

Linear Forwarders

Philippa Gardner^a, Cosimo Laneve^{b,*}, Lucian Wischik^c

^a*Imperial College, London, United Kingdom*

^b*University of Bologna, Italy*

^c*Microsoft, Redmond, USA*

Abstract

A *linear forwarder* is a process that receives one message on a channel and sends it on a different channel. We use linear forwarders to provide a distributed implementation of Milner’s asynchronous pi calculus. Such a distributed implementation is known to be difficult due to *input capability*, where a received name is used as the subject of a subsequent input. This allows the dynamic creation of large input processes in the wrong place, thus requiring comparatively large code migrations in order to avoid consensus problems. Linear forwarders constitute a small atom of *input capability* that is easy to move.

We show that the full input capability can be simply encoded using linear forwarders. We also design a distributed machine, demonstrating the ease with which we can implement the pi calculus using linear forwarders. We also show that linear forwarders allow for a simple encoding of distributed choice and have “clean” behaviour in the presence of failures.

Key words: Pi-calculus, distributed implementation, input capability, linear forwarders.

1 Introduction

Distributed interaction has become a necessary part of modern programming languages. The asynchronous pi calculus, $\lambda\pi$ for short, is widely regarded as a foundation for such languages. In this calculus, a program (or process) has

* Corresponding author.

Email addresses: `pg@doc.ic.ac.uk` (Philippa Gardner), `laneve@cs.unibo.it` (Cosimo Laneve), `lu@wischik.com` (Lucian Wischik).

a collection of channels, and it executes through interaction over these channels. A natural, well-studied distributed interpretation is to let each channel belong to a single location: for instance, one location for the channels x, y, z and another for u, v, w . Interaction on channel x only occurs at the location assigned to x : an *input resource* $x(u).P$ resides at location x , waits to receive formal parameter u and then continues with P ; an *output* $\bar{x}v$ knows to go to the location x in order to find the matching input resource. Outputs are small and so have the freedom to move between locations, whereas inputs have possibly large continuations and so their movement is restricted.

This restricted input movement is subtle, since $\lambda\pi$ has a behaviour called *input capability*, which is the ability 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 $\bar{x}w$ it produces the continuation $w(v).Q\{w/u\}$. This continuation is in the wrong place, since it is still at x whereas it should be at w . A key challenge with distributing $\lambda\pi$ is to find sensible ways of restricting this input capability so that inputs only reside at their correct locations.

The join calculus [13], the local pi calculus [25] and the $\pi_{1\ell}$ calculus [1] are examples of *local* calculi that simply disallow input capability: that is, in a term $x(u).P$, the P may not contain any input on channel u . The join calculus achieves localisation with an elegant syntactic constraint in which the same language construct is used both to declare a channel and define its input behaviour. The local pi calculus achieves it with a well-formedness constraint on processes, and the $\pi_{1\ell}$ calculus with a type system. Whilst these calculi and resulting implementations [10] are elegant and highly successful, we believe the argument for *always* removing input capability is not convincing.

An initial reason for removing input capability, used for example in the development of the join calculus and its associated implementation, was just that it was unimplementable. Researchers then argued that input capability was also unnecessary, since much can be done without it and $\lambda\pi$ can be encoded in the join calculus. This is a compelling yet incomplete reasoning. There remains the open problem of whether it is in fact possible to provide a distributed implementation of input capability, and whether such an implementation is useful. In this paper, we show that such a distributed implementation is theoretically possible. The PiDuce project [7] is a substantial prototype for assessing whether such an implementation is useful.

To help introduce our ideas, we first recall the encoding of input capability into the join calculus. A key point of the encoding is to split channels into two parts: the *input part* and the *output part*. The input $x(u).P$ is encoded as (we write it in $\lambda\pi$ syntax rather than using the join calculus syntax) $(x')(\bar{x}_i x' \mid !x'(u_i, u_o).Q)$, which sends an *input proxy* x' to the input part x_i of x . When

an output on x is available – encoded as a message on the output part x_o of x – the proxy x' will be activated and the translated continuation Q of P be triggered. To obtain a tight correspondence with $\Lambda\pi$ – a full abstraction result – the join calculus encodings must be wrapped by terms protecting the access to the channels (enforcing the protocol splitting names into input and output parts). Therefore, whilst this is an interesting expressivity result, it cannot be regarded as a sensible implementation of $\Lambda\pi$; such an implementation should not split channels and should not require the use of wrappers.

We show that a simple and direct distributed implementation of $\Lambda\pi$ is possible, contrary to folklore. Our idea is to introduce a limited form of input capability, the *linear forwarder*, using which it is possible to encode input capability in a straightforward way. A linear forwarder $x \multimap y$ is a process that allows just one x to be turned into a y . It plays a similar role of an input proxy in the join calculus, except that it really does behave like an input and so does not require a firewall. It is like an input in that it reacts with an output on x , by forwarding (renaming it) to y . It is an *input capability* in that, in the term $u(x).P$ for example, a forwarder $x \multimap y$ is allowed in P . Using this limited form of input capability, we provide a simple encoding of general input capability, prove a simple full-abstraction result, and design a distributed abstract machine showing that it is indeed possible to give a simple distributed implementation in $\Lambda\pi$.

One interpretation of $x \multimap y$ is just as the pi process $x(u).\bar{y}u$. In fact, a linear forwarder may be viewed as an input proxy in a point-to-point network such as the Internet. However, we choose to use forwarders as first-class operators in order to abstract away from a particular style of implementation. Other networks might provide different implementations of linear forwarders. In a broadcast network for example, the forwarder $x \multimap y$ might be located at y : when it hears an offer of $\bar{x}\tilde{u}$ being broadcast, the machine at y can take up the offer.

Our results demonstrate that it is possible to implement $\Lambda\pi$ on a distributed machine. An essential next step is study practical applications of our implementation of $\Lambda\pi$, in particular fully contrasting our implementation with the indirect implementation given by the join machine. We are actively applying our ideas to Web services, where channel update is essential: for example, when query patterns must be combined, as argued in [34,16,23], or when existing services must be orchestrated in order to provide new services [21]. In particular, this orchestration of services is currently being studied in University of Bologna as part of the extensive PiDuce project [7]; see conclusions for more details. We also note that in BizTalk, a language for Web services by Microsoft [40], input capability is offered when run over a reliable message service (MSMQ), but not otherwise. This paper provides a formal treatment of one possible implementation of BizTalk (the implementation details of BizTalk

have not been published). A more precise account of our work follows.

We begin by introducing the *linear forwarder calculus*, $\text{AL}\ell\pi$ for short, which is like $\text{A}\pi$ except that input capability is restricted to linear forwarders. We show that $\text{AL}\ell\pi$ has a direct distributed implementation by describing an abstract machine and provide a simple encoding of $\text{A}\pi$ in $\text{AL}\ell\pi$. To illustrate the encoding, consider the $\text{A}\pi$ term $x(u).u(v).Q$. We encode it as

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

where the input $u(v)$ has been turned into a local input $u'(v)$ at the same location as x , and the forwarder allows one output on u to interact with u' instead. Thus, for example, the process $\bar{x}w \mid \bar{w}y \mid \llbracket x(u).u(v).Q \rrbracket$ manifests the following behaviour:

$$\begin{aligned} \bar{x}w \mid \bar{w}y \mid x(u).(u')(u \multimap u' \mid u'(v).\llbracket Q \rrbracket) & \quad (1) \\ \rightarrow \bar{w}y \mid (u')(w \multimap u' \mid u'(v).\llbracket Q \rrbracket) & \quad \text{interact on } x \\ \rightarrow (u')(\bar{u}'y \mid u'(v).\llbracket Q \rrbracket) & \quad \text{forward } \bar{w}y \text{ to } \bar{u}'y \\ \rightarrow (u')\llbracket Q \rrbracket\{y/v\} & \quad \text{interact on } u' \end{aligned}$$

In the encoding there are exactly as many forwarders $u \multimap u'$ as there are available inputs on u' . We remark that the forwarders being linear is crucial; if there were unlimited forwarders for just one input, then any further u forwarded to u' would become inert. Because of linearity, we have a full abstraction result for the encoding using barbed congruence.

Linear forwarders also give an appealing encoding of *distributed choice*. The process $x(u).P + y(v).Q$ allows either a reaction on x , or one on y , but not both. If x and y are at the same location, the choice is easy to implement. If they are remote, we *linearly forward* x and y to co-located channels x' and y' where a local choice can be made:

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

Linearity is again crucial. If the x' branch were taken, then the other forwarder $y \multimap y'$ can be completely undone by releasing a reverse linear forwarder $y' \multimap y$ (or vice versa). With choice, the full abstraction result is weaker, since the correspondence is mediated by coupled simulation [30] rather than barbed congruence. This weaker result is unsurprising, as coupled simulation is also used in the result of Nestmann and Pierce on encoding input-guarded choice in the choice-free $\text{A}\pi$ [29].

We design a distributed abstract machine for $\text{AL}\ell\pi$, called the *linear forwarder machine*. It is a channel-based machine which is distributed in the sense that

the channels are partitioned between locations. All processes must initially be in the “right” place for a machine to be well-formed. This means that every process inputting on free channels, let us say x , are located at the machine of x . Interaction is always local, occurring at the nominated locations of the channels. A property of our machine is that, during interaction, either processes remain in the correct place, or they are outputs and linear forwarders which are small processes with the freedom to move. We show that the linear forwarder machine is a natural implementation of the linear forwarder calculus, proving a full-abstraction result up to barbed congruence. This result, together with our encoding of $\lambda\pi$ into the linear forwarder calculus, shows that it is possible to give a fully-distributed implementation of $\lambda\pi$.

To examine the robustness of the linear forwarder machine, we analyze two failure models usually studied for distributed machines: (1) when failures are due to *message losses* and (2) when failures are due to *crashes of locations*. In (1), since outputs and linear forwarders are the only processes that move over the network, the failure model only accounts for their losses. The effect of one lost output or forwarder is that one input process may deadlock. Our formal account of message failure in the linear forwarder machine amounts to adding two rules that allow the outputs and linear forwarders to disappear. We show that this fallible linear forwarder calculus, $\text{FAL}\pi$ for short, can be simulated via the corresponding machine with failures. This is a weaker result than one using coupled simulation. Coupled simulation cannot be established in this case because, in the machine, local messages can never be lost. On the contrary, in $\text{FAL}\pi$, missing the location information, every message may be lost. For (2), we give a formal account of location crashes in the linear forwarder machine, where processes running on a location are rebooted to some safe initial state. Such rebooted states are assumed to be the empty location for simplicity. We establish similar results to (1).

Related works. Forwarders have already been studied in detail, but for very different reasons to the work presented here. Much work centres around the private pi calculus [35] – a variant of the pi calculus in which only private names may be emitted, as in $(w)\bar{x}w$. Boreale uses forwarders to encode the emission of free names [6]; the reaction $x(u).Q \mid \bar{x}w$ does not perform the substitution $\{w/u\}$, but instead encodes it as a persistent (non-linear) forwarder from u to w . The forwarders must be persistent since names may occur many times in a term. Because of persistence, the number of forwarders increases during an execution of a program. The same technique is used by Merro and Sangiorgi [25] in proofs about the local pi calculus. Both were inspired by Honda’s *equators* [19,24], which are bidirectional forwarders.

In our work, we use linear forwarders in a very different way, to move messages between locations rather than for encoding substitutions. Our forwarders must be linear, and hence their number decreases during execution. In a translated

pi process, if a linear forwarder moves an output message to another location, there will be an input waiting to interact with it. In contrast, this property would not hold if we used persistent forwarders, hence messages would be lost and we would not be able to prove our full-abstraction result. Our proofs are similar in structure to those of Boreale, but are much simpler due to linearity. We have only seen one other use of linear forwarders, in Kobayashi et al.’s work on using linear forwarders to simulate substitutions when channels are used linearly [20]. This work follows Boreale’s agenda, rather than the ideas presented here.

The linear forwarder machine evolved from our previous work on the *fusion machine* [14], which had a similar mechanism for distributing processes. The fusion machine is based on the explicit fusion calculus [17] that uses explicit fusions $x = y$ rather than linear forwarders $x \multimap y$. Fusions yield equivalence classes of names, without giving information about the direction to which outputs should be forwarded – the representatives of the classes. In the fusion machine, this issue was solved using spanning trees. While our solution was a good first attempt, it is in fact a poor implementation strategy since it is not robust to failures of the network. A break of the spanning tree of fused names can result in many messages being lost. This limitation provided our initial motivation for exploring the work presented here.

Other distributed abstract machines for $\lambda\pi$ in the literature include Facile [18], the Jocaml prototype [10], Distributed pi calculus [2], Nomadic Pict [38], the Ambient Calculus [8], the Channel Ambient Machine [31], and the Klaim Machine [11]. Facile uses two classes of distributed entities: (co-)located processes which execute, and channel-managers which mediate interaction. This requires a hand-shake discipline for communication between the locations of the inputs and outputs, and the different locations of the channel managers. Jocaml simplifies the Facile approach by combining input processes with channel-managers, and hence avoiding the hand-shake. However, it uses a quite different form of interaction, which does not relate that directly to pi calculus communication. This is illustrated by the encoding of the pi calculus, which has a strong correctness result but only because it is mediated by firewall processes. In addition, Jocaml forces a coarser granularity, in that every channel must be co-located with at least one other. Like Jocaml, our linear forwarder machine combines processes with channel-managers. Unlike Jocaml, our machine has finer granularity and uses the same form of interaction as the pi calculus. The other machines mentioned are based on a completely different approach to distribution, which adds explicit location constructs to the pi calculus and uses agent migrations for remote interactions. In summary, the linear forwarder machine describes a direct distributed implementation of the asynchronous pi calculus, which is correct in a stronger way than the other proposed implementations.

Summary of paper. Section 2 describes the linear forwarder calculus, and its reference semantics – barbed congruence. Section 3 gives the encoding of the pi calculus (i.e. the encoding of input capability) and proves the encoding correct. Section 4 gives a distributed abstract machine for implementing the linear forwarder calculus and proves the implementation correct. Section 5 extends the pi calculus with choice and the linear forwarder calculus with local choice, and proves a correctness result for the encoding of choice into local choice. Section 6 discusses the extension of the calculi and the abstract machines with failures, and provides correctness results. Section 7 analyzes the issue of loading program codes into the machine. We conclude in Section 8.

This paper extends our conference paper [15] in several ways: we give proofs in full, extend the work to include distributed choice, and analyze failures.

2 The linear forwarder calculus

We assume an infinite set of *names* ranged over by u, v, x, y, \dots . Names represent communication channels; these names can also be transmitted during communication. Write \tilde{x} for a (possibly empty) finite sequence $x_1 \cdots x_n$ of names. *Name substitutions* $\{\tilde{y}/\tilde{x}\}$ are as usual.

Definition 1 (The calculus $\text{AL}\ell\pi$) *The (asynchronous) local linear forwarder calculus, which we abbreviate $\text{AL}\ell\pi$ ¹, is the calculus whose terms P are given by*

$$P ::= \mathbf{0} \quad | \quad \bar{x}\tilde{u} \quad | \quad x(\tilde{u}).P \quad | \quad (x)P \quad | \quad P|P \quad | \quad !P \quad | \quad x \multimap y$$

and which satisfy the no-input-capability constraint: in $x(\tilde{u}).P$, the P has no free occurrence of $u \in \tilde{u}$ as subject of $u(\tilde{v}).Q$.

Free and bound names are standard: x is bound in $(x)P$ and \tilde{u} is bound in $x(\tilde{u}).P$; names are free when they are not bound. Write $\text{fn}(P)$ for the free names of P .

The operators in $\text{AL}\ell\pi$ 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 Definition 3 below. The meaning of the other operators is standard: the term $\mathbf{0}$ is inert; $\bar{x}\tilde{u}$ is a command to send data \tilde{u} over channel x ; $x(\tilde{u}).P$ receives formal arguments \tilde{u} and then continues as P ; restriction $(x)P$ limits the scope of x to P ; parallel composition $P|Q$ allows two terms to run concurrently and

¹ The acronym $\text{AL}\ell\pi$ follows naming conventions in [37]: A stands for “asynchronous”, L for “local” and ℓ for “linear-forwarders”.

interact; and repetition $!P$ behaves like infinitely many copies of P . Write $(x_1 \cdots x_n)P$ for $(x_1) \cdots (x_n)P$.

Following Milner's presentation [27], we first define a structural congruence which equates all agents that have essentially the same structure and which we will never wish to distinguish. We then use structural congruence when giving the operational semantics.

Definition 2 *Structural congruence \equiv is the smallest equivalence relation which satisfies the following axioms and is closed with respect to contexts and alpha-renaming:*

$$\begin{array}{l} P|\mathbf{0} \equiv P \quad P|Q \equiv Q|P \quad P|(Q|R) \equiv (P|Q)|R \quad !P \equiv P!P \\ (x)\mathbf{0} \equiv \mathbf{0} \quad (x)(y)P \equiv (y)(x)P \quad (x)(P|Q) \equiv P | (x)Q \quad \text{if } x \notin \mathbf{fn}(P) \end{array}$$

Definition 3 *The reduction relation \rightarrow is the least relation satisfying the following rules and closed under \equiv , $(x)_-$ and $_ | _$:*

$$x(\tilde{u}).P | \bar{x}\tilde{v} \rightarrow P\{\tilde{v}/\tilde{u}\} \quad \bar{x}\tilde{u} | x\text{-}o y \rightarrow \bar{y}\tilde{u}$$

Notice that the no-input-capability constraint of $\text{ALL}\pi$ is preserved by structural congruence and reaction.

For behavioural equivalence in $\text{ALL}\pi$ we use the standard notion of *barbed bisimulation* [28]. According to this notion, two agents are considered equivalent if their reductions match and they are indistinguishable under global observations:

Definition 4 *The name x is a barb of P , written $P \downarrow x$, when*

$$\begin{array}{ll} \bar{x}\tilde{u} \downarrow x & (y)P \downarrow x, \quad \text{if } P \downarrow x \text{ and } x \neq y \\ !P \downarrow x, \quad \text{if } P \downarrow x & P|Q \downarrow x, \quad \text{if } P \downarrow x \text{ or } Q \downarrow x \end{array}$$

Write \Rightarrow for \rightarrow^* and \Downarrow for $\Rightarrow \downarrow$.

Barbed bisimulation $\overset{\bullet}{\approx}$ is the largest symmetric relation such that whenever $P \overset{\bullet}{\approx} Q$ then (1) $P \downarrow u$ implies $Q \Downarrow u$, and (2) $P \rightarrow P'$ implies $Q \Rightarrow Q'$ and $P' \overset{\bullet}{\approx} Q'$.

Let $C[\]$ be the set of contexts of $\text{ALL}\pi$ generated by

$$C[\] ::= [\] \mid x(\tilde{u}).C[\] \mid (x)C[\] \mid P | C[\] \mid C[\] | P \mid !C[\]$$

The barbed congruence is the largest symmetric relation $\overset{Lc}{\cong}_a$ such that whenever $P \overset{Lc}{\cong}_a Q$ then, for all contexts $C[\]$ such that $C[P]$ and $C[Q]$ are $\text{ALL}\pi$ terms, $C[P] \overset{\bullet}{\approx} C[Q]$.

As examples of barbed congruent terms in $\text{AL}\ell\pi$, we recall a couple of laws.

Proposition 5 (Merro, Sangiorgi [24])

- (1) If P is a term in $\text{AL}\ell\pi$ and u does not occur free in P as subject of an input then $P\{u'/u\} \cong_a^{Lc} (u)(!u \multimap u' \mid P)$.
- (2) $(u)\bar{x}u \cong_a^{Lc} (u)(\bar{x}u \mid u(v).\mathbf{0})$.

We discuss the first law. The rationale behind it is that, in an asynchronous calculus with a semantics that is unsensible to internal moves, input consumptions cannot be tested. Additionally, since u does not occurs as subject of inputs in P , emitting on u' by $P\{u'/u\}$ is the same as outputting on u , then letting the message be consumed by the (persistent linear) forwarder $!u \multimap u'$ and transformed into a message on u' .

3 The asynchronous pi calculus and its encoding

In this section we present the asynchronous pi calculus $\text{A}\pi$ and encode it into $\text{AL}\ell\pi$.

Definition 6 (Pi calculus) *The asynchronous pi calculus $\text{A}\pi$ has the same terms P as for $\text{AL}\ell\pi$ (Definition 1), but with no linear forwarders and no input-capability condition [26]. Structural congruence \equiv and reduction \rightarrow are defined as in Definitions 2 and 3, but in this case the reduction relation \rightarrow only has the rule $x(\tilde{u}).P \mid \bar{x}\tilde{v} \rightarrow P\{\tilde{v}/\tilde{u}\}$. Barbs $P \downarrow u$, barbed bisimulation $\dot{\approx}$, and barbed congruence $\cong_a^{\pi c}$ are as in Definition 4, replacing $\text{AL}\ell\pi$ with $\text{A}\pi$.*

Note that if P and Q are valid terms in $\text{AL}\ell\pi$ and $\text{A}\pi$, then $P \dot{\approx} Q$ in $\text{AL}\ell\pi$ if and only if $P \dot{\approx} Q$ in $\text{A}\pi$. This justifies our use of the same symbol $\dot{\approx}$ for both. Note also that, for congruence \cong_a^{Lc} in $\text{AL}\ell\pi$, the contexts are a subset of those that are used for $\cong_a^{\pi c}$ in $\text{A}\pi$. We will see, however, that $\text{AL}\ell\pi$ contexts and $\text{A}\pi$ contexts are equally discriminating (Theorem 18).

The following law generalizes Proposition 5.1 to $\text{A}\pi$ and, therefore it also holds in $\text{AL}\ell\pi$.

Proposition 7 (Honda, Yoshida [19]) $P\{u'/u\} \cong_a^{\pi c} (u)(!u' \multimap u \mid !u \multimap u' \mid P)$.

We use the law of Proposition 7 in the *context lemma* below. This lemma usually states that contexts gain no additional discriminating power through restriction and input-prefixing [24,37], namely $P \cong_a^{\pi c} Q$ if and only if $R \mid P\sigma \dot{\approx} R \mid Q\sigma$, for every R and renaming σ . However, in asynchronous calculi a simpler statement without renamings is equally strong. (We suspect that this is a standard result, but have not been able to find a reference for it.)

Lemma 8 $R \mid P \overset{\bullet}{\approx} R \mid Q$, for every R , if and only if, for every $C[\]$, $C[P] \overset{\bullet}{\approx} C[Q]$.

PROOF. The reverse direction is clear. For the forward direction, we negate the consequent: i.e. we suppose there exists $C[\]$ such that $C[P] \not\overset{\bullet}{\approx} C[Q]$. By the Context Lemma in [37] there exist R' and σ such that $R' \mid P\sigma \not\overset{\bullet}{\approx} R' \mid Q\sigma$. By Proposition 7, there exist R'' and \tilde{x} such that $(\tilde{x})(R'' \mid P) \cong_a^{Lc} P\sigma$ and $(\tilde{x})(R'' \mid Q) \cong_a^{Lc} Q\sigma$. Henceforth it must be the case that $R' \mid (\tilde{x})(R'' \mid P) \not\overset{\bullet}{\approx} R' \mid (\tilde{x})(R'' \mid Q)$. Without loss of generality, let $\tilde{x} \cap \text{fn}(R') = \emptyset$. This reduces the proof to $(\tilde{x})(R' \mid R'' \mid P) \not\overset{\bullet}{\approx} (\tilde{x})(R' \mid R'' \mid Q)$. Hence $R' \mid R'' \mid P \not\overset{\bullet}{\approx} R' \mid R'' \mid Q$.

Definition 9 (Encoding pi) The encoding $\llbracket \cdot \rrbracket$ maps terms from $\mathcal{A}\pi$ to $\mathcal{A}\mathcal{L}\pi$ by $\llbracket P \rrbracket = \llbracket P \rrbracket_{\emptyset}$, where $\llbracket P \rrbracket_{\tilde{u}}$ for set of names \tilde{u} is an auxiliary function defined inductively by:

$$\begin{aligned} \llbracket x(\tilde{y}).P \rrbracket_{\tilde{u}} &= \begin{cases} x(\tilde{y}).\llbracket P \rrbracket_{\tilde{u}\tilde{y}} & \text{if } x \notin \tilde{u} \\ (x')(x \rightarrow x' \mid x'(\tilde{y}).\llbracket P \rrbracket_{\tilde{u}\tilde{y}}) & \text{if } x \in \tilde{u}, x' \notin \text{fn}(P) \cup \{x, \tilde{u}, \tilde{y}\} \end{cases} \\ \llbracket \mathbf{0} \rrbracket_{\tilde{u}} &= \mathbf{0} \\ \llbracket \bar{x}\tilde{y} \rrbracket_{\tilde{u}} &= \bar{x}\tilde{y} \\ \llbracket (x)P \rrbracket_{\tilde{u}} &= (x)(\llbracket P \rrbracket_{\tilde{u}}) \\ \llbracket P \mid Q \rrbracket_{\tilde{u}} &= \llbracket P \rrbracket_{\tilde{u}} \mid \llbracket Q \rrbracket_{\tilde{u}} \\ \llbracket !P \rrbracket_{\tilde{u}} &= !\llbracket P \rrbracket_{\tilde{u}} \end{aligned}$$

In the input and restriction cases, we assume that the bound names do not clash with \tilde{u} . We write $\tilde{u}\tilde{v}$ for set union and $\tilde{u} \setminus \tilde{v}$ for set difference.

It is worth remarking that we are abusing notation slightly. For input processes, \tilde{u} denotes a sequence of names. In the subscript $\llbracket P \rrbracket_{\tilde{u}}$, \tilde{u} denotes a set of names.

To understand the encoding, note that the subscript contains all names that have been received in input. By the no-input-capability constraint (Definition 1), they cannot therefore be used as the subject of subsequent input. To achieve this, the encoding uses “primed” names to denote local copies of them. So the encoding of $x(u).u(y).P$ will consume a message on x and result in the processes $u'(y).\llbracket P \rrbracket_{u'y}$, where u' is a new channel, and $u \rightarrow u'$ forwarding one message from u to u' . Meanwhile, any output use of u is left unchanged, since it will be forwarded if appropriate using a linear forwarder.

To illustrate the connection between the reactions of a term and its translation, we consider the $\mathcal{A}\pi$ reduction $\bar{x}v \mid x(u).P \rightarrow P\{v/u\}$. By translating we

obtain:

$$\begin{aligned}
\llbracket \bar{x} v \mid x(u).P \rrbracket_x &= \bar{x} v \mid (x')(x \multimap x' \mid x'(u).\llbracket P \rrbracket_{xu}) \\
&\rightarrow (x')(\bar{x}' v \mid x'(u).\llbracket P \rrbracket_{xu}) \\
&\rightarrow (x')(\llbracket P \rrbracket_{xu}\{v/u\}) \\
&\equiv \llbracket P \rrbracket_{xu}\{v/u\}
\end{aligned}$$

Note that the final state of the translated term is subscripted on x and u , not just on x . In addition, the translated term ends up with some garbage that was not present in the original (the restriction $(x')(\dots)$). 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.

Linearity is crucial in the translation. For instance, consider a nonlinear translation where forwarders are replicated:

$$\langle x(u).P \rangle_x = (x')(!x \multimap x' \mid x'(u).\langle P \rangle_{xu})$$

Then consider the example

$$\begin{aligned}
&\langle x(u).P \mid x(u).Q \mid \bar{x} v \mid \bar{x} w \rangle_x \\
&= (x')(!x \multimap x' \mid x'(u).\langle P \rangle_{xu}) \mid (x'')(!x \multimap x'' \mid x''(u).\langle Q \rangle_{xu}) \mid \bar{x} v \mid \bar{x} w \\
&\Rightarrow (x')(!x \multimap x' \mid \langle P \rangle_{xu}\{v/u\}) \mid (x'')(!x \multimap x'' \mid x''(u).\langle Q \rangle_{xu}) \mid \bar{x} w \\
&\rightarrow (x')(!x \multimap x' \mid \langle P \rangle_{xu} \mid \bar{x}' w) \mid (x'')(!x \multimap x'' \mid x''(u).\langle Q \rangle_{xu})
\end{aligned}$$

Here, both outputs were forwarded to the local name x' , even though the resource $x'(u).\langle P \rangle_{xu}$ had already been consumed by the first one. This precludes the second one from reacting with $(x'')(!x \multimap x'' \mid x''(u).\langle Q \rangle_{xu})$, a reaction that would have been possible in the original $\lambda\pi$ term. Linearity prevents the possibility of such dead ends.

Remark 10 *A simpler encoding is also possible, which does not use subscripts and always applies the $x \in \tilde{u}$ case. We chose the current one for the following appealing property: if P has no free input on x then $\llbracket P \rrbracket_{\tilde{z}x} = \llbracket P \rrbracket_{\tilde{z}}$. As an immediate consequence, if a term P in $\lambda\pi$ satisfies the no-input-capability constraint, then the encoding $\llbracket \cdot \rrbracket$ leaves it unchanged.*

3.1 The correctness of the encoding

The correctness of the encoding $\llbracket \cdot \rrbracket$, that $P \cong_a^{\pi^c} Q$ if and only if $\llbracket P \rrbracket \cong_a^{L^c} \llbracket Q \rrbracket$, is not straightforward to prove because of the garbage left by the encoded term. To show that it is indeed garbage, we must prove that $\llbracket P \rrbracket_u$ and $\llbracket P \rrbracket_{ux}$ are congruent. But the barbed semantics offers too weak an induction hypothesis for this proof. A standard alternative technique (used for instance by Boreale [6])

is to use the barbed semantics as the primary definition, but to switch to a labelled transition system and labelled bisimulation in the proofs. Labelled bisimulation is defined coinductively, which makes it easier to use in proofs, and is stronger than \cong_a^{Lc} . It therefore provides a useful proof technique for establishing results about \cong_a^{Lc} .

To help our argument, we introduce a further calculus, called the asynchronous, *non-local* linear forwarder calculus $\mathcal{A}l\pi$ in Definition 11.

Definition 11 *The calculus $\mathcal{A}l\pi$ has terms P and contexts $C[]$ as in $\mathcal{A}Ll\pi$ (Definitions 1 and 4), but without the no-input-capability constraint. Structural congruence \equiv reduction \rightarrow , and barbs are as in $\mathcal{A}Ll\pi$ (Definitions 2, 3, and 4).*

The calculus $\mathcal{A}l\pi$ has both $\mathcal{A}Ll\pi$ and $\mathcal{A}\pi$ as subcalculi. It is $\mathcal{A}Ll\pi$ without the no-input-capability constraint or, equivalently, $\mathcal{A}\pi$ with linear forwarders. We first define the labelled transition system for $\mathcal{A}l\pi$, and therefore for $\mathcal{A}Ll\pi$ and $\mathcal{A}\pi$. To this aim, let μ range over input labels $x(\tilde{u})$, bound output labels $(\tilde{z})\bar{x}\tilde{u}$ where $\tilde{z} \subseteq \tilde{u}$, and the label τ . Let also $\mathbf{fn}(x(\tilde{u})) = \{x\}$, $\mathbf{fn}((\tilde{z})\bar{x}\tilde{u}) = \{x, \tilde{u}\} \setminus \tilde{z}$, $\mathbf{bn}(x(\tilde{u})) = \{\tilde{u}\}$, $\mathbf{bn}((\tilde{z})\bar{x}\tilde{u}) = \{\tilde{z}\}$, and $\mathbf{fn}(\tau) = \mathbf{bn}(\tau) = \emptyset$.

Definition 12 *The transition relation $P \xrightarrow{\mu} Q$, for terms P and Q in $\mathcal{A}l\pi$, is defined inductively on the structure of P :*

$$\begin{array}{c}
x(\tilde{u}).P \xrightarrow{x(\tilde{u})} P \qquad \bar{x}\tilde{u} \xrightarrow{\bar{x}\tilde{u}} \mathbf{0} \qquad x \circ y \xrightarrow{x(\tilde{u})} \bar{y}\tilde{u} \\
\\
\frac{P \xrightarrow{\mu} Q \quad y \notin \mathbf{fn}(\mu)}{(y)P \xrightarrow{\mu} (y)Q} \qquad \frac{P \xrightarrow{(\tilde{z})\bar{x}\tilde{u}} Q \quad y \neq x, y \in \tilde{u} \setminus \tilde{z}}{(y)P \xrightarrow{(y\tilde{z})\bar{x}\tilde{u}} Q} \qquad \frac{P \mid !P \xrightarrow{\mu} Q}{!P \xrightarrow{\mu} Q} \\
\\
\frac{P \xrightarrow{\mu} P' \quad \mathbf{bn}(\mu) \cap \mathbf{fn}(Q) = \emptyset}{P \mid Q \xrightarrow{\mu} P' \mid Q} \qquad \frac{P \xrightarrow{(\tilde{z})\bar{x}\tilde{v}} P' \quad Q \xrightarrow{x(\tilde{u})} Q' \quad \tilde{z} \cap \mathbf{fn}(Q) = \emptyset}{P \mid Q \xrightarrow{\tau} (\tilde{z})(P' \mid Q'\{\tilde{v}/\tilde{u}\})}
\end{array}$$

The transitions of $P \mid Q$ have mirror cases, which we have omitted. We implicitly identify terms up to alpha-renaming \equiv_α : that is, if $P \equiv_\alpha P'$ and $P' \xrightarrow{\mu} P''$ then $P \xrightarrow{\mu} P''$. Write $\xrightarrow{\tau}$ for $\xrightarrow{\tau^*}$, and $\xrightarrow{\mu}$ for $\xrightarrow{\tau} \xrightarrow{\mu} \xrightarrow{\tau}$ when $\mu \neq \tau$.

The reduction and barb relations in Definition 11 and the transition relation in Definition 12 are related by the proposition below.

Proposition 13 (1) $P \rightarrow P'$ if and only if $P \xrightarrow{\tau} \equiv P'$;
(2) $P \downarrow x$ if and only if $P \xrightarrow{\bar{x}\tilde{u}} P'$, for some \tilde{u} .

The calculus $\mathcal{A}l\pi$ is equipped with an observational semantics called *asynchronous bisimulation* [4]. The next definition also introduces the asynchronous simulation, to be used in Section 5.

Definition 14 *An asynchronous simulation is a binary relation \mathcal{R} between $\mathcal{A}l\pi$ processes such that $P \mathcal{R} Q$ implies:*

- (1) if $P \xrightarrow{\tau} P'$, then $Q \xRightarrow{\tau} Q'$ and $P' \mathcal{R} Q'$;
- (2) if $P \xrightarrow{(\tilde{z})\bar{x}\tilde{u}} P'$ and $\tilde{z} \cap \text{fn}(Q) = \emptyset$, then $Q \xRightarrow{(\tilde{z})\bar{x}\tilde{u}} Q'$ and $P' \mathcal{R} Q'$;
- (3) if $P \xrightarrow{x(\tilde{u})} P'$ and $\tilde{u} \cap \text{fn}(Q) = \emptyset$ then
 - (a) either $Q \xrightarrow{x(\tilde{u})} Q'$, and $P' \mathcal{R} Q'$;
 - (b) or $Q \xrightarrow{\tau} Q'$, and $P' \mathcal{R} (Q' \mid \bar{x}\tilde{u})$.

A relation \mathcal{R} is an asynchronous bisimulation when both \mathcal{R} and \mathcal{R}^{-1} are asynchronous simulations. Asynchronous bisimilarity, written \approx_a is the largest asynchronous bisimulation.

Some standard results, from asynchronous bisimilarity for $\Lambda\pi$ (see [37], Chapter 5), extend naturally to $\Lambda\ell\pi$:

- Theorem 15** (1) \approx_a is an equivalence relation;
- (2) $\equiv \subseteq \approx_a$;
 - (3) \approx_a is a congruence.

Clearly \approx_a is a barbed bisimulation; therefore $\approx_a \subseteq \overset{\bullet}{\approx}$. Because it is a congruence, it is closed under all contexts in $\Lambda\ell\pi$, then *a fortiori* it is also closed under $\Lambda\pi$ contexts and $\Lambda L\ell\pi$ contexts, and so

$$\approx_a \subset \cong_a^{\pi c} \quad \text{and} \quad \approx_a \subset \cong_a^{Lc}. \quad (2)$$

These key equations justify our choice to perform our proofs in \approx_a rather than $\cong_a^{\pi c}$ or \cong_a^{Lc} . We recall that asynchronous bisimilarity, although useful as a proof technique, is stronger than desirable for general use. For instance, by Proposition 5.1, $\bar{x}u \cong_a^{Lc} (v)(\bar{x}v \mid !(v \multimap u))$. However this equality is false in \approx_a . In particular, the term on the right hand side may undergo $\xrightarrow{(v)\bar{x}(v)}$, but the term on the left cannot.

The following laws are key to the correctness proof.

Lemma 16 For all P in $\Lambda\ell\pi$ and $y \notin \text{fn}(P)$,

- (1) $x \multimap z \approx_a x(\tilde{u}).\bar{z}\tilde{u}$.
- (2) $P \approx_a (y)(\bar{y}\tilde{u} \mid y(\tilde{u}).P)$.
- (3) $x(\tilde{u}).P \approx_a (y)(x \multimap y \mid y(\tilde{u}).P)$.

PROOF. For item (i), consider the relation $\mathcal{S} = \{(x \multimap z, x(\tilde{u}).\bar{z}\tilde{u})\} \cup \approx_a$. It is easy to show that \mathcal{S} is contained in \approx_a .

For item (ii), this is an instance of a well-known law in $\Lambda\pi$: $P\{\tilde{v}/\tilde{u}\} \approx_a (y)(\bar{y}\tilde{v} \mid y(\tilde{u}).P)$.

Item (iii) is omitted because it is similar to item 2.

The following corollary shows the operational correspondence of the encoding: any transition of P is matched by one in $\llbracket P \rrbracket_{\tilde{u}}$, and vice versa up to asynchronous bisimilarity.

Corollary 17 (Operational correspondence) $P \approx_a \llbracket P \rrbracket_{\tilde{u}}$, for all \tilde{u} .

PROOF. This follows directly from Lemma 16.3 and Definition 9. We spell out the details. Consider the smallest relation \mathcal{S} containing the identity and closed under the following.

- If $P \mathcal{S} Q$ and $Q \approx_a R$ then $P \mathcal{S} R$.
- If $P \mathcal{S} Q$ then also $x(\tilde{u}).P \mathcal{S} x(\tilde{u}).Q$, $(x)P \mathcal{S} (x)Q$, $!P \mathcal{S} !Q$, and $P\sigma \mathcal{S} Q\sigma$ for all substitutions σ .
- If $P \mathcal{S} Q$ and $P' \mathcal{S} Q'$ then $P \mid P' \mathcal{S} Q \mid Q'$.

The proof that \mathcal{S} is a \approx_a -bisimulation is standard. By Lemma 16.3 we derive $\{(P, \llbracket P \rrbracket_{\tilde{u}})\} \subseteq \mathcal{S}$, thus concluding the proof.

We have all the preliminaries in place for demonstrating the correctness of the encoding.

Theorem 18 (Correctness) For P and Q in $\mathsf{A}\pi$,

- (1) $P \dot{\approx} Q$ if and only if $\llbracket P \rrbracket \dot{\approx} \llbracket Q \rrbracket$;
- (2) $P \cong_a^{\pi c} Q$ if and only if $\llbracket P \rrbracket \cong_a^{Lc} \llbracket Q \rrbracket$.

PROOF. Item 1 is a simple consequence of Corollary 17, since $\approx_a \subset \dot{\approx}$. For item 2 we use Lemma 8 and its analogue in $\mathsf{A}\pi$; namely, for P, Q in $\mathsf{A}\pi$ then

$$\forall R : R \mid P \dot{\approx} R \mid Q \Leftrightarrow P \cong_a^{\pi c} Q. \quad (3)$$

To prove item 2, start in the reverse direction and assume the converse that $P \not\cong_a^{\pi c} Q$. By Equation 3 there exists a term R in $\mathsf{A}\pi$ such that $R \mid P \not\dot{\approx} R \mid Q$. The followings results hold:

$$\begin{array}{ll} R \mid P \not\dot{\approx} R \mid Q & \text{in } \mathsf{A}\pi \\ \approx_a \quad \approx_a & \text{by Corollary 17} \\ \llbracket R \rrbracket \mid \llbracket P \rrbracket \not\dot{\approx} \llbracket R \rrbracket \mid \llbracket Q \rrbracket & \text{in } \mathsf{AL}\ell\pi \end{array}$$

Hence $\llbracket P \rrbracket \not\cong_a^{Lc} \llbracket Q \rrbracket$ in $\mathsf{AL}\ell\pi$.

We now prove the forwards direction. Define a translation $\widehat{\cdot}$ from $\mathsf{A}\ell\pi$ to

$\Lambda\pi$ as follows:

$$\begin{aligned}
\widehat{u \dashv \circ v} &= u(\tilde{x}).\bar{v}\tilde{x} \\
\widehat{u(\tilde{x}).P} &= u(\tilde{x}).\widehat{P} \\
\widehat{\mathbf{0}} &= \mathbf{0} \\
\widehat{\bar{u}\tilde{x}} &= \bar{u}\tilde{x} \\
\widehat{(x)P} &= (x)\widehat{P} \\
\widehat{P|Q} &= \widehat{P} | \widehat{Q} \\
\widehat{!P} &= !\widehat{P}
\end{aligned}$$

By Lemma 16.1, for any R in $\Lambda\ell\pi$, we have $R \approx_a \widehat{R}$. Assume $\llbracket P \rrbracket \not\approx_a^{Lc} \llbracket Q \rrbracket$ in $\text{AL}\ell\pi$. Then by Lemma 8 there exists R in $\text{AL}\ell\pi$ such that $R | \llbracket P \rrbracket \not\approx R | \llbracket Q \rrbracket$. By Corollary 17, $P \approx_a \llbracket P \rrbracket$. Therefore by congruence and transitivity of \approx_a we have $R | \llbracket P \rrbracket \approx_a \widehat{R} | P$ and $R | \llbracket Q \rrbracket \approx_a \widehat{R} | Q$. Additionally, since $\approx_a \subseteq \dot{\approx}$ (see the discussion after Theorem 15) then $\widehat{R} | P \not\approx \widehat{R} | Q$. Hence $P \not\approx_a^{\pi c} Q$.

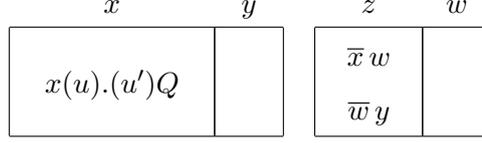
4 The linear forwarder machine

In this section we develop a distributed machine for $\text{AL}\ell\pi$, called the *linear forwarder machine*, which is suitable for a point-to-point network such as the Internet. We start with a set of locations identified by IP numbers. Each location contains several *channel-managers* identified by port numbers, which run in parallel. Each channel-manager contains an unordered set of program-fragments; some fragments may be blocked, waiting to receive on that channel; others may be waiting to be deployed to the right channel manager; others may be waiting to be interpreted locally. Let the names u, v, x, y, \dots range over IP:port pairs, so that each name identifies a particular channel-manager at a particular location.

We first give a diagrammatic overview of the machine. Then we provide a formal syntax and semantics, give a translation from $\text{AL}\ell\pi$ to the machine, and prove a full abstraction result with respect to barbed congruence.

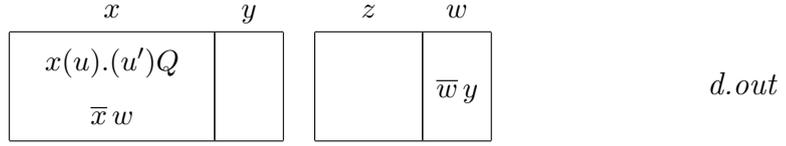
4.1 Machine diagrams

The following diagram shows a machine consisting of two locations, one with channel managers x, y and the other with z, w .

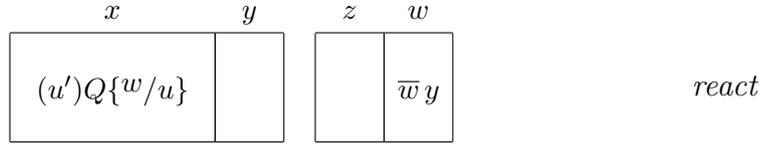


This machine corresponds to the $\text{AL}\ell\pi$ term $\bar{x}w \mid \bar{w}y \mid x(u).(u')Q$. Let $Q = u \multimap u' \mid u'(v).R$, so that the above term also corresponds to the example in the Introduction, which illustrates the encoding of input capability into linear forwarders.

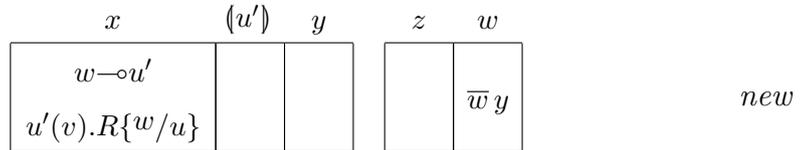
The input fragment at x is currently blocked, waiting for input. The output fragments at z will be sent asynchronously to their intended destinations:



Now the channel-manager x has an available input and output, and so reacts them together. This gives rise to the substitution $\{w/u\}$:

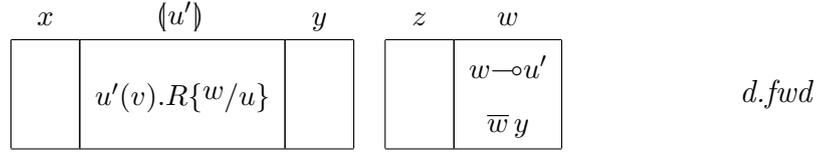


Next the channel-manager x executes the (u') command, by creating a new co-located channel manager. In order to evidence that u' is a channel manager local to the current machine, we wrap u' into the parentheses (\cdot) . Recall that $Q\{w/u\} = w \multimap u' \mid u'(v).R\{w/u\}$.

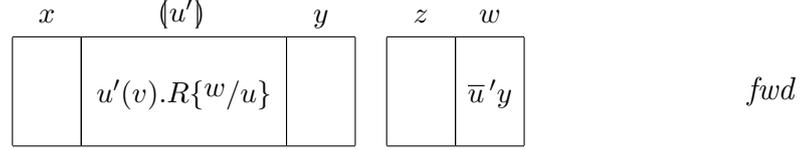


Next the linear forwarder $w \multimap u'$ is sent to its intended destination w . This is easy to implement because the linear forwarder is a small data structure carrying only two channel names. Additionally, the input $u'(z).R\{w/u\}$ is sent

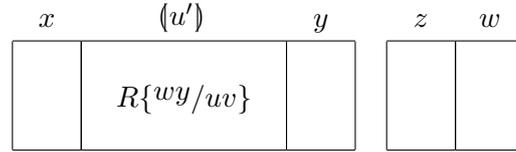
to u' . This is also easy to implement because x and u' are co-located:



The channel-manager at w now contains a linear forwarder and an output, and so allows them to react:



Finally, this output at w is forwarded to channel u' where it can react



A crucial machine step is the deployment of the input $u'(v).R\{w/u\}$ from x to u' . This deployment is only allowed when x and u' are co-located. Otherwise, it would be too unwieldy to move the potentially-large program across the network. We must ensure that programs are never deployed at runtime between channels that are not co-located. We enforce this with well-formedness properties. First, we require that, in a machine $x[P]$, every free input $z(v).Q$ in P is such that the channels x and z are co-located. Second, we require that P satisfies the no-input-capability constraint. Together, these constraints are sufficient to avoid the deployment problem.

4.2 The formal account

Definition 19 (Linear forwarder machine) *The linear forwarder machines M are given by the following grammar, where P ranges over terms in $\text{AL}\ell\pi$ (Definition 1).*

$$M ::= \mathbf{0} \mid \hat{x}[P] \mid (x)[P] \mid M, M$$

The presentation here is similar to that given for the fusion machine [14]. The *basic channel-manager* $\hat{x}[P]$ 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. For example,

if $M = (x)[P], M'$ then x cannot be accessed from the environment without being extruded in advance. We write $\mathbf{lchan}(M)$ for the names of only the *local channel-managers*; we write $\mathbf{chan}(M)$ to denote the set of names of all channel-managers, local and not, in the machine. We also write x to denote \dot{x} or (x) .

In this presentation, we have used arbitrary replication $!P$ for simplicity. In a real machine, we would instead use guarded replication $!x(\tilde{u}).P$ [36], as it is more amenable to implementation.

Assume a *co-location* equivalence relation L on channels. 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: to this end we assume that each equivalence class in L is infinitely large. In the machine calculus, we generally assume L rather than writing it explicitly. Machines meet the following *well-formed conditions*:

Localization. All code is in the right place, and does not need to be moved at runtime. Formally, for every channel machine $x[P]$, then every free input $z(\tilde{u}).Q$ in P satisfies $x@z$. Moreover, P satisfies the no-input-capability constraint.

Single-definition. There is exactly one channel-manager per channel. Formally, a machine $x_1[P_1], \dots, x_n[P_n]$ is singly-defined when $i \neq j$ implies $x_i \neq x_j$.

For example, the machines $x[y().P], y[\bar{y}]$ with $x@y$ and $x[P], x[Q]$ are ill-formed. Note that the localization constraint for machines is stricter than that used for $\text{AL}\ell\pi$ or the $L\pi$ of [37]. This is to be expected, since these calculi do not have explicit locations. Well-formedness is preserved by structural congruence and reactions defined below.

Definition 20 *The structural congruence \equiv is the smallest equivalence relation satisfying the following axioms and closed with respect to alpha-equivalence of local channel managers:*

$$\begin{aligned} M, \mathbf{0} &\equiv M & M_1, M_2 &\equiv M_2, M_1 & M_1, (M_2, M_3) &\equiv (M_1, M_2), M_3 \\ P &\equiv Q & \text{implies} & \dot{x}[P] &\equiv \dot{x}[Q] & \text{and } (x)[P] &\equiv (x)[Q] \end{aligned}$$

The reduction step \rightarrow and the heating step \dashrightarrow are the smallest relations satisfying the rules below, and closed with respect to structural congruence.

$$\begin{aligned} x[\bar{x}\tilde{v} \mid x(\tilde{u}).P \mid R] &\rightarrow x[P\{\tilde{v}/\tilde{u}\} \mid R] && (\text{react}) \\ x[x\multimap z \mid \bar{x}\tilde{v} \mid R] &\rightarrow x[\bar{z}\tilde{v} \mid R] && (\text{fwd}) \end{aligned}$$

$$\begin{array}{l}
x[(y)P \mid R] \rightarrow x[P\{z/y\} \mid R], (z) [\mathbf{0}], \text{ if } z \notin \{x\} \cup \text{fn}(P \mid R), z @ x \\
\hspace{15em} \text{(new)} \\
x[z \multimap y \mid P], z[Q] \rightarrow x[P], z[z \multimap y \mid Q], \hspace{10em} \text{(d.fwd)} \\
x[\bar{z}\tilde{v} \mid P], z[Q] \rightarrow x[P], z[\bar{z}\tilde{v} \mid Q], \hspace{10em} \text{(d.out)} \\
x[z(\tilde{u}).P \mid Q], z[R] \rightarrow x[Q], z[z(\tilde{u}).P \mid R], \text{ if } x @ z \hspace{10em} \text{(d.in)} \\
\hline
M \rightarrow M', \quad \text{chan}(M') \cap \text{chan}(N) = \emptyset \quad M \rightarrow M', \quad \text{chan}(M') \cap \text{chan}(N) = \emptyset \\
M, N \rightarrow M', N \hspace{15em} M, N \rightarrow M', N
\end{array}$$

We write $M \Rightarrow M'$ if $M (\rightarrow \cup \rightarrow)^* M'$.

Alpha-equivalence of local channel managers means that, if $M = (x) [P]$, M' and $z \notin \text{chan}(M)$, then $M \equiv (z) [P\{z/x\}]$, $M'\{z/x\}$. We observe that, in rules *d.fwd*, *d.out*, and *d.in*, $x \neq z$ by well-formedness.

The rationale of splitting steps in reductions and heatings is as follows: reduction steps closely reflect the reductions in $\text{AL}\ell\pi$; heating steps model movements of codes between locations of the distributed machine. We draw attention to steps (*new*) and (*d.in*). The step (*new*) picks a fresh channel-name z and this channel is deemed to be at the location where the command was executed. The step (*d.in*) will only move an input from one channel x to another channel z , if the two channels are co-located; hence, there is no “real” movement of input processes. In other words, the migration of input processes from a location to a *different* one is disallowed in the linear forwarder machine. Actually, such machines are banned by well-formedness. Note that well-formedness is a sufficient but not a *necessary* condition for avoiding such migrations. For instance, the *localization* constraint says that nested free inputs must be co-located. However, if x and u are not co-located then $x[(z)(z().u())]$ violates the localization property, but it is inert and would show no problem at run-time.

Definition 21 A name x is a barb in a machine M , written $M \downarrow x$, when

$$\begin{array}{ll}
\dot{z}[P] \downarrow x & \text{if } P \downarrow x \\
(z) [P] \downarrow x & \text{if } P \downarrow x \text{ and } x \neq z \\
M, N \downarrow x & \text{if } x \notin \mathbf{1}\text{chan}(M, N) \text{ and } (M \downarrow x \text{ or } N \downarrow x)
\end{array}$$

We write $M \Downarrow x$, when $M \Rightarrow M'$ and $M' \downarrow x$.

For example

$$\dot{x}[\bar{x}u] \downarrow x \quad y[\bar{x}u] \downarrow x \quad (x) [\bar{x}u] \not\downarrow x \quad (x) [\mathbf{0}], y[\bar{x}u] \not\downarrow x.$$

Definition 22 Barbed bisimulation $\overset{\bullet}{\approx}$ is defined as in Definition 4, where processes are replaced by machines. Machine equivalence $M \cong_a M'$ holds

when, for all machines N , then $M, N \overset{\circ}{\approx} M', N$ assuming that M, N and M', N are well-formed.

For example $\dot{x}[\bar{x}u] \not\approx_a \dot{x}[\bar{x}v]$ because of the context machine $\dot{z}[x \rightarrow z \mid z(w).\bar{w}]$ or because the machine $\dot{z}[x(w).\bar{w}]$ where $x @ z$. This is an interesting point: in machine contexts, we may always add inputs on remote names either by using linear forwarders or by using co-located names.

The linear forwarder machine is a faithful implementation of $\text{AL}\ell\pi$, as shown by the the translation and results below.

Definition 23 (Translation) Let $\text{calc}(M) = (\tilde{z})(\widehat{M})$ where $\tilde{z} = \text{1chan}(M)$ and

$$\widehat{\mathbf{0}} = \mathbf{0} \quad \widehat{\dot{x}[P]} = P \quad \widehat{(x)[P]} = P \quad \widehat{M, N} = \widehat{M} \mid \widehat{N}.$$

Before demonstrating the correctness of the translation, we give notations for machines and state some preliminary results.

Notation 24 $(\tilde{x})M$, with $\tilde{x} \subseteq \text{chan}(M)$, denotes the machine M where every channel in \tilde{x} has been turned local;

$M \parallel N$, with $\text{1chan}(M) = \text{1chan}(N) = \emptyset$, denotes the machine defined as follows.

$$M \parallel N \stackrel{\text{def}}{=} \begin{cases} N & \text{if } M = \mathbf{0} \\ \dot{x}[P|Q], (M' \parallel (N', N'')) & \text{if } M = \dot{x}[P], M', N = N', \dot{x}[Q], N'' \\ \dot{x}[P], (M' \parallel N) & \text{if } M = \dot{x}[P], M', x \notin \text{chan}(N) \end{cases}$$

$M \mid N$, where $M = (\tilde{x})M'$, $N = (\tilde{y})N'$, $\text{1chan}(M') = \text{1chan}(N') = \emptyset$, and $\tilde{x} \cap \text{chan}(N') = \emptyset = \tilde{y} \cap \text{chan}(M')$, denotes the machine $(\tilde{x}\tilde{y})(M \parallel N)$.

It is easy to verify that $(M \mid M') \mid M'' = M \mid (M' \mid M'')$. Therefore, in the following, we compose machines in parallel without caring about parentheses. The proposition below collects some properties that follow directly from definitions.

Proposition 25 (1) If $M \rightarrow M'$ and $x \notin \text{1chan}(M)$ then $(x)(M \mid x[\mathbf{0}]) \rightarrow (x)(M' \mid x[\mathbf{0}])$. Similarly, if $M \rightarrow M'$ and $x \notin \text{1chan}(M) \cup \text{1chan}(M')$ then $(x)(M \mid x[\mathbf{0}]) \rightarrow (x)(M' \mid x[\mathbf{0}])$.

(2) If $M \rightarrow M'$ and $M \mid N$ is defined then $M \mid N \rightarrow M' \mid N$. Similarly, if $M \rightarrow M'$ and both $M \mid N$ and $M' \mid N$ are defined then $M \mid N \rightarrow M' \mid N$.

Note that Proposition 25 has additional constraints for heating steps. These constraints guarantee the absence of clashes of names created by steps in the context.

Labelled transitions in the calculus and heating steps in the machine are related as detailed by the following lemma. Let $\mathbf{0}_{x_1 \dots x_n}$ denote the machine

$x_1[\mathbf{0}], \dots, x_n[\mathbf{0}]$.

Lemma 26 (1) If $\text{calc}(M) \xrightarrow{x(\tilde{u})} P$ then:

- either $M \rightarrow^* (\tilde{y})(x[x(\tilde{u})].Q) \mid M'$ and $P \equiv \text{calc}((\tilde{y})(x[Q] \mid M'))$,
 - or $M \rightarrow^* (\tilde{y})(x[x \circ z] \mid M')$ and $P \equiv \text{calc}((\tilde{y})(x[\tilde{z}\tilde{u}] \mid M'))$
- according to whether the process that moves in $\text{calc}(M)$ is an input or a linear forwarder;

(2) if $\text{calc}(M) \xrightarrow{(\tilde{z})\tilde{x}\tilde{v}} P$ then $M \rightarrow^* (\tilde{z}\tilde{y})(x'[\tilde{x}\tilde{v}] \mid \mathbf{0}_{\tilde{z}} \mid M')$ with $\tilde{z} \cap \tilde{y} = \emptyset$ and $P \equiv \text{calc}((\tilde{y})(x'[\mathbf{0}] \mid \mathbf{0}_{\tilde{z}} \mid M'))$.

PROOF. The first item is not difficult to prove. We demonstrate the second item by induction on the depth of the derivation of $\text{calc}(M) \xrightarrow{(\tilde{z})\tilde{x}\tilde{v}} P$. The basic case is when $\text{calc}(M) = \tilde{x}\tilde{v}$. Therefore, by definition of calc , there is x' such that $M = x'[\tilde{x}\tilde{v}]$. The statement follows immediately.

The inductive steps regard the rules for parallel, scope, and replication. For parallel we have $\text{calc}(M) = Q \mid R$ with $Q \xrightarrow{(\tilde{z})\tilde{x}\tilde{v}} P$ and $\tilde{z} \notin \text{fn}(R)$. In this case $\text{lchan}(M) = \emptyset$, therefore, by definitions of calc and the operation \mid on machines, there are M' and M'' such that $M = M' \mid M''$, $\text{calc}(M') = Q$, and $\text{calc}(M'') = R$. By inductive hypothesis, $M' \rightarrow^* (\tilde{z}\tilde{y})(x'[\tilde{x}\tilde{v}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1)$ and $P \equiv \text{calc}((\tilde{y})(x'[\mathbf{0}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1))$. Then, by Proposition 25.2, $M \rightarrow^* (\tilde{z}\tilde{y})(x'[\tilde{x}\tilde{v}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1) \mid M''$ and $P \mid R = \text{calc}((\tilde{z}\tilde{y})(x'[\tilde{x}\tilde{v}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1) \mid M'')$ follows by definition of calc . For scope we have $\text{calc}(M) = (u)Q$ and $Q \xrightarrow{(\tilde{z})\tilde{x}\tilde{v}} P$. There are two cases: (a) when $u \notin \tilde{v} \setminus \tilde{z}x$, and (b) when $u \in \tilde{v} \setminus \tilde{z}x$. We discuss subcase (b), the subcase (a) is simpler. By definition of calc , either (b1) $M = (u)[Q'] \mid M'$ or (b2) $M = x[(u)Q]$. We discuss (b2). In this case, $M \rightarrow (u)(x'[Q] \mid u[\mathbf{0}])$ with a (*new*). By inductive hypothesis $x'[Q] \rightarrow^* (\tilde{z}\tilde{y})(x'[\tilde{x}\tilde{v}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1)$ and $P \equiv \text{calc}((\tilde{y})(x'[\mathbf{0}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1))$. Without loss of generality, let $u \notin \tilde{y}$. By Proposition 25, $M \rightarrow^* (u)(\tilde{z}\tilde{y})(x'[\tilde{x}\tilde{v}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1 \mid u[\mathbf{0}]) = (u\tilde{z}\tilde{y})(x'[\tilde{x}\tilde{v}] \mid \mathbf{0}_{u\tilde{z}} \mid M'_1)$. By definition of calc and structural congruence: $\text{calc}((\tilde{y})(x'[\tilde{x}\tilde{v}] \mid \mathbf{0}_{u\tilde{z}} \mid M'_1)) \equiv \text{calc}((\tilde{y})(x'[\mathbf{0}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1)) \equiv P$. The case when the last rule is replication may be reduced to the parallel case.

The theorem about the correctness of the encoding calc follows.

Theorem 27 (Machine correctness) (1) $M \equiv M'$ implies $\text{calc}(M) \equiv \text{calc}(M')$, $M \rightarrow M'$ implies $\text{calc}(M) \equiv \text{calc}(M')$, and $M \rightarrow M'$ implies $\text{calc}(M) \rightarrow \text{calc}(M')$.

(2) $\text{calc}(M) \rightarrow P$ implies there exists N such that $M \Rightarrow N$ and $P \equiv \text{calc}(N)$.

(3) $M \Downarrow x$ if and only if $\text{calc}(M) \Downarrow x$.

(4) $M \dot{\approx} N$ if and only if $\text{calc}(M) \dot{\approx} \text{calc}(M')$.

(5) $M \cong_a M'$ if and only if $\text{calc}(M) \cong_a^{Lc} \text{calc}(M')$.

PROOF. Item 1 is obvious.

For Item 2, the reduction $\text{calc}(M) \rightarrow P$ might have been deduced using structural congruence, which is difficult to match with structural congruence in the machine. Instead, we analyse the reduction via labelled transitions because, by Proposition 13, $\text{calc}(M) \rightarrow P$ implies $\text{calc}(M) \xrightarrow{\tau} P' \equiv P$, for some P' . Therefore it is enough to prove that $\text{calc}(M) \xrightarrow{\tau} P'$. The argument is by induction on the proof of this transition. The basic case is when the last rule is a communication rule:

$$\frac{R \xrightarrow{(\tilde{z})\bar{x}\tilde{v}} R' \quad Q \xrightarrow{x(\tilde{u})} Q' \quad \tilde{z} \cap \text{fn}(Q) = \emptyset}{\text{calc}(M) = R \mid Q \xrightarrow{\tau} (\tilde{z})(R' \mid Q'\{\tilde{v}/\tilde{u}\}) = P'}$$

We discuss the case when $Q \xrightarrow{x(\tilde{u})} Q'$ is due to an input process; the case when the transition is due to a linear forwarder is similar. By Proposition 25, let M', M'' be such that $M = M' \mid M''$ and $\text{calc}(M') = R$ and $\text{calc}(M'') = Q$. By Lemma 26, $M' \rightarrow^* (\tilde{z}\tilde{y})(x'[\bar{x}\tilde{v}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1)$ and $M'' \rightarrow^* (\tilde{y}')(x[x(\tilde{u}).Q''] \mid M''_1)$. Therefore, by assuming that names \tilde{u} and \tilde{z} do not clash with any other name:

$$\begin{aligned} M &\rightarrow^* (\tilde{z}\tilde{y})(x'[\bar{x}\tilde{v}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1) \mid (\tilde{y}')(x[x(\tilde{u}).Q''] \mid M''_1) && \text{by Proposition 25} \\ &\rightarrow (\tilde{z})((\tilde{y})(x'[\mathbf{0}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1) \mid (\tilde{y}')(x[\bar{x}\tilde{v} \mid x(\tilde{u}).Q''] \mid M''_1)) && \text{by } (d.out) \\ &\rightarrow (\tilde{z})((\tilde{y})(x'[\mathbf{0}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1) \mid (\tilde{y}')(x[Q''\{\tilde{v}/\tilde{u}\}] \mid M''_1)) && \text{by } (react) \\ &= (\tilde{z})((\tilde{y})(x'[\mathbf{0}] \mid \mathbf{0}_{\tilde{z}} \mid M'_1) \mid (\tilde{y}')(x[Q''] \mid M''_1)\{\tilde{v}/\tilde{u}\}) \end{aligned}$$

The last machine, say N , is equal to $(\tilde{z})(N' \mid N''\{\tilde{v}/\tilde{u}\})$ and, by Lemma 26, $\text{calc}(N) \equiv (\tilde{z})(R' \mid Q'\{\tilde{v}/\tilde{u}\})$. This concludes the proof of the basic case.

For the inductive step, the rules of interest are the parallel composition, the scope, and the replication. The first two follow by Proposition 25. If the last rule is

$$\frac{P \mid !P \xrightarrow{\tau} P'}{!P \xrightarrow{\tau} P'}$$

and $\text{calc}(M) = !P$, then we observe that there is $M' \equiv M$ such that $\text{calc}(M') = P \mid !P$. Therefore we reduce to reason on M' , thus concluding the proof.

The result about barbs (Item 3) follows directly from Items 1 and 2.

The bisimulation result (Item 4) follows by Items 1, 2, and 3.

For full abstraction (Item 5), by Lemma 8, we may restrict the definition of \cong_a^{Lc} to parallel contexts. The reverse direction of item 5 is straightforward, because machine contexts are partial parallel compositions while calculus contexts allow arbitrary parallel composition. In the forward direction we must show that the partial is as discriminating as the arbitrary. Suppose there ex-

ists a $\text{AL}\ell\pi$ context R such that $R \mid \text{calc}(M) \dot{\not\approx} R \mid \text{calc}(M')$. Let \tilde{z} be the set of free names in R , and suppose without loss of generality that \tilde{z} does not clash with $\text{lchan}(M)$ or $\text{lchan}(M')$. Now define $E_R = u[[R]_{\tilde{z}}]$ for some fresh u . Clearly E_R, M and E_R, M' are well-formed. We may derive the following relations:

$$\begin{array}{ll}
R \mid \text{calc}(M) \dot{\not\approx} R \mid \text{calc}(M') & \text{by assumption} \\
\cong_a^{Lc} & \cong_a^{Lc} & \text{by Corollary 17} \\
[[R]_{\tilde{z}} \mid \text{calc}(M) \dot{\not\approx} [[R]_{\tilde{z}} \mid \text{calc}(M') & \\
\equiv & \equiv & \text{since there are no clashes} \\
(\text{lchan}(M))([R]_{\tilde{z}} \mid \widehat{M}) \dot{\not\approx} (\text{lchan}(M'))([R]_{\tilde{z}} \mid \widehat{M}') &
\end{array}$$

Hence, by Item 4, $E_R, M \dot{\not\approx} E_R, M'$.

We observe that from Theorems 18 and 27 one cannot derive a correspondence between a $\text{A}\pi$ process and its implementation (i.e. $M \cong_a^{Lc} \text{calc}(M)$, or replacing \cong_a^{Lc} with \cong_a). The reason is because the two formalisms are different. However our results relate the formalisms in a very strong way. Practically, their strength means that a program can be debugged purely at source-level rather than at machine-level. This means that the linear forwarder machine is a natural implementation of $\text{A}\pi$, in contrast to other ones, such as Join. Of course, one might establish similar results with weaker semantics, such as *coupled simulation* [32], thus obtaining weaker relationships between source processes and their implementations.

5 The encoding of distributed choice

In this section we study another standard operator in $\text{A}\pi$, the *input-guarded choice*, written $x(\tilde{u}).P + y(\tilde{v}).Q$. This process can either react on x and so discard the y branch, or *vice versa*. It is well known that choice is problematic to implement in a distributed setting because the two channel managers x and y might be remote. In this case, the choice becomes a problem of global consensus.

Nestmann and Pierce [29] have encoded the input-guarded choice in the choice-free $\text{A}\pi$. Their encoding used booleans \mathbf{t} and \mathbf{f} and the conditional statement *if b then P' else P''* . For example, the process $x(\tilde{u}).P + y(\tilde{v}).Q$ is translated

into:

$$(\ell) \left(\bar{\ell} \mathbf{t} \mid x(u).\ell(w).(\bar{\ell} \mathbf{f} \mid \text{if } w \text{ then } P \text{ else } \bar{x} u) \right. \\ \left. \mid y(v).\ell(w).(\bar{\ell} \mathbf{f} \mid \text{if } w \text{ then } Q \text{ else } \bar{y} v) \right)$$

They demonstrated for such encoding a full abstraction result weaker than Theorem 27; the correspondence uses coupled simulation [30] rather than barbed congruence. (Nestmann and Pierce also gave a similar encoding without using booleans and conditionals.)

We design a different encoding of input-guarded choice using linear forwarders. Our encoding is amenable to a distributed implementation because it reduces the choice between remote channels to a choice between co-located channels. We demonstrate a full-abstraction result similar to Nestmann and Pierce's result [29]. We begin by defining the extension of $\text{AL}\ell\pi$ with input-guarded choice.

Definition 28 (Calculi with choice) *The calculi $\text{AL}\ell\pi$ (Definition 1), $\text{Al}\pi$ (Definition 11), and $\text{A}\pi$ (Definition 6) are augmented with choice as follows ²:*

$$P ::= \dots \mid x(\tilde{u}).P + y(\tilde{v}).Q$$

along with the structural congruence law $P+Q \equiv Q+P$ and the reduction rule

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

The $\text{AL}\ell\pi$ calculus with choice has an extended no-input-capability constraint that, in $x(\tilde{u}).P$ occurring alone or as part of a choice, the P has no free occurrence of $v \in \tilde{u}$ as the subject of an input.

Labelled transition semantics of $\text{AL}\ell\pi$ are as in Definition 12 but with two additional rules:

$$x(\tilde{u}).P + y(\tilde{v}).Q \xrightarrow{x(\tilde{u})} P \qquad x(\tilde{u}).P + y(\tilde{v}).Q \xrightarrow{y(\tilde{v})} Q$$

We extend the function $\llbracket \cdot \rrbracket$ of Definition 9 mapping $\text{A}\pi$ processes to $\text{AL}\ell\pi$ processes to account for choice, and we again demonstrate its correctness.

Definition 29 (Encoding choice) *The encoding $\llbracket \cdot \rrbracket$ is as in Definition 9 but*

² A multi-way choice operator $\sum u_i(\tilde{x}_i).P_i$ is often studied, instead of the two-way choice used here. Multi-way choice is used theoretically because it allows an axiomatisation of bisimulation congruence in the replication-less calculus. We stick to the two-way choice for simplicity.

with the additional rule that $\llbracket x(\tilde{u}).P + y(\tilde{v}).Q \rrbracket_{\tilde{z}} =$

$$\begin{array}{ll} x(\tilde{u}).\llbracket P \rrbracket_{\tilde{u}\tilde{z}} + y(\tilde{v}).\llbracket Q \rrbracket_{\tilde{v}\tilde{z}} & \text{if } x \notin \tilde{z}, y \notin \tilde{z} \\ (y')(y \multimap y' \mid x(\tilde{u}).(y' \multimap y \mid \llbracket P \rrbracket_{\tilde{u}\tilde{z}}) + y'(\tilde{v}).\llbracket Q \rrbracket_{\tilde{v}\tilde{z}}) & \text{if } x \notin \tilde{z}, y \in \tilde{z} \\ (x')(x \multimap x' \mid x'(\tilde{u}).\llbracket P \rrbracket_{\tilde{u}\tilde{z}} + y(\tilde{v}).(x' \multimap x \mid \llbracket Q \rrbracket_{\tilde{v}\tilde{z}})) & \text{if } x \in \tilde{z}, y \notin \tilde{z} \\ (x'y')(x \multimap x' \mid y \multimap y' \mid x'(\tilde{u}).(y' \multimap y \mid \llbracket P \rrbracket_{\tilde{u}\tilde{z}}) + y'(\tilde{v}).(x' \multimap x \mid \llbracket Q \rrbracket_{\tilde{v}\tilde{z}})) & \text{if } x \in \tilde{z}, y \in \tilde{z} \end{array}$$

where x' and y' are fresh names and \tilde{u}, \tilde{v} are disjoint from x, y, \tilde{z} .

This encoding is similar to Definition 9, but with the addition of *undo* agents $y' \multimap y$ and $x' \multimap x$. Their function is to undo the forwarder that is not taken, using $(x')(x \multimap x' \mid x' \multimap x) \approx_a \mathbf{0}$ ³.

The correctness of the encoding of distributed choice cannot be established in the same way as for the choice-free fragment. In particular, a basic property of Definition 9 – the Corollary 17 that $P \approx_a \llbracket P \rrbracket_{\tilde{u}}$ – fails in this case because bisimulation is sensitive to the branching structure of processes (*cf. gradual commitment* in [29]). For example $x(u).P + y(v).Q \not\approx_a \llbracket x(u).P + y(v).Q \rrbracket_{xy}$ because $\llbracket x(u).P + y(v).Q \rrbracket_{xy} \xrightarrow{x(u)} (x'y')(\bar{x}'u \mid y \multimap y' \mid x'(u).(y' \multimap y \mid \llbracket P \rrbracket_{x'yu}) + y'(v).(x' \multimap x \mid \llbracket Q \rrbracket_{x'v}))$ can be mimicked by $x(u).P + y(v).Q \xrightarrow{x(u)} P$, but the processes P and $(x'y')(\bar{x}'u \mid y \multimap y' \mid x'(u).(y' \multimap y \mid \llbracket P \rrbracket_{x'yu}) + y'(v).(x' \multimap x \mid \llbracket Q \rrbracket_{x'v}))$ are not bisimilar because the latter has not yet committed to the choice. The asynchronous bisimulation might match this last process with $\bar{x}u \mid x(u).P + y(v).Q$. But this match also fails because $\bar{x}u \mid x(u).P + y(v).Q \xrightarrow{x(u)} \bar{x}u \mid P\{u'/u\}$, while $(x'y')(\bar{x}'u \mid y \multimap y' \mid x'(u).(y' \multimap y \mid \llbracket P \rrbracket_{x'yu}) + y'(v).(x' \multimap x \mid \llbracket Q \rrbracket_{x'v}))$ cannot mimick this transition.

To solve a similar problem, Nestmann and Pierce used a coarser semantics – the *coupled simulation* [29]. We adapt this notion to our case by using asynchronous simulation as defined in Definition 14.

Definition 30 A coupled simulation is a pair $(\mathcal{R}, \mathcal{S})$ of relations such that \mathcal{R} and \mathcal{S}^{-1} are asynchronous simulations and

- (1) $P \mathcal{R} Q$ implies there is Q' such that $Q \xrightarrow{\tau} Q'$ and $P \mathcal{S} Q'$;
- (2) $P \mathcal{S} Q$ implies there is P' such that $P \xrightarrow{\tau} P'$ and $P' \mathcal{R} Q$.

Coupled similarity, written \rightleftharpoons_a , is the largest coupled simulation.

Coupled similarity retains useful properties that we recall from [29].

Theorem 31 (1) \rightleftharpoons_a is an equivalence relation;

- (2) \rightleftharpoons_a is a congruence in $\text{Al}\pi$ (and therefore in $\text{A}\pi$ and in $\text{ALL}\pi$);
- (3) $\approx_a \subseteq \rightleftharpoons_a$.

³ Remark 10 also applies for this extension.

Next we may establish the analogous statement to Corollary 17, this time for coupled simulation.

Lemma 32 (Operational correspondence for choice) $P \rightleftharpoons_a \llbracket P \rrbracket_{\tilde{u}}$, for all \tilde{u} .

PROOF. To simplify the proof, we introduce some notation. Let $\llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y,y'} \stackrel{\text{def}}{=} x'(\tilde{v}).(y' \multimap y \mid \llbracket P \rrbracket_{\tilde{u}\tilde{w}}) + y'(\tilde{w}).(x' \multimap x \mid \llbracket Q \rrbracket_{\tilde{u}\tilde{w}})$ and let $\llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y} \stackrel{\text{def}}{=} x'(\tilde{v}).\llbracket P \rrbracket_{\tilde{u}\tilde{w}} + y(\tilde{w}).(x' \multimap x \mid \llbracket Q \rrbracket_{\tilde{u}\tilde{w}})$, where the names x', y', \tilde{v} , and \tilde{w} are disjoint from the names in \tilde{u} .

Consider the smallest pair of relations $(\mathcal{R}, \mathcal{S})$ such that both \mathcal{R} and \mathcal{S} contain structural congruence and are closed under the following conditions:

•

$$\begin{array}{l} P \mathcal{R} \llbracket P \rrbracket_{\tilde{u}} \\ P\{\tilde{v}'/\tilde{v}\} \mathcal{R} (x', y')(\bar{x}' \tilde{v}' \mid y \multimap y' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y,y'}) \\ \bar{y} \tilde{w}' \mid P\{\tilde{v}'/\tilde{v}\} \mathcal{R} (x', y')(\bar{x}' \tilde{v}' \mid \bar{y}' \tilde{w}' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y,y'}) \\ P\{\tilde{v}'/\tilde{v}\} \mathcal{R} (x')(\bar{x}' \tilde{v}' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y}) \end{array}$$

where x', y' are fresh. Additionally, if $P \mathcal{R} Q$, $P' \mathcal{R} Q'$, and $x' \notin \text{fn}(Q)$ then

$$P \mathcal{R} (x')(x \multimap x' \mid x' \multimap x \mid Q) \quad \text{and} \quad P \mid P' \mathcal{R} Q \mid Q' \quad \text{and} \quad (x)P \mathcal{R} (x)Q$$

•

$$\begin{array}{l} P \mathcal{S} \llbracket P \rrbracket_{\tilde{u}} \\ \bar{x} \tilde{v}' \mid x(\tilde{v}).P + y(\tilde{w}).Q \mathcal{S} (x', y')(\bar{x}' \tilde{v}' \mid y \multimap y' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y,y'}) \\ \bar{x} \tilde{v}' \mid \bar{y} \tilde{w}' \mid x(\tilde{v}).P + y(\tilde{w}).Q \mathcal{S} (x', y')(\bar{x}' \tilde{v}' \mid \bar{y}' \tilde{w}' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y,y'}) \\ \bar{x} \tilde{v}' \mid x(\tilde{v}).P + y(\tilde{w}).Q \mathcal{S} (x')(\bar{x}' \tilde{v}' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y}) \end{array}$$

where x', y' are fresh. Additionally, if $P \mathcal{S} Q$, $P' \mathcal{S} Q'$, and $x' \notin \text{fn}(Q)$ then

$$P \mathcal{S} (x')(x \multimap x' \mid x' \multimap x \mid Q) \quad \text{and} \quad P \mid P' \mathcal{S} Q \mid Q' \quad \text{and} \quad (x)P \mathcal{S} (x)Q$$

It is worth observing that \mathcal{R} also pairs the processes

$$\begin{array}{l} - Q\{\tilde{v}'/\tilde{v}\} \text{ and } (x', y')(\bar{y}' \tilde{v}' \mid x \multimap x' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y,y'}); \\ - \bar{x} \tilde{w}' \mid Q\{\tilde{v}'/\tilde{v}\} \text{ and } (x', y')(\bar{y}' \tilde{v}' \mid \bar{x}' \tilde{w}' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x,x',y,y'}) \end{array}$$

because it includes structural congruence. Similarly for \mathcal{S} .

The proof that \mathcal{R} and \mathcal{S}^{-1} are asynchronous simulations is entirely standard. We therefore focus on the proof that $(\mathcal{R}, \mathcal{S})$ is a coupled simulation. Since \mathcal{R} and \mathcal{S} are inductively defined, the argument is by induction on the number of rules used for deriving PRQ and PSQ .

• relation \mathcal{R} :

- (1) P is \mathcal{R} -paired to Q since $P \equiv Q$. This follows by $\equiv \subseteq \rightleftarrows_a$ by using Theorem 15.2 and Theorem 31.3.
- (2) P is \mathcal{R} -paired to $\llbracket P \rrbracket_{\tilde{u}}$. Immediate because these two processes are also \mathcal{S} -paired.
- (3) $P\{\tilde{v}'/\tilde{v}\}$ is \mathcal{R} -paired with $(x', y')(\overline{x'} \tilde{v}' \mid y \circ y' \mid \llbracket P, Q \rrbracket_{\tilde{u}\tilde{v}}^{x, x', y, y'})$. There is a transition $(x', y')(\overline{x'} \tilde{v}' \mid y \circ y' \mid \llbracket P, Q \rrbracket_{\tilde{u}\tilde{v}}^{x, x', y, y'}) \xrightarrow{\tau} (x', y')(y \circ y' \mid y' \circ y \mid \llbracket P \rrbracket_{\tilde{u}\tilde{v}}^{\tilde{v}'/\tilde{v}}) \equiv (y')(y \circ y' \mid y' \circ y \mid \llbracket P\{\tilde{v}'/\tilde{v}\} \rrbracket_{\tilde{u}\tilde{v}'})$ since $x' \notin \text{fn}(\llbracket P\{\tilde{v}'/\tilde{v}\} \rrbracket_{\tilde{u}\tilde{v}'})$ and \tilde{v} are disjoint from \tilde{u} . This last process is \mathcal{S} -paired with $P\{\tilde{v}'/\tilde{v}\}$.
- (4) $\overline{y} \tilde{w}' \mid P\{\tilde{v}'/\tilde{v}\}$ is \mathcal{R} -paired with $(x', y')(\overline{x'} \tilde{v}' \mid \overline{y'} \tilde{w}' \mid \llbracket P, Q \rrbracket_u^{x, x', y, y'})$. There is a derivation $(x', y')(\overline{x'} \tilde{v}' \mid \overline{y'} \tilde{w}' \mid \llbracket P, Q \rrbracket_u^{x, x', y, y'}) \xrightarrow{\tau} (x', y')(\overline{y'} \tilde{w}' \mid \llbracket P \rrbracket_{\tilde{u}\tilde{v}}^{\tilde{v}'/\tilde{v}}) \mid y' \circ y \xrightarrow{\tau} (x', y')(\overline{y} \tilde{w}' \mid \llbracket P \rrbracket_{\tilde{u}\tilde{v}}^{\tilde{v}'/\tilde{v}}) \equiv \overline{y} \tilde{w}' \mid \llbracket P\{\tilde{v}'/\tilde{v}\} \rrbracket_{\tilde{u}\tilde{v}'}$. This last process is \mathcal{S} -paired with $\overline{y} \tilde{w}' \mid P\{\tilde{v}'/\tilde{v}\}$ because of parallel closure.
- (5) $P\{\tilde{v}'/\tilde{v}\}$ is \mathcal{R} -paired with $(x')(\overline{x'} \tilde{v}' \mid \llbracket P, Q \rrbracket_u^{x, x', y})$. Similar to 3.
- (6) The inductive cases (addition of a forward and backward linear forwarder, parallel closure, and restriction closure) follow by the inductive hypotheses.

• relation \mathcal{S} :

- (1) P is \mathcal{S} -paired to Q because $P \equiv Q$. As for \mathcal{R} , this case follows since $\equiv \subseteq \rightleftarrows_a$.
- (2) P is \mathcal{S} -paired to $\llbracket P \rrbracket_{\tilde{u}}$. Immediate because these two processes are also \mathcal{R} -paired.
- (3) $\overline{x} \tilde{v}' \mid x(\tilde{v}).P + y(\tilde{w}).Q$ is \mathcal{S} -paired with $(x', y')(\overline{x'} \tilde{v}' \mid y \circ y' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x, x', y, y'})$. The coupling condition follows by $\overline{x} \tilde{v}' \mid x(\tilde{v}).P + y(\tilde{w}).Q \xrightarrow{\tau} P\{\tilde{v}'/\tilde{v}\}$ and the process $P\{\tilde{v}'/\tilde{v}\}$ is \mathcal{R} -paired with $(x', y')(\overline{x'} \tilde{v}' \mid y \circ y' \mid \llbracket P, Q \rrbracket_{\tilde{u}}^{x, x', y, y'})$.
- (4) $\overline{x} \tilde{v}' \mid \overline{y} \tilde{w}' \mid x(\tilde{v}).P + y(\tilde{w}).Q$ is \mathcal{S} -paired with $(x', y')(\overline{x'} \tilde{v}' \mid \overline{y'} \tilde{w}' \mid \llbracket P, Q \rrbracket_u^{x, x', y, y'})$. The coupling condition follows in two ways: either (a) $\overline{x} \tilde{v}' \mid \overline{y} \tilde{w}' \mid x(\tilde{v}).P + y(\tilde{w}).Q \xrightarrow{\tau} \overline{y} \tilde{w}' \mid P\{\tilde{v}'/\tilde{v}\}$ or (b) $\overline{x} \tilde{v}' \mid \overline{y} \tilde{w}' \mid x(\tilde{v}).P + y(\tilde{w}).Q \xrightarrow{\tau} \overline{x} \tilde{v}' \mid Q\{\tilde{w}'/\tilde{w}\}$. The final processes are both \mathcal{R} -paired with $(x', y')(\overline{x'} \tilde{v}' \mid \overline{y'} \tilde{w}' \mid \llbracket P, Q \rrbracket_u^{x, x', y, y'})$.
- (5) $\overline{x} \tilde{v}' \mid x(\tilde{v}).P + y(\tilde{w}).Q$ is \mathcal{S} -paired with $(x')(\overline{x'} \tilde{v}' \mid \llbracket P, Q \rrbracket_u^{x, x', y})$. There is a transition $\overline{x} \tilde{v}' \mid x(\tilde{v}).P + y(\tilde{w}).Q \xrightarrow{\tau} P\{\tilde{v}'/\tilde{v}\}$ and this last process is \mathcal{R} -paired with $(x')(\overline{x'} \tilde{v}' \mid \llbracket P, Q \rrbracket_u^{x, x', y})$.
- (6) The inductive cases are immediate consequences of the inductive hypotheses.

Our full abstraction result for the encoding of choice is an immediate consequence of Lemma 32 and the transitivity of \rightleftharpoons_a .

Theorem 33 For P and Q in $\Lambda\pi$,

- (1) $P \rightleftharpoons_a \llbracket P \rrbracket$.
- (2) $P \rightleftharpoons_a Q$ if and only if $\llbracket P \rrbracket \rightleftharpoons_a \llbracket Q \rrbracket$.

5.1 Choice in the machine

Our machine will only implement local choice, $x(\tilde{u}).P + y(\tilde{v}).Q$ where x and y are co-located. The low-level implementation is easy: a single thread can manage both x and y and so it chooses one branch atomically; or two threads could manage x and y but the choice is protected by a mutex. The machine reduction rule is

$$\begin{array}{c}
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} u \\ u(x).P \oplus v(x).Q \\ \bar{u}y \end{array} & \begin{array}{c} v \\ \end{array} \\
 \hline
 \end{array}
 \quad \rightarrow \quad
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} u \\ P\{y/x\} \end{array} & \begin{array}{c} v \\ \end{array} \\
 \hline
 \end{array}
 \end{array}$$

The following diagrams show the execution of $\llbracket u(x).P + v(x).Q \rrbracket$ on the machine:

$$\begin{array}{c}
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} u \\ u \multimap u' \\ \bar{u}y \end{array} & \begin{array}{c} v \\ v \multimap v' \\ \bar{v}y \end{array} \\
 \hline
 \end{array}
 \quad \Rightarrow \quad
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} (u') \\ u'(x).P \oplus v'(x).Q \end{array} & \begin{array}{c} (v') \\ \end{array} \\
 \hline
 \end{array}
 \quad \text{react, d.out} \\
 \Rightarrow \quad
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} u \\ \end{array} & \begin{array}{c} v \\ \end{array} \\
 \hline
 \end{array}
 \quad \Rightarrow \quad
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} (u') \\ u'(x).P \oplus v'(x).Q \\ \bar{u}'y \end{array} & \begin{array}{c} (v') \\ \bar{v}'y \end{array} \\
 \hline
 \end{array}
 \quad \text{choice} \\
 \rightarrow \quad
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} u \\ \end{array} & \begin{array}{c} v \\ \end{array} \\
 \hline
 \end{array}
 \quad \rightarrow \quad
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} (u') \\ P\{y/x\} \\ v' \multimap v \end{array} & \begin{array}{c} (v') \\ \bar{v}'y \end{array} \\
 \hline
 \end{array}
 \quad \text{choice} \\
 \rightarrow \rightarrow \quad
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} u \\ \end{array} & \begin{array}{c} v \\ \bar{v}y \end{array} \\
 \hline
 \end{array}
 \quad \rightarrow \rightarrow \quad
 \begin{array}{|c|c|}
 \hline
 \begin{array}{c} (u') \\ P\{y/x\} \end{array} & \begin{array}{c} (v') \\ \end{array} \\
 \hline
 \end{array}
 \quad \text{react, d.out}
 \end{array}$$

We now give a formal treatment of choice in the machine.

Definition 34 The machine with choice is as before (Definition 19), but where terms P now come from the $\text{ALL}\pi$ calculus with choice (Definition 28).

The machine with choice has the same locality constraint as before: in $x[P]$ and $(x)[P]$, then every free input $y(\tilde{v}).Q$ in P satisfies $x@y$. But this constraint now also applies to inputs that are parts of summations.

The machine dynamics are as before (Definition 20), plus two additional reaction rules:

$$\begin{aligned} x[y(\tilde{u}).P+z(\tilde{v}).R \mid S], y[\bar{y}\tilde{w} \mid T] &\rightarrow x[S], y[P\{\tilde{w}/\tilde{u}\} \mid T] \quad \text{if } x@y \quad (\text{choice-1}) \\ x[x(\tilde{u}).P+z(\tilde{v}).R \mid \bar{x}\tilde{w} \mid S] &\rightarrow x[P\{\tilde{w}/\tilde{u}\} \mid S] \quad (\text{choice-2}) \end{aligned}$$

Rule (choice-1) reduces a machine even if the subject y of the chosen branch is different from the channel manager x performing the reduction. In this case we constrain the two channels to be co-located. This seems at odds with (*d-in*) that allowed migration of input processes on co-located machines by means of an heating step. We prefer the reduction rule in order to avoid divergency of heating steps. The reduction (choice-2) is similar to the choice rule in the calculus.

We remark that all the machine operations preserve locality. Now we have to prove again the bisimulation and full abstraction results, analogous to those in Theorem 27. The new work is in re-proving the deduction from $\text{calc}(M) \rightarrow P'$ to $M \Rightarrow M'$ and $P' \equiv \text{calc}(M')$. If the calculus reaction involves (choice-1), we rely on the well-formedness of M to satisfy the (choice-1) side condition that $x@y$.

Proposition 35 (1) $M \dot{\approx} M'$ if and only if $\text{calc}(M) \dot{\approx} \text{calc}(M')$.
(2) $M \cong_a M'$ if and only if $\text{calc}(M) \cong_a^{Lc} \text{calc}(M')$.

PROOF. The proofs are substantially the same to those given earlier (from Proposition 25 to Theorem 27). The only difference is in the analogous result to Lemma 26.1 which, as well as the two cases given previously, requires the additional case:

- $M \rightarrow^* (\tilde{y})(M', z[x(\tilde{u}).Q+y(\tilde{v}).Q' \mid R], M''), z@x$ and $P \equiv \text{calc}((\tilde{y})(M', x[Q \mid R], M''))$.

Apart from this difference, the proof carries over unmodified from Theorem 27.

6 Failures

In this section we study two extensions of the linear forwarder machine to incorporate failure information. The first extension considers failures that are due to message losses or to software exceptions that deviate the normal flow of

execution. The timescale of such failures is milliseconds. The second extension considers failures due to crashes of locations. This model, whose timescale is between minutes and hours, is similar to those discussed in [13,33,3]. For each model of failure, we establish correctness results between the fallible versions of the pi calculus, of the linear forwarder calculus and the machine.

6.1 Message loss failures

When failures such as message losses occur there are a number of problems (for instance, consensus) that cannot be solved in an asynchronous framework as $\text{AL}\ell\pi$. In order to avoid these limitations, a calculus should ultimately use timeout-based failure detectors. At the low-level (eg. TCP and UDP programming), these will have the form $x(\tilde{u})^{\mathfrak{t}}.P?Q$ and meaning: “if a message arrives on channel x within \mathfrak{t} milliseconds then continue with P ; but if nothing arrives within a given timeout then execute Q ”. It is beyond the scope of the current contribution to deal with the *compensations* Q and with timed issues (the reader is referred to [5,22] for details). We therefore fix $\mathfrak{t} = \infty$ and $Q = \mathbf{0}$ in what follows, and write $\alpha.P?\mathbf{0}$ as just $\alpha.P$.

Definition 36 *The calculus $\text{FA}\pi$ is as in Definitions 6 and 12, but with the additional reactions:*

$$\bar{x}\tilde{u} \xrightarrow{\tau} \mathbf{0} \quad x(\tilde{u}).P \xrightarrow{\tau} \mathbf{0}$$

*The calculus $\text{FA}\ell\pi$ is as in Definitions 11 and 12, but with these additional reactions*⁴.

$$x\text{-}\circ y \xrightarrow{\tau} \mathbf{0} \quad \bar{x}\tilde{u} \xrightarrow{\tau} \mathbf{0}$$

The fallible local linear forwarder machine is as in Definition 20 but with these additional reaction steps:

$$\begin{array}{llll} x[z\text{-}\circ y \mid P] & \rightarrow & x[P] & \text{if } x\text{-}\circ z \quad (f.fwd) \\ x[\bar{z}\tilde{v} \mid P] & \rightarrow & x[P] & \text{if } x\text{-}\circ z \quad (f.out) \end{array}$$

Note that, in the fallible machine, only forwarders and outputs can be deployed to remote locations. That is why we have added failure of forwarders (*f.fwd*) and failure of outputs (*f.out*), but no rule for failure of input such as $x[z(\tilde{y}).Q \mid P] \xrightarrow{\tau} x[P]$. Similarly, in the fallible version of $\text{AL}\ell\pi$, failures only concern linear forwarders and outputs, which are the unique objects that are intended

⁴ For simplicity, in this section and the next one, we discuss failures in the choice-free fragment of $\text{AL}\pi$.

to move. However, *every* forwarder and output may fail in $\text{FAL}\pi$ since the calculus does not specify any location information.

As regards the relationship between $\text{FA}\pi$ and the fallible machine, consider the $\text{FA}\pi$ term $P = \bar{x}u \mid x(w).Q$. When translated into $\text{FAL}\pi$, the term $\llbracket P \rrbracket_x$ admits two reactions, one of which is

$$\bar{x}u \mid (x')(x \dashv\!\!\!\dashv x' \mid x'(w)).\llbracket Q \rrbracket_{xw} \quad \xrightarrow{\tau} \quad (x')(\bar{x}'u \mid x'(w)).\llbracket Q \rrbracket_{xw}.$$

At this point, $\llbracket P \rrbracket_x$ has committed to consuming the $\bar{x}u$, but it remains in an intermediate state where it is open either to succeed and yield $\llbracket Q \rrbracket_{wx}\{u/w\}$, or to fail and yield a process that is \approx_a -equivalent to $\mathbf{0}$. Meanwhile, the original term P in $\text{FA}\pi$ admits no such intermediate state. This is an instance of a phenomenon similar to the one discussed in Section 5, namely that of an implementation having intermediate states even though they are not observable. As before, the behavioural equivalence we use is coupled simulation.

Theorem 37 *For P and Q in $\text{FA}\pi$,*

- (1) $P \rightleftharpoons_a \llbracket P \rrbracket$.
- (2) $P \rightleftharpoons_a Q$ if and only if $\llbracket P \rrbracket \rightleftharpoons_a \llbracket Q \rrbracket$.

PROOF. The proof is similar to the one of Theorem 33. We therefore focus on the proof of $P \rightleftharpoons_a \llbracket P \rrbracket_{\tilde{u}}$, for all \tilde{u} , by providing a coupled simulation.

Let $(\mathcal{R}, \mathcal{S})$ be the smallest pair such that both \mathcal{R} and \mathcal{S} contain \approx_a and are closed under the following ($x' \notin \text{fn}(P)$).

- For every \tilde{u} :

$$\begin{array}{c} P \quad \mathcal{R} \quad \llbracket P \rrbracket_{\tilde{u}} \\ P\{\tilde{v}'/\tilde{v}\} \quad \mathcal{R} \quad (x')(\bar{x}'\tilde{v}' \mid x'(\tilde{v})).\llbracket P \rrbracket_{\tilde{u}\tilde{v}} \end{array}$$

Additionally, if $P \mathcal{R} Q$ and $P' \mathcal{R} Q'$ then both $P \mid P' \mathcal{R} Q \mid Q'$ and $(x)P \mathcal{R} (x)Q$.

- For every \tilde{u} :

$$\begin{array}{c} P \quad \mathcal{S} \quad \llbracket P \rrbracket_{\tilde{u}} \\ \bar{x}\tilde{v}' \mid x(\tilde{v}).P \quad \mathcal{S} \quad (x')(\bar{x}'\tilde{v}' \mid x'(\tilde{v})).\llbracket P \rrbracket_{\tilde{u}\tilde{v}} \end{array}$$

Additionally, if $P \mathcal{S} Q$ and $P' \mathcal{S} Q'$ then both $P \mid P' \mathcal{S} Q \mid Q'$ and $(x)P \mathcal{S} (x)Q$.

The proof that $(\mathcal{R}, \mathcal{S})$ is a coupled simulation is omitted because similar to that of Lemma 32.

Regarding fallible machines and their associated processes, we do not have an analogous result to Theorem 27, not even for coupled simulation. This is

because M manifests less failures than $\text{calc}(M)$ since the latter misses the co-location information, thus making every $x \multimap y$ and $\bar{x} \tilde{u}$ fail. For example, if $y @ x$, $(x) [\bar{x} \mid x(\cdot).\bar{y}], \dot{y}[\]$ never fails and always emits y , while $\text{calc}((x) [\bar{x} \mid x(\cdot).\bar{y}], \dot{y}[\])$ may fail without emitting anything.

It turns out that the relationship one may establish between a M and $\text{calc}(M)$ is a so-called *barbed simulation*, defined as in Definitions 4 and 22 without the constraint of symmetry. Let $P \leq Q$ if Q simulates P . The proof of the following statement is omitted because it is straightforward.

Proposition 38 $M \leq \text{calc}(M)$.

Remark 39 *It is worth remarking that Proposition 38 may be strengthened because the machine implementation does not introduce deadlocks. Deadlock-avoidance has been demonstrated for the implementation in Pict and in join calculus of pi calculus [39,12].*

6.2 Location crash failures

The second model we analyze considers locations that may crash and be rebooted into some safe initial state. This style of failure is studied in detail in [3,33]. As in these papers, we assume that the rebooted state is the empty location and that channel managers located at crashed locations cannot be used anymore. The usual failure detector has the form $\text{ping}^t x. P?Q$ with the meaning: “if the location x is active (not crashed) then continue with P ; but if nothing arrives within a given timeout t then execute Q ”. For the sake of simplicity we will not explicitly use any ping command. Instead, in $\text{ping}^t x. P?Q$ we let t be infinite, $Q = \mathbf{0}$ and shorten $\text{ping}^t x. P?\mathbf{0}$ into $\text{ping} x. P$. Additionally, the process $\text{ping} x. P$ is encoded by $(z)(\overline{\text{ping}_x} z \mid z.P)$, where ping_x is a channel co-located with x whose behaviour is $!\text{ping}_x(z).\bar{z}$.

The machine in Definition 20 is extended with terms representing *crashed channel managers*:

$$M ::= \dots \mid x\langle\mathbf{0}\rangle$$

The machine $x\langle\mathbf{0}\rangle$ represents a location that cannot perform any action. This means that every other machine co-located with x is crashed as well. The following semantics enforces this constraint. We write $\{x_1 \cdots x_n\}\langle\mathbf{0}\rangle$ for the (crashed) machine $x_1\langle\mathbf{0}\rangle, \dots, x_n\langle\mathbf{0}\rangle$. We use function $\text{fchan}(M)$ to collect the names of the failed channel managers in M and extend the function $\text{chan}(M)$ to include both failed and not-failed channel managers; we have $\text{fchan}(M) \subseteq \text{chan}(M)$.

Definition 40 *The calculus $\text{CFA}\pi$ is as in Definitions 11 and 12, but with*

these additional reactions.

$$x \circ y \xrightarrow{\tau} \mathbf{0} \quad \bar{x} \tilde{u} \xrightarrow{\tau} \mathbf{0} \quad x(\tilde{u}).P \xrightarrow{\tau} \mathbf{0}$$

The machine with crash-failures is as in Definition 20 but with the additional reaction steps:

$$\frac{\text{for every } x, y \in \text{chan}(M). x@y}{M \rightarrow \text{chan}(M)\langle \mathbf{0} \rangle} \quad (\text{crash})$$

$$\frac{M \rightarrow M', \quad \text{for every } x \in \text{fchan}(M'), y \in \text{chan}(N). x@y}{M, N \rightarrow M', N}$$

With respect to $\text{FAL}\pi$, the calculus $\text{CFAL}\pi$ also admits that an input process may fail. This is because the corresponding channel manager may unexpectedly crash: see the rule (*crash*) of the machine with crash-failures. Rule (*crash*) constraints co-located channel managers to crash simultaneously. In fact, the inference rule has a premise verifying that failed channel managers in M' are not co-located with machines in the environment.

We extend $\text{calc}(\cdot)$ as follows:

$$\text{calc}(x\langle \mathbf{0} \rangle) = \mathbf{0}$$

By definition of $\text{CFAL}\pi$, the statements corresponding to Theorem 37 also hold for this calculus. As regards the correspondence between the machine with crash failures and $\text{CFAL}\pi$, we observe that, in contrast to the previous failure model, M manifests *more* failures than $\text{calc}(M)$ because the co-location information of the formers causes failures of a group of processes at once. For instance, take $M = x[x(u).P, x(v).Q]$ then $\text{calc}(M) \xrightarrow{\tau} x(v).Q$ but there is no M' yielded from M such that $\text{calc}(M') = x(v).Q$. However $x(v).Q \xrightarrow{\tau} \mathbf{0}$ and $M \rightarrow x\langle \mathbf{0} \rangle$ with $\text{calc}(x\langle \mathbf{0} \rangle) = \mathbf{0}$. Henceforth, Proposition 38 may be also established for location crash failures.

7 Loadings of program codes

Due to the localisation constraint in $\text{AL}\pi$, loading a calculus program onto a machine is not straightforward. Let us briefly discuss this topic. A given program typically can be loaded in several ways: for instance, $\bar{x}u$ may be loaded onto a machine either at x or at u (or indeed at any existing channel manager). If all the channels are co-located – said otherwise, the machine is a multiprocessor – then every loading of a $\text{AL}\pi$ program is satisfactory and they are all equivalent (by rule (*d.in*)). However, when the machine is distributed,

there are $\text{ALL}\pi$ programs that *cannot* be loaded directly onto it. This is the case when the program has nested free inputs but they are not co-located. Following Remark 10 it is straightforward to use the encoding $\llbracket \cdot \rrbracket$ to avoid such nested free inputs, so making every calculus term loadable. For instance, if $\tilde{z} = \text{fn}(P)$, then $\llbracket P \rrbracket_{\tilde{z}}$ is loadable.

Alternatively, one may achieve the same result by replacing *on-the-fly* the input on a remote name with an input on a new co-located name and a linear forwarder:

$$x[z(\tilde{u}).P \mid Q], z[R] \rightarrow x[Q], (x') [x'(\tilde{u}).P], z[z \multimap x' \mid R] \quad \text{if } x' \text{ fresh, } x @ x', x @ z$$

(d.din)

We actually use the on-the-fly solution in practice in PiDuce [9] since it avoids the need for the encoding. The theory developed in this paper sustains the correctness of this solution.

8 Conclusions

In this paper we have implemented the asynchronous pi calculus, and in particular its problematic feature of input capability on a distributed machine. We have shown that a basic form of input capability, the linear forwarders, is sufficient to express the general form. We have demonstrated this by introducing a calculus with linear forwarders, providing a correct encoding of the asynchronous pi calculus in the linear forwarder calculus, and presenting a distributed abstract machine with a correct implementation of the linear forwarder calculus. Distributed choice and failures have also been studied.

Laneve now has a large project in Bologna which provides a distributed implementation of PiDuce [9], a language that combines pi calculus processes and XML datatypes, whose implementation relies on linear forwarders. This machine, as in the Join Calculus Machine, groups processes at their channels (or locations). In contrast with the Join Calculus Machine that does not support messages carrying channels with input capability, the PiDuce Machine supports input capability using linear forwarders as illustrated in this paper. The ease of implementation of linear forwarders is demonstrated in the PiDuce Machine: the programs managing channels and linear forwarders are implemented by a code of one thousand lines in C#.

It still remains for us to fully assess the practical significance of linear forwarders. One promising application is the so-called orchestration and choreography languages in Web services (WS-BPEL, WSCDL, BizTalk, etc.). These

languages use input capability to allow services to change dynamically the behaviour of other services, or to design new services out of existing ones. For example the service

$$w(x, y, z).x(u).y(v).\bar{z}uv$$

takes three services addresses (URI) x , y , and z , then catches one message to x and to y , and finally packages the two messages and send them to z . The theory of these simple patterns have been recently studied in [21]. It turns out that, combining linear forwarders and join-patterns, it is possible to describe several workflow patterns in Web services. The next aim is to extend these ideas to develop an expressive calculus for Web services, where every primitive is easily amenable to a distributed implementation.

Acknowledgments. We thank the anonymous referees for providing relevant comments on the submitted draft. In particular, one of them pointed out a significant flaw in a previous version of Lemma 32.

References

- [1] Roberto Amadio. An asynchronous model of locality, failure, and process mobility. In D. Garlan and D. Le Métayer, editors, *COORDINATION 1997*, volume 1282 of *Lecture Notes in Computer Science*, pages 374–391. Springer-Verlag, 1997.
- [2] Roberto Amadio, Gérard Boudol, and Cédric Lhoussaine. The receptive distributed pi-calculus (extended abstract). In C. Pandu Rangan, V. Raman, and R. Ramanujam, editors, *FSTTCS 1999*, volume 1738 of *Lecture Notes in Computer Science*, pages 304–315. Springer-Verlag, 1999.
- [3] Roberto Amadio and Sanjiva Prasad. Localities and failures (extended abstract). In *Proceedings of the 14th Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 880 of *Lecture Notes in Computer Science*, pages 205–216. Springer-Verlag, 1994.
- [4] Roberto M. Amadio, Iliaria Castellani, and Davide Sangiorgi. On bisimulations for the asynchronous π -calculus. *Theoretical Computer Science*, 195(2):291–324, 1998. An extended abstract appeared in *Proceedings of CONCUR '96*, LNCS 1119: 147–162.
- [5] Martin Berger. Basic theory of reduction congruence for two timed asynchronous π -calculi. In *CONCUR '04: Proceedings of the 15th International Conference on Concurrency Theory*, volume 3170 of *Lecture Notes in Computer Science*, pages 115–130. Springer-Verlag, 2004.
- [6] Michele Boreale. On the expressiveness of internal mobility in name-passing calculi. *Theoretical Computer Science*, 195(2):205–226, 1998.

- [7] A. Brown, C. Laneve, and L.G. Meredith. PiDuce: A process calculus with native xml datatypes. In *Proceedings of 2nd International Workshop on Web Services and Formal Methods*, volume 3670 of *LNCS*, pages 18–34. Springer-Verlag, 2005.
- [8] Luca Cardelli and Andrew D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, 2000.
- [9] Samuele Carpineti, Cosimo Laneve, and Paolo Milazzo. The bopi machine: a distributed machine for experimenting web services technologies. In *Fifth International Conference on Application of Concurrency to System Design (ACSD 2005)*, pages 202–211. IEEE Computer Society Press, 2005.
- [10] Silvain Conchon and Fabrice Le Fessant. Jocaml: Mobile agents for objective-caml. In *First International Symposium on Agent Systems and Applications Third International Symposium on Mobile Agents*, pages 22–29. IEEE Computer Society Press, October, 1999.
- [11] Rocco De Nicola, Gian Luigi Ferrari, and Pugliese Pugliese. KLAIM: A kernel language for agents interaction and mobility. *IEEE Transactions on Software Engineering*, 24(5):315–330, May 1998.
- [12] Cédric Fournet. *The Join-Calculus: A Calculus for Distributed Mobile Programming*. PhD thesis, École Polytechnique, Paris, France, 1998.
- [13] Cédric Fournet and Georges Gonthier. The reflexive chemical abstract machine and the join-calculus. In *POPL 1996*, pages 372–385. ACM, ACM Press, 1996.
- [14] Philippa Gardner, Cosimo Laneve, and Lucian Wischik. The fusion machine. In M. Kretinsky A. Kucera L. Brim, P. Jancar, editor, *CONCUR 2002*, volume 2421 of *Lecture Notes in Computer Science*, pages 418–433. Springer-Verlag, 2002.
- [15] Philippa Gardner, Cosimo Laneve, and Lucian Wischik. Linear forwarders. In D. Lugiez R. Amadio, editor, *CONCUR 2002*, volume 2761 of *Lecture Notes in Computer Science*, pages 415–430. Springer-Verlag, 2003.
- [16] Philippa Gardner and Sergio Maffei. Modelling dynamic Web data. In Georg Lausen and Dan Suciu, editors, *Proc. of DBPL'03*, volume 2921 of *LNCS*, pages 130–146, September 2003.
- [17] Philippa Gardner and Lucian Wischik. Explicit fusions. In Mogens Nielsen and Branislav Rován, editors, *MFCS 2000*, volume 1893 of *Lecture Notes in Computer Science*, pages 373–382. Springer-Verlag, 2000.
- [18] Alessandro Giacalone, Prateek Mishra, and Sanjiva Prasad. Facile: A symmetric integration of concurrent and functional programming. *International Journal of Parallel Programming*, 18(2):121–160, 1989.
- [19] Kohei Honda and Nobuko Yoshida. On reduction-based process semantics. *Theoretical Computer Science*, 152(2):437–486, 1995.

- [20] Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. Linearity and the pi-calculus. *ACM Transactions on Programming Languages and Systems*, 21(5):914–947, 1999.
- [21] Cosimo Laneve and Luca Padovani. Smooth orchestrators. In *Proceedings of Foundations of Software Science and Computation Structures (FOSSACS 2006)*, volume 3921 of *LNCS*, pages 32–46. Springer-Verlag, 2006.
- [22] Cosimo Laneve and Gianluigi Zavattaro. Foundations of web transactions. In *FOSSACS 2005: Proceedings of Foundations of Software Science and Computation Structure*, volume 3441 of *Lecture Notes in Computer Science*, pages 282–298. Springer Verlag, 2005.
- [23] Sergio Maffei. Dynamic web data: A process-algebraic approach. PhD thesis, Imperial College London, 2006.
- [24] Massimo Merro. On equators in asynchronous name-passing calculi without matching. In *EXPRESS 1999*, volume 27 of *Electronic Notes in Theoretical Computer Science*, pages 57–70. Elsevier Science Publishers, 1999.
- [25] Massimo Merro and Davide Sangiorgi. On asynchrony in name-passing calculi. In Kim G. Larsen, Sven Skyum, and Glynn Winskel, editors, *ICALP 1998*, volume 1443 of *Lecture Notes in Computer Science*, pages 856–867. Springer-Verlag, 1998.
- [26] Robin Milner. Functions as processes. *Journal of Mathematical Structures in Computer Science*, 2(2):119–141, 1992.
- [27] Robin Milner. The polyadic pi calculus: A tutorial. In Friedrich L. Bauer, Wilfried Brauer, and Helmut Schwichtenberg, editors, *Logic and Algebra of Specification, NATO 1991*, volume 94 of *Series F*. Springer-Verlag, 1993.
- [28] Robin Milner and Davide Sangiorgi. Barbed bisimulation. In W. Kuich, editor, *ICALP 1992*, volume 623 of *Lecture Notes in Computer Science*, pages 685–695. Springer-Verlag, 1992.
- [29] Uwe Nestmann and Benjamin C. Pierce. Decoding choice encodings. *Information and Computation*, 163(1):1–59, 2000.
- [30] Joachim Parrow and Peter Sjödin. Multiway synchronization verified with coupled simulation. In W. R. Cleaveland, editor, *CONCUR 1992*, volume 630 of *Lecture Notes in Computer Science*, pages 518–533. Springer-Verlag, 1992.
- [31] A. Phillips, N. Yoshida, and S. Eisenbach. A distributed abstract machine for boxed ambient calculi. In *Proceedings of the European Symposium on Programming (ESOP 2004)*, LNCS. Springer-Verlag, April 2004.
- [32] Benjamin C. Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. *Mathematical Structures in Computer Science*, 6(5):409–454, 1996.
- [33] James Riely and Matthew Hennessy. Distributed processes and location failures. *Theoretical Computer Science*, 226:693–735, 2001.

- [34] Arnaud Sahuguet and Val Tannen. Resource Sharing Through Query Process Migration. University of Pennsylvania Technical Report MS-CIS-01-10, 2001.
- [35] Davide Sangiorgi. π -calculus, internal mobility and agent-passing calculi. *Theoretical Computer Science*, 167(1,2):235–274, 1996.
- [36] Davide Sangiorgi. On the bisimulation proof method. *Mathematical Structures in Computer Science*, 8(5):447–479, 1998.
- [37] Davide Sangiorgi and David Walker. *The Pi-calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
- [38] Peter Sewell, Pawel Wojciechowski, and Benjamin Pierce. Location independence for mobile agents. In H. E. Bal, B. Belkhouche, and L. Cardelli, editors, *ICCL 1998*, volume 1686 of *Lecture Notes in Computer Science*, pages 1–31. Springer-Verlag, 1999.
- [39] David N. Turner. *The Polymorphic Pi-Calculus: Theory and Implementation*. PhD thesis, LFCS, University of Edinburgh, June 1996. CST-126-96 (also published as ECS-LFCS-96-345).
- [40] URL. Microsoft biztalk server. www.microsoft.com/biztalk.