

Using *linear forwarders*
to implement *input capability*
for implementing the pi calculus

Cosimo Laneve, Philippa Gardner, Lucian Wischik,

Concur, Marseilles September 2003
www.wischik.com/lu/research/linfwd.html

The previous year, at CONCUR, we presented work on our distributed implementation of the pi calculus. It used something called “fusions”. A fusion is something that, when two channels are fused, they can be used interchangeably.

Lots of people found fusions difficult to swallow... also, it seemed hard to make them robust in the presence of failure. So we went back to the drawing board, dissected our implementation, and tried to find out what the core idea inside it was. This talk is about that core idea, “linear forwarders”. We think that linear forwarders make a more simple and elegant story, and one that’s easier to implement.

channel-oriented distribution

```
new  $\overline{tunnel@mar}.$ (           // Create a fresh channel, at Marseilles
   $\overline{globtunnel}$            // extrude channel  $tunnel$ 
  |  $tunnel(svr).svr(c).\bar{c}$  // find the current server  $svr$ , ask it what
)                               // service  $c$  it provides, and use the service
|  $glob(t).\bar{t}server$       // Receive  $t$  and tell it our server
|  $\overline{server}alert$       // (this server allows people to alert)
```

Each channel ($tunnel$, $glob$, $server$, $alert$) has a fixed location.

To receive on a channel, must either

- already be at the location (eg. Join), or
- migrate to the location (eg. Fusion Machine), or
- delegate an **input-proxy** to that location (eg. Facile, current work)

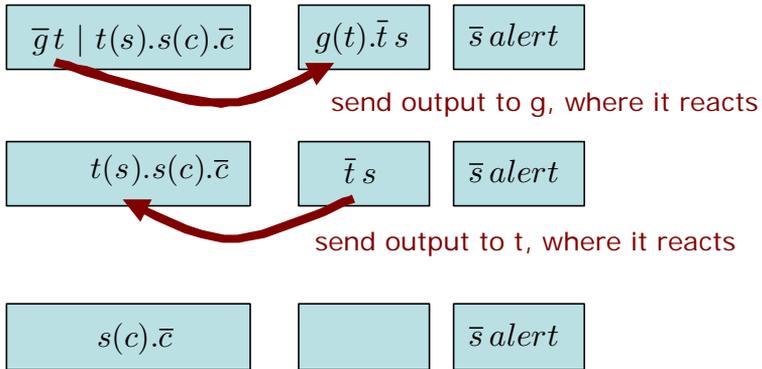
This example program does none, hence it fails at input on $svr...$

There are two philosophies for implementing concurrent calculi.

- (1) Mobile Ambients, Nomadic Pict, this group says that pi-calculus rendezvous is not appropriate for a distributed implementation. These have processes grouped together at locations.
- (2) Facile, Join, Fusion machine, Linear Forwarders, MS Biztalk, Highwire – this group says that pi-calculus rendezvous is a good primitive to have. In these, each channel belongs to a single location.

Within the second camp, we face the problem of input capability: as in this example, when a received name is used as the subject of subsequent input.

input capability



Fails because $s(c)$ is at the wrong location to react further.

This is **input capability**, where a received name s is used as the subject of subsequent input.

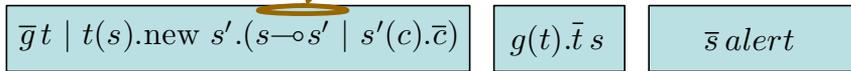
We'll pre-compile, replacing *input-capability* with *input-proxies*...

This is the step-by-step execution trace of the previous program.

encoding input capability

Pre-compile pi into an intermediate, implementable language:
the **localised linear forwarder calculus**

the linear forwarder will be sent to s
it will transform just one output from s to s'



... then, after a few outputs/reactions...



linear forwarder goes to s



then will transform the output to s'...

(Step-by-step execution trace after we pre-compile the program into linear forwarders, avoiding input mobility).

Facile was a bit similar. Facile was a synchronous calculus – ie. it had continuations after SEND as well as after RECEIVE. That meant that neither SEND nor RECEIVE could move around in the network. Instead, Facile used output-proxies as well as input-proxies. But it got stuck a bit, because all these proxies made for a difficult protocol that was hard to make robust in the presence of failures.

How do linear forwarders behave in the presence of failure? Well, if a linear forward fails to be delivered to its intended target, that's gives the same results as if it had reached the target and had caused an output to be transformed but then the output was subsequently lost.

linear forwarder calculus

The **linear forwarder** calculus:

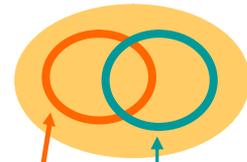
$$P ::= \mathbf{0} \mid \bar{x}\tilde{y} \mid x(\tilde{y}) \mid \nu x.P$$

$$P|P \mid !P \mid x \multimap y$$

reaction:

$$\bar{u}\tilde{x} \mid u(\tilde{y}).P \rightarrow P\{\tilde{x}/\tilde{y}\}$$

$$\bar{u}\tilde{x} \mid u \multimap u' \rightarrow \bar{u}'\tilde{x}$$



pi subcalculus:
subset without
linear forwarders
 $t(s).s(c).\bar{c}$

Localised subcalculus:
subset without input
capability, ie.
 $u(x).P$ never inputs on x
 $t(s).\text{new } s'.(s \multimap s' \mid s'(c).\bar{c})$

We define the entire linear forwarder calculus first, and then identify two subsets within it.

The linear forwarder calculus is just like the standard pi calculus, but with the addition of a “linear forwarder” term $x \multimap y$, and a new reduction rule for it.

linear forwarder calculus

The **linear forwarder** calculus:

$$P ::= \mathbf{0} \mid \bar{x}\tilde{y} \mid x(\tilde{y}) \mid \nu x.P$$

$$P|P \mid !P \mid x \multimap y$$

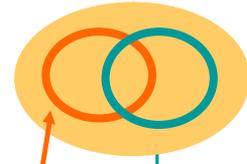
We encode **pi** into **Localised**:

$$\llbracket x(\tilde{y}).P \rrbracket = x(y).\llbracket P \rrbracket_y$$

The 'y' subscript means that this name has been received, and so can't be used as subject of subsequent input. Hence, in full...

$$\llbracket x(\tilde{y}).P \rrbracket_{\tilde{u}} = \begin{cases} x(\tilde{y}).\llbracket P \rrbracket_{\tilde{u}\tilde{y}} & \text{if } x \notin \tilde{u} \\ \nu x'.(x \multimap x' \mid x'(\tilde{y}).\llbracket P \rrbracket_{\tilde{u}\tilde{y}}) & \text{if } x \in \tilde{u} \end{cases}$$

$$\llbracket - \rrbracket_{\tilde{u}} = -$$



pi subcalculus:
subset without
linear forwarders
 $t(s).s(c).\bar{c}$

Localised subcalculus:
subset without input
capability

$$t(s).\text{new } s'.(s \multimap s' \mid s'(c).\bar{c})$$

As discussed previously, we pre-compile from pi calculus into a localised linear calculus which lacks input mobility.

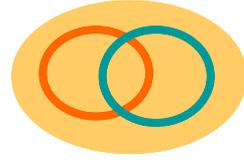
ground bisimulation \approx_ℓ

It's all standard, but with one extra rule:

$$u(\tilde{x}).P \xrightarrow{u(\tilde{x})} P$$

$$\bar{u}\tilde{x} \xrightarrow{\bar{u}\tilde{x}} \mathbf{0}$$

$$u \dashv\vdash v \xrightarrow{u(\tilde{x})} \bar{v}\tilde{x}$$



$$\frac{P \xrightarrow{\mu} P' \quad y \notin \mu}{(y)P \xrightarrow{\mu} (y)P'}$$

$$\frac{P \xrightarrow{(\tilde{z})\bar{u}\tilde{x}} P' \quad y \neq u, y \in \tilde{x} \setminus \tilde{z}}{(y)P \xrightarrow{(y\tilde{z})\bar{u}\tilde{x}} P'}$$

$$\frac{P|!P \xrightarrow{\mu} P'}{!P \xrightarrow{\mu} P'}$$

$$\frac{P \xrightarrow{\mu} P' \quad \text{bn}(\mu) \cap \text{fn}(Q) = \emptyset}{P|Q \xrightarrow{\mu} P'|Q}$$

$$\frac{P \xrightarrow{(\tilde{z})\bar{u}\tilde{y}} P' \quad Q \xrightarrow{u(\tilde{x})} Q' \quad \tilde{z} \cup \text{fn}(Q) = \emptyset}{P|Q \xrightarrow{\tau} (\tilde{z})(P'|Q'\{\tilde{y}/\tilde{x}\})}$$

Let *Ground bisimulation* be the largest symmetric relation \approx_ℓ such that $P \approx_\ell Q$ and $P \xrightarrow{\mu} P'$ implies $Q \xrightarrow{\mu} Q'$.
Common result: in asynchronous calculi, \approx_ℓ is a congruence.

The labelled transition system for the linear forwarder calculus.

- Pi contexts are exactly as discriminating as (localised) linear-forwarder contexts.

Or, by limiting general input capability to only input capability of linear forwarders, we keep the same expressive power of terms and the same distinguishing power of contexts.

$$P \approx_{\pi} Q \quad \text{iff} \quad \llbracket P \rrbracket \approx \llbracket Q \rrbracket$$

- Ground bisimulation helps as a proof technique

$$P \approx_{\ell} \llbracket P \rrbracket \approx_{\ell} \llbracket P \rrbracket_{\tilde{u}} \quad \approx \supset \approx_{\ell}|_{\text{localised}}$$

but it's too strict in general, eg. for this standard result from equators:

$$u(x).P \approx u(x').\nu x.(!x \multimap x' \mid !x' \multimap x \mid P)$$

$$\not\approx_{\ell}$$

ground bisimulation \approx_{ℓ}

standard pi barbed congruence \approx_{π}

localised barbed congruence \approx

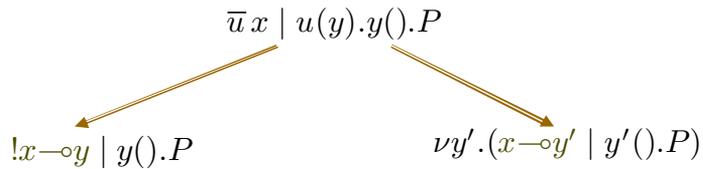
The first result, “full abstraction”, is what justifies our work.

We really have three separate equivalence relations that we’re juggling with. The ground bisimulation from the previous slide is helpful in proofs – it’s as though it were tailor-made specially to prove our encoding correct! But it’s too strict for general use. So instead we use barbed congruence.

A “congruence” is a relation that’s closed under contexts. Obviously you have to specify what those contexts should be. The standard barbed congruence is closed under pi-contexts. Our localised barbed congruence is closed under localised-contexts. We might also have defined a “linear forwarder congruence” that was closed under all linear forwarder contexts, but there didn’t seem much use. In any case, thanks to the ground-bisimulation results about the encoding, it’d be just the same.

It’s worth being careful about the statement of ground bisimulation. “If two terms are in the localised subcalculus, and they are judged equivalent according to ground bisimulation, then they are also judged equivalent by localised-barbed-congruence”.

linearity of forwarders



A **non-linear** use of forwarders:

Reaction gives a substitution $\{x/y\}$.
This can be implemented by forwarding all future x s to y .

(similar to equator work by Mero,
pi-1 work by Boreale, our earlier
work on fusions)

Our **linear** use of forwarders:

Reaction gives a substitution $\{x/y\}$.
In our subsequent source code we
have only one input on x , so we need
forward only one x to our input.

Non-linearity here would be an error:



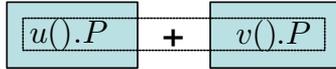
even after our input $y'()$ had been
used up, it'd keep forwarding to y' , so
denying other inputs their chance.

The interesting thing is about linearity. Really, we have given a calculus that uses forwarders for the code P that's inside a term. There's a finite amount of code, and so we have LINEAR forwarders.

Other people have used forwarders for the possible data that might be changed subsequently. There's a potentially infinite amount of external data that might have to be changed, so these people have to use NON-LINEAR forwarders.

(Actually, the equation given here doesn't quite work as it is. It needs either a non-linear forwarder in the reverse direction, giving equators, or it needs some kind of protocol around it as used by Boreale, or it needs a more complicated mechanism as we did with fusions.)

distributed choice



Distributed choice is hard: any proposed reaction at u would have to ask permission from v before proceeding.

"Where is the $+$ implemented?"

So we encode it with linear forwarders:

forward u and v both to a common location.

(and after a choice has been made, **undo** the other forwarder)

$$\begin{aligned} & \llbracket u().P + v().Q \rrbracket \\ & = \nu u' v'. (u \multimap u' \mid v \multimap v' \mid u'().(P \mid v' \multimap v) + v'().(Q \mid u' \multimap u)) \end{aligned}$$

$u \multimap u'$

$v \multimap v'$

$u'().\dots + v'().\dots$

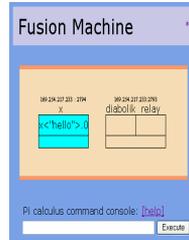
the fusion project at bologna

www.cs.unibo.it/fusion

Cosimo Laneve
Lucian Wischik
Manuel Mazzara
Fabrizio Bisi
Paolo Milazzo



prototype
implementation
in Java



distributed
implementation
in C++/Win32

```
spawn |
  channel f =
    new channel(Result);
  mailer.send(email,
    sub,body, response);
  K.recv(r: Result);
  if (r!="ok")
    pending.send(email);
}
```

Java impl.
with XML
data-types

The first two implementations, by Wischik, used fusions. The third implementation, by Paolo and Fabrizio, uses linear forwarders.