

A roundtrip from π -calculus to synchronous programming

Roberto AMADIO

Université de Paris 7

Laboratoire Preuves, Programmes et Systèmes

Groupe de travail Concurrency: Chevaleret, Thursday 2 pm.

<http://www.pps.jussieu.fr/~amadio/cc/>

Introduction

In truth, there is nothing canonical about our choice of basic combinators, even though they were chosen with great attention to economy. What characterises our calculus is not the exact choice of combinators, but rather the choice of interpretation and of *mathematical framework*.

R. Milner, Communication and Concurrency, Prentice-Hall, 1989, page 195.

We now wish to discuss a calculus [...] It arose from the author's attempt to relate *asynchrony* to *synchrony*. The contrast between these terms may be understood in more than one way. Here, we mean the contrast between the assumption which we have hitherto made that concurrent agents proceed at indeterminate relative speeds (asynchrony), and the alternative assumption that they proceed in lockstep - i.e. that at every instant each agent performs a single action (synchrony).

Op. cit.

Some references

- R. Milner, *Calculi for synchrony and asynchrony*, TCS, 25, 1983.
- D. Austry and G. Boudol, *Algèbre de processus et synchronisation*, TCS, 30, 1984.
- G. Berry and G. Gonthier, *The Esterel synchronous programming language*. Science of computer programming, 19, 1992.
- F. Boussinot and R. De Simone, *The SL Synchronous Language*. IEEE Trans. on Software Engineering, 22, 1996.
- R. Amadio, *The SL synchronous language, revisited*. Technical Report, Université Paris 7, Laboratoire PPS, November 2005.

NB There is a parallel thread of synchronous programming languages that is based on data flow ideas (Lustre, Signal, synchronous Kahn networks,...) that we will not be discussed here. Refer to the course on *Synchronous Programming*.

Goal of this course

- Introduce and revisit the *SL synchronous programming model* (a relaxation of the ESTEREL model).
- Apply the *mathematical framework* that has been developed for CCS/ π -calculus to the SL model.
- Highlight a *research problem* arising when extending the model with data values and name mobility.

Some general ideas

Synchrony:

- Computation is regulated by a notion of *instant*.
- Threads can proceed at different speeds but periodically they go through a *global synchronisation*.

NB An instant can turn out to be considerably more complex than in Milner's description where "at every instant each agent performs a *single* action".

Reactivity:

- Each instant should *terminate* in reasonable time/space.
- If a thread loops the whole system loops!

Parallelism

- Parallelism might be useful to exploit a parallel machine.
- Parallelism is also useful to have a *modular programming style* (*e.g.* think of a graphical user interface).

Determinism

- Debugging/testing/verifying non-deterministic programs is difficult.
- Whenever possible try to work with *deterministic* programs.
E.g. Kahn networks are an early example of deterministic and parallel model.

Cooperative concurrency

- In preemptive concurrency a thread can be *interrupted at any point*.
- In cooperative concurrency a thread must *explicitly return the control* to the scheduler.
- In a synchronous framework, it makes sense to adopt cooperative concurrency. This may help in making the program look more deterministic.

The SL model

- A *program* is a *multi-set of threads* interacting via *shared signals*.
- A *signal* is emitted or not, and once it is emitted it *persists* within the current instant.

Basic principle of the SL model

Reaction to the *absence* of a signal within an instant can only happen at the *next instant*.

- This is a relaxation of the ESTEREL model (Berry-Gonthier 1992) where the reaction to the absence can be *instantaneous* (SL designed to avoid ESTEREL's 'causality problem').

- For instance,

νs present s then 0 else (emit s)

- How does one ‘run’ this statement?
- If terms of boolean equations, this is like

$$s = \neg s$$

which has no solution...

- SL avoids the paradox by assuming that the branch **else** is only computed at the end of the instant, once the presence or absence of the signal s is settled.

Evolution

The SL model has evolved towards a *programming language* for concurrent applications:

- More *control operators, thread spawning,...*
- *Data types.*
- *Programming environments:* C, JAVA, SCHEME, ML,...
(Boussinot et al. 95-05).
- *Multi-processing* (Boussinot 04).
- *Migration* (Boudol 04).

Philosophy

- Synchronous models like ESTEREL or LUSTRE are designed to allow a *static scheduling* and a compilation to *finite state machines*.
- The SL model:
 - more relaxed view on *efficiency*,
 - more *flexibility* in programming and compiling,
 - preserve *determinism* and *reactivity*.

Typical applications (for the basic model)

- Graphical user interfaces.
- Simulations.
- Web services.
- Network games.

A programming example (informal)

- Suppose you want to program a generic *cellular automaton*.
- Assign a *thread* to each *cell*.
- The parameters of each thread/cell are: its *current state*, its *signal*, the signals of its *neighbours*.

- Each thread performs the following actions periodically:
 1. It *emits* its current state in the signals of its neighbours.
 2. It *waits the end of the instant*.
 3. At the end of the instant it *collects* the states of its neighbours on its signal and it computes its new state.

Remarks

- Similar examples exist for the simulation of mechanical phenomena such as the n -body problem. See <http://www-sop.inria.fr/mimoso/rp> for demos.
- In these examples, signals carry *values* (e.g. the state of the cell).
- In the following, we will focus on a calculus where signals are *pure* (cf. CCS).
- In the end, we will consider the problems that arise when this hypothesis is dropped and we move to a kind of *synchronous* π -calculus.

The SL model (revisited)

Signals and Environments

Signal names $S = \{s, s', \dots\}$

Interface signals

$$Int = Input \cup Output \subset S$$

assume Int finite.

Environment

$$E : S \rightarrow \{f, t\}$$

assume $dom(E)$ finite and containing Int .

Threads and Programs

$$\begin{aligned} T ::= & 0 \mid \\ & T; T \mid \\ & A(\mathbf{s}) \\ & (\nu s T) \mid \\ & (\text{thread } T) \mid \\ & (\text{emit } s) \mid \\ & (\text{await } s) \mid \\ & (\text{watch } s T) \end{aligned}$$

A *program* P is a *multi-set* of threads.

Evaluation (informal)

To execute a program during an instant, proceed as follows:

1. Schedule the execution of the threads *non-deterministically*.
2. When no thread can progress then all threads are either *terminated* or *waiting* on an await.
3. If this happens, evaluate the watching conditions in the order of their insertion.

Redexes

- Assume that sequential composition ‘;’ associates to the right.
- A *redex* Δ is defined by the grammar:

$$\Delta ::= 0;T \mid (\text{emit } s) \mid (\nu s T) \mid (\text{thread } T) \mid (\text{await } s) \mid (\text{watch } s 0) \mid A(\mathbf{s}) .$$

Evaluation context

An *evaluation context* C is defined by the grammar:

$$C ::= [] \mid []; T \mid (\text{watch } s C) \mid (\text{watch } s C); T .$$

Unique decomposition

A thread $T \neq 0$ admits a *unique decomposition* $T = C[\Delta]$ into an evaluation context C and a redex Δ . Moreover, if $T = 0$ then no decomposition exists.

Reduction rules

$$(0; T, E) \xrightarrow{\emptyset} (T, E)$$

$$(\text{emit } s, E) \xrightarrow{\emptyset} (0, E[t/s])$$

$$(\text{watch } s \ 0, E) \xrightarrow{\emptyset} (0, E)$$

$$(\nu s \ T, E) \xrightarrow{\emptyset} (T, E[f/s]) \quad \text{if } s \notin \text{dom}(E)$$

$$(A(\mathbf{s}), E) \xrightarrow{\emptyset} ([\mathbf{s}/\mathbf{x}]T, E) \quad \text{if } A(\mathbf{x}) = T$$

$$(\text{await } s, E) \xrightarrow{\emptyset} (0, E) \quad \text{if } E(s) = \mathbf{t}$$

$$(\text{thread } T, E) \xrightarrow{\{T\}} (0, E)$$

$$(C[\Delta], E) \xrightarrow{P} (C[T'], E') \quad \text{if } (\Delta, E) \xrightarrow{P} (T', E')$$

Program evaluation

$$(P \cup \{T\}, E) \rightarrow (P \cup \{T'\} \cup P'', E') \quad \text{if} \quad (T, E) \xrightarrow{P''} (T', E') .$$

Suspension = End of instant

- $(T, E) \downarrow$ if T cannot be reduced in the environment E according to the rules above.
- $(T, E) \downarrow$ if and only if $T = 0$ or $T = C[(\text{await } s)]$ with $E(s) = \text{f}$.
- A program is suspended if all its thread are suspended.

Computation at the end of the instant

We transform all $(\text{watch } s \ T)$ instructions where the signal s is present into the terminated thread 0. Formally,

$$\begin{aligned} [P]_E &= \{[T]_E \mid T \in P\} \\ [0]_E &= 0 \\ [T; T']_E &= [T]_E; T' \\ [\text{await } s]_E &= (\text{await } s) \\ [\text{watch } s \ T]_E &= \begin{cases} 0 & \text{if } E(s) = \text{t} \\ (\text{watch } s \ [T]_E) & \text{otherwise} \end{cases} \end{aligned}$$

Some definable instructions

(loop T) = A where: $A = T; A$

(now T) = νs (emit s); (watch s T) $s \notin sig(T)$

pause = νs (now (await s)) $s \notin sig(T)$

$(\text{present } s \ T_1 \ T_2) = \nu s' \ (\text{thread}$
 $(\text{now } (\text{await } s); (\text{thread } T_1; (\text{emit } s'))),$
 $(\text{watch } s \ \text{pause}; (\text{thread } T_2; (\text{emit } s')))) \);$
 $(\text{await } s')$

$$\begin{aligned}
(T_1 \parallel T_2) &= \nu s_1, s_2, s'_1, s'_2 \text{ (thread} \\
&\quad \text{(watch } s'_1 T_1; \text{(loop (emit } s_1); \text{pause))),} \\
&\quad \text{(watch } s'_2 T_2; \text{(loop (emit } s_2); \text{pause)))}; \\
&\quad \text{(await } s_1); \text{(emit } s'_1); \text{(await } s_2); \text{(emit } s'_2)
\end{aligned}$$

Remark

- The original SL model has `loop`, `present`, and `||` rather than recursion and thread spawning.
- This has the consequence that the number of threads and the size of the evaluation contexts is bound by a constant.
- Also the `present` and `||` constructs in SL have a different behaviour with respect to the evaluation context.

Monotonicity and determinacy

- During an instant the collection of signals emitted can only grow and all that can be tested is the presence (not the absence) of a signal.
- Technically, the reduction relation is *strongly confluent* up to some renaming.
- This implies that programs are *deterministic*, *i.e.*, for a given input they always react in the same way and recursively the property is satisfied by the continuation in the following instant (if any).

Reactivity

- It is easy to write *looping programs*.
- In the programming practice, reactivity is ensured by *instrumenting the code* with `pause` statements that force the computation to suspend for the current instant.

- For instance,

$$A = (\text{watch } s_1 \ B); (\text{emit } s_4); A$$
$$B = (\text{await } s_2); (\text{emit } s_3); \text{pause}; B$$

- It is feasible to design *static analysis techniques* that guarantee reactivity.

How do we justify the choice of the synchronisation operators?

- The present operator reflects directly the design principle of the SL language.
- But why having watch and await?
- What about other synchronisation operators such as

(when s T)

that runs T when s is present ?

A tail recursive model

$t ::= 0 \mid A(s) \mid \text{emit } s.t \mid \nu s t \mid \text{thread } t.t \mid \text{present } s t b$

$b ::= t \mid \text{ite } s b b$

- The reduction semantics is even simpler than for the previous language (no evaluation context).
- The basic reduction rule is

$$(\text{present } s \ t \ b, E) \xrightarrow{\emptyset} (t, E) \quad \text{if } E(s) = t$$

- The basic rule at the end of the instant is

$$\langle \text{ite } s \ b_1 \ b_2 \rangle_E = \begin{cases} \langle b_1 \rangle_E & \text{if } E(s) = t \\ \langle b_2 \rangle_E & \text{if } E(s) = f \end{cases}$$

Continuation passing style translation

- There is a *semantic preserving* CPS translation of the SL language into the tail recursive SL language.
- The translation $\llbracket _ \rrbracket$ is parameterised on a pair (t, τ) :
 - t is the *default continuation*.
 - τ is a *list of pairs* $(s_1, t_1) \cdots (s_n, t_n)$ representing the signals that are under the control of a watch statement and the relative continuations.

CPS-translation

$$\llbracket 0 \rrbracket(t, \tau) = t$$

$$\llbracket T_1; T_2 \rrbracket(t, \tau) = \llbracket T_1 \rrbracket(\llbracket T_2 \rrbracket(t, \tau), \tau)$$

$$\llbracket \text{emit } s \rrbracket(t, \tau) = \text{emit } s.t$$

$$\llbracket \nu s T \rrbracket(t, \tau) = \nu s \llbracket T \rrbracket(t, \tau), \quad \text{where: } s \notin \text{sig}(t) \cup \text{sig}(\tau)$$

$$\llbracket \text{thread } T \rrbracket(t, \tau) = \text{thread } \llbracket T \rrbracket(0, \epsilon).t$$

$$\llbracket \text{watch } s \ T \rrbracket(t, \tau) = \llbracket T \rrbracket(t, \tau \cdot (s, t))$$

$$\llbracket \text{await } s \rrbracket(t, \tau) = \text{present } s \ t \ b,$$

$$\text{where: } \tau = (s_1, t_1) \cdots (s_m, t_m), \quad b \equiv (\text{ite } s_1 \ t_1 \ \dots (\text{ite } s_m \ t_m \ A) \ \dots),$$

$$A = \text{present } s \ t \ b$$

$$\llbracket A(\mathbf{s}) \rrbracket(t, \tau) = A^{(t, \tau)}(\mathbf{s}, \mathbf{s}'),$$

$$\text{where: } \text{sig}(t, \tau) = \{\mathbf{s}'\}, \quad A(\mathbf{x}) = T,$$

$$\{\mathbf{x}\} \cap \{\mathbf{s}'\} = \emptyset, \quad A^{(t, \tau)}(\mathbf{x}, \mathbf{s}') = \llbracket T \rrbracket(t, \tau) .$$

Example

- Consider the program A defined by:

$$A = (\text{watch } s_1 B); (\text{emit } s_4); A$$

$$B = (\text{await } s_2); (\text{emit } s_3); \text{pause}; B$$

- To keep the translation compact, assume:

$$\llbracket \text{pause} \rrbracket (t, (s_1, t_1) \cdots (s_n, t_n)) = \text{pause.ite } s_1 t_1 (\cdots (\text{ite } s_n t_n t) \cdots)$$

- Then the program after CPS translation becomes:

$$\begin{aligned}
 A^{(0,\epsilon)} &= B^{(t_1,\tau_1)} \\
 t_1 &= \text{emit } s_4.A^{(0,\epsilon)} \\
 \tau_1 &= (s_1, t_1) \\
 B^{(t_1,\tau_1)} &= \text{present } s_2 t_2 (\text{ite } s_1 t_1 B^{(t_1,\tau_1)}) \\
 t_2 &= \text{emit } s_3.\text{pause.ite } s_1 t_1 B^{(t_1,\tau_1)} .
 \end{aligned}$$

Towards a process algebraic treatment

Some criticism

The formalisation of the SL model we have considered so far is close to an abstract machine.

- Ad hoc α -renaming.
- A global notion of environment.
- Threads compose but do not reduce.
- Programs reduce but do not compose.

A revised syntax

We introduce a notion of program that subsumes, threads, programs, and environments.

$$P ::= 0 \mid \text{emit } s \mid \text{present } s \ P \ B \mid P \mid P \mid \nu s \ P \mid A(\mathbf{s})$$

$$B ::= P \mid \text{ite } s \ B \ B$$

NB We omit ‘emit $s.P$ ’ and ‘thread $P'.P$ ’. They are syntactic sugar for $(\text{emit } s) \mid P$ and $P' \mid P$, respectively.

A labelled transition system

Actions

$$\alpha ::= \tau \mid s \mid \bar{s}$$

Transitions

$$\frac{}{\text{emit } s \xrightarrow{\bar{s}} \text{emit } s}$$

$$\frac{}{\text{present } s \ P \ B \xrightarrow{s} P \mid (\text{emit } s)}$$

$$\frac{P_1 \xrightarrow{s} P'_1 \quad P_2 \xrightarrow{\bar{s}} P'_2}{P_1 \mid P_2 \xrightarrow{\tau} P'_1 \mid P'_2}$$

$$\frac{P_1 \xrightarrow{\alpha} P'_1}{P_1 \mid P_2 \xrightarrow{\alpha} P'_1 \mid P_2}$$

$$\frac{P \xrightarrow{\alpha} P' \quad s \notin \alpha}{\nu s P \xrightarrow{\alpha} \nu s P'}$$

$$\frac{A(\mathbf{x}) = P}{A(\mathbf{s}) \xrightarrow{\tau} [\mathbf{s}/\mathbf{x}]P}$$

Remarks

- Emission of a signal is *persistent*.
- When a signal is present, the continuation is guaranteed to see the signal as present.
- This has the effect that:

$$(\text{emit } s) \mid (\text{present } s \ 0 \ 0) \xrightarrow{\tau} (\text{emit } s) \mid 0 \mid (\text{emit } s)$$

- However this is not a problem, because emitting a signal *once* is the *same* as emitting it *twice*.

End of the instant

The idea is always the same:

- Collect all emitted signals (restricted or not).
- Reset all emissions.
- Compute the continuations B of the suspended present statements.

Towards equivalence

- Can we define *a manageable compositional equivalence* based on a *labelled bisimulation*?
- Can we *justify its definition* via a notion of *contextual bisimulation*?

Worm up: mini-‘asynchronous’ CCS (aCCS)

$$P ::= \bar{a} \mid a.P \mid (P \mid P) \mid \nu a P$$

- No output prefix.
- Communication is asynchronous, *i.e.* based on *unordered and unbounded buffers* rather than on *rendez vous* (Internet vs. telephone).
- If an observer sends a message it has no way to see if it has been received. In other terms, we can observe the outputs but not the inputs.

Barbed bisimulation for aCCS

A relation R is a *barbed bisimulation* if it is symmetric and whenever $P R Q$ then

1. $P \xrightarrow{\tau} P'$ then $Q \xRightarrow{\tau} Q'$ and $P' R Q'$.
2. $P \xrightarrow{\bar{a}} \cdot$ then $Q \xRightarrow{\tau} Q'$, $Q' \xrightarrow{\bar{a}} \cdot$, and $P R Q'$.

Denote with \approx_B the largest relation that satisfies this property.

Contextual bisimulation for aCCS

A static context C is defined by

$$C ::= [] \mid C \mid P \mid \nu a C$$

A relation R is a *contextual bisimulation* if

1. It is a barbed bisimulation.
2. If $P R Q$ then $C[P] R C[Q]$.

Denote with \approx_C the largest contextual bisimulation.

Problem

- One can prove that:

$$a.\bar{a} \approx_C 0$$

- How can we account for such equivalence in a labelled bisimulation?

Idea

If $P R Q$ and $P \xrightarrow{a} P'$ then

- Either $Q \xrightarrow{a} Q'$ and $P' R Q'$ (as usual!)
- or $Q \xrightarrow{\tau} Q'$ and $P' R (Q' \mid \bar{a})$ (what's that?).

Ref. A., Castellani, Sangiorgi. *On bisimulations for the asynchronous π -calculus*. *Theor. Comp. Sci.* 1998.

Example

$$a.\bar{a} \approx_L 0$$

- If $a.\bar{a} \xrightarrow{a} \bar{a}$,
- then $0 \xrightarrow{\tau} 0$,
- and $\bar{a} \approx_L (0 \mid \bar{a})$.

Labelled bisimulation

So the full definition becomes. A relation R is a *labelled bisimulation* if it is symmetric and whenever $P R Q$ then if

1. $P \xrightarrow{\tau} P'$ then $Q \xRightarrow{\tau} Q'$ and $P' R Q'$.
2. $P \xrightarrow{\bar{a}} P'$ then $Q \xRightarrow{\bar{a}} Q'$, and $P' R Q'$.
3. $P \xrightarrow{a} P'$ then either $Q \xRightarrow{a} Q'$ and $P' R Q'$ or $Q \xRightarrow{\tau} Q'$ and $P' R (Q' \mid \bar{a})$.

Congruence properties of labelled bisimulation

Proposition If $P_1 \approx_L P_2$ then $P_1 \mid Q \approx_L P_2 \mid Q$
(and $\nu a P_1 \approx_L \nu a P_2$).

Proof outline

The proof structure is *not* quite the usual one.

1. We introduce a structural equivalence \equiv including equations such as $P \mid 0 \equiv P$.
2. We check that the proof technique *up to structural equivalence* is sound.
3. We show that $P_1 \approx_L P_2$ implies $(P_1 \mid \bar{a}) \approx_L (P_2 \mid \bar{a})$.

Exercise Analyse the case where $(P_1 \mid \bar{a})$ does a τ or an input move.

4. We show that \approx_L is transitive.

Exercise Analyse the case where $P_1 \approx_L P_2 \approx_L P_3$ and P_1 does an input move.

5. Only then can we show that

$$\{(P_1 \mid Q, P_2 \mid Q) \mid P_1 \approx_L P_2\}$$

is a labelled bisimulation up to \equiv .

Exercise Analyse the case where $P_1 \mid Q$ does a τ move or an input move.

6. It follows that

$$P_1 \approx_L P_2 \text{ implies } P_1 \approx_C P_2$$

Exercise Why?

Building contexts

Proposition If $P_1 \approx_C P_2$ then $P_1 \approx_L P_2$.

- We show that

$$\{(P_1, P_2) \mid P_1 \approx_C P_2\}$$

is a labelled bisimulation. This amounts to build suitable contexts.

Exercise What happens if P_1 does an input move?

Back to SL

So why are we looking at aCCS?

- It is an interesting variation of the model to the point that *programming languages* based on the π -calculus have adopted it.
- The *reception of a signal* in SL is not directly observable just as the *reception of a message* in aCCS.
- However there are two additional complications to consider in SL:
 1. The *end of the instant* should be observable.
 2. Once a program is closed and running, an *emission* is only observable at the end of the instant.

End of the instant

- Consider the following two *suspended* programs:

$$P = \text{present } s_1 \ 0 \ (\text{ite } s_2 \ (\text{emit } s_3) \ 0) \quad \text{and} \quad Q = \text{present } s_2 \ 0 \ 0 .$$

- If we plug them in the context $[] \mid (\text{emit } s_2)$ they exhibit a different behaviour.
- However a definition such as:

$$\frac{P \ R \ Q \quad P \downarrow}{Q \xRightarrow{\tau} Q', \quad Q' \downarrow, \quad P \ R \ Q', \quad [P] \ R \ [Q']}$$

is not enough to distinguish them.

- What will do is the following stronger definition:

$$\frac{P \ R \ Q \ S = (\text{emit } s_1) \mid \cdots \mid (\text{emit } s_n) \ (P \mid S) \downarrow}{
 \begin{array}{l}
 (Q \mid S) \xrightarrow{\tau} (Q' \mid S), \quad (Q' \mid S) \downarrow, \\
 (P \mid S) \ R \ (Q' \mid S), \quad [P \mid S] \ R \ [Q' \mid S]
 \end{array}
 }$$

Emission and end of the instant

- Suppose Ω is a looping program.
- Then in $(\text{emit } s) \mid \Omega$, the emission on s should not be observable because the instant *never terminates* (no matter who runs in parallel).
- So we look for a predicate $Pred$ such that the clause for emission is formulated as follows:

$$\frac{P R Q, \quad Pred(P), \quad P \xrightarrow{\bar{s}} P'}{Q \xRightarrow{\bar{s}} Q' \quad \text{and} \quad P' R Q'}$$

NB Incidentally, in SL if $P \xrightarrow{\bar{s}} P'$ then $P = P'$.

**Two candidates:
Weak suspension and Labelled suspension**

$P \Downarrow$ if $\exists P' P \xrightarrow{\tau} P'$ and $P' \Downarrow$ (weak suspension)

$P \Downarrow_L$ if $P \xrightarrow{\alpha_1} P_1 \cdots \xrightarrow{\alpha_n} P_n$, $n \geq 0$, and $P_n \Downarrow$ (L-suspension)

Characterisation of Labelled suspension

Let P be a program and let C denote a static context made of parallel composition and restriction. The following are equivalent:

1. $P \Downarrow_L$.
2. There is a program Q such that $(P \mid Q) \Downarrow$.
3. There is a static context C such that $C[P] \Downarrow_L$.

- An important property of L-suspension is that

$$(P \mid Q) \Downarrow_L \text{ implies } P \Downarrow_L$$

- This *may fail* for weak suspension in *non-deterministic* extensions of the language.
- However, in SL the property holds for weak suspension too because thanks to *confluence* we can show:

$$P \Downarrow_L \text{ iff } P \Downarrow$$

Full definition of labelled bisimulation

If $P R Q$ then

- $P \xrightarrow{\tau} P'$ implies $Q \xRightarrow{\tau} Q'$ and $P' R Q'$.
- $P \Downarrow_L$ and $P \xrightarrow{\bar{s}} P'$ implies $Q \xRightarrow{\bar{s}} Q'$ and $P' R Q'$.
- $P \xrightarrow{s} P'$ implies either $Q \xRightarrow{s} Q'$ and $P' R Q'$ or $Q \xRightarrow{\tau} Q'$ and $P' R (Q' \mid (\text{emit } s))$.
- $S = (\text{emit } s_1) \mid \cdots \mid (\text{emit } s_n)$, $(P \mid S) \Downarrow$ implies $(Q \mid S) \xRightarrow{\tau} (Q' \mid S)$, $(Q' \mid S) \Downarrow$, $(P \mid S) R (Q' \mid S)$, and $[P \mid S] R [Q' \mid S]$.

Denote with \approx_L the largest one.

Proof of congruence

The proof outline is similar to the one for asynchronous bisimulation (but of course the details are different...):

1. We introduce a *structural equivalence*.
2. Develop *proof technique up to*.
3. Show that $P_1 \approx_L P_2$ implies $(P_1 \mid (\text{emit } s)) \approx_L (P_2 \mid (\text{emit } s))$.
4. Show that \approx_L is *transitive*.
5. Show that

$$\{(P_1 \mid Q, P_2 \mid Q) \mid P_1 \approx_L P_2\}$$

is a labelled bisimulation up to \equiv .

Justifying labelled bisimulation

It remains to justify labelled bisimulation via a notion of contextual bisimulation.

Definition of barbed bisimulation If $P R Q$ then

- $P \xrightarrow{\tau} P'$ implies $Q \xRightarrow{\tau} Q'$ and $P' R Q'$.
- $P \Downarrow_L$ and $P \xrightarrow{\bar{s}} \cdot$ implies $Q \xRightarrow{\tau} Q' \xrightarrow{\bar{s}}$ and $P' R Q'$.
- $P \downarrow$ implies $Q \xRightarrow{\tau} Q'$, $Q' \downarrow$, $P R Q'$, and $[P] R [Q']$.

Definition of contextual bisimulation A barbed bisimulation such that $P R Q$ implies $C[P] R C[Q]$. Denote with \approx_C the largest one.

Theorem

$$P \approx_L Q \text{ iff } P \approx_C Q.$$

NB Again, to show the direction (\Leftarrow), we have to build suitable contexts ...

A quick look at ongoing research

- Practical languages that have been developed on top of the model have *data types*.
- In particular, signals carry *values*. E.g. think of primitives:

(emit s e) present $s(x).P, K$

- This introduces two sources of *non-determinism* in the language.

1. Two signals compete to be received within an instant:

$$(\text{emit } s \ 0) \mid (\text{emit } s \ 1) \mid \text{present } s(x).P, K$$

To avoid this situation one needs to make sure (static analysis) that *at most one value* is emitted on a signal at each instant.

2. Alternatively, we could wait the end of the instant to collect the *set of values* emitted within the instant (cf. Reactive ML).

In practice, the set has to be represented by some data structure such as a list. We have to make sure that the behaviour of the continuation *does not depend* on the particular representation chosen, *i.e.*, it has to be invariant under permutations of the list.

Determinacy vs. observational equivalence

The last example shows that the notion of *determinacy* depends on the notion of *observational equivalence*. We propose the following challenge:

Extend the theory of observational equivalence developed for SL to a kind of *synchronous* π -calculus.

- This seems interesting by itself since so far there is no analogous of Synchronous CCS for the π -calculus.
- Moreover, it may be useful to analyse the notion of determinacy.