receipt) k
- $k : a\langle P \rangle \div C$ alternative C to the message $a\langle P \rangle$
- $k \prec (k_1, k_2)$ thread k is divided into two sub-threads k_1 and k_2

$$(\text{COM}) \frac{\mu = (k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q)}{(k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q) \rightarrow \nu k . (k : Q\{P, k / X, \gamma\}) \mid [\mu; k]}$$

$$(\text{TAGP}) k : P \mid Q \rightarrow \nu k_1 k_2 . k \prec (k_1, k_2) \mid k_1 : P \mid k_2 : Q$$

- Communication side-effects:
 - tags the new instance of Q with the new key k
 - γ is replaced by k in the instance of Q
 - stores the configuration that generated the even k in a memory
- Partial order among process identifiers

$$(\text{COM}) \frac{\mu = (k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q)}{(k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q) \rightarrow \nu k . (k : Q\{P, k / X, \gamma\}) \mid [\mu; k]}$$

$$(\text{TAGP}) k : P \mid Q \rightarrow \nu k_1 k_2 . k \prec (k_1, k_2) \mid k_1 : P \mid k_2 : Q$$

- Communication side-effects:
 - tags the new instance of Q with the new key k
 - γ is replaced by k in the instance of Q
 - stores the configuration that generated the even k in a memory
- Partial order among process identifiers

$$(\text{COM}) \frac{\mu = (k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q)}{(k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q) \rightarrow \nu k . (k : Q\{P, k / X, \gamma\}) \mid [\mu; k]}$$

$$(\text{TAGP}) k : P \mid Q \rightarrow \nu k_1 k_2 . k \prec (k_1, k_2) \mid k_1 : P \mid k_2 : Q$$

- Communication side-effects:
 - tags the new instance of Q with the new key k
 - γ is replaced by k in the instance of Q
 - stores the configuration that generated the even k in a memory
- Partial order among process identifiers

$$(\text{COM}) \frac{\mu = (k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q)}{(k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q) \rightarrow \nu k . (k : Q\{P, k / X, \gamma\}) \mid [\mu; k]}$$

$$(\text{TAGP}) k : P \mid Q \rightarrow \nu k_1 k_2 . k \prec (k_1, k_2) \mid k_1 : P \mid k_2 : Q$$

- Communication side-effects:
 - tags the new instance of Q with the new key k
 - γ is replaced by k in the instance of Q
 - stores the configuration that generated the even k in a memory
- Partial order among process identifiers

$$\text{(COM)} \quad \frac{\mu = (k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q)}{(k_1 : a\langle P \rangle \div C) \mid (k_2 : a(X) \triangleright_{\gamma} Q) \rightarrow \nu k . (k : Q\{P, k / X, \gamma\}) \mid [\mu; k]}$$

$$\text{(TAGP)} \quad k : P \mid Q \rightarrow \nu k_1 k_2 . k \prec (k_1, k_2) \mid k_1 : P \mid k_2 : Q$$

- Communication side-effects:
 - tags the new instance of Q with the new key k
 - γ is replaced by k in the instance of Q
 - stores the configuration that generated the even k in a memory
- Partial order among process identifiers

$$(S.ROLL) \quad \frac{k <: N \quad \text{complete}(N \mid [\mu; k] \mid (k_r : \text{roll } k)) \quad \mu' = \text{xtr}(\mu)}{N \mid [\mu; k] \mid (k_r : \text{roll } k) \rightarrow \mu' \mid N \not\downarrow_k}$$

How the rule works:

- 1 **collects** all the processes caused by k
- 2 **deletes** them and frees resources consumed by the computation to be rolled-back
- 3 **substitutes** the message in μ with its alternative (xtr function)

$$\text{xtr}(k_1 : a\langle P_1 \rangle \div C) \mid (k_2 : a(X) \triangleright_\gamma Q) = k_1 : C \mid (k_2 : a(X) \triangleright_\gamma Q)$$

$$(S.ROLL) \quad \frac{k <: N \quad \text{complete}(N \mid [\mu; k] \mid (k_r : \text{roll } k)) \quad \mu' = \text{xtr}(\mu)}{N \mid [\mu; k] \mid (k_r : \text{roll } k) \rightarrow \mu' \mid N \not\downarrow_k}$$

How the rule works:

- 1 **collects** all the processes caused by k
- 2 **deletes** them and frees resources consumed by the computation to be rolled-back
- 3 **substitutes** the message in μ with its alternative (xtr function)

$$\text{xtr}(k_1 : a\langle P_1 \rangle \div C) \mid (k_2 : a(X) \triangleright_\gamma Q) = k_1 : C \mid (k_2 : a(X) \triangleright_\gamma Q)$$

Roadmap

- 1 Reversibility
- 2 A flexible reversible language: $\text{croll}\pi$
- 3 Results

Simple messages as alternatives are powerful enough to express:

- Different kind of alternatives:
 - General alternatives: e.g. $a\langle P \rangle \div Q$
 - Triggers with alternatives: $(a(X) \triangleright_{\gamma} P) \div C$
- More sophisticated patterns:
 - Finite retry: $a\langle P \rangle \div_n C$ (try $a\langle P \rangle$ **n times** and then C)
 - Infinite retry: $a\langle P \rangle$

Some encodings in $\text{croll}\pi$

$$\langle a\langle P \rangle \div Q \rangle_{aa} = \nu c. a\langle \langle P \rangle_{aa} \rangle \div c\langle \langle Q \rangle_{aa} \rangle \div \mathbf{0} \mid c(X) \triangleright X$$

$$\langle a\langle P \rangle \rangle_{er} = \nu t. Y \mid a\langle \langle P \rangle_{er} \rangle \div t\langle Y \rangle \quad Y = t(Z) \triangleright Z \mid a\langle \langle P \rangle_{er} \rangle \div t\langle Z \rangle$$

Encodings Correctness

$$P \approx_c \langle P \rangle_{aa} \text{ for any closed process } P$$

$$P \approx_c \langle P \rangle_{er} \text{ for any closed process } P$$

\approx_c weak barbed congruence

- A transaction is a computation that may:
 - **succeed**, making results permanent
 - **abort**, undoing all its effect
 - have a compensation to be executed upon abort
- Interacting transactions are allowed to interact with the environment during their execution.
 - no **isolation**

TransCCS

$$\text{(COM)} \quad \bar{a} \mid a.P \rightarrow P \qquad \text{(CO)} \quad \llbracket P \mid \text{co } k \triangleright_k Q \rrbracket \rightarrow P$$

$$\text{(AB)} \quad \llbracket P \triangleright_k Q \rrbracket \rightarrow Q \qquad \text{(EMB)} \quad \llbracket P \triangleright_k Q \rrbracket \mid R \rightarrow \llbracket P \mid R \triangleright_k Q \mid R \rrbracket$$

- \rightarrow closed under structural congruence (π one) and under the contexts: $\bullet \mid P$, $\llbracket \bullet \triangleright_k Q \rrbracket$ and $\nu a. \bullet$
- $\llbracket P \triangleright_k Q \rrbracket$ transaction with name k , body P and compensation Q
- rule Emb allows an arbitrary process to get into the transactional scope
- abort is spontaneous, commit is explicit

Encoding

$$\langle \nu a. P \rangle = \nu a. \langle P \rangle \quad \langle P \mid Q \rangle = \langle P \rangle \mid \langle Q \rangle \quad \langle \bar{a} \rangle = \bar{a}$$

$$\langle a.P \rangle = a \triangleright P \quad \langle \text{co } k \rangle = l(X) \triangleright \mathbf{0} \quad \langle \mathbf{0} \rangle = \mathbf{0}$$

$$\langle \llbracket P \triangleright_l Q \rrbracket \rangle = [\nu l. \langle P \rangle \mid l(\text{roll } \gamma) \mid l(X) \triangleright X, \langle Q \rangle]_\gamma$$

$$\langle P, Q \rangle_\gamma = \nu a, c. \bar{a} \div \bar{c} \div \mathbf{0} \mid (a \triangleright_\gamma P) \mid (c \triangleright Q)$$

- abort is modelled as rollback
- a transaction is started by a message with alternative
- compensation Q is started upon rollback

Theorem (Encoding Correctness)

For each TransCCS process P , $P \approx_{\text{croll}\pi} \nu k. k : (P)$

- $\approx_{\text{croll}\pi}$ ad-hoc weak barbed bisimulation between transCCS processes and croll π processes
- croll π causality tracking mechanism is more precise than Hennessy embedding
- we avoid spurious rollbacks they have (thread independence)

Conclusion

- we presented a calculus with explicit rollback and core facilities for alternatives
- simple messages are enough to built more complex alternatives
- rollback and alternatives can encode transactional constructs
- messages with alternatives increase the expressive power of the calculus (**not shown here**)
- simple Maude interpreter (**not shown here**)
 - some primitive for errors handling (try/catch , stabilizer)
 - <http://proton.inrialpes.fr/~mlienhar/croll-pi/implem/>
- context lemma to help proofs (**not shown here**)

- try to encode more transactional constructs (e.g. STM)
- study relationships between alternatives and compensations in long-running transactions
- improve the Maude interpreter (e.g. garbage collection of unreachable memories)
- test the interpreter against more complex case studies