MPRI – Course on Concurrency

Lecture 17

Probabilistic asynchronous pi-calculus

Catuscia Palamidessi
INRIA Futurs and LIX
catuscia@lix.polytechnique.fr
www.lix.polytechnique.fr/~catuscia

Page of the course:
http://mpri.master.univ-paris7.fr/C-2-3.html

Motivations

- A language as expressive as the π-calculus with mixed-choice, and suitable for a distributed implementation
  - Used also as an intermediate language for the fully distributed implementation of the π-calculus with mixed-choice. The mixed choice mechanism of the π-calculus cannot be implemented in a fully distributed way deterministically, but can be done in a randomized way. Correctness is achieved with probability 1
- Applications in distributed computing: Some problems can only be solved with the use of randomization
  - Dining Philosophers, Leader Election, ...
- Applications in Security: some protocols use randomization
  - Anonymity: Crowds, Onion Routing, FreeNet, Dining Cryptographers, ...
  - Fair exchange: Contract signing, Non-repudiation, Partial-Secret exchange, ...

Plan of the lecture

- Overview of the basic notions of Probability theory and Measure theory
- Probabilistic automata
- Probabilistic π-calculus
- Examples

Probability and measure theory

References
- Prakash Panangaden, Measure and probability for Concurrency Theorists. *TCS*. 253(2): 287-309. www.lix.polytechnique.fr/~catuscia/teaching/papers_and_books/panangaden.ps
- Prakash Panangaden, Stochastic techniques in Concurrency. Lecture notes. www.lix.polytechnique.fr/~catuscia/teaching/papers_and_books/notes.ps

- Probability in the finite case
  - An experiment with a finite set S of possible results
  - Event: a subset A of S
  - Assuming that all outcomes are equally likely, the probability of event A is defined by pb(A) = $|A| / |S|$
- Example: tossing a fair dice
  - Set of possible results $S = \{ 1,2,3,4,5,6 \}$
  - Event “the result is even”: $A = \{ 2,4,6 \}$, pb(A) = 1/2
  - Event “the result is at least 5” : $B = \{ 5,6 \}$, pb(B) = 1/3

Probability and measure theory

- Example: tossing 3 three times a fair coin
  - Event “all coins are H”:
    pb(HHH) = p(H--) x pb(-H-) x pb(--H) independent
    = 1/2 x 1/2 x 1/2 = 1/8
  - “all coins are H or all coins are T” :
    pb(HHH or TTT) = pb(HHH) + pb(TTT) = 1/4 disjoint
  - “at least one coin is H”:
    pb(not TTT) = 1 – pb(TTT) = 7/8
  - “at least one coin is H and at least one coin is T”
    pb(not TTT and not HHH) = 6/8 not independent
  - “at least one coin is H or at least one coin is T”
    pb(not TTT or not HHH) not disjoint
    = pb(not TTT) + pb(not HHH) – pb(not TTT and not HHH)
    = 7/8 + 7/8 – 6/8 = 1

Probability and measure theory

- The need for measure theory
  - Some “experiments” have infinitary nature
  - Example: tossing infinitely many time a fair coin
    - The set S of all infinite sequences of H/T is infinite (uncountable)
    - The probability of each sequence is 0, so we cannot expect that the single result will be enough as “building blocks” (i.e. we cannot expect to be able to define the probability of every event by summing up the probability of the singletones)
    - When *S* is uncountable, we cannot expect of being able to define the probability of every set of results.

Probability and measure theory

We would like to preserve the laws of the finite case, notaby:

- Countable union of disjoint sets

$$pb(\bigcup_{i\in I} A_i) = \sum_{i\in I} pb(A_i)$$

if I is countable and $\forall i, j. i \neq j \Rightarrow A_i \cap A_j = \emptyset$

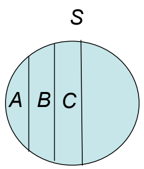

- Complementation:

$$pb(\bar{A}) = 1 - pb(A)$$

Probability and measure theory

Definition: A measurable space is a pair (S, Σ) where S is a set and Σ is a family of subsets of S, called σ-field or σ-algebra, satisfying the following axioms:

- $\emptyset \in \Sigma$
- if $A \in \Sigma$ then $\bar{A} \in \Sigma$
- if I is countable and $\forall i \in I.\ A_i \in \Sigma$, then $\bigcup_{i\in I} A_i \in \Sigma$

Probability and measure theory

Proposition The intersection of an arbitrary collection of σ-fields on S is a σ-field on S
Proof: Exercise

Corollary Given a set S and a set $\mathcal{B}$ of subsets of S, there exists a least σ-field $(S, \Sigma_{\mathcal{B}})$ containing $\mathcal{B}$

We call $\mathcal{B}$ *base* of $(S, \Sigma_{\mathcal{B}})$ and we say that $(S, \Sigma_{\mathcal{B}})$ *is generated by* $\mathcal{B}$

Probability and measure theory

Examples Given a set S

The set of all subsets of S, $\mathcal{P}(\mathcal{S})$, is a σ-field

The set $\{\emptyset, S\}$ is a σ-field

If S is countable, the σ-field generated by the singletons is $\mathcal{P}(\mathcal{S})$

If S is uncountable, the σ-field generated by the singletons is the set of all countable and cocountable elements of $\mathcal{P}(\mathcal{S})$

Probability and measure theory

Definition (Measure) Given a measurable space (S, Σ), a measure is a function $\mu : \sigma \to [0, \infty]$ such that for every countable I and for every family $\{A_i \in \Sigma\}_{i \in I}$, if $\forall i, j \in I.\ A_i \cap A_j = \emptyset$, then

$$\mu(\bigcup_{i \in I} A_i) = \Sigma_{i \in I} \mu(A_i)$$

Definition (Probability measure) A probability measure is like a measure (the elements of Σ are called *events*) with the difference that $\mu : \sigma \to [0, 1]$ and that we additionally require

$$\mu(S) = 1$$

Proposition $\mu(\bar{A}) = 1 - \mu(A)$

Probability and measure theory

Proposition Given a measurable space $(S, \Sigma_{\mathcal{B}})$ generated by a base $\mathcal{B}$ containing S, and given $f : \mathcal{B} \to [0, \infty]$ which satisfies the countable disjoint union property, there exists a unique measure $\mu_f : \Sigma_{\mathcal{B}} \to [0, \infty]$ which coincides with f on the elements of $\mathcal{B}$.

We say that μ_f *is induced by* f. μ_f can be constructed inductively from f in the same way as $(S, \Sigma_{\mathcal{B}})$ can be constructed from $\mathcal{B}$.

We have a similar reslt for probability measure, except that we require also $f(S) = 1$.

Probability and measure theory

Example

- S = the set of all infinite sequences of a fair coin tossing
- From elementary finite probability theory: Each finite sequence of H/T $x_0, x_1, \ldots, x_{n-1}$ (x_i=H or x_i=T) has probability $1/2^n$ (independence)
- Each infinite sequence has probability 0
- **Cone**: given a sequence s = $x_0, x_1, \ldots, x_{n-1}$ [,…], the set A of all sequences which have s as prefix is called **cone**
  - We assign to A the probability measure of its prefix: $m(A) = 1/2^n$
- Define B (base) as the set of all cones. Note that they are All disjoint

Probability and measure theory

Consider the space (S,S_{B}) generated by S and the set B of all cones, with probability measure induced by m

What is the probability that a sequence has infinitely many H?

Probability of exactly one H in any position: 0 (countable disjoint union of sets with measure 0)

Probability of exactly n H in any position: 0 (same reason)

Probability of finitely many H in any position: 0 (same reason)

Probability of infinitely many H:

1 – pb(finitely many H) = 1 – 0 (complementation property)

The probabilistic π-calculus

References:

- Roberto Segala and Nancy Lynch. Probabilistic simulations for probabilistic processes. Nordic Journal of Computing, 2(2):250--273, 1995.
- C. Palamidessi, O.M. Herescu. A Randomized Distributed Encoding of the π-Calculus with Mixed Choice. Theoretical Computer Science 335(2-3): 373-404, 2005 www.lix.polytechnique.fr/~catuscia/papers/prob_enc/report.ps

The probabilistic π-calculus: syntax

Similar to the asynchronous p-calculus of Amadio,Castellani and Sangiorgi, the only difference is that the input-guarded choice is probabilistic

The prefixes

$$\alpha \ ::= \ x(y) \mid \tau \qquad \text{input | silent action}$$

The processes

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

where $\Sigma_i p_i = 1$

The probabilistic π-calculus: operational sem

- Based on the probabilistic automata of Segala and Lynch
- nondeterministic and probabilistic behavior
- nondeterminism associated to a scheduler (adversary)
- probabilistic behavior associated to the choice of the process
  - groups, probabilitic distributions, steps

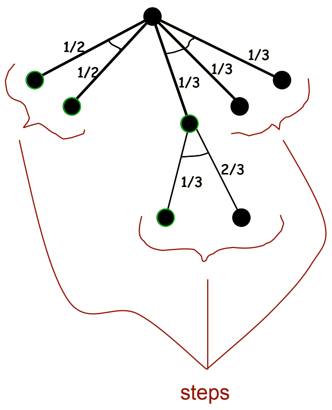

The probabilistic π-calculus: operational sem

Steps will b represented as follows:

$$P\ \{\xrightarrow[p_i]{\mu_i} P_i \mid i \in I\}$$

where the μ_i's are actions (n the sense of the π-calculus) and $\Sigma_{i \in I}\, p_i = 1$

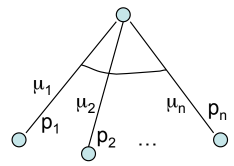

The probabilistic π-calculus: operational sem

The steps are defined in a SOS style as follows

Sum $\sum_i p_i \alpha_i.P_i\ \{\xrightarrow[p_i]{\alpha_i} P_i\}_i$

Out $\bar{x}y\ \{\xrightarrow[1]{\bar{x}y} \mathbf{0}\}$

Res $$\frac{P\ \{\xrightarrow[p_i]{\mu_i} P_i\}_i}{\nu y P\ \{\xrightarrow[p_i']{\mu_i} \nu y P_i\}_{i : y \not\in fn(\mu_i)}}$$

$\exists i.\ y \not\in fn(\mu_i)$ and
$\forall i.\ p_i' = p_i / \sum_{j : y \not\in fn(\mu_j)} p_j$

The probabilistic π-calculus: operational sem

Open $$\frac{P\ \{\xrightarrow[1]{\bar{x}y} P'\}}{\nu y P\ \{\xrightarrow[1]{\bar{x}(y)} P'\}}\quad x \neq y$$

Par $$\frac{P\ \{\xrightarrow[p_i]{\mu_i} P_i\}_i}{P \mid Q\ \{\xrightarrow[p_i]{\mu_i} P_i \mid Q\}_i}$$

Com $$\frac{P\ \{\xrightarrow[1]{\bar{x}y} P'\} \qquad Q\ \{\xrightarrow[p_i]{\mu_i} Q_i\}_i}{P \mid Q\ \{\xrightarrow[p_i]{\tau} P' \mid Q_i[y/z_i]\}_{i : \mu_i = x(z_i)} \cup \{\xrightarrow[p_i]{\mu_i} P \mid Q_i\}_{i : \mu_i \neq x(z_i)}}$$

The probabilistic π-calculus: operational sem

$$\text{Close}\quad \frac{P\ \{\xrightarrow[1]{\bar{x}(y)} P'\} \qquad Q\ \{\xrightarrow[p_i]{\mu_i} Q_i\}_i}{P \mid Q\ \{\xrightarrow[p_i]{\tau} \nu y(P' \mid Q_i[y/z_i])\}_{i:\mu_i = x(z_i)} \cup \{\xrightarrow[p_i]{\mu_i} P \mid Q_i\}_{i:\mu_i \neq x(z_i)}}$$

$$\text{Cong}\quad \frac{P \equiv P' \qquad P'\ \{\xrightarrow[p_i]{\mu_i} Q'_i\}_i \qquad \forall i.\ Q'_i \equiv Q_i}{P\ \{\xrightarrow[p_i]{\mu_i} Q_i\}_i}$$

The probabilistic π-calculus: operational sem

Structural congruence:

We assume that structural congruence satisfies the standard rules: associative monoid rules for $|$, the commutativity of the summands for Σ, the alpha-conversion, and the following:

$(\nu x\, P) \mid Q \equiv \nu x\, (P \mid Q)$ if $x \notin fn(Q)$

$!\, P = P \mid !\, P.$

Example: D.C. in the probabilistic asynchronous π-calculus

$$\boxed{\begin{aligned} Master &= \textstyle\sum_{i=0}^{2} \tau\,.\,\overline{m}_i \mathsf{p}\,.\,\overline{m}_{i\oplus 1}\mathsf{n}\,.\,\overline{m}_{i\oplus 2}\mathsf{n}\,.\,0 \\ &+\ \ \tau.\overline{m}_0\mathsf{n}\,.\,\overline{m}_1\mathsf{n}\,.\,\overline{m}_2\mathsf{n}\,.\,0 \end{aligned}}$$

Nondeterministic choice

$$\begin{aligned} Crypt_i &= m_i(x)\,.\,c_{i,i}(y)\,.\,c_{i,i\oplus 1}(z)\,. \\ &\qquad \text{if } x = \mathsf{p} \\ &\qquad\quad \text{then } \overline{pay}_i\,.\,\text{if } y = z \\ &\qquad\qquad \text{then } \overline{out}_i\, disagree \\ &\qquad\qquad \text{else } \overline{out}_i\, agree \\ &\qquad\quad \text{else if } y = z \\ &\qquad\qquad \text{then } \overline{out}_i\, agree \\ &\qquad\qquad \text{else } \overline{out}_i\, disagree \end{aligned}$$

Anonymous actions

Observables

$$\boxed{Coin_i = p_h\tau\,.\,Head_i + p_t\tau\,.\,Tail_i}$$

Probabilistic choice

$$\begin{aligned} Head_i &= \overline{c}_{i,i} head\,.\,\overline{c}_{i\ominus 1,i} head\,.\,0 \\ Tail_i &= \overline{c}_{i,i} tail\,.\,\overline{c}_{i\ominus 1,i} tail\,.\,0 \\ DCP &= (\nu\vec{m})(Master \\ &\qquad \mid\ (\nu\vec{c})(\Pi_{i=0}^{2} Crypt_i\ \mid\ \Pi_{i=0}^{2} Coin_i)\,) \end{aligned}$$

Probabilistic automaton associated to the probabilistic π program for the D.C.