

MPRI C.2.3 - Concurrency

Probabilistic models and applications Lecture 3

Kostas Chatzikokolakis

Jan 10, 2012

Outline of the lectures

- Dec 13
- Dec 20
- Jan 10
- Jan 17
- Jan 24

Outline of the lectures

- The need for randomization
- Probabilistic automata
- Probabilistic bisimulation
- Probabilistic calculi
- Encoding of the pi-calculus into the asynchronous fragment
- Introduction to probabilistic model checking and PRISM
- Verification of anonymity protocols: Dining Cryptographers, Crowds

Exercises from the last lecture

Exercise 1: Show that probabilistic bisimulation is a generalization of traditional bisimulation

Puzzle from the last lecture

- I select two real numbers in some arbitrary way
- I put them in two envelopes, you select one of them (in any way you want)
- You **see** the number and you have 2 options: **keep** it, or **exchange** it with the other envelope
- Your goal is to select the bigger number
- Is there any strategy that **guarantees** winning this game with pb higher than $1/2$?

Outline

Encoding of π -calculus in the asynchronous fragment

Encoding of π -calculus in the probabilistic asynchronous π

Non-deterministic transition systems and CTL

Probabilistic asynchronous π -calculus

The input guarded choice is probabilistic.

The prefixes

$\alpha ::= x(y) \mid \tau$ input | silent action

The processes

$P ::=$	0	inaction
	$\sum_i p_i \alpha_i . P_i$	probabilistic choice
	$\bar{x}y$	output
	$P \mid P$	parallel
	$(\nu x)P$	new name
	$!P$	replication

where $\sum_i p_i = 1$

Expressive power of π_a wrt π

- Clearly π is at least as expressive as π_a
- The latter is practically a subset of the former:
 $\bar{x}y$ (in π_a) can be seen as $\bar{x}y.0$ (in π)
- What about the **opposite direction**? We need to encode:
 - the **output prefix**
 - the choice operator

Three types of choice: internal, separate, mixed

- In general, in order to compare the expressive power of two languages, we look for the **existence/non existence of an encoding** with certain properties among these languages
- What is a **good notion of encoding** to be used as basis to measure the relative expressive power?

A “good” notion of encoding

In general we would be happy with an encoding $\llbracket \cdot \rrbracket : \pi \rightarrow \pi_a$ being:

- Compositional wrt the operators $\llbracket P \text{ op } Q \rrbracket = \text{Cop}[\llbracket P \rrbracket, \llbracket Q \rrbracket]$
- (Preferably) homomorphic wrt $|$ (distribution-preserving) $\llbracket P | Q \rrbracket = \llbracket P \rrbracket | \llbracket Q \rrbracket$
- Preserving some kind of semantics. Here there are several possibilities
 - Preserving observables $\text{Obs}(P) = \text{Obs}(\llbracket P \rrbracket)$
 - Preserving equivalence

$$\llbracket P \rrbracket \text{ equiv } \llbracket Q \rrbracket \Rightarrow P \text{ equiv}' Q \text{ (soundness)}$$

$$\llbracket P \rrbracket \text{ equiv } \llbracket Q \rrbracket \Leftarrow P \text{ equiv}' Q \text{ (completeness)}$$

$$\llbracket P \rrbracket \text{ equiv } \llbracket Q \rrbracket \Leftrightarrow P \text{ equiv}' Q \text{ (full abstraction, correctness)}$$

Testing semantics

- A **test** O is a process with a distinct success action ω
- A process P **may** pass O iff there is a computation of $[P|O]$ where ω is enabled
eg. $a.b + a$ may pass $\bar{a}.\bar{b}.\omega$
- A process P **must** pass O iff all computations of $[P|O]$ reach a state where ω is enabled
eg. $a.b + a$ must pass $\bar{a}.\omega$

Testing semantics

- $P \sqsubseteq_{\mathbf{may}} Q$ iff $\forall O : P \text{ may } O \Rightarrow Q \text{ may } O$
- $P \sqsubseteq_{\mathbf{must}} Q$ iff $\forall O : P \text{ must } O \Rightarrow Q \text{ must } O$
- Exercise: are $\sqsubseteq_{\mathbf{may}}, \sqsubseteq_{\mathbf{must}}$ pre-congruences for CCS, π ?
- We would like the encodings to satisfy:
 - $P \text{ may pass } O$ iff $\llbracket P \rrbracket \text{ may pass } \llbracket O \rrbracket$
 - $P \text{ must pass } O$ iff $\llbracket P \rrbracket \text{ must pass } \llbracket O \rrbracket$

The encoding of Boudol

Encodes the output prefix (but without choice). Idea: we proceed only when it is sure that the communication can take place, by using a sort of rendez-vous protocol.

- $\llbracket \bar{x}y.P \rrbracket = (\nu z)(\bar{x}z \mid (z(w)(\bar{w}y \mid \llbracket P \rrbracket)))$

- $\llbracket x(y).Q \rrbracket = x(z).(\nu w)(\bar{z}w \mid w(y).\llbracket Q \rrbracket)$

$\llbracket \cdot \rrbracket$ is homomorphic for all the other operators

- $\llbracket 0 \rrbracket = 0$

- $\llbracket P \mid Q \rrbracket = \llbracket P \rrbracket \mid \llbracket Q \rrbracket$

- $\llbracket (\nu x)P \rrbracket = (\nu x)\llbracket P \rrbracket$

- $\llbracket !P \rrbracket = !\llbracket P \rrbracket$

The encoding satisfies P may pass O iff $\llbracket P \rrbracket$ may pass $\llbracket O \rrbracket$

Encoding of Honda-Tokoro

A more compact encoding, it takes two steps instead than three.
The idea is to let the receiver take the initiative.

- $\llbracket \bar{x}y.P \rrbracket = x(z).(\bar{z}y \mid \llbracket P \rrbracket)$
- $\llbracket x(y).Q \rrbracket = (\nu z)(\bar{x}z \mid z(y).\llbracket Q \rrbracket)$

$\llbracket \cdot \rrbracket$ is homomorphic for all the other operators

- $\llbracket 0 \rrbracket = 0$
- $\llbracket P \mid Q \rrbracket = \llbracket P \rrbracket \mid \llbracket Q \rrbracket$
- $\llbracket (\nu x)P \rrbracket = (\nu x)\llbracket P \rrbracket$
- $\llbracket ! P \rrbracket = ! \llbracket P \rrbracket$

The encoding satisfies P may pass O iff $\llbracket P \rrbracket$ may pass $\llbracket O \rrbracket$

Encoding of the output prefix

- The encodings of Boudol and Honda-Tokoro do **not** satisfy P must pass O iff $\llbracket P \rrbracket$ must pass $\llbracket O \rrbracket$
- This is a problem of **fairness**
- must testing is preserved if we restrict to fair computations only
- The encodings preserve a version of testing called “fair must testing”

Encoding of internal choice

The blind choice (or internal choice) construct $P \oplus Q$ has the following semantics

$$\frac{}{P \oplus Q \xrightarrow{\tau} P} \qquad \frac{}{P \oplus Q \xrightarrow{\tau} Q}$$

In π this operator can be represented by the construct $\tau.P + \tau.Q$

Exercise: Let π be π where the $+$ operator can only occur as a blind choice.

Give an encoding $\llbracket \cdot \rrbracket : \pi^{\oplus} \longrightarrow \pi_a$ such that $\forall P \llbracket P \rrbracket \sim P$

Encoding of input-guarded choice

Input-guarded choice is a construct of the form: $\sum_{i \in I} x_i(y_i).P_i$


Let π^i be π where $+$ can only occur in an input-guarded choice. The following encoding of π^i into π_a was defined by Nestmann and Pierce [1996]

$$\llbracket \sum_{i \in I} x_i(y_i)P_i \rrbracket = (\nu l)(\bar{\ell} \text{ true} \mid \prod_{i \in I} \text{Branch}_{\ell_i})$$

$$\text{Branch}_{\ell_i} = x_i(z_i).\ell(w).(if \quad w \\ \text{then } (\bar{\ell} \text{ false} \mid \llbracket P_i \rrbracket) \\ \text{else } (\bar{\ell} \text{ false} \mid \bar{x}_i z_i))$$

Nestmann and Pierce proved that his encoding is fully abstract wrt a notion of equivalence called coupled bisimulation, and it does not introduce divergences.

The π -calculus hierarchy



π_a : asynchronous π

π_{ic} : asynchronous π + input-guarded choice

π_{op} : asynchronous π + output prefix

π_s : asynchronous π + separate choice

π_I : π with internal mobility (Sangiorgi)

CCS_{vp} : value-passing CCS

\longrightarrow : Language inclusion

\longrightarrow (blue) : Encoding

\longrightarrow (pink with slash) : Non-encoding

The separation between π and π_s

This separation result is based on the fact that it is not possible to solve the symmetric leader election problem in π_s , while it is possible in π

Leader Election Problem (LEP): All the nodes of a distributed system must agree on who is the leader. This means that in every possible computation, all the nodes must eventually output the name of the leader on a special channel

- No deadlock
- No livelock
- No conflict (only one leader must be elected, every process outputs its name and only its name)

The separation between π and π_s

Theorem


It is impossible to write in π_s a symmetric (having an automorphism with a single orbit) solution to the LEP.

Crucial point: **Diamond** lemma: when a node P_i performs an action, any other node P_j can perform the same action returning to a symmetric state. (Note: this does not hold in π)

Corollary: in a symmetric π_s network trying to solve the LEP, there is at least one diverging computation.

The separation between π and π_s

Remark: In π (in π with mixed choice) we can easily write a symmetric solution for the LEP in a network of two nodes:



$$P_0 = x.\overline{out} 0 + \bar{y}.\overline{out} 1$$

$$P_1 = y.\overline{out} 1 + \bar{x}.\overline{out} 0$$

The separation between π and π_s

Corollary: there does not exist an encoding of π in π_s which is homomorphic wrt $|$ and renaming, and preserves the observables on every computation.

Proof (sketch): An encoding homomorphic wrt $|$ and renaming transforms a symmetric solution to the LEP in the source language into a symmetric solution to the LEP in the target language.

Outline

Encoding of π -calculus in the asynchronous fragment

Encoding of π -calculus in the probabilistic asynchronous π

Non-deterministic transition systems and CTL

Probabilistic testing semantics

- Test O : same as before (but can be probabilistic)
- $\text{sexec}([P|O])$ the set of **successful executions** of $[P|O]$ (those containing ω)
- Note: $\text{sexec}([P|O])$ can be obtained as a **countable union of disjoint cones**
- $\mu_\sigma(\text{sexec}([P|O]))$ the **probability of success** under scheduler σ

Probabilistic testing semantics

- P **may** pass O iff $\exists \sigma: \mu_{\sigma}(\text{sexec}([P|O])) > 0$
- P **must** pass O iff $\forall \sigma: \mu_{\sigma}(\text{sexec}([P|O])) = 1$
- $\sqsubseteq_{\text{may}}, \sqsubseteq_{\text{must}}$: same as before
 - $P \sqsubseteq_{\text{may}} Q$ iff $\forall O : P \text{ may } O \Rightarrow Q \text{ may } O$
 - $P \sqsubseteq_{\text{must}} Q$ iff $\forall O : P \text{ must } O \Rightarrow Q \text{ must } O$

Probabilistic testing semantics

Exercises:

- Give alternative definitions that consider the exact probabilities of success. Are they equivalent to the ones of the previous slide?
- Show that for all probabilistic CCS processes P, Q :
 - $P +_p Q \sqsubseteq_{\text{may}} \tau.P + \tau.Q$
 - $\tau.P + \tau.Q \sqsubseteq_{\text{must}} P +_p Q$

Encoding of π into π_{ap} (high level idea)

- Similarly to the encoding of Nestmann-Pierce, the branches of the choice are put in parallel together with a lock
- An input process will try to acquire **both its own lock and its partner's lock**
- If a lock is not available it aborts and tries again
- **Crucial** point: the locks are tried in random order, similarly to the solution of the Dining Philosophers problem
- This ensures that **divergences will have probability 0**

Encoding of π into π_{ap} (high level idea)

- This encoding satisfies:
 - P may pass O iff $\llbracket P \rrbracket$ may pass $\llbracket O \rrbracket$
 - P must pass O iff $\llbracket P \rrbracket$ must pass $\llbracket O \rrbracket$
- under a weak assumption on the schedulers (weaker than fairness)
- For details: C. Palamidessi, M. O. Herescu. A randomized encoding of the π -calculus with mixed choice. Theoretical Computer Science. 335(2-3): 373-404, 2005.

Outline

Encoding of π -calculus in the asynchronous fragment

Encoding of π -calculus in the probabilistic asynchronous π


Non-deterministic transition systems and CTL

Model Checking

- Techniques to automatically **verify systems** (software, hardware, protocols, . . .) using **automata** and **temporal logics**
- We give a formal **model** M of the system (typically using some kind of automaton)
- We give a formal **property** φ to verify (typically a formula in some temporal logic)
- An algorithm decides whether **M satisfies the formula φ** , written

$$M \models \varphi$$

Non-deterministic Transition Systems




A tuple (S, s_0, \rightarrow) where

- S is a **finite** set of **states**
- $s_0 \in S$ is the **initial state**
- $\rightarrow \subseteq S \times S$ is a **transition relation**.

We write $s_1 \rightarrow s_2$ when there $(s_1, s_2) \in \rightarrow$

Non-deterministic Transition Systems

Example: a coffee machine



We want to verify properties like “the machine always goes back to its initial state”.

Computation Tree Logic (CTL)

- Formulas are **evaluated** on a transition system M
- On each state s we assign a set $L(s)$ of **atomic propositions**. These are the propositions that we consider to be **true** on this state.
- Two types of formulas:
 - **state** formulas are evaluated on states
 - **path** formulas are evaluated on infinite sequences of states

Computation Tree Logic (CTL)

Syntax:

$\varphi ::= p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \neg \varphi \mid A\psi \mid E\psi$ state formulas

$\psi ::= \circ\varphi \mid \Box\varphi \mid \Diamond\varphi \mid \varphi U\varphi'$ path formulas

Path quantifiers:

$A\psi$ for all paths starting from this state ψ holds

$E\psi$ there exists a path starting from this state s.t. ψ holds

Temporal operators:

$\circ\varphi$ in the next state φ holds ($X\varphi$)


$\Box\varphi$ (always) in all states φ holds ($G\varphi$)

$\Diamond\varphi$ (eventually) in some future state φ holds ($F\varphi$)

$\varphi U\varphi'$ φ holds in all states until a state where φ' holds


Computation Tree Logic (CTL)

finally P




$AF P$

globally P




$AG P$

next P




$AX P$


P until q




$A[P U q]$




$EF P$



$EG P$



$EX P$



$E[P U q]$