

CO406H: Concurrent Processes

- π -calculus: dynamic reconfiguration of communication links
- Session types for distributed protocols

Plan of Lectures and Tutorials

- Lectures, Tutorials and Coursework
 - Room 342, Monday, 16.00–18.00 Lectures and Tutorials
 - Room 145, Monday, 13.00–14.00 (from Week 4)
 - **October 19th: no lecture**
 - 16th November or 30th November: Feedback and Revision
 - π -calculus
 - * **Coursework** (Publish 26th October, Deadline 9th November)
 - * Assessed and Examined
 - **Applications are highlight of this course:**
 - * *" Very pragmatic and really teaches from experience rather than a textbook. It's how all other lecturers should teach as well!"* (SOLE comments)
 - Web services and distributed protocols: 9th and 16th November
 - * Lecturers: Nobuko Yoshida & Romyana Neykova
 - * Not Assessed but Examined

Plan of Lectures and Tutorials

- Please print the lecture notes on Cate **before** the next lecture.
- Tutorial sheets and their answers will be available from Cate.
- **Coursework** π -calculus (Electronic submission).
- **Examination** Between 14th and 18th December
 - Question 1: π -calculus Similar with Coursework
 - Question 2: π -calculus and Distributed protocols (Session Calculus/Types)
 - Question 3: π -calculus and Distributed protocols (Session Calculus/Types)

The π -calculus

The π -calculus was first presented in 1989 by Milner, Parrow and Walker, based on an extension of CCS by Engberg and Nielsen.

It is useful for building models of concurrent/distributed/mobile systems and study their properties, like Turing Machines and the λ -calculus are useful for studying sequential computation.

The impact of π -calculus on industry and academia:

- Message-passing programming (SJ, Session C, Occam-Pi, ...).
- Distributed programming (JBoss Red Hat Scribble and Savara projects and the Ocean Observatories Initiative)
- **Web service orchestration and choreography** (e.g. W3C Web Service Choreography Description Languages) and Financial protocols
- Systems Biology and Security: the asynchronous π -calculus is used to specify and verify security protocols
- Active field of research, especially in Europe, in the UK, and at Imperial

References


- What you need for the exam is on these slides (and the tutorial sheets)
- **Introductory book:**
 - Communicating and Mobile Systems: the π -calculus* (Milner 1999)
- Advanced books:
 - The π -calculus: a Theory of Mobile Processes* (Sangiorgi, Walker 2001)
 - Distributed Pi-Calculus* (Hennessy 2007)

About the π -calculus

- There is no “canonical” π -calculus. For each application domain, there are alternative notations and lots of specialised variants:
 - we fix one notation and stick to it (once you know one notation, it’s easy to understand the others)
 - we start with the simplest variant (the asynchronous π -calculus) and explore various extensions
- An interaction happens by *message passing* rather than *synchronisation*.
- In the asynchronous π -calculus, the communication is *asynchronous* rather than *synchronous*.
- The π -calculus evolved from (value-passing) CCS, but it is more expressive:
 - channel mobility: send and receive *channel names* as messages
 - restriction is interpreted as new (private, secret) channel generation

Synchrony and Asynchrony

➤ Synchrony



➤ Asynchrony



cf. E-mail.
News.

Overview of the π -calculus lectures

- A simple version of asynchronous π -calculus: syntax and semantics
- The asynchronous π -calculus: syntax and semantics
- Joyful Hacking in asynchronous π -calculus
- Protocols and Session-calculus

The asynchronous π -calculus

The asynchronous π -calculus is a subset of the π -calculus presented independently by Honda and Tokoro (1991), and by Boudol (1992).

Communication is asynchronous: the output process $\bar{a}\langle b \rangle$ represents a message which is in the communication layer waiting to be picked up by a receiver (it does not have a continuation P like the synchronous process $\bar{a}\langle b \rangle.P$). Several messages can be in the communication layer at the same time, and their order is not preserved. Asynchronous communication is common in distributed systems, and can be used to simulate synchronous communication (handshake) when needed.

The asynchronous π -calculus is easier and more efficient to implement than the full π -calculus. It is widely used as the basis for building more complicated calculi, by adding primitives for distributed or object-oriented programming.

As a start, we consider a subset of asynchronous π -calculus, called *asynchronous CCS*.

Syntax of CCS

$a, b, c, ..$	name
$P, Q ::=$	processes
$\mathbf{0}$	nil process
$P \mid Q$	parallel composition of P and Q
$(\nu a)P$	restriction of a (scope) in P
\bar{a}	output on channel a
$a.P$	input on channel a , with continuation P

Notation: $\left\{ \begin{array}{l} \tilde{a} \stackrel{\text{df}}{=} a_1, \dots, a_n \quad \text{when we don't care about each } a_i \\ (\nu a_1, \dots, a_n)P \stackrel{\text{df}}{=} (\nu a_1) \dots (\nu a_n)P \end{array} \right.$

Abbreviation: $\left\{ \begin{array}{l} \bar{a} \stackrel{\text{df}}{=} \bar{a}.0 \quad \text{when we don't have a continuation} \\ a \stackrel{\text{df}}{=} a.0 \quad \text{when we don't have a continuation} \end{array} \right.$

Examples of CCS

- 0 means nothing and $0|0$ is as same as one 0 , hence nothing
- \bar{a} is one message to a ; and $\bar{a}|\bar{a}$ means two messages to a .
- $\bar{a}|\bar{b}|\bar{c}$
One message to a , one message to b and One message to c
- $\bar{a}|\bar{b}$ is as the same as $\bar{b}|\bar{a}$.
- $a.b.c$ inputs from a , then inputs from b and then inputs from c
- $a.b$ does not mean $b.a$
- $a.\bar{b}$ inputs from a then outputs to b , and $a.(\bar{b}|\bar{c}|\bar{d})$ inputs from a , then outputs to b , c and d .
- $a.b.(\bar{c}|\bar{a})$
- $a.(b.\bar{c}|d.\bar{e})$

Bad Syntax of CCS

- $\bar{a}.\bar{b}$ and $\bar{a}.b$
- $a.(b|c).d$ and $a.(b|c).\bar{d}$
- **0.0**

Reduction of CCS: Informally

We write $P \longrightarrow Q$ if P reduces to Q , like $8 + 2 \longrightarrow 10$.

- $\bar{a} | a.0 \longrightarrow 0$

- $\bar{a} | a.\bar{b} \longrightarrow \bar{b}$

- $\bar{a} | a.\bar{b} | \bar{c} | c.\bar{d}$ (hint: $(10 + 2) - (2 + 4)$)

- $\bar{a} | a.\bar{b} | a.\bar{c}$ reduces to either $\bar{b} | a.\bar{c}$ or $a.\bar{b} | \bar{c}$, that is

$$\bar{a} | a.\bar{b} | a.\bar{c} \longrightarrow \bar{b} | a.\bar{c} \quad \text{or} \quad \bar{a} | a.\bar{b} | a.\bar{c} \longrightarrow a.\bar{b} | \bar{c}$$

- $\bar{a} | \bar{a} | a.\bar{b} | a.\bar{c}$

- $\bar{a} | \bar{b} | a.b.\bar{e} | b.a.\bar{d}$

Name Restriction of CCS: Informally

- $(\nu a)(\bar{a} | a.\bar{b}) | a.\bar{c}$ means $(\nu d)(\bar{d} | d.\bar{b}) | a.\bar{c}$
Similar with $f(x) = x + 2$ means $f(y) = y + 2$
- $(\nu a)(\bar{a} | a.\bar{b}) | a.\bar{c} \longrightarrow (\nu a)\bar{b} | a.\bar{c} \equiv \bar{b} | a.\bar{c}$
- but $(\nu a)(\bar{a} | a.\bar{b}) | a.\bar{c} \not\longrightarrow a.\bar{b} | \bar{c}$

Names and Variables

We separate channel names, which are like constants in a programming language, from variables, which are used to instantiate messages received in input. Consequently, we have two sorts:

$$a, b, c \in \mathcal{N}$$

Channel Names

$$x, y, z \in \mathcal{V}$$

Variables

This distinction is not mandatory to build the theory, but is convenient when the calculus is used to model actual systems.

Syntax of asynchronous π -calculus

$u, v ::=$	identifiers
$a, b, c, ..$	name
$x, y, z, ..$	variable
$P, Q ::=$	processes
0	nil process
$P Q$	parallel composition of P and Q
$(\nu a)P$	generation of a with scope P (also called <i>restriction</i>)
$!P$	replication of P , i.e. infinite parallel composition $P P P \dots$
$\bar{u}\langle v \rangle$	output of v on channel u
$u(x).P$	input of <i>distinct</i> variables x on u , with continuation P

Notation: $\left\{ \begin{array}{ll} \tilde{u} \stackrel{\text{df}}{=} u_1, \dots, u_n & \text{when we don't care about each } u_i \\ (\nu a_1, \dots, a_n)P \stackrel{\text{df}}{=} (\nu a_1) \dots (\nu a_n)P \end{array} \right.$

Later on, we will consider other operators such as choice, output continuation, recursive definitions.

Free variables and free names

In order to understand the formal semantics of the π -calculus, it is important to know exactly what are the free variables fv and the free names fn of each term.

$$\begin{array}{ll} fv(x) = \{x\} & fn(x) = \emptyset \\ fv(a) = \emptyset & fn(a) = \{a\} \\ fv(\mathbf{0}) = \emptyset & fn(\mathbf{0}) = \emptyset \\ fv(P \mid Q) = fv(P) \cup fv(Q) & fn(P \mid Q) = fn(P) \cup fn(Q) \\ fv((\nu a)P) = fv(P) & fn((\nu a)P) = fn(P) \setminus \{a\} \\ fv(!P) = fv(P) & fn(!P) = fn(P) \\ fv(\bar{u}\langle v \rangle) = fv(u) \cup fv(v) & fn(\bar{u}\langle v \rangle) = fn(u) \cup fn(v) \\ fv(u(x).P) = fv(u) \cup (fv(P) \setminus \{x\}) & fn(u(x).P) = fn(u) \cup fn(P) \end{array}$$

Both $u(-)$ and $(\nu -)$ are called *binders*. A term is *closed* if it has no free variables, it is *open* otherwise.

In the process $P = (\nu b)a(x).(\bar{x}\langle z \rangle \mid \bar{x}\langle b \rangle)$, we have highlighted all the *free occurrences* of names and variables. In particular, $fv(P) = \{z\}$ and $fn(P) = \{a\}$. Above, the first occurrences of b and x are called *binding occurrences*, whereas the second occurrence of b and the second and third occurrences of x are called *bound*.

α -conversion

α -conversion is the meta-operation of renaming consistently (i.e. avoiding clashes) the bound names or variables of a process. If P is obtained from Q by α -conversion, we say that P and Q are α -equivalent, and we write $P =_\alpha Q$. For example:

$$(\nu a)(\bar{a}\langle b \rangle \mid (\nu c)\bar{c}\langle a \rangle) =_\alpha (\nu d)(\bar{d}\langle b \rangle \mid (\nu c)\bar{c}\langle d \rangle)$$

But we cannot replace a with b :

$$(\nu a)(\bar{a}\langle b \rangle \mid (\nu c)\bar{c}\langle a \rangle) \neq_\alpha (\nu b)(\bar{b}\langle b \rangle \mid (\nu c)\bar{c}\langle b \rangle)$$

Can we replace a with c ?

$$(\nu a)(\bar{a}\langle b \rangle \mid (\nu c)\bar{c}\langle a \rangle) \neq_\alpha (\nu c)(\bar{c}\langle b \rangle \mid (\nu c)\bar{c}\langle c \rangle)$$

Not naively! We can if, for example, we first α -convert the name c to e .

$$(\nu a)(\bar{a}\langle b \rangle \mid (\nu c)\bar{c}\langle a \rangle) =_\alpha (\nu c)(\bar{c}\langle b \rangle \mid (\nu e)\bar{e}\langle c \rangle)$$

The intuition is that α -conversion preserves each difference between names. We will use α -conversion very often, without explicit mention.

Substitution

A *substitution* $\{a/x\}$ applied to a process P (denoted by $P\{a/x\}$), has the effect of replacing all the free occurrences of x in P with a .

The substitution avoids clashes with any bound names by implicitly using α -conversion (*capture-avoiding substitution*). For example:

$$((\nu d)(\bar{a}\langle b \rangle \mid \bar{a}\langle d \rangle \mid \bar{a}\langle x \rangle))\{d/x\} = (\nu c)(\bar{a}\langle b \rangle \mid \bar{a}\langle c \rangle \mid \bar{a}\langle d \rangle)$$

where we have avoided the capture of the external d by α -converting (νd) to (νc) .

Given a substitution σ and an open process P , if $P\sigma$ is closed (i.e. $fv(P\sigma) = \emptyset$) we say that σ is a *closing substitution*. For example, $\{b/x\}$ is closing for $a(y).\bar{x}\langle y \rangle$:

$$a(y).\bar{x}\langle y \rangle\{b/x\} = a(y).\bar{b}\langle y \rangle \quad fv(a(y).\bar{b}\langle y \rangle) = \emptyset$$

As a shorthand, we use the notation $\{v_1, \dots, v_n/x_1, \dots, x_n\} \stackrel{\text{df}}{=} \{v_1/x_1\} \dots \{v_n/x_n\}$.

Quiz: α -conversion and substitution

1. Correct?

$$(\nu a)(\bar{a}\langle b \rangle \mid \bar{c}\langle a \rangle) =_{\alpha} (\nu d)(\bar{d}\langle b \rangle \mid \bar{c}\langle d \rangle)$$

2. Correct?

$$(\nu a)(\bar{a}\langle b \rangle \mid (\nu a)\bar{c}\langle a \rangle) =_{\alpha} (\nu d)(\bar{d}\langle b \rangle \mid (\nu a)\bar{c}\langle a \rangle)$$

3. Correct?

$$(\nu a)(\bar{a}\langle b \rangle \mid (\nu a)\bar{c}\langle a \rangle) =_{\alpha} (\nu d)(\bar{d}\langle b \rangle \mid (\nu a)\bar{c}\langle d \rangle)$$

4. Correct?

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

5. Correct?

$$a(y).(\bar{x}\langle y \rangle \mid (\nu a)\bar{c}\langle a \rangle)\{a/x\} = a(y).(\bar{a}\langle y \rangle \mid (\nu a)\bar{c}\langle a \rangle)$$

6. Correct?

$$a(y).(\bar{x}\langle y \rangle \mid (\nu a)\bar{c}\langle x \rangle)\{a/x\} = a(y).(\bar{a}\langle y \rangle \mid (\nu a)\bar{c}\langle a \rangle)$$

Quiz: Answers: α -conversion and substitution

1. Yes

$$(\nu a)(\bar{a}\langle b \rangle \mid \bar{c}\langle a \rangle) =_{\alpha} (\nu d)(\bar{d}\langle b \rangle \mid \bar{c}\langle d \rangle)$$

2. Yes

$$(\nu a)(\bar{a}\langle b \rangle \mid (\nu a)\bar{c}\langle a \rangle) =_{\alpha} (\nu d)(\bar{d}\langle b \rangle \mid (\nu a)\bar{c}\langle a \rangle)$$

3. No.

$$(\nu a)(\bar{a}\langle b \rangle \mid (\nu a)\bar{c}\langle a \rangle) =_{\alpha} (\nu a)(\bar{a}\langle b \rangle \mid (\nu d)\bar{c}\langle d \rangle) \neq_{\alpha} (\nu d)(\bar{d}\langle b \rangle \mid (\nu a)\bar{c}\langle d \rangle)$$

4. No. Because

$$a(y).\bar{x}\langle y \rangle \{b/y\} =_{\alpha} a(z).\bar{x}\langle z \rangle \{b/y\} = a(z).\bar{x}\langle z \rangle$$

5. Yes. Because

$$a(y).(\bar{x}\langle y \rangle \mid (\nu a)\bar{c}\langle a \rangle) \{a/x\} = a(y).(\bar{a}\langle y \rangle \mid (\nu a)\bar{c}\langle a \rangle) =_{\alpha} a(y).(\bar{a}\langle y \rangle \mid (\nu d)\bar{c}\langle d \rangle)$$

6. No. Because

$$a(y).(\bar{x}\langle y \rangle \mid (\nu a)\bar{c}\langle x \rangle) =_{\alpha} a(y).(\bar{a}\langle y \rangle \mid (\nu d)\bar{c}\langle x \rangle)$$

Hence

$$a(y).(\bar{x}\langle y \rangle \mid (\nu a)\bar{c}\langle x \rangle) \{a/x\} =_{\alpha} a(y).(\bar{a}\langle y \rangle \mid (\nu d)\bar{c}\langle a \rangle)$$

Structural congruence (1)

The reduction semantics of the π -calculus is inspired by the Chemical Abstract Machine (CHAM) of Berry and Boudol. Processes "float around" like molecules in a solution using structural congruence (\equiv) and "react" using a reduction relation (\longrightarrow).

The intuition is that if $P \equiv Q$ then we consider P and Q completely interchangeable.

Structural congruence is an equivalence relation, is preserved by all the syntactic operators, and contains α -equivalence.

$P \equiv P$ (Eq Reflexivity)

$P \equiv Q \implies Q \equiv P$ (Eq Symmetry)

$P \equiv R$ and $R \equiv Q \implies P \equiv Q$ (Eq Transitivity)

$P \equiv Q \implies (\nu a)P \equiv (\nu a)Q$ (Cong Res)


$P \equiv Q \implies P \mid R \equiv Q \mid R$ (Cong Par)

$P \equiv Q \implies u(x).P \equiv u(x).Q$ (Cong In)



$P \equiv Q \implies !P \equiv !Q$ (Cong Rep)

$P =_{\alpha} Q \implies P \equiv Q$ (α -equivalence)

Structural congruence satisfies also additional rules specific to the π -calculus. . .



$$a(x).P \mid \bar{a}\langle v \rangle \mid \bar{a}\langle w \rangle$$



$$(a(x).P \mid \bar{a}\langle w \rangle \mid \bar{a}\langle v \rangle)$$

Structural congruence (2)

The set of all processes, with the operation of parallel composition, and with the nil process as the neutral element are a *commutative monoid*.

$$P \mid (Q \mid Q') \equiv (P \mid Q) \mid Q' \quad (\text{Associativity})$$

$$P \mid Q \equiv Q \mid P \quad (\text{Commutativity})$$

$$P \mid \mathbf{0} \equiv P \quad (\text{Zero})$$

That is, we can exchange freely the position of parallel processes and we can insert or delete the $\mathbf{0}$ process at will. For example:

$$\mathbf{0} \mid \mathbf{0} \mid P \mid Q \mid R \equiv Q \mid \mathbf{0} \mid R \mid P$$

Replication can be folded or unfolded as many times as needed.

$$!P \equiv P \mid !P \quad (\text{Rep})$$

Thanks to this rule, the π -calculus can express infinite computations.

Structural congruence (3)

The last rules that complete the definition of \equiv state that the scope of a restricted name can be extended (or contracted) across terms which do not contain free occurrences of that name.

$$(\nu a)0 \equiv 0 \quad (\text{Res Nil})$$

$$(\nu a)(\nu b)P \equiv (\nu b)(\nu a)P \quad (\text{Res Res})$$

$$a \notin \text{fn}(P) \implies P \mid (\nu a)Q \equiv (\nu a)(P \mid Q) \quad (\text{Res Par})$$

For example:

$$(\nu a, b)(a(x).\bar{x}\langle c \rangle \mid \bar{a}\langle b \rangle) \equiv (\nu a)(a(x).\bar{x}\langle c \rangle \mid (\nu b)(\bar{a}\langle b \rangle))$$

$$(\nu a, b)(c(x).\bar{c}\langle x \rangle \mid \bar{c}\langle d \rangle) \equiv c(x).\bar{c}\langle x \rangle \mid \bar{c}\langle d \rangle$$

It is very important to gain familiarity with structural congruence, we will use it very often during the rest of the course.

Question 1 Prove $(\nu c)P \equiv P$ if $c \notin \text{fn}(P)$.

Question 2 Prove if $P \equiv Q$ then $\text{fn}(P) = \text{fn}(Q)$ and $\text{fv}(P) = \text{fv}(Q)$.

Reduction relation

Reduction describes how processes interact by exchanging messages. It is the smallest *partial relation* between processes satisfying the rules given below.

$$\bar{a}\langle v \rangle \mid a(x).P \longrightarrow P\{v/x\} \quad (\text{Comm})$$

$$\frac{P \longrightarrow P'}{P \mid Q \longrightarrow P' \mid Q} \quad (\text{Par})$$

$$\frac{P \longrightarrow P'}{(\nu a)P \longrightarrow (\nu a)P'} \quad (\text{Res})$$

$$\frac{P \equiv Q \longrightarrow Q' \equiv P'}{P \longrightarrow P'} \quad (\text{Struct})$$

As a shorthand notation, we write $P \xrightarrow{*} Q$ if $P \equiv Q$ or $P \longrightarrow \dots \longrightarrow Q$.

Reduction Examples

Communication:

$$a(x).\bar{x}\langle c \rangle \mid \bar{a}\langle b \rangle \longrightarrow \bar{b}\langle c \rangle \quad \text{by (Struct), (Commutativity) and (Comm)}$$

$$\begin{aligned} & a(x).\bar{x}\langle c \rangle \mid \bar{a}\langle b \rangle \\ \equiv & \bar{a}\langle b \rangle \mid a(x).\bar{x}\langle c \rangle && \text{by (Struct), (Commutativity)} \\ \longrightarrow & \bar{b}\langle c \rangle && \text{by (Comm)} \end{aligned}$$

Scope extrusion:

$$(\nu b)(\bar{a}\langle b \rangle) \mid a(x).\bar{c}\langle x \rangle \longrightarrow (\nu b)\bar{c}\langle b \rangle \quad \text{by (Struct), (Res Par), (Res) and (Comm)}$$

$$\begin{aligned} & (\nu b)(\bar{a}\langle b \rangle) \mid a(x).\bar{c}\langle x \rangle && \text{note } b \notin \text{fn}(a(x).\bar{c}\langle x \rangle) \\ \equiv & (\nu b)(\bar{a}\langle b \rangle \mid a(x).\bar{c}\langle x \rangle) && \text{by (Struct), (Res Par)} \\ \longrightarrow & (\nu b)\bar{c}\langle b \rangle && \text{by (Res) and (Comm)} \end{aligned}$$

Infinite behaviour and Sorts

We use the macros for process definition $\mathbf{A}(\tilde{x})$ and usage $\mathbf{A}\langle\tilde{v}\rangle$ *only informally*. To represent infinite behaviour, we use the *replication* operator $!P$, stating that there are as many copies of P as needed, all running in parallel. For example:

$$a(x).P \mid a(x).Q \mid !\bar{a}\langle b \rangle \longrightarrow P\{b/x\} \mid a(x).Q \mid !\bar{a}\langle b \rangle \longrightarrow P\{b/x\} \mid Q\{b/x\} \mid !\bar{a}\langle b \rangle$$

Replication is very simple, yet combined with channel generation it can encode recursive definitions (later on we will see how).

Reduction Examples

Infinite behaviour:

$$\bar{a}\langle b \rangle \mid !a(x).\bar{a}\langle x \rangle \longrightarrow \bar{a}\langle b \rangle \mid !a(x).\bar{a}\langle x \rangle \quad \text{by (Struct), (Rep), (Par) and (Comm)}$$

$$\begin{aligned} & \bar{a}\langle b \rangle \mid !a(x).\bar{a}\langle x \rangle \\ \equiv & \bar{a}\langle b \rangle \mid a(x).\bar{a}\langle x \rangle \mid !a(x).\bar{a}\langle x \rangle && \text{by (Struct), (Rep)} \\ \longrightarrow & \bar{a}\langle b \rangle \mid !a(x).\bar{a}\langle x \rangle && \text{by (Par) and (Comm)} \end{aligned}$$

Nondeterminism:

$$\bar{a}\langle b \rangle \mid \bar{a}\langle d \rangle \mid a(x).\bar{c}\langle x \rangle \begin{array}{l} \nearrow \bar{a}\langle b \rangle \mid \bar{c}\langle d \rangle \\ \searrow \bar{a}\langle d \rangle \mid \bar{c}\langle b \rangle \end{array} \quad \begin{array}{l} \text{by (Par) and (Comm)} \\ \text{using also (Struct), (Commutativity)} \end{array}$$

Atoms for Name Passing

Small Agents (1)

Forwarder $\mathbf{FW}(a, b) \stackrel{\text{df}}{=} a(z).\bar{b}\langle z \rangle$

A *forwarder* from channel a to channel b , is a process that forwards a message for a on b .

$$\mathbf{FW}\langle a, b \rangle \mid \bar{a}\langle d \rangle \longrightarrow \bar{b}\langle d \rangle$$

$$(\nu b)(\mathbf{FW}\langle a, b \rangle \mid \mathbf{FW}\langle b, c \rangle) \mid \bar{a}\langle d \rangle \longrightarrow \longrightarrow \bar{c}\langle d \rangle$$

Duplicator $\mathbf{D}(a, b, c) \stackrel{\text{df}}{=} a(z).(\bar{b}\langle z \rangle \mid \bar{c}\langle z \rangle)$

A *duplicator* from channel a to b and c , is a process that duplicates a message for a to b and c .

$$\mathbf{D}\langle a, b, c \rangle \mid \bar{a}\langle d \rangle \longrightarrow (\bar{b}\langle d \rangle \mid \bar{c}\langle d \rangle)$$

$$(\nu b)(\mathbf{D}\langle a, b, c_1 \rangle \mid \mathbf{D}\langle b, c_2, c_3 \rangle) \mid \bar{a}\langle d \rangle \longrightarrow \longrightarrow (\bar{c}_1\langle d \rangle \mid \bar{c}_2\langle d \rangle \mid \bar{c}_3\langle d \rangle)$$

Killer $\mathbf{K}(a) \stackrel{\text{df}}{=} a(z).0$

A *killer* at channel a is a process that kills a message from a .

$a(z).(P \mid Q)$ can be decomposed as $(\nu c_1, c_2)(\mathbf{D}\langle a, c_1, c_2 \rangle \mid c_1(z).P \mid c_2(z).Q)$.

For example, $a(z).(\bar{b}\langle z \rangle \mid 0)$ is the same as $(\nu c_1, c_2)(\mathbf{D}\langle a, c_1, c_2 \rangle \mid \mathbf{FW}\langle c_1, b \rangle \mid \mathbf{K}\langle c_2 \rangle)$

Small Agents (3)

Identity Receptor $\mathbf{I}(a) \stackrel{\text{df}}{=} !\mathbf{FW}\langle a, a \rangle$

An *identity receptor* at channel a is a process that forwards messages for a on a .

$$\bar{a}\langle d \rangle \mid \mathbf{I}\langle a \rangle \longrightarrow \bar{a}\langle d \rangle \mid \mathbf{I}\langle a \rangle \longrightarrow \bar{a}\langle d \rangle \mid \mathbf{I}\langle a \rangle \longrightarrow \bar{a}\langle d \rangle \mid \mathbf{I}\langle a \rangle \dots$$

Equator $\mathbf{EQ}(a, b) \stackrel{\text{df}}{=} !\mathbf{FW}\langle a, b \rangle \mid !\mathbf{FW}\langle b, a \rangle$.

An *equator* between two channels a and b , is a process that forwards all messages for a on b , and viceversa, making a and b in some sense “equivalent”. For example,

$$\bar{a}\langle d \rangle \mid \mathbf{EQ}\langle a, b \rangle \longrightarrow \bar{b}\langle d \rangle \mid \mathbf{EQ}\langle a, b \rangle \longrightarrow \bar{a}\langle d \rangle \mid \mathbf{EQ}\langle a, b \rangle \longrightarrow \dots$$

Omega $\Omega \stackrel{\text{df}}{=} (\nu a)(!\mathbf{FW}\langle a, a \rangle \mid \bar{a}\langle a \rangle)$

An *omega* is a process that continues infinite reductions by himself.

$$\Omega \longrightarrow \Omega \longrightarrow \Omega \longrightarrow \dots$$

New Name Generator $\mathbf{NN}(a) \stackrel{\text{df}}{=} !a(x).(\nu b)\bar{x}\langle b \rangle$. A *new name generator* is a process that creates a new name infinitely when it is asked.

$$\bar{a}\langle c \rangle \mid \bar{a}\langle d \rangle \mid \mathbf{NN}\langle a \rangle \longrightarrow (\nu b)\bar{c}\langle b \rangle \mid \bar{a}\langle d \rangle \mid \mathbf{NN}\langle a \rangle \longrightarrow (\nu b)\bar{c}\langle b \rangle \mid (\nu b')\bar{d}\langle b' \rangle \mid \mathbf{NN}\langle a \rangle$$

Quiz: Name Generator

1. Is $\mathbf{NN}(a) \stackrel{\text{df}}{=} !a(x).(\nu b)\bar{x}\langle b \rangle$ structural congruent with $\mathbf{NN}_1(a) = (\nu b)!a(x).\bar{x}\langle b \rangle$?
2. Is $\mathbf{NN}(a) \stackrel{\text{df}}{=} !a(x).(\nu b)\bar{x}\langle b \rangle$ structural congruent with $\mathbf{NN}_2(a) = !(\nu b)a(x).\bar{x}\langle b \rangle$?

Quiz: Answer: Name Generator

1. Is $\mathbf{NN}(a) \stackrel{\text{df}}{=} !a(x).(\nu b)\bar{x}\langle b \rangle$ structural congruent with $\mathbf{NN}_1(a) = (\nu b)!a(x).\bar{x}\langle b \rangle$?

No: since

$$\begin{aligned} \mathbf{NN}(a) &= !a(x).(\nu b)\bar{x}\langle b \rangle \\ &\equiv a(x).(\nu b)\bar{x}\langle b \rangle \mid a(x).(\nu b)\bar{x}\langle b \rangle \mid \cdots \mid a(x).(\nu b)\bar{x}\langle b \rangle \mid !a(x).(\nu b)\bar{x}\langle b \rangle \\ \mathbf{NN}_1(a) &= (\nu b)!a(x).\bar{x}\langle b \rangle \\ &\equiv (\nu b)(a(x).\bar{x}\langle b \rangle \mid a(x).\bar{x}\langle b \rangle \mid \cdots \mid a(x).\bar{x}\langle b \rangle \mid !a(x).\bar{x}\langle b \rangle) \end{aligned}$$

2. Is $\mathbf{NN}(a) \stackrel{\text{df}}{=} !a(x).(\nu b)\bar{x}\langle b \rangle$ structural congruent with $\mathbf{NN}_2(a) = !(\nu b)a(x).\bar{x}\langle b \rangle$?

No: since we do **not** have the rule $(\nu b)a(x).P \equiv a(x).(\nu b)P$ but they are observationally equivalent as we shall see soon.

Quiz: Small Agents

1. Is $I\langle a \rangle$ like 0 ?
2. Is Ω like 0 ?
3. Is $\mathbf{EQ}\langle a, b \rangle \mid \bar{a}\langle d \rangle$ like $\mathbf{EQ}\langle a, b \rangle \mid \bar{b}\langle d \rangle$?
4. Is $\mathbf{EQ}\langle a, b \rangle \mid \bar{d}\langle a \rangle$ like $\mathbf{EQ}\langle a, b \rangle \mid \bar{d}\langle b \rangle$?
5. Is $\mathbf{EQ}\langle a, b \rangle \mid \bar{d}\langle a \rangle$ like $\mathbf{EQ}\langle a, b \rangle \mid \bar{d}\langle b \rangle$?

Joyful Hacking in π -calculus

Channel mobility example

Internet connection: a **Client** and a **Server** talk on dedicated communication ports, set up using the channel a :

$$\mathbf{Client}(a, c) \stackrel{\text{df}}{=} (\bar{a}\langle c \rangle \mid c(x).\mathbf{Client}_1\langle c, x \rangle)$$

$$\mathbf{Server}(a, s) \stackrel{\text{df}}{=} a(y).(\bar{y}\langle s \rangle \mid \mathbf{Server}_1\langle y, s \rangle)$$

Difference with CCS: the variable y is used as the name of a channel.

$$\begin{aligned} & \mathbf{Client}\langle a, c \rangle \mid \mathbf{Server}\langle a, s \rangle \\ \longrightarrow & c(x).\mathbf{Client}_1\langle c, x \rangle \mid \bar{c}\langle s \rangle \mid \mathbf{Server}_1\langle c, s \rangle \\ \longrightarrow & \mathbf{Client}_1\langle c, s \rangle \mid \mathbf{Server}_1\langle c, s \rangle \end{aligned}$$

After the two communication steps, client and server know each other's port.

Channel generation example

We can make the Internet connection example more flexible. To avoid external interferences, the client and server exchange newly generated port names:

$$\mathbf{Client}(a) \stackrel{\text{df}}{=} (\nu c)(\bar{a}\langle c \rangle \mid c(x).\mathbf{Client}_1\langle c, x \rangle)$$

$$\mathbf{Server}(a) \stackrel{\text{df}}{=} a(y).(\nu s)(\bar{y}\langle s \rangle \mid \mathbf{Server}_1\langle y, s \rangle)$$

No other process knows about c and s because they are inside the name restriction: you can imagine that they will be created at run-time by the ν operator (pronounced “new”).

$$\mathbf{Client}\langle a \rangle \mid \mathbf{Server}\langle a \rangle$$

$$\longrightarrow (\nu c)(c(x).\mathbf{Client}_1\langle c, x \rangle \mid (\nu s)(\bar{c}\langle s \rangle \mid \mathbf{Server}_1\langle c, s \rangle))$$

$$\longrightarrow (\nu c, s)(\mathbf{Client}_1\langle c, s \rangle \mid \mathbf{Server}_1\langle c, s \rangle)$$

Note that \mathbf{Client}_1 and \mathbf{Server}_1 are now within the scope of the ν operator. This phenomenon, called *scope extrusion*, is a distinguishing feature of the π -calculus.

Secure client-server communication

We can now represent secure client-server communication. The client creates a new (secret) channel c before contacting the server on the public channel a :

$$\mathbf{SClient}_i(a) \stackrel{\text{df}}{=} (\nu c)(\bar{a}\langle c \rangle \mid c(x).\mathbf{PClient}_i\langle c, x \rangle)$$

The server is a replicated process, ready to spawn a *session* for each client request:

$$\mathbf{SServer}(a) \stackrel{\text{df}}{=} !a(y).(\nu s)(\bar{y}\langle s \rangle \mid \mathbf{PServer}\langle y, s \rangle)$$

The server can interact with multiple clients at the same time:

$$\begin{aligned} & \mathbf{SClient}_1\langle a \rangle \mid \mathbf{SClient}_2\langle a \rangle \mid \mathbf{SServer}\langle a \rangle \xrightarrow{*} \\ & (\nu s_1, c_1)(\mathbf{PClient}_1\langle c_1, s_1 \rangle \mid \mathbf{PServer}\langle c_1, s_1 \rangle) \\ & \mid (\nu s_2, c_2)(\mathbf{PClient}_2\langle c_2, s_2 \rangle \mid \mathbf{PServer}\langle c_2, s_2 \rangle) \\ & \mid \mathbf{SServer}\langle a \rangle \end{aligned}$$

Each session is protected from the external environment by the restriction on the client and server communication ports, which can be used to exchange data without external interferences.