

Concurrency theory

types to reason about processes: subtyping, receptiveness
from process calculi to programming languages
some slides to navigate through the literature

Francesco Zappa Nardelli

INRIA Paris-Rocquencourt, MOSCOVA research team

`francesco.zappa_nardelli@inria.fr`

together with

Frank Valencia (INRIA Futurs)



Roberto Amadio (PPS)



Emmanuel Haucourt (CEA)



Simple types for pi-calculus

Objective: avoid run-time errors

$$\bar{x}\langle\text{true}\rangle.P \parallel x(y).\bar{y}\langle 4 \rangle \rightarrow P \parallel \overline{\text{true}}\langle 4 \rangle \rightarrow \text{error}$$

Ideas: associate a type to each channel: $x : \text{ch}(\text{bool})$. The free names of each process are stored in an **environment** Γ , that represents a *contract* between the process and the environment about the use of the channels. Processes willing to interact *must agree* on the contract.

$$\begin{array}{ll} \bar{x}\langle\text{true}\rangle.P & \Gamma_1 = x : \text{ch}(\text{bool}) \\ x(y).\bar{y}\langle 4 \rangle & \Gamma_2 = y : \text{ch}(\text{int}); x : \text{ch}(\text{ch}(\text{int})) \end{array}$$

Γ_1 and Γ_2 disagree on the type of x : the process above is not well typed.

Simply-typed pi-calculus: the type rules (excerpt)

$\Gamma \vdash M : T$ value M has type T under the type assignment for names Γ ;

$$\begin{array}{c} \mathfrak{z} : \text{int} \\ \frac{\Gamma(x) = T}{\Gamma \vdash x : T} \quad \frac{\Gamma \vdash M_1 : T_1 \quad \Gamma \vdash M_2 : T_2}{\Gamma \vdash (M_1, M_2) : T_1 \times T_2} \end{array}$$

$\Gamma \vdash P$ process P respects the type assignment for names

$$\begin{array}{c} \Gamma \vdash \mathbf{0} \\ \frac{\Gamma \vdash P_1 \quad \Gamma \vdash P_2}{\Gamma \vdash P_1 \parallel P_2} \quad \frac{\Gamma, x:T \vdash P}{\Gamma \vdash (\nu x : T)P} \\ \frac{\Gamma \vdash x : \text{ch}(T) \quad \Gamma, y:T \vdash P}{\Gamma \vdash x(y : T).P} \quad \frac{\Gamma \vdash x : \text{ch}(T) \quad \Gamma \vdash M : T \quad \Gamma \vdash P}{\Gamma \vdash \bar{x}\langle M \rangle.P} \end{array}$$

Soundness

Extend the syntax with the `wrong` process, and add reduction rules to capture runtime errors:

$$\frac{\text{where } x \text{ is not a name}}{\bar{x}\langle M \rangle.P \xrightarrow{\tau} \text{wrong}}$$

$$\frac{\text{where } x \text{ is not a name}}{x(y:T).P \xrightarrow{\tau} \text{wrong}}$$

1. Prove that if $\Gamma \vdash P$, with Γ closed, and $P \rightarrow^* P'$, then P' does not have `wrong` as a subterm.
2. If P has `wrong` as a subterm, then P is not typable.

Types to reason about processes

Specification and an implementation of the factorial function:

Spec = $!f(x, r).\bar{r}\langle \text{fact}(x) \rangle$

Imp = $!f(x, r).\text{if } x = 0 \text{ then } \bar{r}\langle 1 \rangle \text{ else } (\nu r')\bar{f}\langle x - 1, r' \rangle.r'(m).\bar{r}\langle x * m \rangle$

In general, Spec $\not\approx$ Imp. (Why?)

Idea: the channel f should be *write-only* for the environment.

Subtyping

Idea: refine the type of channels $\text{ch}(T)$ into

$i(T)$	input (read) capability
$o(T)$	output (write) capability

Example: a context that interacts with the term

$$x : \text{ch}(o(T)) \vdash (\nu f : \text{ch}(T)) \bar{x}\langle f \rangle.\text{Imp}$$

and respects the environment $x : \text{ch}(o(T))$ can only write at f .

Problem: f has type $\text{ch}(T)$, not $o(T)$.

However, in some sense $\text{ch}(T)$ and $o(T)$ are *compatible*.

Subtyping

We say that a type T_1 is a *subtype* of a type T_2 , denoted $T_1 <: T_2$ if it is safe to use a value of type T_1 in all contexts that expect a value of type T_2 . For instance:

- $\text{int} <: \text{float}$: the value 3 can safely be used in the context $\text{sqr}(-)$.
- $\text{float} \not<: \text{int}$: the value 3.14 cannot be used in the context $- \bmod 3$ because mod is not defined over floats.
- $\{a:\text{int}; b:\text{int}\} <: \{a : \text{int}\}$: a context accepting records containing the label a will just ignore the existence of the label b .
- $\text{ch}(T) <: \text{o}(T)$: a context expecting a channel of type $\text{o}(T)$ will just ignore the read capability available.

The subtyping relation, formally

– *is a preorder*

$$T <: T$$
$$\frac{T_1 <: T_2 \quad T_2 <: T_3}{T_1 <: T_3}$$

– *capabilities can be forgotten*

$$\text{ch}(T) <: \text{i}(T) \qquad \text{ch}(T) <: \text{o}(T)$$

– *i is a covariant type constructor, o is contravariant, ch is invariant*

$$\frac{T_1 <: T_2}{\text{i}(T_1) <: \text{i}(T_2)} \qquad \frac{T_2 <: T_1}{\text{o}(T_1) <: \text{o}(T_2)} \qquad \frac{T_2 <: T_1 \quad T_1 <: T_2}{\text{ch}(T_1) <: \text{ch}(T_2)}$$

Subtyping, ctd.

Intuition: if $x : o(T)$ then it is safe to send along x values of of a subtype of T .
Dually, if $x : i(T)$ then it is safe to assume to assume that values received along x belong to a supertype of T .

Type rules must be updated as follows:

$$\frac{\Gamma \vdash x : i(T) \quad \Gamma, y:T \vdash P}{\Gamma \vdash x(y : T).P}$$

$$\frac{\Gamma \vdash x : o(T) \quad \Gamma \vdash M : T \quad \Gamma \vdash P}{\Gamma \vdash \bar{x}\langle M \rangle.P}$$

$$\frac{\Gamma \vdash M : T_1 \quad T_1 <: T_2}{\Gamma \vdash M : T_2}$$

Exercises

Show that:

1. $a : \text{ch}(\text{int}), b : \text{ch}(\text{real}) \vdash \bar{a}\langle 5 \rangle \parallel a(x).\bar{b}\langle x \rangle$, assuming $\text{int} <: \text{real}$;
2. $x : \text{ch}(\text{o}(T)) \vdash (\nu y:\text{ch}(T))(\bar{x}\langle y \rangle.!\!y(z:T))$
3. $x : \text{ch}(\text{o}(T)), z : \text{ch}(\text{i}(T)) \vdash (\nu y:\text{ch}(T))(\bar{x}\langle y \rangle \parallel \bar{z}\langle y \rangle)$
4. $b : \text{ch}(S), x : \text{ch}(\text{i}(S)), a : \text{ch}(\text{o}(\text{i}(S))) \vdash \bar{a}\langle x \rangle \parallel x(y).y(z) \parallel a(x).\bar{x}\langle b \rangle$

Remarks on i/o types

– *different processes may have different visibility of a name:*

$$\begin{aligned} & (\nu x:\text{ch}(T)) \bar{y}\langle x \rangle. \bar{z}\langle x \rangle. P \parallel y(a:\text{i}(T)). Q \parallel z(b:\text{o}(T)). R \quad \rightarrow \rightarrow \\ & (\nu x:\text{ch}(T)) (P \parallel Q\{x/a\} \parallel R\{x/b\}) \end{aligned}$$

Q can only read from x , R can only write to x .

– *acquiring the o and i capabilities on a name is different from acquiring ch:*
the term

$$(\nu x:\text{ch}(\text{unit})) \bar{y}\langle x \rangle. \bar{z}\langle x \rangle \parallel y(a:\text{i}(\text{unit})). z(b:\text{o}(\text{unit})). \bar{a}\langle \rangle$$

is not well-typed.

Types for reasoning

Types can be seen as *contracts* between a process and its environment: the environment *must respect* the constraints imposed by the typing discipline.

In turn, *types reduce the number of legal contexts* (and give us more process equalities).

Example: an observer whose typing is

$$\Gamma = a : \text{o}(T), b : T, c : T' \qquad T \text{ and } T' \text{ unrelated}$$

- can offer an output $\bar{a}\langle b \rangle$;
- cannot offer an output $\bar{a}\langle c \rangle$, or an input at a .

A typed contextual equivalence, informally

Definition (informal): The processes P and Q are equivalent in Γ , denoted

$$P \cong_{\Gamma} Q$$

iff $\Gamma \vdash P, Q$ and they are equivalent in all the testing contexts that respect the types in Γ .

To formalize this equivalence we need to type contexts: a context $C[-]$ is a Δ/Γ -context if $\Delta \vdash C$ is a valid type judgement when the hole is typed as

$$\frac{\Omega \text{ extends } \Gamma}{\Omega \vdash -}$$

Main property: if C is a Δ/Γ -context and $\Gamma \vdash P$ then $\Delta \vdash C[P]$.

Semantic consequences of i/o types

Example: the processes

$$\begin{aligned} P &= (\nu x)\bar{a}\langle x\rangle.\bar{x}\langle\rangle \\ Q &= (\nu x)\bar{a}\langle x\rangle.\mathbf{0} \end{aligned}$$

are different in the untyped or simply-typed pi-calculus.

With i/o types, it holds that

$$P \cong_{\Gamma} Q \quad \text{for } \Gamma = a : \text{ch}(\text{o}(\text{unit}))$$

because the residual $\bar{x}\langle\rangle$ of P is deadlocked (the context cannot read from x).

Semantic consequences of i/o types, ctd.

Back to the factorial function:

Spec = $!f(x, r).\bar{r}\langle \text{fact}(x) \rangle$

Imp = $!f(x, r).\text{if } x = 0 \text{ then } \bar{r}\langle 1 \rangle \text{ else } (\nu r')\bar{f}\langle x - 1, r' \rangle.r'(m).\bar{r}\langle x * m \rangle$

With i/o types, we can protect the input end of the function, obtaining

$$(\nu f)\bar{a}\langle f \rangle.\text{Spec} \cong_{\Gamma} (\nu f)\bar{a}\langle f \rangle.\text{Imp}$$

for $\Gamma = a : \text{ch}(\text{o}(\text{int}) \times \text{o}(\text{int}))$.

Which is the type that I omitted after the restriction of r' ?

Semantic consequences of i/o types, ctd.

$$P = (\nu x, y)(\bar{a}\langle x \rangle \parallel \bar{a}\langle y \rangle \parallel !x().R \parallel !y().R)$$

$$Q = (\nu x)(\bar{a}\langle x \rangle \parallel \bar{a}\langle x \rangle \parallel !x().R)$$

In the untyped calculus $P \not\approx Q$: a context that tells them apart is

$$- \parallel a(z_1).a(z_2).(z_1().\bar{c}\langle \rangle \parallel \bar{z_2}\langle \rangle) .$$

With i/o types

$$P \cong_{\Gamma} Q \quad \text{for } \Gamma = a : \text{ch}(\text{o}(\text{unit})) .$$

Notation: I will often omit redundant type informations.

Challenge: a labelled equivalence

Problem: only a subset of the actions of the tested process is observable; in general the typings of the observer and of the tested process do not coincide.

Example: in an initial type environment $\Gamma = a : \text{ch}(\text{o}(T))$, we have

$$(\nu b : \text{ch}(T))\bar{a}\langle b \rangle.P \xrightarrow{(\nu b : \text{ch}(T))\bar{a}\langle b \rangle} P .$$

The final typing for P is $\Gamma, b : \text{ch}(T)$, while for the observer it is $\Gamma, b : \text{o}(T)$.

Solution: separate the observer and the process points of view on the types of the names.

The typings should however be typewise compatible (that is, the types of the same name should have a common subtype).

Challenge: a labelled equivalence, ctd.

Problem: aliasing! Going back to the example on slide 15:

$$P = (\nu x, y)(\bar{a}\langle x \rangle \parallel \bar{a}\langle y \rangle \parallel !x().R \parallel !y().R)$$

$$Q = (\nu x)(\bar{a}\langle x \rangle \parallel \bar{a}\langle x \rangle \parallel !x().R)$$

It holds $P \cong_{\Gamma} Q$ for $\Gamma = a : \text{ch}(\text{o}(\text{unit}))$. Consider this interaction with a tester:

$$a(p).a(q).R \parallel P \rightarrow^* R\{x/p\}\{y/q\} \parallel P'$$

$$\underbrace{a(p).a(q).R}_{\text{the tester}} \parallel Q \rightarrow^* R\{x/p\}\{x/q\} \parallel Q'$$

At the beginning, the tester/observer is the same but after two reductions the observers are different. *In the bisimulation game the observer should be unique.*

Aliasing, ctd.

In other terms, equivalent terms may realise different sequence of transitions:

$$(\nu x, y)(\bar{a}\langle x \rangle \parallel \bar{a}\langle y \rangle \parallel !x().R \parallel !y().R) \xrightarrow{(\nu x)\bar{a}\langle x \rangle} \xrightarrow{(\nu y)\bar{a}\langle y \rangle} \xrightarrow{y(v)} \dots$$

must be matched by the sequence

$$(\nu x)(\bar{a}\langle x \rangle \parallel \bar{a}\langle x \rangle \parallel !x().R) \xrightarrow{(\nu x)\bar{a}\langle x \rangle} \xrightarrow{\bar{a}\langle x \rangle} \xrightarrow{x(v)} \dots$$

Solution: separate the observer's view on the identity of names from their real identity.

Digression

Aliasing is not related to types. The same problem arises if we consider a dialect of pi-calculus *without matching*.

In fact, the transition $P \xrightarrow{\bar{x}\langle y \rangle} P'$ can be read as:

1. P is interacting with a context ready to receive a name over x ;
2. there exists a context that can test that P is ready to send the name y over x .

Interpretation 2. requires that the context can test if the name sent over x is y or not (eg, matching is required). *Completeness of the standard bisimulation wrt contextual equivalence is lost if matching is omitted.*

Exercise

1. Extend the syntax, the reduction semantics, and the type rules of pi-calculus with i/o types with the nondeterministic sum operator, denoted $+$;
2. Show that the terms

$$P = \bar{b}\langle x \rangle . a(y) . (y() \mid \bar{x}\langle \rangle)$$

$$Q = \bar{b}\langle x \rangle . a(y) . (y() . \bar{x}\langle \rangle + \bar{x}\langle \rangle . y())$$

are not equivalent in the untyped calculus. Propose a i/o typing such that $P \simeq_{\Gamma} Q$.

Receptiveness

A local environment:

$$(\nu x)(!x(y).R \parallel Q)$$

(Q and R have only the output capability on x).

The name x is **stateless**, and its input end is always available: x is *uniformly receptive*.

Uniform receptiveness occurs when modelling functions, objects, RPC protocols, etc. It is a common idiom in programming languages based on pi-calculus (Pict, Join, Blue): `def $x(y) = R$ in Q` . Variant: *linear receptiveness*: x is used at most once.

It is possible to impose receptiveness using *syntactical constraints*, or a *type system*.

Some properties of receptiveness

If x is *uniformly receptive* in Γ , then it holds that

$$\bar{y}\langle x \rangle.P \approx_{\Gamma} (\nu z)\bar{y}\langle z \rangle.(!z(a).\bar{x}\langle a \rangle \parallel P)$$

for z fresh.

(an output of a global name becomes an output of a fresh name)

Not true in the ordinary pi-calculus because of contexts like

$$y(u).\bar{u}\langle c \rangle.\bar{p}\langle \rangle \parallel -$$

But this context is not receptive on u .

Some properties of receptiveness, ctd.

If x is *uniformly receptive* in Γ , then it holds that

$$\bar{x}\langle v \rangle.P \approx_{\Gamma} \bar{x}\langle v \rangle \parallel P$$

(makes a synchronous communication into an asynchronous one).

If $P \rightarrow P'$ by means of a communication at x , then

$$P \approx_{\Gamma} P'$$

(insensitiveness to internal reductions)

Remark: all names uniformly receptive \Rightarrow **confluence**.

Factorial function again

$$S1 = !f(x:\text{int}, y:\text{int}, r:\text{lrec}(\text{int})). \text{ if } x = 0 \text{ then } \bar{r}\langle y \rangle \\ \text{ else } (\nu r')(\bar{f}\langle x - 1, x * y, r' \rangle.r'(m).\bar{r}\langle m \rangle)$$
$$S2 = !f(x:\text{int}, y:\text{int}, r:\text{lrec}(\text{int})). \text{ if } x = 0 \text{ then } \bar{r}\langle y \rangle \\ \text{ else } \bar{f}\langle x - 1, x * y, r \rangle$$

Exercise: Give a (non-well typed) context that shows that S_1 and S_2 are not equivalent.

However, taking into account the receptiveness of f, r and r' , S_1 and S_2 are behaviourally equivalent (*tail-call optimisation*).

References

Milner: *The polyadic pi-calculus - a tutorial*, ECS-LFCS-91-180.

Pierce, Sangiorgi: *Typing and subtyping for mobile processes*, LICS '93.

Boreale, Sangiorgi: *Bisimulation in name-passing calculi without matching*, LICS '98.

Sangiorgi, Walker: *The pi-calculus*, CUP.

...there is a large literature on the subject. The articles above have been reported because they are explicitly mentioned in this lecture.

From process languages to programming languages

Implementations of the pi-calculus semantics include:

- Pict : statically typed programming language based on the pi-calculus (mostly local computation)
- Nomadic Pict: communication is local to each runtime + migration between runtimes

Implementing in a distributed setting the pi-calculus semantics of channels requires solving the *distributed consensus*.

Erlang

An dynamically typed language *a la Lisp*;

each thread has a unique identifier;

each thread has *one* channel, identified by the thread id;

```
send: id!msg
```

```
receive:
```

```
receive
```

```
    patt1 -> action1
```

```
    patt2 -> action2
```

```
    ...
```

```
end.
```

The join calculus

The join-idea: receptors and channels defined at the same time.

In informal syntax:

```
let C1() | C2() = P1
or  C2()          = P2
```

defines channels C1, C2, and reactions P1 and P2.

Consequence: all receptors are known *statically*.

- Synchronizations are solved locally;
- Static resolution of many problems (automata for synchronization, implicit polymorphic typing, etc.)

Join as a process language

The definition

```
let Count(n) | Tick() = Count(n+1)
or  Count(n) | Show() = Count(n) | Print(n)
```

translates to rules

```
Count(n) || Tick()  ▷  Count(n+1)
Count(n) || Show()  ▷  Count(n) || Print(n)
```

The semantics performs rewriting modulo equivalence (similarly to what we did when we defined the reduction semantics using structural equivalence):

$$\text{Count}(2) \parallel \text{Tick}() \parallel \text{Show}() \longrightarrow \text{Count}(3) \parallel \text{Show}()$$

(the new) JoCaml

- An extension of OCaml 3.10;
- distributed asynchronous channels (synchronous channels by CPS);
- polymorphic typing à la ML;
- easy encoding of concurrent programming constructs.

For instance, what does the code below do?

```
type 'a buffer = { get : unit -> 'a ; put : 'a -> unit ; }
```

```
let create_buff () =  
  def some(v) & get() = none() & reply v to get  
  or none() & put(v) = some(v) & reply () to put in  
  spawn none() ;  
  { get = get ; put = put ; }
```

A more useful buffer

A classic algorithm that represents state by a pair of lists.

```
def state(xs,ys) & put(x) =  
    state(x::xs,ys) & reply () to put  
or state(xs,y::ys) & get() =  
    state(xs,ys) & reply y to get  
or state(_::_ as xs,[]) & get() =  
    state([], List.rev xs) & reply get() to get
```

Try it:

<http://jocaml.inria.fr>

Navigating through the literature

Pi-calculus literature describes **zillions** of slightly different languages, semantics, equivalencies.

Some slides for not getting lost.

Barbed congruence vs. reduction-closed barbed congruence

Let *barbed equivalence*, denoted \cong^\bullet , be the largest symmetric relation that is barb preserving and reduction closed. Barbed equivalence is not preserved by context, so define *barbed congruence*, denoted \cong^c , as

$$\{(P, Q) : C[P] \cong^\bullet C[Q] \text{ for every context } C[-].\}$$

- Barbed congruence is *more natural* and *less discriminating* than reduction-closed barbed congruence (for pi-calculus processes).
- Completeness of bisimulation for image-finite processes holds with respect to barbed congruence, but its proof requires transfinite induction.

Late bisimulation

Change the definition of the LTS:

$$x(y).P \xrightarrow{x(y)} P \qquad \frac{P \xrightarrow{\bar{x}\langle v \rangle} P' \quad Q \xrightarrow{x(y)} Q'}{P \parallel Q \xrightarrow{\tau} P' \parallel Q'\{v/y\}}$$

and extend the definition of bisimulation with the clause: if $P \approx_l Q$ and $P \xrightarrow{x(y)} P'$, then there is Q' such that $Q \xrightarrow{x(y)} Q'$ and for all v it holds $P'\{v/y\} \approx_l Q'\{v/y\}$.

- Late bisimulation differs (slightly) from (early) bisimulation. More importantly, *the label $x(y)$ does not denote an interacting context.*

Ground bisimulation

Idea: play a standard bisimulation on the late LTS. Or,

Let *ground bisimulation* be the largest symmetric relation, \approx_g , such that whenever $P \approx_g Q$, there is $z \notin \text{fn}(P, Q)$ such that if $P \xrightarrow{\alpha} P'$ where α is $\bar{x}\langle y \rangle$ or $x(z)$ or $(\nu z)\bar{x}\langle z \rangle$ or τ , then $Q \xrightarrow{\hat{\alpha}} \approx_g P'$.

Contrast it with bisimilarity: to establish $x(z).P \approx x(z).Q$ it is necessary to show that $P\{v/z\} \approx Q\{v/z\}$ for all v . Ground bisimulation requires to test only *a single, fresh, name*.

However, ground bisimilarity is less discriminating than bisimilarity, and it is not preserved by composition (still, it is a reasonable equivalence for sublanguages of pi-calculus).

Open bisimulation

Full bisimilarity is the closure of bisimilarity under substitutions, and is a congruence with respect to all contexts. Unfortunately, full bisimilarity is not defined co-inductively.

Question: can we give a co-inductive definition of a useful congruence?

Yes, with *open bisimulation*.

Idea: (on the restriction free calculus) let \bowtie be the largest symmetric relation such that whenever $P \bowtie Q$ and σ is a substitution, $P\sigma \xrightarrow{\alpha} P'$ implies $Q\sigma \xrightarrow{\hat{\alpha}} \bowtie P'$.

It is possible to avoid the σ quantification by means of an appropriate LTS.

Subcalculi

Idea: In pi-calculus contexts have a great discriminating power. It may be useful to consider other languages in which contexts "*observe less*", so that we have *more equations*.

Asynchronous pi-calculus: no continuation after an output prefix.

Localized pi-calculus: given $x(y).P$, the name y is not used as subject of an input prefix in P .

Private pi-calculus: only output of new names.

Distribution, action at distance, and mobility

The parallel composition operator of CCS and pi-calculus does not specify whether the concurrent threads are running on the same machine, or on different machines connected by a network.

Some phenomena typical of distributed systems require a finer model, that explicitly keeps track of the spatial distribution of the processes.

We will briefly sketch two models that have been proposed: *DPI* (Hennessy and Riely, 1998) and *Mobile Ambients* (Cardelli and Gordon, 1998).

The aim of this section is to get a glimpse of more complex process languages, and to rediscover the idea of “transitions in an LTS characterise the interactions a term can have with a context” in this setting.

DPI, design choices

- add explicit locations to pi-calculus processes: $\ell[P]$;
- locations are identified by their name: $\ell[P] \parallel \ell[Q] \equiv \ell[P \parallel Q]$;
- communication is local to a location:

$$\ell[\bar{x}\langle y \rangle.P] \parallel \ell[x(u).Q] \rightarrow \ell[P] \parallel \ell[Q\{y/u\}] ;$$

- add explicit migration: $\ell[\text{goto } k.P] \rightarrow k[P]$.

We also include the restriction and match operators, subject to the usual pi-calculus semantics.

Behavioural equivalence for DPI

Again, we apply the standard recipe:

- define the suitable contexts:

$$C[-] ::= - \mid C[-] \parallel \ell[P] \mid (\nu n)C[-] .$$

- define the observation:

$$M \downarrow x@l \text{ iff } P \equiv (\nu \tilde{n})(\ell[x(u).P'] \parallel P'') \text{ for } x, l \notin \tilde{n} .$$

Can we characterise this equivalence with a labelled bisimulation?

Labelled bisimulation for DPI

$$\frac{P \rightarrow P'}{P \xrightarrow{\tau} P'} \quad \frac{P \equiv (\nu \tilde{n})(\ell \llbracket x(u).P' \rrbracket \parallel P'') \quad x, \ell \notin \tilde{n}}{P \xrightarrow{x(y)@l} (\nu \tilde{n})(\ell \llbracket P' \{y/u\} \rrbracket \parallel P'')}$$

$$\frac{P \equiv (\nu \tilde{n})(\ell \llbracket \bar{x}\langle y \rangle.P' \rrbracket \parallel P'') \quad x, y, \ell \notin \tilde{n}}{P \xrightarrow{\bar{x}\langle y \rangle@l} (\nu \tilde{n})(\ell \llbracket P' \rrbracket \parallel P'')}$$

$$\frac{P \equiv (\nu \tilde{n})(\ell \llbracket \bar{x}\langle y \rangle.P' \rrbracket \parallel P'') \quad x, \ell \notin \tilde{n} \quad y \in \tilde{n}}{P \xrightarrow{(\nu y)\bar{x}\langle y \rangle@l} (\nu \tilde{n} \setminus y)(\ell \llbracket P' \rrbracket \parallel P'')}$$

Labelled bisimulation for DPI, ctd.

The standard bisimulation on top of the LTS below coincides with reduction barbed congruence.

Remark: the LTS is written in an *unconventional* style, which precisely characterises the interactions a term can have with a context.

Questions:

- 1- every label should correspond to a (minimal) interacting context: can you spell out these contexts?
- 2- why there are no explicit labels for the "goto" action?

Mobile Ambients, design choices

Objective: build a process language on top of the concepts of barriers (administrative domains, firewalls, ...) and of barrier crossing.

A graphical representation of the syntax and of the reduction semantics of Mobile Ambients can be found here:

`http://research.microsoft.com/Users/luca/Slides/
2000-11-10%20Wide%20Area%20Computation%20\(Valladolid\).pdf`

Mobile Ambients syntax (in ISO 10646)

Processes:

$$P, Q, R ::= \mathbf{0}$$
$$P_1 \parallel P_2$$
$$(\nu n)P$$
$$n[P]$$
$$C.P$$
$$!P$$

Capabilities:

$$C ::= \text{in}_n$$
$$\text{out}_n$$
$$\text{open}_n$$

Mobile Ambients: interaction

- Locations migrate under the control of the processes located at their inside:

$$\begin{aligned}n[\text{in}_m.P \parallel Q] \parallel m[R] &\rightarrow m[n[P \parallel Q] \parallel R] \\m[n[\text{out}_m.P \parallel Q] \parallel R] &\rightarrow n[P \parallel Q] \parallel m[R]\end{aligned}$$

- a location may be opened:

$$\text{open}_n.P \parallel n[Q] \rightarrow P \parallel Q$$

Hint about an LTS for Mobile Ambients

Consider the term $M \equiv (\nu \tilde{m})(k[\text{in}_n.P \parallel Q] \parallel R)$ where $k \notin \tilde{m}$. It can interact with the context $n[T] \parallel -$, where T is an arbitrary process, yielding $O \equiv (\nu \tilde{m})(n[T \parallel k[P \parallel Q]] \parallel R)$. This interaction can be captured with a transition $M \xrightarrow{k.\text{enter}_n} O$.

Remark that, contrarily to what happens in CCS and pi-calculus, a bit of the interacting context is still visible in the outcome!

Along these lines (asynchrony is needed too!) it is possible to characterise reduction barbed congruence using a labelled bisimilarity.

References

James Riely, Matthew Hennessy: *Distributed Pprocesses and location failures*. Theoretical Computer Science, 2001. An extended abstract appeared in ICALP 97.

Luca Cardelli, Andrew Gordon: *Mobile Ambients*. Theoretical Computer Science, 2000. An extended abstract appeared in FOSSACS 1998.

Summary

- syntax and semantics of CCS:
 - non-determinism, parallel composition: from automata to CCS
 - LTS for CCS: compositional definition of automata
 - reduction semantics for CCS
- equivalences:
 - traces, completed traces, failures, simulation, bisimulation
 - from strong to weak equivalences
 - proof techniques for bisimulation (up-to bisimulation, up-to context)
 - axiomatisation (proof of soundness and completeness in the the strong case)
 - Hennessy-Milner logic (proof of soundness and completeness)
- name passing:
 - syntax and reduction semantics of pi-calculus
 - data structures as processes

-
- contextual equivalence:
 - relationship between contextual equivalences and labelled equivalence (proof of soundness and completeness for CCS in the weak case)
 - derivation of the early LTS for pi-calculus
 - proof of soundness of bisimilarity for pi-calculus, and counter-example to completeness
 - asynchronous interaction
 - asynchronous contextual equivalences
 - how to build an LTS for asynchronous pi-calculus (ACS and HT)
 - functions as processes:
 - cbv and cbn CPS transform
 - encoding of cbv and cbn lambda calculus in pi-calculus
 - types:
 - types to avoid errors (simple types, proof of type preservation)
 - types to reason about processes: subtyping, receptiveness, types and contextual equivalence

Partiel

20/11/07 — 12h45-15h45

Salle 1C12 — Chevaleret

Personal notes and lecture notes authorised.

Thank you for your attention

Feel free to contact me if you want to know more about research in concurrency.

`francesco.zappa_nardelli@inria.fr`