Linear Forwarders[*]

Philippa Gardner   Cosimo Laneve   Lucian Wischik

June 2003

Abstract

A *linear forwarder* is a process which receives one message on a channel and sends it on a different channel. Such a process allows for a simple implementation of the asynchronous pi calculus, by means of a direct encoding of the pi calculus' *input capability* (that is, where a received name is used as the subject of subsequent input). This encoding is fully abstract with respect to barbed congruence.

Linear forwarders are actually the basic mechanism of an earlier implementation of the pi calculus called the *fusion machine*. We modify the fusion machine, replacing fusions by forwarders. The result is more robust in the presence of failures, and more fundamental.

1   Introduction

Distributed interaction has become a necessary part of modern programming languages. We regard the asynchronous pi calculus as a basis for such a language. In the pi calculus, a program (or process) has a collection of channels, and it runs through interaction over these channels. A possible distributed implementation is to let each channel belong to a single location. For instance, there is one location for the channels u, v, w and another for x, y, z, and the input resource $u(a).P$ goes in the first location. If an output $\overline{u}\,x$ should arise anywhere else in the system, it knows where it can find a matching input. This basic scheme is used in the join calculus [7], in the $\pi_{1\ell}$ calculus [3], and in the fusion machine [8]. (A different approach is taken in Dπ [2], in nomadic pict [19], and in the ambient calculus [6], where agent migration is used for remote interaction.)

We immediately face the problem of **input capability**, which is the ability in the pi calculus to receive a channel name and subsequently accept input on it. Consider the example $x(u).u(v).Q$. This program is located at (the location of) x, but upon reaction with $\overline{x}\,w$ it produces the continuation $w(v).Q\{w/u\}$ – and this continuation is still at x, whereas it should actually be at w. Solving the problem of input capability is the key challenge in distributing the pi calculus.

---

[*]Available at www.wischik.com/lu/research/linfwd.html. This paper appears in the proceedings of CONCUR 2003, where one proof and some remarks were omitted for lack of space.

The point of this paper is to solve the problem of input capability with a language that is "just right" – it neither disallows more features than necessary (as does the join calculus), nor adds more implementation work than is necessary (as does the fusion machine). One measure of our solution is that we obtain full abstraction with the asynchronous pi calculus, up to weak barbed congruence.

First of all, let us consider in more detail the other solutions to input capability. The join calculus and localised pi calculus [13] simply disallow it: that is, in a term $x(u).P$, the P may not contain any inputs on channel u. The problem now is how to encode input capability into such a *localised* calculus. An encoding is possible, but awkward: when the term $x(u).u(v).Q \mid \overline{x} w$ is encoded and then performs the reaction, it does not perform the substitution $\{w/u\}$, but rather encodes this substitution as a persistent forwarder between w and u. Next, a firewall is needed to protect the protocol used by these forwarders. (The forwarder is called a "merged proxy pair" in the join calculus).

The fusion machine instead implements input capability through the runtime migration of code. In our example, $w(v).Q\{w/u\}$ would migrate from x over to w after the interaction. The migration is costly however when the continuation Q is large. In addition, code migration requires an elaborate infrastructure. To mitigate this, a large amount of the work on the fusion machine involved an encoding of arbitrary programs into *solos* programs (ones which have only simple continuations) without incurring a performance penalty. But the encoding used fusions, implemented through persistent trees of forwarders, which seem awkward and fragile in the presence of failures.

The solution presented in this paper is to disallow general input capability, and to introduce instead a limited form of input, the **linear forwarder**. A linear forwarder $x \multimap y$ is a process which allows just one x to be turned into a y. The essential point is that this limited form can be used to easily encode general input capability. For example, consider the pi calculus term $x(u).u(v).Q$. We will encode it as

$$x(u).(u')(u \multimap u' \mid u'(v).Q')$$

where the input $u(v)$ has been turned into a local input $u'(v)$ at the same location as x, and where the forwarder allows one output on u to interact with u' instead. The encoding has the property that if the forwarder $u \multimap u'$ exists, then there is guaranteed to be an available input on u'. We remark that linearity is crucial: if the forwarder persisted, then the guarantee would be broken; any further u turned into u' would become inert since there are no other inputs on u'.

One might think of a linear forwarder $x \multimap y$ as the pi calculus agent $x(u).\overline{y} u$ located at x. This agent would be suitable for a point-to-point network such as the Internet. But we have actually turned forwarders into first-class operators in order to abstract away from any particular implementation. This is because other kinds of networks benefit from different implementations of linear forwarders. In a broadcast network, $x \multimap y$ might be located at y; whenever it hears an offer of $\overline{x} u$ being broadcast, the machine at y can take up the offer. Another possibility is to use a shared tuple-space such as Linda [9], and ignore

2

all linearity information. (The fusion machine also amounts to a shared state which ignores linearity).

In this paper we show how to encode the pi calculus into a linear forwarder calculus. Conversely, we also show how linear forwarders can be encoded into the pi calculus. We therefore obtain full abstraction with respect to barbed congruence.

We also describe a *linear forwarder machine*. It is a simplified form of our earlier fusion machine, and more robust with respect to failures. This machine gives an implementation of distributed rendezvous which can be performed locally. In this respect it is different from Facile [10], which assumes a three-party handshake. This handshake is a protocol for interaction, and so prevents full abstraction. We prove full abstraction between the machine and the linear forwarder calculus, with respect to barbed congruence.

Related work. Forwarders have already been studied in detail by the pi community. Much work centres around the πI calculus [17] – a variant of the pi calculus in which only private names may be emitted, as in $(x)\overline{u}\,x$. Boreale uses forwarders to encode the emission of free names [4]: the reaction $u(a).Q \mid \overline{u}\,x$ does not perform the substitution $\{x/a\}$, but instead encodes it as a persistent forwarder between a and x. The same technique is used by Merro and Sangiorgi [13] in proofs about the localised pi calculus; and both are inspired by Honda's *equators* [11], which are bidirectional forwarders. Something similar is also used by Abadi and Fournet [1]. When channels are used linearly, Kobayashi et al. [12] show that a linear forwarder can simulate a substitution.

We remark upon some differences. If substitutions are encoded as persistent forwarders, then the ongoing execution of a program will create steadily more forwarders. In contrast, we perform substitution directly, and in our setting the number of forwarders decreases with execution. More fundamentally, the πI calculus uses forwarders to effect the substitution of data, and they must be persistent (nonlinear) since the data might be used arbitrarily many times by contexts. We use forwarders to effect the input capability of code, and this is linear because a given piece of source code contains only finitely many input commands. Our proofs are similar in structure to those of Boreale, but are much simpler due to linearity.

Structure. The structure of this paper is as follows. Section 2 gives the linear forwarder calculus, and shows how to encode the pi calculus (with its input mobility) into this calculus. Section 3 gives bisimulations for the linear forwarder calculus, and Section 4 proves full abstraction of the pi calculus encoding. Section 5 describes a distributed abstract machine for implementing the linear forwarder calculus, and Section 6 proves full abstraction for this implementation. We outline future developments in Section 7.

3

2 The Linear Forwarder Calculus

We assume an infinite set $\mathcal{N}$ of *names* ranged over by $u, v, x, \ldots$. Names represent communication channels, which are also the values being transmitted in communications. We write $\widetilde{x}$ for a (possibly empty) finite sequence $x_1 \cdots x_n$ of names. *Name substitutions* $\{\widetilde{y}/\widetilde{x}\}$ are as usual.

Definition 1 (Linear forwarder calculus) *Terms are given by*

$$P \ ::= \ \mathbf{0} \ \mid \ \overline{x}\,\widetilde{y} \ \mid \ x(\widetilde{y}).P \ \mid \ (x)P \ \mid \ P|P \ \mid \ !P \ \mid \ x{\multimap}y$$

Structural congruence $\equiv$ is the smallest equivalence relation satisfying the following and closed with respect to contexts and alpha-renaming:

$$P|\mathbf{0} \equiv P \qquad P|Q \equiv Q|P \qquad P|(Q|R) \equiv (P|Q)|R \qquad !P \equiv P|!P$$
$$(x)(y)P \equiv (y)(x)P \qquad (x)(P|Q) \equiv P \mid (x)Q \quad \textit{if } x \notin \text{fn}\, P$$

Reaction is the smallest equivalence satisfying the following and closed under $\equiv$, $(x)_{\_}$ and $_{\_} \mid {_{\_}}$:

$$u(\widetilde{x}).P \mid \overline{u}\,\widetilde{y} \ \rightarrow \ P\{\widetilde{y}/\widetilde{x}\} \qquad \overline{x}\,\widetilde{u} \mid x{\multimap}y \ \rightarrow \ \overline{y}\,\widetilde{u}$$

The operators in the syntax are all standard apart from the linear forwarder $x{\multimap}y$. This allows one output on x to be transformed into one on y, through the second reaction rule. In the output $\overline{x}\,\widetilde{y}$ and the input $x(\widetilde{y}).P$, the name x is called the *subject* and the names $\widetilde{y}$ are the *objects*. In the restriction $(x)P$, the name x is said to be *bound*. Similarly, in $x(\widetilde{y}).P$, the names $\widetilde{y}$ are bound in P. The *free names* in P, denoted $\text{fn}(P)$, are the names in P with a non-bound occurrence. We write $(x_1 \cdots x_n)P$ for $(x_1)\cdots(x_n)P$.

Next we make a *localised* sub-calculus, by adding the *no-input-capability* constraint. It is standard from the π_L calculus [13] and the join calculus that such a constraint makes a calculus amenable to distributed implementation.

Definition 2 (Localised calculus) *The* localised *linear forwarder calculus, which we abbreviate $L\ell$, is the sub-calculus of the linear forwarder calculus which satisfies the* no-input-capability *constraint: in $x(\widetilde{u}).P$, the P has no free occurrence of $\widetilde{u}$ as the subject of an input.*

We remark that the no-input-capability constraint is preserved by structural congruence and by reaction.

The asynchronous pi calculus [5] is a sub-calculus of the linear forwarder calculus, obtained by dropping linear forwarders. We give an encoding of the asynchronous pi calculus into the *localised* linear forwarder calculus $L\ell$, showing that the input capability can be expressed using forwarders and local inputs. Henceforth, when we refer to the pi calculus, we mean the asynchronous pi calculus.

Definition 3 (Encoding pi) *The encoding $\llbracket \cdot \rrbracket$ maps terms in the pi calculus into terms in the $L\ell$ calculus as follows. (In the input and restriction cases, assume that the bound names do not clash with $\widetilde{u}$.) Define $\llbracket P \rrbracket = \llbracket P \rrbracket_\emptyset$, where*

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

$$\llbracket (x)P \rrbracket_{\widetilde{u}} \quad = \quad (x)(\llbracket P \rrbracket_{\widetilde{u}})$$

$$\llbracket P|Q \rrbracket_{\widetilde{u}} \quad = \quad \llbracket P \rrbracket_{\widetilde{u}} \mid \llbracket Q \rrbracket_{\widetilde{u}}$$

$$\llbracket !P \rrbracket_{\widetilde{u}} \quad = \quad !\llbracket P \rrbracket_{\widetilde{u}}$$

$$\llbracket \overline{x}\,\widetilde{y} \rrbracket_{\widetilde{u}} \quad = \quad \overline{x}\,\widetilde{y}$$

$$\llbracket \mathbf{0} \rrbracket_{\widetilde{u}} \quad = \quad \mathbf{0}$$

To understand the encoding, note that we use "primed" names to denote local copies of names. So the encoding of $x(u).u(y).P$ will use a new channel u' and a process $u'(y).P$, both at the same location as x. It will also create exactly one forwarder $u \multimap u'$, from the argument passed at runtime to u'. Meanwhile, any output use of u is left unchanged.

To illustrate the connection between the reactions of a pi term and of its translation, we consider the pi calculus reduction $\overline{u}\,y \mid u(x).P \;\rightarrow\; P\{y/x\}$. By translating we obtain:

$$\llbracket\, \overline{u}\,y \mid u(x).P \,\rrbracket_u \quad = \quad \overline{u}\,y \mid (u')(u \multimap u' \mid u'(x).\llbracket P \rrbracket_{xu})$$
$$\rightarrow \quad (u')(\overline{u}'y \mid u'(x).\llbracket P \rrbracket_{xu})$$
$$\rightarrow \quad (u')(\llbracket P \rrbracket_{xu}\{y/x\})$$
$$\equiv \quad \llbracket P \rrbracket_{xu}\{y/x\}$$

Note that the final state of the translated term is subscripted on x and u, not just on u. In effect, the translated term ends up with some garbage that was not present in the original. Because of this garbage, it is not in general true that $Q \rightarrow Q'$ implies $\llbracket Q \rrbracket \rightarrow^* \llbracket Q' \rrbracket$; instead we must work up to some behavioural congruence. The following section deals with barbed congruence.

We remark that linearity is crucial in the translation. For instance, consider a non-linear translation where forwarders are replicated:

$$\llbracket u(x).P \rrbracket_u \quad = \quad (u')(!u \multimap u' \mid u'(y).P)$$

Then consider the example

$$\llbracket u().P \mid u().Q \mid \overline{u} \mid \overline{u} \rrbracket_u \quad = \quad (u')(!u \multimap u' \mid u'().P) \mid (u'')(!u \multimap u'' \mid u''().Q) \mid \overline{u} \mid \overline{u}$$
$$\rightarrow \quad (u')(P \mid !u \multimap u') \mid (u'')(!u \multimap u'' \mid u''().Q) \mid \overline{u}$$
$$\rightarrow \quad (u')(P \mid \overline{u}' \mid !u \multimap u') \mid (u'')(!u \multimap u'' \mid u''().Q)$$

Here, both outputs were forwarded to the local name u', even though the resource $u'().P$ had already been used up by the first one. This precludes the second one from reacting with Q — a reaction that would have been possible in the original pi calculus term. We need linearity to prevent the possibility of such dead ends.

3 Bisimulation and Congruence

We use barbed congruence [15] as our semantics for the $L\ell$ calculus.

Definition 4 (Barbed congruence) *The observation relation $P \downarrow u$ is the smallest relation generated by*

$$\overline{u}\,\widetilde{x} \downarrow u \qquad\qquad\qquad P|Q \downarrow u \quad \text{if } P \downarrow u \text{ or } Q \downarrow u$$
$$(x)P \downarrow u \quad \text{if } P \downarrow u \text{ and } u \neq x \qquad !P \downarrow u \quad \text{if } P \downarrow u$$

We write $\Downarrow$ for $\rightarrow^* \downarrow$ and $\Rightarrow$ for $\rightarrow^*$. A symmetric relation $\mathcal{R}$ is a weak barbed bisimulation *if whenever $P \mathrel{\mathcal{R}} Q$ then*

1. *$P \Downarrow u$ implies $Q \Downarrow u$*

2. *$P \rightarrow P'$ implies $Q \Rightarrow Q'$ such that $P' \mathrel{\mathcal{R}} Q'$*

Let $\overset{\cdot}{\approx}$ be the largest weak barbed bisimulation. Two terms P and Q are weak barbed congruent *in the $L\ell$ calculus when, for every C, then $C[P] \overset{\cdot}{\approx} C[Q]$, where $C[P]$ and $C[Q]$ are assumed to be terms in the $L\ell$ calculus. Let $\approx$ be the least relation that relates all congruent terms.*

We remark that barbed bisimulation $\overset{\cdot}{\approx}$ is defined for the linear forwarder calculus. However, the weak barbed *congruence* $\approx$ is predicated upon the $L\ell$ sub-calculus. Similar definitions may be given for the pi calculus, and, with abuse of notation, we keep $\overset{\cdot}{\approx}$ and $\approx$ denoting the corresponding semantic relations.

As an example of $\approx$ congruent terms in the $L\ell$ calculus, we remark that

$$u(x).P \quad \approx \quad u(x').(x)(!x\!\multimap\!x' \mid !x'\!\multimap\!x \mid P). \tag{1}$$

This is a straightforward variant of a standard result for equators [14], and we use it in Lemma 9.

Our overall goal is to prove that the encoding $[\![\cdot]\!]$ preserves the $\approx$ congruence. The issue, as described near the end of the previous section, is that an encoded term may leave behind garbage. To show that it is indeed garbage, we must prove that $[\![P]\!]_u$ and $[\![P]\!]_{ux}$ are congruent. But the barbed semantics offer too weak an induction hypothesis for this proof. A standard alternative technique (used for example by Boreale [4]) is to use barbed semantics as the primary definition, but then to use in the proofs a labelled transition semantics and its corresponding bisimulation — which is stronger than barbed congruence. The remainder of this section is devoted to the labelled semantics.

Definition 5 (Labelled semantics) *The labels, ranged over by μ, are the standard labels for interaction $\overset{\tau}{\longrightarrow}$, input $\overset{u(\widetilde{x})}{\longrightarrow}$ and possibly-bound output $\overset{(\widetilde{z})\overline{u}\,\widetilde{x}}{\longrightarrow}$ where $\widetilde{z} \subseteq \widetilde{x}$. The bound names $\mathrm{bn}(\mu)$ of these input and output labels are $\widetilde{x}$*

and $\widetilde{z}$ respectively.

$$u(\widetilde{x}).P \xrightarrow{u(\widetilde{x})} P \qquad \overline{u}\,\widetilde{x} \xrightarrow{\overline{u}\,\widetilde{x}} \mathbf{0} \qquad u{\multimap}v \xrightarrow{u(\widetilde{x})} \overline{v}\,\widetilde{x}$$

$$\frac{P \xrightarrow{\mu} P' \quad y \notin \mu}{(y)P \xrightarrow{\mu} (y)P'} \qquad \frac{P \xrightarrow{(\widetilde{z})\overline{u}\,\widetilde{x}} P' \quad y \neq u, \ y \in \widetilde{x} \backslash \widetilde{z}}{(y)P \xrightarrow{(y\widetilde{z})\overline{u}\,\widetilde{x}} P'}$$

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

$$\frac{P \xrightarrow{(\widetilde{z})\overline{u}\,\widetilde{y}} P' \quad Q \xrightarrow{u(\widetilde{x})} Q' \quad \widetilde{z} \cup \mathrm{fn}(Q) = \emptyset}{P|Q \xrightarrow{\tau} (\widetilde{z})(P'|Q'\{\widetilde{y}/\widetilde{x}\})}$$

The transitions of $P|Q$ have mirror cases, which we have omitted. We implicitly identify terms up to alpha-renaming $\equiv_\alpha$: that is, if $P \equiv_\alpha \xrightarrow{\mu} P'$ then $P \xrightarrow{\mu} P'$. We write $\overset{\mu}{\Longrightarrow}$ for $\xrightarrow{\tau}{}^* \xrightarrow{\mu}$.

A symmetric relation $\mathcal{R}$ is a weak labelled bisimulation if whenever $P \mathcal{R} Q$ then $P \xrightarrow{\mu} P'$ implies $Q \overset{\mu}{\Longrightarrow} Q'$. Let $\approx_\ell$ be the largest labelled bisimulation.

This definition is given for terms in the full linear forwarder calculus. It is a standard result that $\approx_\ell$ is a congruence with respect to contexts in the full calculus, and hence also with respect to contexts in the $L\ell$ calculus and the pi calculus. The connection between labelled and barbed semantics is also standard:

Lemma 6 *In the $L\ell$ calculus,*

1. *$P \to P'$ iff $P \xrightarrow{\tau} \equiv P'$.*

2. *$P \downarrow u$ iff $P \xrightarrow{(\widetilde{z})\overline{u}\,\widetilde{x}} P'$.*

3. *$\approx_\ell \subset \approx$.*

The bisimulation $\approx_\ell$ allows for some congruence properties to be proved trivially: (the first will be used in Proposition 10)

$$\begin{aligned} u{\multimap}v \ &\approx_\ell \ u(\widetilde{x}).\overline{v}\,\widetilde{x} \ \approx_\ell \ (u')(u{\multimap}u' \mid u'(\widetilde{x}).\overline{v}\,\widetilde{x}). \qquad (2) \\ u(x).P \ &\approx_\ell \ (u')(u{\multimap}u' \mid u'(x).P). \end{aligned}$$

We remark that Equation 1 does not hold when we use labelled bisimulation $\approx_\ell$ instead of barbed congruence. In particular, in the case $P = \mathbf{0}$, the right hand side undergoes $\xrightarrow{u(x')x'(\widetilde{v})}$ but the left does not. This illustrates why we chose barbed congruence as our primary measure process equivalence, and left labelled bisimulation as a technical tool for proofs.

4 Full Abstraction for the Pi Calculus Encoding

The $L\ell$ calculus is fully abstract with respect to the pi calculus encoding: $P \approx Q$ in pi if and only if $[\![P]\!] \approx [\![Q]\!]$ in $L\ell$. Informally, this is because the pi calculus input capability can be encoded with linear forwarders (as in Definition 3); and conversely a linear forwarder $x {\multimap} y$ can be encoded as the pi calculus term $x(\widetilde{u}).\overline{y}\,\widetilde{u}$. This section builds up to a formal proof of the result.

The structure of the proof follows that of Boreale ([4], Definition 2.5 to Proposition 3.6). However, the proofs are significantly easier in our setting. We begin with a basic lemma about the encoding $[\![\cdot]\!]$.

Lemma 7 *In the linear forwarder calculus,*

1. $[\![P]\!]_{\widetilde{z}} \approx_\ell [\![P]\!]_{x\widetilde{z}}$.

2. $[\![P]\!]_{\widetilde{z}}\{y/x\} \approx_\ell [\![P\{y/x\}]\!]_{\widetilde{z}}$.

3. $[\![P]\!]_{\widetilde{z}}\{\widetilde{y}/\widetilde{x}\} \approx_\ell (u')(\overline{u}\,'\widetilde{y} \mid u'(\widetilde{x}).[\![P]\!]_{\widetilde{z}\widetilde{x}})$.

Proof. The first two are trivial inductions on P. The last one follows directly.□

We draw attention to the first part of Lemma 7. This is an important simplifying tool. It means that, even though the encoding $[\![u(x).P]\!]_u = (u')(u'{\multimap}u \mid u'(x).[\![P]\!]_{ux})$ involves progressively more subscripts, they can be ignored up to behavioural equivalence. Thus, although a context C might receive names $\widetilde{x}$ in input, we can ignore this fact: $C\big[[\![P]\!]_{\widetilde{x}}\big] \approx_\ell C\big[[\![P]\!]\big]$ in the linear forwarder calculus. (Notice that, given a localised term $C\big[[\![P]\!]_{\widetilde{x}}\big]$, it is not necessary the case that $C\big[[\![P]\!]\big]$ is also localised; hence the result does not carry over to $\approx$, which is only defined for $L\ell$ contexts).

Boreale [4] uses more complicated proofs than those used in this paper. In particular, Part 1 does not hold for Boreale, and so his equivalent of Part 3 uses a significantly longer (5-page) alternative proof. The difference is that our encoding $[\![P]\!]_x$ merely indicates that each input use of x must be encoded via a local channel $(x')(x{\multimap}x' \mid \ldots)$. In contrast, Boreale uses a protocol instead of linear forwarders, and his encoding $\{|P|\}_{x'/x}$ indicates that inputs on x will use the protocol on the free channel x', and so the protocol is observable in $\{|P|\}_{x'/x}$ but not in $\{|P|\}_\emptyset$. This is why Part 1 does not hold. Additionally, our encoding causes forwarders to be consumed after reaction; Boreale's causes them to be produced by reaction, so requiring the use of an expansion preorder rather than our simpler labelled bisimulation.

Note that this section has even simpler proofs, using the property $P \approx_\ell [\![P]\!]$ which is deduced directly from Lemma 7 and the definition of $[\![\cdot]\!]$. However, this property relates terms from two different sub-calculi of the linear forward calculus, which some readers found inelegant – so we have avoided it.

Proposition 8 *For P, Q in the pi calculus, $P \mathbin{\dot\approx} Q$ if and only if $[\![P]\!] \mathbin{\dot\approx} [\![Q]\!]$.*

Proof. This is similar to Boreale's Proposition 3.5 and 3.6. We need two lemmas:

1. $P \xrightarrow{\mu} P'$ implies $[\![P]\!]_{\widetilde{z}} \overset{\mu}{\Longrightarrow} \approx_{\ell} [\![P']\!]_{\widetilde{z}}$

2. $[\![P]\!]_{\widetilde{z}} \xrightarrow{\mu} P'_1$ implies $P \overset{\mu}{\Longrightarrow} P'$ and $P'_1 \approx_{\ell} [\![P']\!]_{\widetilde{z}}$

For Part 1 the cases of $\mu = (\widetilde{v})\overline{u}\,\widetilde{y}$ and $\mu = u(\widetilde{x})$ are straightforward inductions on the derivation of $\xrightarrow{\mu}$. We draw attention to the case $u(x).P \xrightarrow{u(\widetilde{x})} P$, since this relates specifically to linear forwarders. Suppose $u \in \widetilde{z}$. Then

$$[\![u(\widetilde{x}).P]\!]_{\widetilde{z}} \quad = \quad (u')(u{-}\!\circ u' \mid u'(\widetilde{x}).[\![P]\!]_{\widetilde{z}\widetilde{x}}) \qquad\qquad (3)$$
$$\xrightarrow{u(\widetilde{x})} \quad (u')(\overline{u}\,'\widetilde{x} \mid u'(\widetilde{x}).[\![P]\!]_{\widetilde{z}\widetilde{x}})$$

and by Lemma 7, this is $\approx_{\ell} [\![P]\!]_{\widetilde{z}}$. For $\mu = \tau$ we will write out the interesting case, of communication. Suppose $P|Q \xrightarrow{\tau} (\widetilde{v})(P'\{\widetilde{y}/\widetilde{x}\}|Q'$ with $P \xrightarrow{u(\widetilde{x})} P'$ and $Q \xrightarrow{(\widetilde{z})\overline{u}\,\widetilde{y}} Q'$. By the induction hypothesis and Lemma 7.1, $[\![P]\!]_{\widetilde{z}} \overset{u(\widetilde{x})}{\Longrightarrow} P'_1 \approx_{\ell} [\![P']\!] \approx_{\ell} [\![P']\!]_{\widetilde{z}}$ and $[\![Q]\!]_{\widetilde{z}} \overset{(\widetilde{v})\overline{u}\,\widetilde{y}}{\Longrightarrow} Q'_1 \approx_{\ell} [\![Q']\!]_{\widetilde{z}}$. Therefore,

$$[\![P|Q]\!]_{\widetilde{z}} \overset{\tau}{\Longrightarrow} \quad (\widetilde{v})(P'_1\{\widetilde{y}/\widetilde{x}\} \mid Q'_1)$$
$$\approx_{\ell} \quad (\widetilde{v})\big([\![P']\!]_{\widetilde{z}}\{\widetilde{y}/\widetilde{x}\} \mid [\![Q']\!]_{\widetilde{z}}\big)$$
$$= \quad [\![(\widetilde{v})\big(P'\{\widetilde{y}/\widetilde{x}\}|Q'\big)]\!]_{\widetilde{z}}.$$

Part 2 is similar. Again, we draw attention to the input case $[\![u(\widetilde{x}).P]\!]_{\widetilde{z}}$. Suppose $u \in \widetilde{z}$. Then $[\![u(\widetilde{x}).P]\!]_{\widetilde{z}} \xrightarrow{u(x)} \approx_{\ell} [\![P]\!]_{\widetilde{z}}$ as in Equation 3. The original term also has a matching transition $u(\widetilde{x}).P \xrightarrow{u(\widetilde{x})} P$, so completing the proof. $\qquad\square$

We will also need the following lemma. It generalises Lemma 7.1 to apply to barbed congruence rather than just labelled bisimulation. Effectively, it implies that a non-localised context can be transformed into a localised one.

Lemma 9 *For P,Q in the pi calculus, $[\![P]\!] \approx [\![Q]\!]$ implies $[\![P]\!]_{\widetilde{z}} \approx [\![Q]\!]_{\widetilde{z}}$.*

Proof. From Lemma 7.1 we get $[\![P]\!]_{\widetilde{z}} \approx_{\ell} [\![P]\!] \approx [\![Q]\!] \approx_{\ell} [\![Q]\!]_{\widetilde{z}}$. The result follows by Lemma 6.3 and the transitivity of $\approx$. (We thank an anonymous reviewer for this simpler proof.) $\qquad\square$

We are now ready to establish full abstraction for the encoding of the pi calculus into the $L\ell$ calculus.

Theorem 10 (Full abstraction) *For P,Q in the pi calculus, $P \approx Q$ if and only if $[\![P]\!] \approx [\![Q]\!]$ in the $L\ell$ calculus.*

Proof. We show that (1) $P \not\approx Q$ implies $[\![P]\!] \not\approx [\![Q]\!]$ and (2) $[\![P]\!] \not\approx [\![Q]\!]$ implies $P \not\approx Q$. We write C_{π} and $C_{L\ell}$ to range over contexts such that $C_{\pi}[P], C_{\pi}[Q]$ are terms in the pi calculus, and $C_{L\ell}\big[[\![P]\!]\big], C_{L\ell}\big[[\![Q]\!]\big]$ are terms in the $L\ell$ calculus.

To establish (1), extend the translation $\llbracket \cdot \rrbracket$ to contexts in the obvious way. Since the translation $\llbracket \cdot \rrbracket$ is compositional, we get $\llbracket C[P] \rrbracket = \llbracket C \rrbracket \big[\llbracket P \rrbracket_{\widetilde{z}} \big]$ and $\llbracket C[Q] \rrbracket = \llbracket C \rrbracket \big[\llbracket Q \rrbracket_{\widetilde{z}} \big]$ for some $\widetilde{z}$ determined by C. Next, we reason by contradiction: we prove that $P \not\approx Q$ and $\llbracket P \rrbracket \approx \llbracket Q \rrbracket$ is false. Assuming $P \not\approx Q$, there exists a context $C_\pi[\,]$ such that $C_\pi[P] \not\approx C_\pi[Q]$. By $\llbracket P \rrbracket \approx \llbracket Q \rrbracket$ and Lemma 9, we also have $\llbracket P \rrbracket_{\widetilde{z}} \approx \llbracket Q \rrbracket_{\widetilde{z}}$. Therefore, in particular $\llbracket C_\pi \rrbracket \big[\llbracket P \rrbracket_{\widetilde{z}} \big] \dot{\approx} \llbracket C_\pi \rrbracket \big[\llbracket Q \rrbracket_{\widetilde{z}} \big]$ and, by the above equalities, $\llbracket C_\pi[P] \rrbracket \dot{\approx} \llbracket C_\pi[Q] \rrbracket$. By Proposition 8, this latter bisimulation contradicts $C_\pi[P] \not\approx C_\pi[Q]$.

To establish (2), we show that pi contexts are as expressive as linear forwarder contexts, by exhibiting a pi implementation of linear forwarders. To this end, we define $\widehat{\cdot}$, which translates $x \multimap y$ into $x(\widetilde{u}).\overline{y}\,\widetilde{u}$ and leaves all else unchanged. Similarly to (1), we prove that $\llbracket P \rrbracket \not\approx \llbracket Q \rrbracket$ and $P \approx Q$ are contradictory. We are given a context $C_{L\ell}[\,]$ such that $C_{L\ell}[\llbracket P \rrbracket] \not\approx C_{L\ell}[\llbracket Q \rrbracket]$. Consider the agent $\llbracket \widehat{C_{L\ell}[P]} \rrbracket$, which by definition is equal to $\llbracket \widehat{C_{L\ell}} \rrbracket \big[\llbracket P \rrbracket_{\widetilde{z}} \big]$ for some $\widetilde{z}$. By Lemma 7 this is $\approx_\ell$-bisimilar to $\llbracket \widehat{C_{L\ell}} \rrbracket \big[\llbracket P \rrbracket \big]$. Now we consider the double translation $\llbracket \widehat{\cdot} \rrbracket$; it will convert each forwarder $u \multimap v$ into either $u(\widetilde{x}).\overline{v}\,\widetilde{x}$ or $(u')(u \multimap u' \mid u'(\widetilde{x}).\overline{v}\,\widetilde{x})$. Thanks to Equation 2, $\llbracket \widehat{C_{L\ell}} \rrbracket \big[\llbracket P \rrbracket \big] \approx_\ell C_{L\ell}\big[\llbracket P \rrbracket \big]$. And, with similar reasoning, the same holds for Q. The proof follows from these results. From the assumption that $P \approx Q$ we get $\widehat{C_{L\ell}[P]} \dot{\approx} \widehat{C_{L\ell}[Q]}$. By Proposition 8, $\llbracket \widehat{C_{L\ell}[P]} \rrbracket \dot{\approx} \llbracket \widehat{C_{L\ell}[Q]} \rrbracket$. Now we focus on P. $\llbracket \widehat{C_{L\ell}[P]} \rrbracket = \llbracket \widehat{C_{L\ell}} \rrbracket \big[\llbracket P \rrbracket_{\widetilde{z}} \big]$ (by definition of $\llbracket \cdot \rrbracket$); $\approx_\ell \llbracket \widehat{C_{L\ell}} \rrbracket \big[\llbracket P \rrbracket \big]$ (by Lemma 7.1); $\approx_\ell C_{L\ell}\big[\llbracket P \rrbracket \big]$ (by Equation 2). Doing the same to Q, we obtain $C_{L\ell}\big[\llbracket P \rrbracket \big] \dot{\approx} C_{L\ell}\big[\llbracket Q \rrbracket \big]$, contradicting $\llbracket P \rrbracket \not\approx \llbracket Q \rrbracket$ and so proving the result. $\qquad\square$

5 A Linear Forwarder Machine

In this section we develop a distributed machine for the $L\ell$ calculus, suitable for a point-to-point network such as the Internet. This machine is actually very similar to the fusion machine [8], but with linear forwarders instead of fusions (trees of persistent forwarders). We first give a diagrammatic overview of the machine. Then we provide a formal syntax, and prove full abstraction with respect to barbed congruence.

We assume a set of locations. Each channel belongs to a particular location. For instance, channels u, v, w might belong to ℓ_1 and x, y to ℓ_2. The structure of a channel name u might actually be the pair (IP:TCP), giving the IP number and port number of a channel-manager service on the Internet. Every input process is at the location of its subject channel. Output processes may be anywhere. For example,

$\ell_1{:}uvw$	$\ell_2{:}xy$
$u(x).(x')(x \multimap x' \mid x'(z).P)$	$\overline{u}\,y \mid \overline{y}\,w$

In a point-to-point network such as the Internet, the output message $\overline{u}\,y$ would be sent to u to react; in a broadcast network such as wireless or ethernet, the offer of output would be broadcast and then ℓ_1 would accept the offer. In both cases, the result is a reaction and a substitution $\{y/x\}$ as follows:

$\ell_1{:}uvw\,x'$	$\ell_2{:}xy$
$\rightarrow$ $y{-}\!\circ x' \mid x'(z).P\{y/x\}$	$\overline{y}\,w$

The overall effect of the linear forwarder $y{-}\!\circ x'$ will be to turn the $\overline{y}\,w$ into $\overline{x}\,'w$. In a point-to-point network this can be implemented by migrating the $y{-}\!\circ x'$ to ℓ_2, there to *push* the $\overline{y}\,w$ to x', as shown below. (In the following diagrams, some steps are shown as heating transitions $\rightharpoonup$; these steps were abstracted away in the $L\ell$ calculus).

$\ell_1{:}uvw\,x'$	$\ell_2{:}xy$
$\rightharpoonup$ $x'(z).P\{y/x\}$	$y{-}\!\circ x' \mid \overline{y}\,w$
$\rightarrow$ $x'(z).P\{y/x\}$	$\overline{x}\,'w$
$\rightharpoonup$ $\overline{x}\,'w \mid x'(z).P\{y/x\}$	
$\rightarrow$ $P\{y/x\}\{w/z\}$	

In a broadcast network, the linear forwarder $y{-}\!\circ x'$ would instead stay at ℓ_1; later, when the offer of $\overline{y}\,a$ is broadcast, the linear forwarder can *grab* the offer.

We remark that the above machine basically avoids code migration: after an input, the continuation remains in the same place (with the minor exception that forwarders $x{-}\!\circ x'$ and outputs $\overline{x}\,y$ may migrate, which is easy to implement). Because code does not migrate, there is no need for a run-anywhere infrastructure such as Java, and it is possible to compile into CPU-specific machine code.

Distributed choice. A well-known operator in process calculi is the input-guarded choice $x(u).P + y(v).Q$. In the case where x and y are at separate locations, the choice is awkward to implement, but we can compile it into an (easily implementable) localised choice as follows:

$$\llbracket x(u).P + y(v).Q \rrbracket \;=\; (x'y')\big(\; x{-}\!\circ x' \mid y{-}\!\circ y'$$
$$\mid\; x'(u).(P \mid y'{-}\!\circ y) + y'(v).(Q \mid x'{-}\!\circ x) \;\big)$$

To understand this encoding, note that the new names x' and y' will be created at the same location, and so the choice between x' and y' will a local one. Next, $\overline{x}\,v$ may be forwarded to x', or $\overline{y}\,v$ to y', or both. If the reaction with x' is taken, this yields $y'\!\multimap\! y$ to "undo" the effect of the forwarder that was not taken. (It undoes it up to weak barbed congruence, in the sense that $(y')(y\!\multimap\! y' \mid y'\!\multimap\! y) \approx \mathbf{0}$.) Note that even if the location of x should fail, then the y option remains open, and vice versa.

Failure. We briefly comment on the failure model for the linear forwarder machine. It is basically the same as failure in the join calculus: either a message can be lost, or an entire location can fail. If a linear forwarder $x\!\multimap\! y$ should fail, the effect is the same as if a single future message $\overline{x}\,\widetilde{u}$ should fail. A command *'iffail(u) then P'* might be added to determine if u's location is unresponsive. The current theory was largely prompted by criticisms of the fragility of our earlier fusion machine.

5.1 The Machine Calculus

Definition 11 (Linear forwarder machine) *Localised linear forwarder machines M are given by the following grammar, where P ranges over terms in the $L\ell$ calculus (Definition 2):*

$$M \quad ::= \quad \mathbf{0} \quad \big| \quad x[P] \quad \big| \quad (\!x\!)[P] \quad \big| \quad M, M$$

The presentation here is similar to that given for the fusion machine [8]. The *basic channel-manager* $x[P]$ denotes a channel-manager at channel x containing a *body P*. The *local channel-manager* $(\!x\!)[P]$ denotes a channel-manager where the name x is not visible outside the machine. We write chan M to denote the set of names of all channel-managers in the machine, and Lchan M for the names of only the local channel-managers.

We assume a *co-location* equivalence relation L on channels. We write $x@y$ to mean that $(x, y) \in L$, with the intended meaning that the two channels are at the same location. It is always possible to create a fresh channel at an existing location: therefore let each equivalence class in L be infinitely large. In the machine calculus, we generally assume L rather than writing it explicitly.

There are a number of well-formedness conditions on machines:

(1) *Localised.* All code is in the right place, and does not need to be moved at runtime. Formally, in every channel $u[P]$, every free input $v(\widetilde{x}).Q$ satisfies $u@v$. (Also, no received name is used as the subject of input; this already holds from the $L\ell$ calculus).

(2) *Singly-defined.* There is exactly one channel-manager per channel. Formally, a machine $x_1[B_1], \cdots, x_n[B_n]$ is singly-defined when $i \neq j$ implies $x_i \neq x_j$ (x_i or x_j may be local).

(3) *Complete.* It does not make sense to write a program that refers to channels which do not exist. We say that a machine is complete when it has

no such references. Formally, the free names of a machine must be contained in $\mathsf{chan}\,M$.

A machine is *well-formed* when it is localised, singly-defined and complete. In the following, we consider only well-formed machines.

Definition 12 (Dynamics) *The* structural congruence *for well-formed machines* $\equiv$ *is the smallest equivalence relation satisfying the following laws:*

$$M, \mathbf{0} \equiv M \qquad M_1, M_2 \equiv M_2, M_1 \qquad M_1, (M_2, M_3) \equiv (M_1, M_2), M_3$$

$$P \equiv Q \quad implies \quad u[P] \equiv u[Q] \ and \ (\!u\!)[P] \equiv (\!u\!)[Q]$$

The reduction step $\rightarrow$ *and the heating step* $\rightharpoonup$ *are the smallest relations satisfying the rules below, and closed with respect to structural congruence. Each rule addresses generically both free and local channel-managers.*

$$u[\overline{u}\,\widetilde{y} \mid u(\widetilde{x}).P \mid R] \quad \rightarrow \quad u[P\{\widetilde{y}/\widetilde{x}\} \mid R] \qquad\qquad \text{(react)}$$
$$u[u\!\multimap\! v \mid \overline{u}\,\widetilde{x} \mid R] \quad \rightarrow \quad u[\overline{v}\,\widetilde{x} \mid R] \qquad\qquad \text{(fwd)}$$

$$u[(x)P \mid R] \quad \rightharpoonup \quad u[P\{x'/x\} \mid R],\ (\!x'\!)[\,], \quad x' \ fresh,\ x'@u \quad \text{(dep.new)}$$
$$u[x\!\multimap\! y \mid R_1],\ x[R_2] \quad \rightharpoonup \quad u[R_1],\ x[x\!\multimap\! y \mid R_2], \quad if\ u \neq x \qquad \text{(dep.fwd)}$$
$$u[\overline{x}\,\widetilde{y} \mid R_1],\ x[R_2] \quad \rightharpoonup \quad u[R_1],\ x[\overline{x}\,\widetilde{y} \mid R_2], \quad if\ u \neq x \qquad \text{(dep.out)}$$
$$u[x(\widetilde{y}).P \mid R_1],\ x[R_2] \quad \rightharpoonup \quad u[R_1],\ x[x(\widetilde{y}).P \mid R_2], \quad if\ u \neq x,\ u@x \qquad \text{(dep.in)}$$

For every transition rule above, we close it under contexts:

$$\frac{M \rightarrow M', \quad \mathsf{chan}\,M' \cap \mathsf{chan}\,N = \emptyset}{M, N \rightarrow M', N} \qquad \frac{M \rightharpoonup M', \quad \mathsf{chan}\,M' \cap \mathsf{chan}\,N = \emptyset}{M, N \rightharpoonup M', N}$$

We draw attention to two of the rules. The rule *(dep.new)* picks a fresh channel-name x' and this channel is deemed to be at the location where the command was executed. The rule *(dep.in)* will only move an input command from one channel u to another channel x, if the two channels are co-located; hence, there is no "real" movement. In the current presentation we have used arbitrary replication $!P$, but in a real machine we would instead use guarded replication [18], as is used in the fusion machine. All rules preserve well-formedness.

Definition 13 (Bisimulation) *The observation* $M \downarrow u$ *is the smallest relation satisfying* $u[P] \downarrow u$ *if* $P \downarrow u$, *and* $M_1, M_2 \downarrow u$ *if* $M_1 \downarrow u$ *or* $M_2 \downarrow u$. *Write* $\Rightarrow$ *for* $(\rightarrow^* \rightharpoonup^*)^*$, *and* $M \Downarrow u$ *for* $M \Rightarrow\downarrow u$. *A weak barbed bisimulation* $\mathcal{R}$ *between machines is a symmetric relation such that if* $M\ \mathcal{S}\ N$ *then*

1. *$M \Downarrow u$ implies $N \Downarrow u$*

2. *$M \Rightarrow M'$ implies $N \Rightarrow N'$ such that $M'\ \mathcal{R}\ N'$*

Let $\dot{\approx}$ *be the largest barbed bisimulation. Two machines* M_1 *and* M_2 *are weak barbed equivalent, written* $M_1 \simeq M_2$, *when for every machine* N, *then* $N, M_1 \dot{\approx} N, M_2$. *(Note that* N, M_1 *and* N, M_2 *are assumed to be well-formed, so* $\mathsf{chan}\,N \cap \mathsf{chan}\,M_1 = \mathsf{chan}\,N \cap \mathsf{chan}\,M_2 = \emptyset$.)*

We will prove correctness using a translation $\mathsf{calc}\, M = (\mathsf{Lchan}\, M)\widehat{M}$ from machines to terms in the $L\ell$ calculus, where

$$\widehat{\mathbf{0}} = \mathbf{0} \qquad \widehat{u[P]} = P \qquad \widehat{(u)[P]} = P \qquad \widehat{M_1, M_2} = \widehat{M_1} \mid \widehat{M_2}$$

One might prefer to prove correctness of a "compiling" translation, which takes a term and compiles it into a machine — rather than the reverse translation calc. However, different compilers are possible, differing in their policy for which location to upload code into. We note that all correct compilers are contained in the inverse of calc, so our results are more general.

The correctness of the forwarder machine relies on the following lemma

Lemma 14 (Correctness) $M_1 \stackrel{.}{\approx} M_2$ *if and only if* $\mathsf{calc}\, M_1 \stackrel{.}{\approx} \mathsf{calc}\, M_2$

Proof. It is clear that machine operations are reflected in the calculus: $M \equiv M'$ implies $\mathsf{calc}\, M \equiv \mathsf{calc}\, M'$, and $M \rightharpoonup M'$ implies $\mathsf{calc}\, M \equiv \mathsf{calc}\, M'$, and $M \to M'$ implies $\mathsf{calc}\, M \to \mathsf{calc}\, M'$, and $M \downarrow u$ implies $\mathsf{calc}\, M \downarrow u$.

The reverse direction is more difficult. We wish to establish that

1. $\mathsf{calc}\, M \downarrow u$ implies $M \rightharpoonup^* \downarrow u$, and

2. $\mathsf{calc}\, M \to P'$ implies $\exists M' : M \rightharpoonup^* \to M'$ and $P' \equiv \mathsf{calc}\, M'$.

Both parts share a similar proof; we focus on part 2. Given the machine M, there is also a *fully-deployed* machine M' such that $M \rightharpoonup^* M'$ and M' has no heating transitions: that is, all unguarded restrictions have been used to create fresh channels, and all outputs and forwarders are at the correct location. Therefore $\mathsf{calc}\, M \equiv \mathsf{calc}\, M' \to P'$. The structure of $\mathsf{calc}\, M'$ has the form $(\mathsf{Lchan}\, M')\widehat{M'}$. The reaction must have come from an output $\overline{u}\,\widetilde{y}$ and an input $u(\widetilde{x}).P$ in M' (or from an output and a forwarder). Because M' is fully-deployed, it must contain $u[\overline{u}\,\widetilde{y} \mid u(\widetilde{x}).P]$. Therefore it too must allow the corresponding reaction. The bisimulation result follows directly from the above. $\qquad \square$

We now prove full abstraction: that two machines are barbed *equivalent* if and only if their corresponding $L\ell$ calculus terms are barbed *congruent*. There is some subtlety here. With the $L\ell$ calculus, $\approx$ is closed under restriction, input prefix and parallel contexts. With the abstract machine, $\simeq$ is only closed under the partial machine composition: contexts can add new channel managers to the machine, but not additional programs to existing ones. We defined machine equivalence around this partial closure, because this is the most natural equivalence in the machine setting. (It is not surprising that contexts gain no additional discriminating power through restriction and input-prefixing, since the same holds in the asynchronous pi calculus without matching [14]: $P \approx Q$ if and only if $R|P \stackrel{.}{\approx} R|Q$ for every R.)

It has been suggested that a weaker simulation result would suffice. But we believe that full abstraction shows our machine to be a natural implementation of the pi calculus, in contrast to Facile and Join. Practically, full abstraction means that a program can be debugged purely at source-level rather than machine-level.

Theorem 15 (Full abstraction) $M_1 \simeq M_2$ *if and only if* $\mathsf{calc}\, M_1 \approx \mathsf{calc}\, M_2$.

Proof. The reverse direction is straightforward, because machine contexts are essentially parallel compositions. In the forwards direction, it suffices to prove $R |\, \mathsf{calc}\, M_1 \,\dot{\approx}\, R |\, \mathsf{calc}\, M_2$ for every R. By contradiction suppose the contrary: namely, there exists an R such that the two are not barbed bisimilar. Expanding the definition of calc we obtain that $R | (\mathrm{lchan}\, M_1) \widehat{M_1} \,\dot{\not\approx}\, R | (\mathrm{lchan}\, M_2) \widehat{M_2}$.

We now show how to construct a machine context M_R such that $M_R, M_1 \,\dot{\not\approx}\, M_R, M_2$, thus demonstrating a contradiction. Without loss of generality, suppose that R does not clash with the local names $\mathrm{lchan}\, M_1$ or $\mathrm{lchan}\, M_2$. This gives $(\mathrm{lchan}\, M_1)(R | \widehat{M_1}) \,\dot{\not\approx}\, (\mathrm{lchan}\, M_2)(R | \widehat{M_2})$. In order to ensure well-formedness of M_R, M_1 and M_R, M_2, let $\widetilde{z} = \mathrm{chan}\, M_1 \cup \mathrm{chan}\, M_2$. By Lemma 7 we get $[\![R]\!]_{\widetilde{z}} \approx_\ell R$, and by definition $[\![R]\!]_{\widetilde{z}}$ contains no inputs on $\widetilde{z}$, so satisfying the localised property. Now assume without loss of generality that R contains no top-level restrictions. Let $M_R = u[R]$ for a fresh name u such that, for every free input $u_i(\widetilde{x}).R'$ in R, then $u_i @ u$. Hence $\mathrm{lchan}\, M_R = \emptyset$ and $\widehat{M_R} = [\![R]\!]_{\widetilde{z}}$. This yields $\mathsf{calc}\, M_R, M_1 = (\mathrm{lchan}\, M_1)([\![R]\!]_{\widetilde{z}} | \widehat{M_1}) \approx_\ell (\mathrm{lchan}\, M_1)(R | \widehat{M_1})$, and similarly for M_2. And finally, by construction, both M_R, M_1 and M_R, M_2 are singly-defined and complete. □

6 Further Issues

The point of this paper is to provide a distributed implementation of the input capability of the pi calculus. We have shown that a limited form of input capability (linear forwarders) is enough to easily express the full input capability. We have expressed this formally through a calculus with linear forwarders, and a proof of its full abstraction with respect to the pi calculus encoding.

The calculus in this paper abstracts away from certain details of implementation (such as the choice between a point-to-point or broadcast network). Nevertheless, thanks to its *localisation* property, it remains easy to implement.

Coupled bisimulation. There is an interesting subtlety in the encoding of input capability. Our first attempt at an encoding gave, for example,

$$[\![u(x).(x().P \mid x().Q)]\!] \quad = \quad u(x).(x')(x \multimap x' \mid x'().P \mid x \multimap x' \mid x'().Q)$$

That is, we tried to reuse the same local name x' for all bound inputs. But then, a subsequent reaction $x \multimap x' \mid \overline{x}\, z \to \overline{x}\,'z$ would be a commitment to react with one of $x'().P$ or $x'().Q$, while ruling out any other possibilities. This *partial* commitment does not occur in the original pi calculus expression, and so the encoding does not even preserve behaviour. An equivalent counterexample in the pi calculus is that $\tau.P | \tau.Q | \tau.R$ and $\tau.P | \tau.(\tau.Q | \tau.R)$ are not weak bisimilar. We instead used an encoding which has a fresh channel for each bound input:

$$[\![u(x).(x().P \mid x().Q)]\!] \quad = \quad u(x).\big((x')(x \multimap x' \mid x'().P) \mid (x'')(x \multimap x'' \mid x''().Q)\big)$$

Now, any reaction with a forwarder is a *complete* rather than a partial commitment. In fact, both encodings are valid. The original encoding, although not a bisimulation, is still a *coupled bisimulation* [16]. (Coupled bisimulation is a less-strict form of bisimulation that is more appropriate for an implementation, introduced for the same reasons as given here.) In this paper we chose the normal bisimulation and the repaired encoding, because they are simpler.

The join calculus and forwarders. We end with some notes on the difference between the join calculus [7] and the $L\ell$ calculus. The core join calculus is

$$P \ ::= \ \mathbf{0} \ \mid \ \overline{x}\,\widetilde{u} \ \mid \ P|P \ \mid \ \mathrm{def}\ x(\widetilde{u})|y(\widetilde{v}) \rhd P \ \mathrm{in}\ Q$$

The behaviour of the *def* resource is, when two outputs $\overline{x}\,\widetilde{u}'$ and $\overline{y}\,\widetilde{v}'$ are available, then it consumes them to yield a copy $P\{\widetilde{u}'\widetilde{v}'/\widetilde{u}\widetilde{v}\}$ of P. Note that x and y are bound by *def*, and so input capability is disallowed by syntax. The core join calculus can be translated into the pi calculus (and hence $L\ell$) as follows [7]:

$$\llbracket \mathbf{0} \rrbracket \ = \ \mathbf{0} \qquad \llbracket \overline{x}\,\widetilde{u} \rrbracket \ = \ \overline{x}\,\widetilde{u} \qquad \llbracket P|Q \rrbracket \ = \ \llbracket P \rrbracket \mid \llbracket Q \rrbracket$$
$$\llbracket \mathrm{def}\ x(\widetilde{u})|y(\widetilde{v}) \rhd P \ \mathrm{in}\ Q \rrbracket \ = \ (xy)(\,\llbracket Q \rrbracket \mid !x(\widetilde{u}).y(\widetilde{v}).\llbracket P \rrbracket\,)$$

If a join program is translated into the linear forwarder machine and then executed, then the result has exactly the same runtime behaviour (i.e. same number of messages) as the original join program. Additionally, we can provide the same distribution of channels through the co-location operator discussed above.

A reverse translation is more difficult, because of linear forwarders. One might try to translate $x{\multimap}y \mid R$ into def $x(\widetilde{u}) \rhd \overline{y}\,\widetilde{u}$ in $\llbracket R \rrbracket$, analogous with the translation of a forwarder into the pi calculus that was used in Proposition 10. But the $L\ell$ calculus allows a received name to be used as the source of a forwarder, as in $u(x).(x{\multimap}y \mid P)$, and the same is not possible in the join calculus. Therefore contexts in the $L\ell$ calculus are strictly more discriminating than contexts in the join calculus. (As an example, def $x(u) \rhd \overline{y}\,u$ in $\overline{z}\,x$ is equivalent to $\overline{z}\,y$ in the join calculus, but the context $z(a).(a{\multimap}b \mid \overline{a} \mid \_)$ can distinguish them in the $L\ell$ calculus.)

References

[1] M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *POPL 2001*, pages 104–115. ACM Press.

[2] R. Amadio, G. Boudol, and C. Lhoussaine. The receptive distributed pi-calculus (extended abstract). In *FSTTCS 1999*, LNCS 1738:304–315.

[3] R. Amadio. An asynchronous model of locality, failure, and process mobility. In *COORDINATION 1997*, LNCS 1282:374–391.

[4] M. Boreale. On the expressiveness of internal mobility in name-passing calculi. *Theoretical Computer Science*, 195(2):205–226, 1998.

[5] G. Boudol. Asynchrony and the π-calculus (note). Rapport de Recherche 1702, INRIA Sophia-Antipolis, 1992.

[6] L. Cardelli and A. D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, 2000.

[7] C. Fournet and G. Gonthier. The reflexive chemical abstract machine and the join-calculus. In *POPL 1996*, pages 372–385. ACM Press.

[8] P. Gardner, C. Laneve, and L. Wischik. The fusion machine. In *CONCUR 2002*, LNCS 2421:418–433.

[9] D. Gelernter, N. Carriero, S. Chandran, and S. Chang. Parallel programming in Linda. In *ICPP 1985*, pages 255–263. IEEE.

[10] A. Giacalone, P. Mishra, and S. Prasad. Facile: A symmetric integration of concurrent and functional programming. *International Journal of Parallel Programming*, 18(2):121–160, 1989.

[11] K. Honda and N. Yoshida. On reduction-based process semantics. *Theoretical Computer Science*, 152(2):437–486, 1995.

[12] N. Kobayashi, B. C. Pierce, and D. N. Turner:. Linearity and the pi-calculus. *ACM Transactions on Programming Languages and Systems*, 21(5):914–947, 1999.

[13] M. Merro and D. Sangiorgi. On asynchrony in name-passing calculi. In *ICALP 1998*, LNCS 1443:856–867.

[14] M. Merro. On equators in asynchronous name-passing calculi without matching. In *EXPRESS 1999*, volume 27 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science Publishers.

[15] R. Milner and D. Sangiorgi. Barbed bisimulation. In *ICALP 1992*, LNCS 623:685–695.

[16] J. Parrow and P. Sjödin. Multiway synchronization verified with coupled simulation. In *CONCUR 1992*, LNCS 630:518–533.

[17] D. Sangiorgi. Pi-calculus, internal mobility and agent-passing calculi. *Theoretical Computer Science*, 167(1–2):235–275, 1996.

[18] D. Sangiorgi. On the bisimulation proof method. *Mathematical Structures in Computer Science*, 8(5):447–479, 1998.

[19] P. Sewell, P. Wojciechowski, and B. Pierce. Location independence for mobile agents. In *ICCL 1998*, LNCS 1686:1–31.