

The Pi-Calculus

2 March 2021

0. What motivated the search for a process calculus?
1. What did prehistoric models look like?
2. How did the pi-calculus improve upon them?
3. What is The Truth for a process calculus?
4. How do you add discipline to the world?

Pre-History

In the late 70's / early 80's, researchers had a model of computation (λ -calculus), but no model for concurrent communication.

Researchers wanted to reason about processes running independently and in communication with each other.

Many 'process calculi' emerged:

- Actors
- CSP
- Petri Nets
- ...

Many models, little agreement.

1. Milner 1980 - CCS

Calculus of Communicating Systems (CCS).

Two Ideas:

1. A process can be defined based on how an Observer interacts with it. To build complex systems, you can take an observer and model it as a process that can be observed.
2. Processes communicate over channels pathways over which data can be sent between processes.

→ NOT a new idea.

Intro to CCS

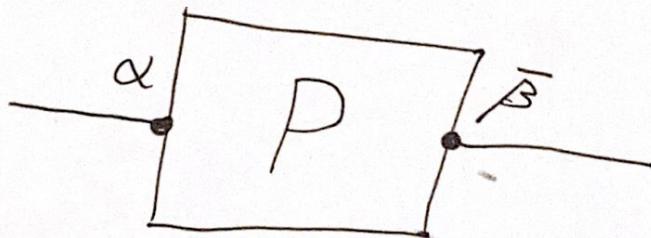
Preliminaries:

channels are defined by a set of names, unique symbols that differ from other values. We use $\alpha, \beta, \gamma, \delta, \dots$ and range over channels with μ .

Denote processes with P, Q, R .

We define processes via their interactions with processes over channels.

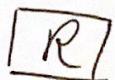
Channels are a static property of a process.



Example: Checkout at grocery store.

1. nil

- a closed checkout aisle.



2. receive

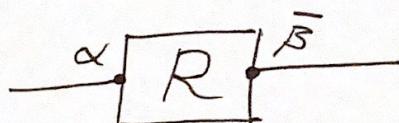
- the cashier receives \$5 from customer.



$$d(x). R' \xrightarrow{\alpha \$5} R'[x \leftarrow \$5]$$

3. send

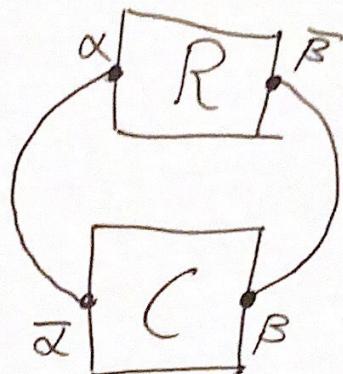
- the cashier gives the customer what they purchased.



$$\bar{\beta}_{\text{gum}}. R' \xrightarrow{\beta_{\text{gum}}} R'$$

4. Parallel (Par)

- A cashier and customer communicating.



$(\alpha(x). \bar{\beta} \text{ gum}. R' / \bar{\alpha} \$5. \beta(y). C')$

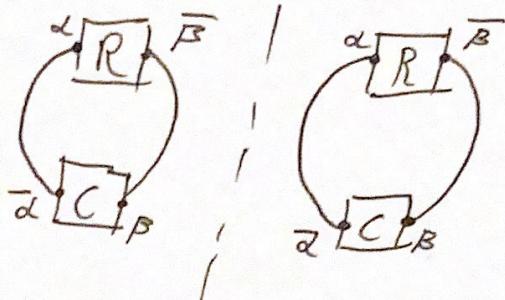
$\xrightarrow{\alpha \$5} (\bar{\beta} \text{ gum}. R'[x \leftarrow \$5] / \beta(y). C')$

$\xrightarrow{\beta \text{ gum}} (R'[x \leftarrow \$5] / C'[y \leftarrow \text{gum}])$

Note: From the perspective of an outside observer, they don't interact with the larger system. These 'unobserved' actions are called 'silent' actions labeled with τ -transitions ($\xrightarrow{\tau}$). This distinction is mostly useful for theoretical purposes.

5. Restriction

- Cashiers & customers in different aisles.



$$\frac{(\alpha \& \beta) \dots}{\text{Not the same!}} \quad | \quad (\alpha \& \beta) \dots$$

6. Summation / Non-Determinism

- The cashier might ask if you have a rewards card.

$$\frac{\alpha \& \beta}{\begin{array}{c} \xrightarrow{\alpha} \\ \bullet \end{array} \boxed{R} \xrightarrow{\beta}}$$

$$\xrightarrow{\alpha \& \beta} \alpha(x). R_1 + \alpha(x). R_2$$

$$\rightarrow R_2[x \leftarrow \$5]$$

7. Relabelling

- Channels must be renamed in the surface syntax.

$$\xrightarrow{\alpha \bullet \boxed{R} \xrightarrow{\beta}} \xrightarrow{[\alpha \leftarrow \gamma, \beta \leftarrow \delta]} \xrightarrow{\gamma \bullet \boxed{R} \xrightarrow{\delta}}$$

$$\begin{aligned} & (\alpha(x) \cdot \overline{\beta}_{\text{gum}}. R')[\alpha \leftarrow \gamma, \beta \leftarrow \delta] \\ \rightarrow & (\gamma(x) \cdot \overline{\delta}_{\text{gum}}. R'[\dots]) \end{aligned}$$

Problem: where did the Truth go?

In the λ -calculus, the Truth is Observational Equivalence.
Its meaning is derived from expressions reducing to values.

Processes don't do so; what do we do?

Insight: two processes are equivalent if they simulate each other.

A process simulates another process if it has the same interface (channels) and, after a transition, continues to simulate the other process.

This is called a bisimulation / bisimilarity relation.

Milner defines three such relations for CCS.

Example:

The 'strong equivalence' relation:

$P \sim Q$ iff

1. If $P \xrightarrow{\mu V} P'$, then for some Q' , $Q \xrightarrow{\mu V} Q'$ and $P' \sim Q'$.
2. If $Q \xrightarrow{\mu V} Q'$, then for some P' , $P \xrightarrow{\mu V} P'$ and $Q' \sim P'$.

Examples:

1.	$\bar{B} S. P'$	$\bar{x} y z. Q'$	X
2.	$\bar{B} S. P'$	$\bar{B} S. P'$	✓
3.	$\bar{B} S(\bar{x} b. \text{nil} d(y), \bar{y} y. \text{nil})$	$\bar{B} S. \bar{x}. \bar{y} b. \text{nil}$	✓
4.	$\bar{B} S(\bar{x} b. \text{nil} d(y), \bar{y} y. \text{nil})$	$\bar{B} S. \bar{y} b. \text{nil}$	X

We need a new relation to relate 4.

$$\begin{aligned}\xrightarrow{\mu V} &\stackrel{\text{def}}{=} \xrightarrow{\tau} \xrightarrow{\tau} \xrightarrow{\tau} \dots \xrightarrow{\mu V} \xrightarrow{\tau} \dots \xrightarrow{\tau} \\ \xrightarrow{\mu V} &= \xrightarrow{\tau^k} \xrightarrow{\mu V} \xrightarrow{\tau^k}\end{aligned}$$

$P \approx Q$ iff

1. if $P \xrightarrow{\mu V} P'$, then for some Q' , $Q \xrightarrow{\mu V} Q'$ and $P' \approx Q'$.
2. if $Q \xrightarrow{\mu V} Q'$, then for some P' , $P \xrightarrow{\mu V} P'$ and $Q' \approx P'$.

Flaw: Not a congruence relation.

Reflections on CCS

1. Channels are static. In the real world, communication links change dynamically, with time.
2. "...its primitive constructs deserve re-examination."
 - Are there other operators that make the calculus more expressive?
 - Are there unneeded operators?

Intermediate Steps - Ulfberg & Nielsen, 1986

Solution: allow processes to pass channels over channels.
Extended CCS (ECCS).

Two disjoint sets of variables:

- 'value' variables: $x, y, z \dots$
- name / channel variables: $a, b, c \dots$

Still use α, β to represent actual channels.

Example : Walkie - Talkies.

$$(\nu\beta)(\alpha\beta. \beta(x). P') \mid \alpha(talk). \overline{talk} ``hi". Q' \mid \overline{\gamma} ``hello?". R'$$

$$\xrightarrow{\perp\beta} (\nu\beta)(\beta(x). P' \mid \overline{\beta} ``hi". Q[talk \leftarrow \beta]) \mid \overline{\gamma} ``hello?". R'$$

$$\xrightarrow{\beta ``hi"} (\nu\beta)(P'[x \leftarrow ``hi"] \mid Q[talk \leftarrow \beta]) \mid \overline{\gamma} ``hello?". R'$$

ECCS Reflection

The capacity for the links between processes to change dynamically is called 'mobility' (CP channels).

↳ Channels no longer unidirectional.

This also means shrinking the calculus.

We no longer need relabelling, now that we've defined substitution over channels.

2. Milner, Parrow, Walker 1992 - π -calculus

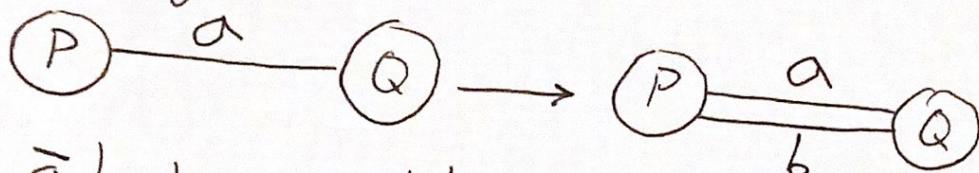
Small improvement: distinguishing between channel and value variables is unnecessary to achieve the same expressiveness.

The relatively small change of removing the distinction is what produces the π -calculus.

Characteristic Examples

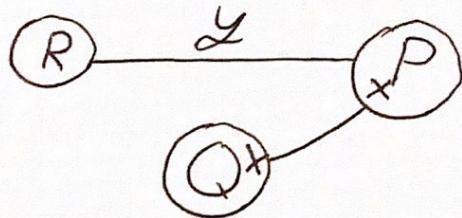
We do not use greek letters for channels anymore

1. Passing Channels



$$\begin{array}{l} \bar{a}b. b(x). P' | \alpha(\text{chan}). \text{chan } s. Q' \\ \xrightarrow{ab} b(x). P' | \bar{b}s. Q' [\text{chan} \leftarrow b] \\ \xrightarrow{bs} P'[x \leftarrow s] | Q' [\text{chan} \leftarrow b] \end{array}$$

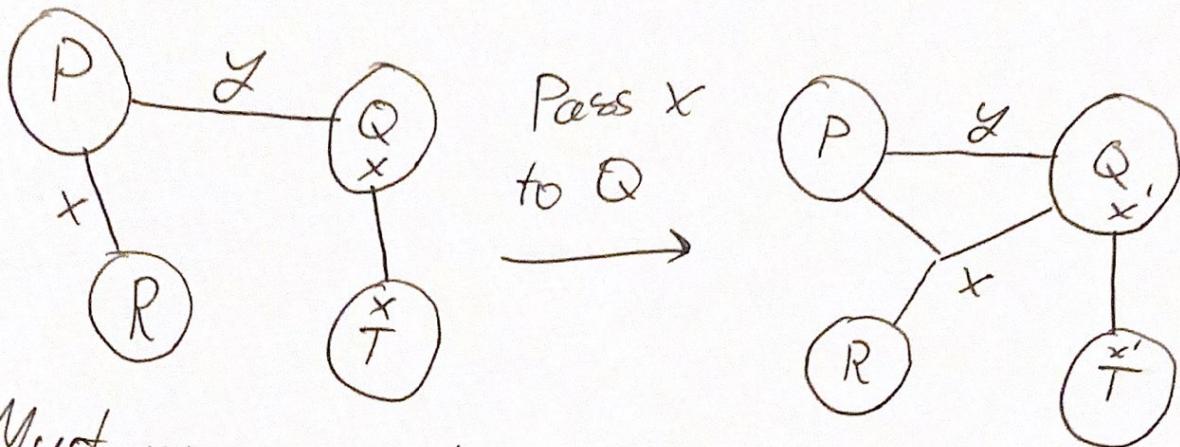
2. Restricted Channels



$$\begin{array}{l} (\forall x)(\bar{x}s. P' | x(y). Q') | R \\ \xrightarrow{x s} (\forall x)(P' | Q' [y \leftarrow s]) | R \end{array}$$

Restriction introduces channels like lambdas introduce variables.

3. Scope Intrusion



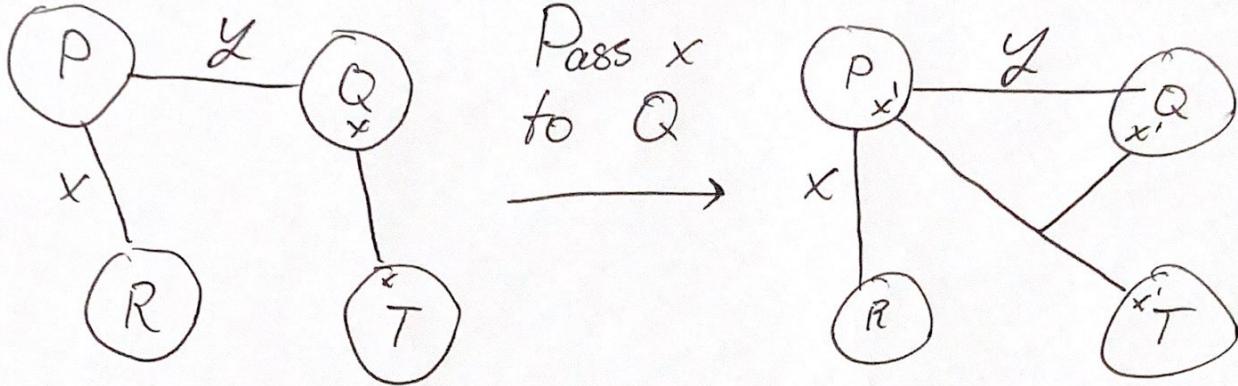
Must rename restricted channels.

$$\begin{array}{l}
 (\bar{y}x.P' \\
 | R \\
 | (\forall x) \\
 (\bar{y}(z).Q' \\
 | T))
 \end{array}$$

$\bar{y}x \rightarrow$

$$\begin{array}{l}
 (P' \\
 | R \\
 | (\forall x) \\
 (Q'[x \leftarrow x'][z \leftarrow x] \\
 | T[x \leftarrow x']))
 \end{array}$$

4o Scope Extrusion



Must rename restricted channels.

$(y(z).P'$

$| R$

$| (\nu x)$

$(\bar{y}x.Q'$

$| T))$

$\cancel{y x'}$

$((\nu x')$

$(P'[z \leftarrow x'])$

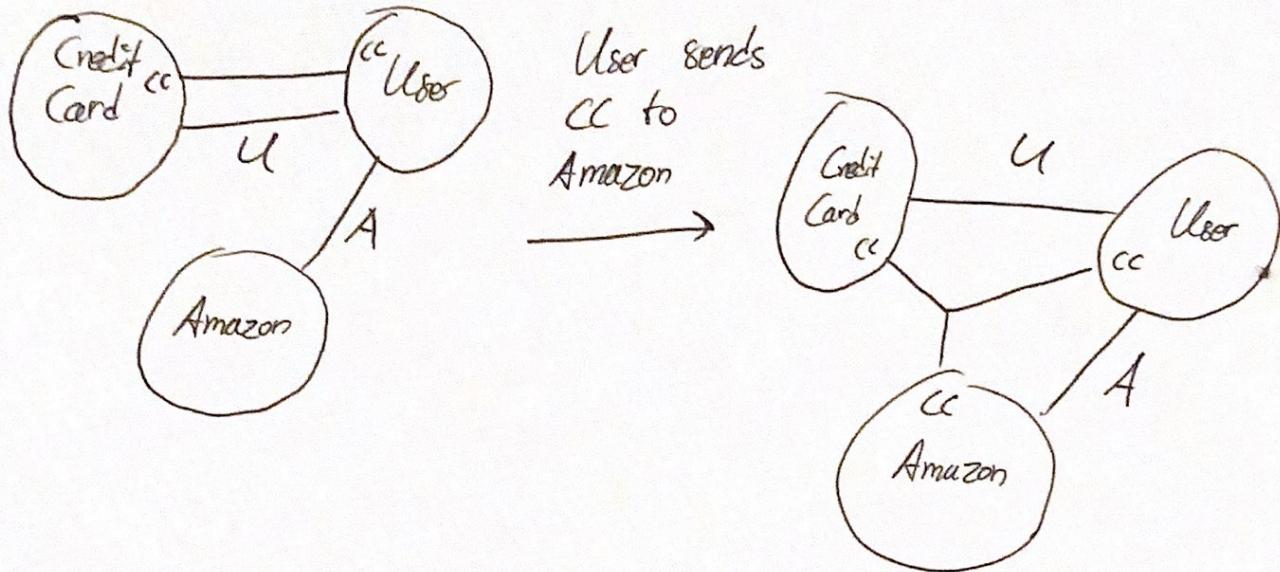
$| Q'[x \leftarrow x']$

$| T[x \leftarrow x'])$

$| R)$

Must hoist the
channel restriction up
so that x' is in
scope for P.

Another Example: Amazon Purchase



$(\text{CreditCard} = (\lambda)(\lambda)(\lambda)(\text{User}(\text{cc}), \text{User}(\text{charge}), \text{CreditCard}'))$

$\text{User} = \text{User}(\text{cc}), \text{User}(\text{book}), \text{User}(\text{cc}), \text{User}'$

$\text{Amazon} = \text{Amazon}(\text{product}), \text{Amazon}(\text{cc}), \text{Amazon}(\text{price_of_product}), \text{Amazon}'$

Final Note on The Truth:

Milner et al. introduce Five bisimilarity relations.

- strong bisimilarity
- strong equivalence
- strong D-Equivalence
- early bisimilarity
- late bisimilarity

And there's more where that came from:

- barbed congruence
- strong barbed congruence
- open bisimilarity

• • •

If there's 25 ways for characterizing
the Truth . . . do we really have the Truth?

Π-Calculus Reflection

A relatively small step puts the Π -calculus over the edge.

The Π -calculus goes on to dominate the world of process calculi.

Flaw: the world of concurrent communication is chaotic, and the Π -calculus offers no way of remedying that.

- Race Conditions
- Non-Determinism
- Dead-locking

Example: the molecular action:

$$(\nu x)(\bar{x}5.\bar{x}6.P') \mid x(y).x(z).Q' \mid x(y).x(z).R'$$

Values are NOT broadcast over channels.

$$\xrightarrow{x5} (\nu x)(\bar{x}6.P') \mid x(z).Q[y \leftarrow 5] \mid x(y).x(z).R'$$

$$\xrightarrow{x6} (\nu x)(P') \mid x(z).Q[y \leftarrow 5] \mid x(z).R[y \leftarrow 6]$$

We've failed! A race condition has occurred.

One possible solution: send a channel of communication first.

$$P = (\nu w)(\bar{x}w.\bar{w}5.\bar{w}6.P')$$

$$Q = x(w).w(y).w(z).Q'$$

$$R = x(w).w(y).w(z).R'$$

The π -calculus is expressive and powerful, but intricate.

3. Milner 1991 - A Sorting Discipline

Solution: add a sorting discipline to the π -calculus (analogous to a typing discipline).

A sorting discipline lets the programmer specify how a process must interact with a channel.

Intermediate step: Polyadic Π -calculus

Instead of sending single values, send and receive tuples of arbitrary arity.

$$P = \bar{x}(5, 6). P'$$

$$Q = x(y, z). Q'$$

$$R = x(y, z). R'$$

- Good: can pass multiple values at a time.
- Bad: Arity errors.

The Sorting Discipline

Intuition: we express the arities of each channel
and (recursively) the arities of the values passed
over it.

A channel has a sort S . Our sorting
contains the sorts of channels.

Example sorting:

$$\{ \text{SENDER} : [\text{LEFT}, \text{RIGHT}], \text{LEFT} : [], \text{RIGHT} : [] \}$$

Add type annotation:

$$(\forall x : [\text{LEFT}, \text{RIGHT}]) (\dots)$$

Note: we need names because our types are
recursive.

Inference Rules:

We show that a process P obeys the sorts of the channels it interacts with.

Preliminary:

- $[\tilde{x}] \stackrel{\text{def}}{=} [x_1, x_2 \dots x_n]$
- A judgement $\Gamma \vdash P$ should be read as: "Process P is well-behaved under Γ ."

$$\frac{\Gamma(x) = [\tilde{s}] \quad \Gamma[\tilde{b}]:[\tilde{s}] + P}{\Gamma \vdash x(\tilde{b}; \tilde{s}).P} \quad [T-IN]$$

$$\frac{\Gamma(x) = [\Gamma(\tilde{b})] \quad \Gamma \vdash P}{\Gamma \vdash \bar{x}\langle \tilde{b} \rangle.P} \quad [T-OUT]$$

$$\frac{\Gamma \vdash P \quad \Gamma \vdash Q}{\Gamma \vdash (P|Q)} \quad [T-PAR]$$

Reflection

The sorting guarantees no only runtime errors.

Is this what we want?

This discipline is satisfactory. We've learned relatively little about how to use channels.

Imagine if modern programming languages only let the programmer specify the arities of functions!

Example: office printer

job, paper \in channels.

Printer = job(x). paper(x). Printer

If another process reads on the job channel, they can steal a printing job!

Pierce & Sangiorgi: 1994 - I/O Types

Solution: specify the read-write capabilities of each channel.

Preliminaries:

$$\text{Polarities} = \{i, o\}$$

$$r = \{i\}$$

$$w = \{o\}$$

$$b = \{i, o\}$$

Lowercase p, q range over polarities

A sort now looks like: $p[s_1, s_2, \dots, s_n]$

Office Printer:

$(\forall \text{job}: r[s], \text{paper}: w[s]) (\text{job}(x). \overline{\text{paper}} x. \text{Printer})$

A client will receive the job channel with only write access:

Client = $\text{recv}(\text{job}: w[s]). \overline{\text{job}} \text{ my-talk. Client'}$

We need a way to show that a sort obeys the interface of a channel. For that, we introduce subsorts. We specify that:

$b \lessdot r$

$b \lessdot w$

Intuition: a sort is a subsort of another if the interface is supported and the children obey the relationships.

Example subsort rule:

Judgement $\Sigma \vdash S \lessdot T$ means "Sort S is a subsort of Sort T under Σ ."

$p \lessdot r$

For each i , $\Sigma \vdash s_i \lessdot t_i$

$\Sigma \vdash p[\tilde{s}] \lessdot r[\tilde{t}]$

Sorting rules with I/O & subtyping

Two rules are modified:

$$\frac{\Gamma(x) \lessdot r[\tilde{s}] \quad \Gamma[b] : [\tilde{s}] \vdash P}{\Gamma \vdash x(b : \tilde{s}) . P} \quad [T-IN]$$

$$\frac{\Gamma(x) \lessdot w[\Gamma(b)] \quad \Gamma \vdash P}{\Gamma \vdash x \langle b \rangle . P} \quad [T-OUT]$$

Reflection on I/O Types

We've specified a critical component of a channel's interface!

But, we're still missing another frequent part of a protocol.

Consider a web server:

- it receives any number of requests and a channel on which to reply (in π -calculus).
- it always sends one response per request.

It would be great if our typechecker could let the implementer know that they forgot to send a reply in a particular case.

Kobayashi, Pierce, Turner 1998 - Linear Types

Objective: allow the programmer to specify the usage of channels on top of the ~~other~~ I/O interface and the ~~others~~.

Linear type system will be leveraged to achieve this.

Introduction to Linear Types

Linear types let us specify & enforce the usage of resources. We assign to each resource a 'multiplicity' that indicates how many times it must be used.

This prevents two other features in type systems:

1. Weakening (Dropping):

(first $x:y$) $\Rightarrow x$ NOT ALLOWED

$\Gamma, x:T \vdash x:T$ X

$x:T \vdash x:T$ ✓

We say ' Γ unlimited' to specify an environment with no limited resources.

2. Contraction (Duplication):

(define double x) $(x \ x)$ NOT ALLOWED

$\Gamma \vdash v \in V \quad \Gamma \vdash w \in W$

$\frac{\Gamma \vdash v \in V \quad \Gamma \vdash w \in W}{\Gamma + \langle v, w \rangle \in V \times W}$ X

$\frac{\Gamma_1 \vdash v \in V \quad \Gamma_2 \vdash w \in W}{\Gamma_1 + \Gamma_2 \in \langle v, w \rangle \quad V \times W}$ ✓

Inference Rules:

m is either 1 or ω .
Write a type $p[s_1 \dots s_n]^m$.

$$\frac{\Gamma, z : \tilde{S} \vdash P \quad \Gamma^{\text{unlimited}}}{\Gamma + x : r[\tilde{S}]^m \vdash x(z). P} \quad [\Gamma\text{-IN}]$$

$$\frac{\Gamma^{\text{unlimited}}}{\Gamma + x : w[\tilde{S}] + z : \tilde{S} \vdash \bar{x}(z). P} \quad [\Gamma\text{-OUT}]$$

$$\frac{\Gamma_1 \vdash P \quad \Gamma_2 \vdash Q}{\Gamma_1 + \Gamma_2 \vdash P / Q} \quad [\Gamma\text{-PAR}]$$

How do we combine types from a split context?

The capabilities of our channels must be split:
a linear process can be read from one process
and written from one process, but no more.

$$p[S_1 \dots S_n]^\omega + q[S_1 \dots S_n]^\omega = (p \cup q)[S_1 \dots S_n]^\omega$$

$$p[S_1 \dots S_n]^1 + q[S_1 \dots S_n]^1 = (p \cup q)[S_1 \dots S_n]^1$$

$$\text{if } p \cap q = \emptyset$$

Reflection on Linear Types

Returning to the web server: we've begun to encode something more powerful than just the interface of a channel - we're beginning to specify protocols.

The element of typing in the π -calculus plays a big role in the future of work on communication and mobile processes (Andreas' talk)

Summary

What have we achieved?

- Our original goal — one process calculus to rule them all.
- The key to the model is mobility of channels.
- We've defined a new notion of the Truth for processes.
- We've introduced discipline into our world.

Shortcomings:

- There is no notion of the definitive Π -calculus.
- Nobody programs in the Π -calculus.

