

A theory of communicating transactions

Matthew Hennessy

joint work with Edsko de Vries, Vasileios Koutavas

CHoCoLa, PPS Paris December 2012



TRINITY COLLEGE DUBLIN
COLÁISTE NA TRÍONÓIDE, BAILE ÁTHA CLIATH



Outline

Introduction

TransCCS

Liveness and safety properties

Compositional semantics



Standard Transactions

- ▶ Transactions provide *an abstraction for error recovery* in a concurrent setting.
- ▶ Guarantees:
 - ▶ **Atomicity**: Each transaction either runs in its entirety (commits) or not at all
 - ▶ **Consistency**: When faults are detected the transaction is automatically rolled-back
 - ▶ **Isolation**: The effects of a transaction are concealed from the rest of the system until the transaction commits
 - ▶ **Durability**: After a transaction commits, its effects are permanent
- ▶ **Isolation**:
 - ▶ good: provides coherent semantics
 - ▶ bad: limits concurrency
 - ▶ bad: limits co-operation between transactions and their environments

Communicating/Co-operating Transactions

- ▶ We *drop isolation to increase concurrency*
 - ▶ There is no limit on the communication between a transaction and its environment
- ▶ These new transactional systems guarantee:
 - ▶ **Atomicity**: Each transaction will either run in its entirety or not at all
 - ▶ **Consistency**: When faults are detected the transaction is automatically rolled-back, *together with all effects of the transaction on its environment*
 - ▶ **Durability**: After *all transactions that have interacted* commit, their effects are permanent (coordinated checkpointing)

Example: three-way rendezvous

$$P_1 \parallel P_2 \parallel P_3 \parallel P_4$$

Problem:

- ▶ P_n process/transaction subject to failure
- ▶ Some three P_n should decide to collaborate

Result:

- ▶ Each P_n in the coalition outputs id of its partners on channel out_n

Example: programming a three-way rendezvous

$$P_1 \parallel P_2 \parallel P_3 \parallel P_4$$

Algorithm for P_n :

- ▶ Broadcast id n randomly to two arbitrary partners
 $b!\langle n \rangle \mid b!\langle n \rangle$
- ▶ Receive ids from two random partners $b?(y) . b?(z)$
- ▶ Propose coalition with these partners $s_y!\langle n, z \rangle . s_z!\langle n, y \rangle$
- ▶ Confirm that partners are in agreement:
 - ▶ if YES, **commit** and report
 - ▶ if NO, **abort&retry**

Example: programming a three-way rendezvous

$$P_1 \parallel P_2 \parallel P_3 \parallel P_4$$

$$P_n \Leftarrow b!\langle n \rangle \mid b!\langle n \rangle \mid$$

$$\text{atomic} \llbracket b?(y) . b?(z) .$$

$$s_y!\langle n, z \rangle . s_z!\langle n, y \rangle . \quad \text{proposing}$$

$$s_n?(y_1, z_1) . s_n?(y_2, z_2) . \quad \text{confirming}$$

$$\text{if } \{y, z\} = \{y_1, z_1\} = \{y_2, z_2\}$$

$$\text{then } \textit{commit} \mid \text{out}_n!\langle y, z \rangle$$

$$\text{else } \textit{abrt\&retry} \rrbracket$$

Communicating Transactions: Issues

- ▶ Language Design
 - ▶ Transaction Synchronisers (Luchangco et al 2005)
 - ▶ Transactional Events for ML (Fluet, Grossman et al. ICFP 2008)
 - ▶ Communication Memory Transactions (Lesani, Palsberg PPOPP 2011)
 - ▶ ...
- ▶ Implementation strategies
 - ▶ See above
- ▶ Semantics Behavioural theory: what should happen when programs are run
 - ▶ TransCCS (Concur 2010, APLAS 2010)

Communication Memory Transactions Lesani Palsberg

- ▶ Builds on optimistic semantics of memory transactions O'Herlihy et al 2010
- ▶ Adds asynchronous channel-based message passing as in Actors CML etc
- ▶ Formal reduction semantics
- ▶ Formal properties of semantics proved
- ▶ Implementation as a Scala library
- ▶ Performance evaluation using benchmarks



TransCCS

An extension of CCS with communicating transactions.

1. **Simple language**: 2 additional language constructs and 3 additional reduction rules.
2. **Intricate concurrent and transactional behaviour**:
 - ▶ encodes nested, restarting, and non-restarting transactions
 - ▶ does not limit communication between transactions
3. **Simple behavioural theory**: based on properties of systems:
 - ▶ *Safety* property: nothing bad happens
 - ▶ *Liveness* property: something good happens



TransCCS

Syntax:	$P, Q ::= \sum \mu_i.P_i$	guarded choice
	$ P Q$	parallel
	$ \nu a.P$	hiding
	$ \mu X.P$	recursion
	$ \llbracket P \triangleright_k Q \rrbracket$	transaction (k bound in P)
	$ \text{co } k$	commit

Transaction $\llbracket P \triangleright_k Q \rrbracket$

- ▶ execute P to completion ($\text{co } k$)
- ▶ subject to random aborts
- ▶ if aborted roll back all effects of P and initiate Q
- ▶ roll back includes ... **environmental impact of P**



Rollbacks and Commits

Co-operating actions: $a \leftarrow \text{needs co-operation of } \rightarrow \bar{a}$

$$T_a | T_b | T_c | P_d | P_e$$

where

$$T_a = \llbracket \bar{d}.\bar{b}.\text{co } k_1 | a \triangleright_{k_1} \mathbf{0} \rrbracket$$

$$T_b = \llbracket \bar{c}.\text{co } k_2 | b \triangleright_{k_2} \mathbf{0} \rrbracket$$

$$T_c = \llbracket \bar{e}.c.\text{co } k_3 \triangleright_{k_3} \mathbf{0} \rrbracket$$

$$P_d = d.R_d$$

$$P_e = e.R_e$$

- ▶ if T_c aborts, what roll-backs are necessary?
- ▶ When can action a be considered permanent?
- ▶ When can code R_d be considered permanent?



Reduction semantics main rules

R-COMM

$$\frac{a_i = \bar{b}_j}{\sum_{i \in I} a_i.P_i \mid \sum_{j \in J} b_j.Q_j \rightarrow P_i \mid Q_j}$$

Communication

R-Co

$$\frac{}{\llbracket P \mid \text{co } k \triangleright_k Q \rrbracket \rightarrow P}$$

Commit

R-AB

$$\frac{}{\llbracket P \triangleright_k Q \rrbracket \rightarrow Q}$$

Random abort

R-EMB

$$\frac{k \notin R}{\llbracket P \triangleright_k Q \rrbracket \mid R \rightarrow \llbracket P \mid R \triangleright_k Q \mid R \rrbracket}$$

Embed



Simple Example

Convention:

- ▶ ω : I am happy
- ▶ \circlearrowleft : I am sad

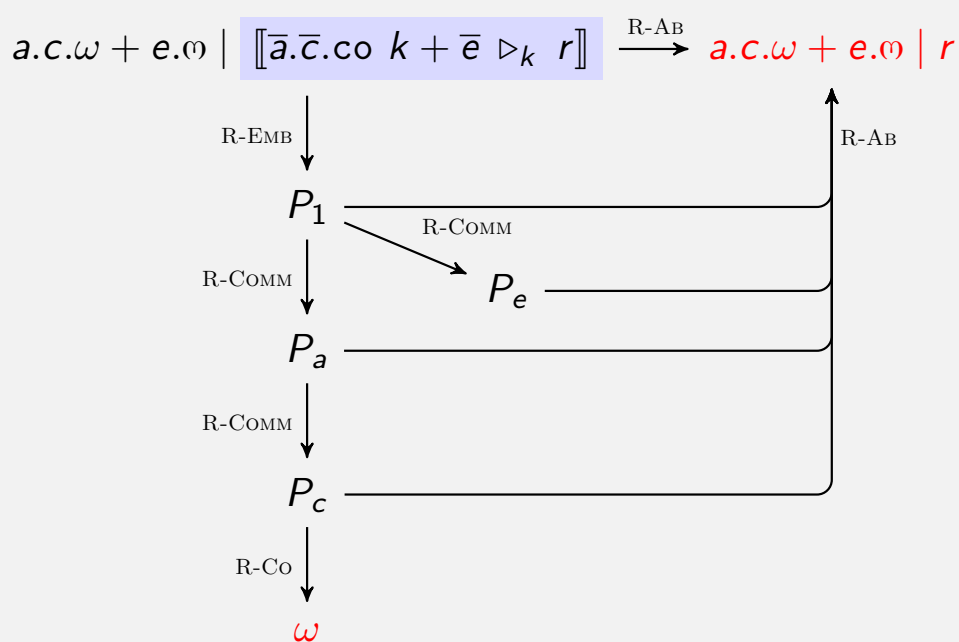
$$\begin{aligned}
 & a.c.\omega + e.\circlearrowleft \mid \llbracket \bar{a}.\bar{c}.\text{co } k + \bar{e} \triangleright_k r \rrbracket \\
 \xrightarrow{\text{R-EMB}} & \llbracket a.c.\omega + e.\circlearrowleft \mid \bar{a}.\bar{c}.\text{co } k + \bar{e} \triangleright_k a.c.\omega + e.\circlearrowleft \mid r \rrbracket \\
 \xrightarrow{\text{R-COMM}} & \llbracket c.\omega \mid \bar{c}.\text{co } k \triangleright_k a.c.\omega + e.\circlearrowleft \mid r \rrbracket \\
 \xrightarrow{\text{R-COMM}} & \llbracket \omega \mid \text{co } k \triangleright_k a.c.\omega + e.\circlearrowleft \mid r \rrbracket \\
 \xrightarrow{\text{R-Co}} & \omega
 \end{aligned}$$



Simple Example (a second trace)

$$\begin{aligned}
 & a.c.\omega + e.\emptyset \mid \llbracket \bar{a}.\bar{c}.co\ k + \bar{e} \triangleright_k r \rrbracket \\
 \xrightarrow{\text{R-EMB}} & \llbracket a.c.\omega + e.\emptyset \mid \bar{a}.\bar{c}.co\ k + \bar{e} \triangleright_k a.c.\omega + e.\emptyset \mid r \rrbracket \\
 \xrightarrow{\text{R-COMM}} & \llbracket \emptyset \triangleright_k a.c.\omega + e.\emptyset \mid r \rrbracket \quad (\text{Deadlocked}) \\
 \xrightarrow{\text{R-AB}} & a.c.\omega + e.\emptyset \mid r \quad (\text{The environment is restored})
 \end{aligned}$$

Simple Example (all traces)

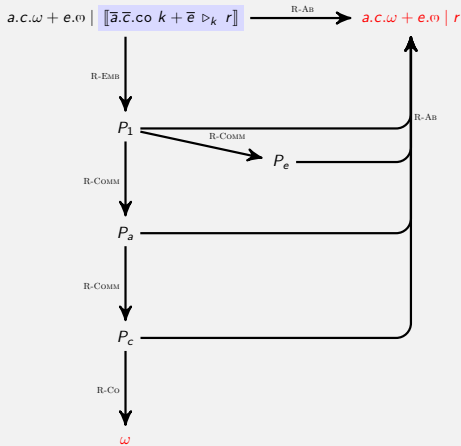


Will never be sad:

\emptyset

assuming r does not contain \bar{e}

Aborting transactions



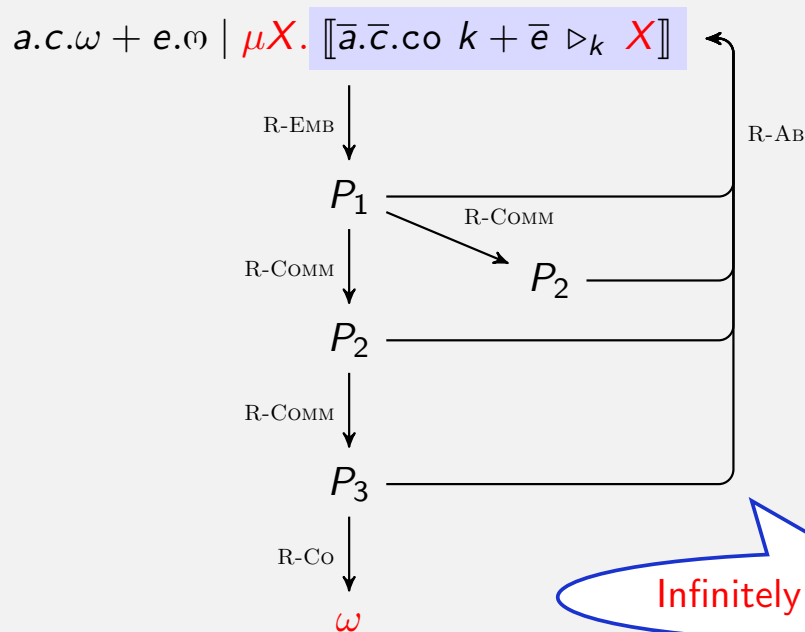
A commit step makes the effects of the transaction permanent (**Durability**)

An abort step:

- ▶ restarts the transaction
- ▶ rolls-back embedded processes to their state before embedding (**Consistency**)
- ▶ does not roll-back actions that happened before embedding
- ▶ does not affect non-embedded processes

The behavioural theory will show the **Atomicity** property.

Restarting transactions



Infinitely aborting loop

Will never be sad:

⊙

Double Embedding

$$\begin{array}{l}
 \llbracket a.co\ k \mid b \triangleright_k \mathbf{0} \rrbracket \mid \llbracket \bar{a}.co\ l \mid c \triangleright_l \mathbf{0} \rrbracket \\
 \xrightarrow{\text{R-EMB}} \llbracket a.co\ k \mid b \mid \llbracket \bar{a}.co\ l \mid c \triangleright_l \mathbf{0} \rrbracket \triangleright_k \llbracket \bar{a}.co\ l \mid c \triangleright_l \mathbf{0} \rrbracket \rrbracket \\
 \xrightarrow{\text{R-EMB}} \llbracket b \mid \llbracket a.co\ k \mid \bar{a}.co\ l \mid c \triangleright_l a.co\ k \rrbracket \triangleright_k \llbracket \bar{a}.co\ l \mid c \triangleright_l \mathbf{0} \rrbracket \rrbracket \\
 \xrightarrow{\text{R-COMM}} \llbracket b \mid \llbracket co\ k \mid co\ l \mid c \triangleright_l a.co\ k \rrbracket \triangleright_k \llbracket \bar{a}.co\ l \mid c \triangleright_l \mathbf{0} \rrbracket \rrbracket \\
 \xrightarrow{\text{R-Co}} \llbracket b \mid co\ k \mid c \triangleright_k \llbracket \bar{a}.co\ l \mid c \triangleright_l \mathbf{0} \rrbracket \rrbracket \\
 \xrightarrow{\text{R-Co}} b \mid c
 \end{array}$$

Safety properties

Safety: “Nothing bad will happen” [Lamport’77]

- ▶ A safety property can be formulated as a *safety test* T° which signals on channel \circ when it detects the bad behaviour

Examples:

- ▶ $\mu X.(a.X + e.\circ)$ can not perform e while performing any sequence of a s
- ▶ $T^\circ = e.\circ \mid \bar{a}.b$ can not perform e when a followed by b is offered.
- ▶ P passes the safety test T° when $P \mid T^\circ$ **can not** output on \circ
 - ▶ This is the negation of passing a “may test” [DeNicola-Hennessy’84]

Examples:

- ▶ $l_3 = \mu X. \llbracket a.b.co\ k + \bar{e} \triangleright_k X \rrbracket$ passes safety test T°
- ▶ $l_4 = \mu X. \llbracket a.b.co\ k \mid \bar{e} \triangleright_k X \rrbracket$ does **not** pass safety test T°

Safety

Definition (Basic Observable)

$P \Downarrow_{\mathfrak{m}}$ iff there exists P' such that $P \rightarrow^* P' \mid \mathfrak{m}$

- ▶ Basic observable actions are *permanent*
- ▶ **True:** $\llbracket a.b.co \ k \mid \bar{e} \triangleright_k \mathbf{0} \rrbracket \mid (e.\mathfrak{m} \mid \bar{a}.\bar{b}) \Downarrow_{\mathfrak{m}}$
- ▶ **False:** $\llbracket a.b.co \ k + \bar{e} \triangleright_k \mathbf{0} \rrbracket \mid (e.\mathfrak{m} \mid \bar{a}.\bar{b}) \Downarrow_{\mathfrak{m}}$

Definition (P Passes safety test $T^{\mathfrak{m}}$)

P cannot $T^{\mathfrak{m}}$ when $P \mid T^{\mathfrak{m}} \not\Downarrow_{\mathfrak{m}}$

Definition (Safety Preservation)

$S \sqsubseteq_{\text{safe}} I$ when $\forall T^{\mathfrak{m}}. S \text{ cannot } T^{\mathfrak{m}} \text{ implies } I \text{ cannot } T^{\mathfrak{m}}$



Safety preservation: Examples

$$S_{ab} = \mu X. \llbracket a.b.co \ k \triangleright_k X \rrbracket$$

$$I_3 = \mu X. \llbracket a.b.co \ k + \bar{e} \triangleright_k X \rrbracket$$

$$I_4 = \mu X. \llbracket a.b.co \ k \mid \bar{e} \triangleright_k X \rrbracket$$

- ▶ $S_{ab} \not\sqsubseteq_{\text{safe}} I_4$ use test $T^{\mathfrak{m}} = e.\mathfrak{m} \mid \bar{a}.\bar{b}$
- ▶ $S_{ab} \sqsubseteq_{\text{safe}} I_3$ – proof techniques required
- ▶ $\tau.P + \tau.Q \sqsubseteq_{\text{safe}} \llbracket P \triangleright_k Q \rrbracket$, for any P, Q – proof techniques reqd



Liveness

Liveness: “Something good will eventually happen” [Lampert'77]

- ▶ A liveness property can be formulated as a *liveness test* T^ω which detects and reports good behaviour on ω .

Examples:


- ▶ $T^\omega = \bar{a}.\bar{b}.\omega$ can do an a then a b
 - ▶ $\mu X. \llbracket \bar{a}.\bar{b}.\omega \mid \text{co } I \rrbracket \triangleright_I X$ can eventually do an a, b uninterrupted?
 - ▶ $a.\mu X. \llbracket \bar{b}.\bar{c}.\omega \mid \text{co } I \rrbracket \triangleright_I X$ English?
- ▶ P passes the liveness test T^ω when ω is eventually guaranteed

Dilemma: What does this mean?



Dilemma

Does $\mu X. \llbracket a.b.\text{co } k \rrbracket \triangleright_k X$ pass liveness test $T_{ab}^\omega = \bar{a}.\bar{b}.\omega$?



- ▶ **must-testing:** NO because of infinite loop
- ▶ **should-testing:** YES



Liveness testing

Definition (P Passes liveness test T^ω [Rensink-Vogler'07])

$$P \text{ shd } T^\omega \text{ when } \forall R. P \mid T^\omega \rightarrow^* R \text{ implies } R \downarrow_\omega$$

Examples:

- ▶ $\mu X. \llbracket a.b.co \ k \triangleright_k \ X \rrbracket$ passes liveness test $T_{ab}^\omega = \bar{a}.\bar{b}.\omega$
- ▶ $\llbracket a.b.co \ k \triangleright_k \ \mathbf{0} \rrbracket$ does not pass test T_{ab}^ω

Definition (Liveness preservation)

$$S \sqsubseteq_{\text{live}} I \text{ when } \forall T^\omega. S \text{ shd } T^\omega \text{ implies } I \text{ shd } T^\omega$$



Liveness preservation: Examples

$$S_{ab} = \mu X. \llbracket a.b.co \ k \triangleright_k \ X \rrbracket$$

$$I_2 = \mu X. \llbracket a.b.\mathbf{0} \triangleright_k \ X \rrbracket$$

$$I_3 = \mu X. \llbracket a.b.co \ k + \bar{e} \triangleright_k \ X \rrbracket$$

- ▶ $S_{ab} \not\sqsubseteq_{\text{live}} I_2$ use test $T^\omega = \bar{a}.\bar{b}.\omega$
- ▶ $S_{ab} \sqsubseteq_{\text{live}} I_3$ – proof techniques required
- ▶ $\mu X. \llbracket P \mid co \ k \triangleright_k \ X \rrbracket \approx_{\text{live}} P$, for any P – proof techniques reqd


Proof techniques:

Require characterisations using “traces” and “refusals”



Compositional Semantics

- ▶ The embedding rule is simple but entangles the processes
- ▶ We need to reason about the behaviour of $P|Q$ in terms of P and Q
- ▶ We introduce a compositional Labelled Transition System that uses *secondary transactions*: $\llbracket P \triangleright_k Q \rrbracket^\circ$



Compositional Semantics: safe-testing

The behaviour of processes in TransCCS can be understood by a *simple subset of the LTS traces*:

- ▶ where *all actions are eventually committed*
- ▶ that *ignore transactional annotations* on the traces

$$\text{Tr}_{\text{clean}} \left(\llbracket a.c.co \ k \triangleright_k e \rrbracket \right) = \{\epsilon, \mathbf{a c}, \mathbf{e}\}$$

$$\text{Tr}_{\text{clean}} \left(\mu X. \llbracket a.c.co \ k \triangleright_k X \rrbracket \right) = \{\epsilon, \mathbf{a c}\}$$

- ▶ Set of clean traces not prefix closed: **atomicity**

Characterisation of Safe Testing:

$$P \sqsubseteq_{\text{safe}} Q \quad \text{iff} \quad \text{Tr}_{\text{clean}}(P) \subseteq \text{Tr}_{\text{clean}}(Q)$$

- ▶ To understand the safe-testing behaviour of P we only need to consider the clean traces $\text{Tr}_{\text{clean}}(P)$.

Compositional semantics: should-testing

Tree Failures: [Rensink-Vogler'07]

(t, Ref) where

- ▶ t is a clean trace
- ▶ Ref is a set of clean traces

can be non-prefixed closed

Tree failures of a process:

(t, Ref) is a **tree failure** of P when

$$\exists P'. P \xrightarrow{t}_{CL} P' \text{ and } \mathcal{L}(P') \cap Ref = \emptyset$$

$$\mathcal{F}(P) = \{(t, Ref) \text{ tree failure of } P\}$$



Characterisation of should-testing:

$$S \sqsubseteq_{\text{live}} I \text{ iff } \mathcal{F}(S) \supseteq \mathcal{F}(I)$$

Simple Examples

$$\text{Let } S_{ab} = \mu X. \llbracket a.b.co \ k \triangleright_k \ X \rrbracket \quad \mathcal{L}(S_{ab}) = \{\epsilon, ab\}$$

$$\mathcal{F}(S_{ab}) = \{(\epsilon, S \setminus ab), (ab, S) \mid S \subseteq A^*\}$$

$$\begin{aligned} \text{▶ } S_{ab} &\approx_{\text{safe}} I_1 = \llbracket a.b.co \ k \triangleright_k \ \mathbf{0} \rrbracket & \mathcal{L}(I_1) &= \{\epsilon, ab\} \\ S_{ab} &\not\approx_{\text{live}} I_1 & \mathcal{F}(I_1) &= \{(\epsilon, S), (ab, S) \mid S \subseteq A^*\} \end{aligned}$$

$$\begin{aligned} \text{▶ } S_{ab} &\approx_{\text{safe}} I_2 = \mu X. \llbracket a.b.co \ k + e \triangleright_k \ X \rrbracket & \mathcal{L}(I_2) &= \mathcal{L}(S_{ab}) \\ S_{ab} &\approx_{\text{live}} I_2 & \mathcal{F}(I_2) &= \mathcal{F}(S_{ab}) \end{aligned}$$

Summary

- ▶ TransCCS: a language for communicating/co-operative transactions
- ▶ simple reduction semantics using an *embedding* rule
- ▶ behavioural theories for preservation of
 - ▶ safety properties
 - ▶ liveness properties
- ▶ characterisations which allow
 - ▶ proofs of equivalences
 - ▶ equational laws

References:

- ▶ *Communicating Transactions*, Concur 2010
- ▶ *Liveness of Communicating Transactions*, APLAS 2010

Current work:

- ▶ Extension to Haskell/CML
- ▶ prototype implementation
- ▶ Proof techniques based on traces, refusal trees, co-induction



THANK YOU!



Workshop announcement

1st Workshop on Optimistic Cooperation in Concurrent Programming (OCCP 2013)

- ▶ Location: Rome, Italy (co-located with ETAPS 2013)
- ▶ Date: Saturday March 16th, 2013
- ▶ Submissions: 14th Dec (abstracts) 21st Dec (Papers)

Details: <http://www.cs.tcd.ie/Vasileios.Koutavas/occp-workshop>