



HAL
open science

Asynchronous pi-calculus at Work: The Call-by-Need Strategy

Davide Sangiorgi

► **To cite this version:**

Davide Sangiorgi. Asynchronous pi-calculus at Work: The Call-by-Need Strategy. The Art of Modelling Computational Systems: A Journey from Logic and Concurrency to Security and Privacy, Nov 2019, Paris, France. pp.33-49, 10.1007/978-3-030-31175-9_3 . hal-02399695

HAL Id: hal-02399695

<https://inria.hal.science/hal-02399695>

Submitted on 9 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

This is a post-peer-review, pre-copyedit version of an article published in Lecture Notes in Computer Science, vol 11760. The final authenticated version is available online at: https://doi.org/10.1007/978-3-030-31175-9_3

This version is subjected to Springer Nature terms for reuse that can be found at: <https://www.springer.com/gp/open-access/authors-rights/aam-terms-v1>

Asynchronous π -calculus at work: the call-by-need strategy

Davide Sangiorgi

Focus Team, University of Bologna and INRIA

Abstract. In a well-known and influential paper [17] Palamidessi has shown that the expressive power of the Asynchronous π -calculus is strictly less than that of the full (synchronous) π -calculus. This gap in expressiveness has a correspondence, however, in sharper semantic properties for the former calculus, notably concerning algebraic laws. This paper substantiates this, taking, as a case study, the encoding of call-by-need λ -calculus into the π -calculus. We actually adopt the *Local* Asynchronous π -calculus, that has even sharper semantic properties. We exploit such properties to prove some instances of validity of β -reduction (meaning that the source and target terms of a β -reduction are mapped onto behaviourally equivalent processes). Nearly all results would fail in the ordinary synchronous π -calculus. We show that however the full β -reduction is not valid. We also consider a refined encoding in which some further instances of β -validity hold. We conclude with a few questions for future work.

1 Introduction

A lot of effort has been devoted to the comparison of the π -calculus with the λ -calculus, beginning with Milner's seminar work on functions as processes [14]. The attention has gone mostly to *call-by-name* and *call-by-value* λ -calculi [19], and the main results concern operational correspondence, validity of β -reduction, characterisation of the equivalence induced on λ -terms by the π -calculus encoding [14, 21, 22, 27, 6]. In particular, the call-by-name encoding, for its simplicity, is often presented as *the* π -calculus representation of functions.

In a call-by-name reduction, the redex contracted is the leftmost one; the reduction occurs regardless of whether the argument of the function is a value (as in call-by-value). As a consequence, if the argument is not a value and will be used several times, its evaluation will be repeated the same number of times. In implementation of programming languages following call-by-name, this repetition of evaluation is avoided: evaluation occurs only once, the first time the term is used, and the value so obtained is recorded for future uses. This implementation technique is referred to as *call-by-need* evaluation (or strategy) [28]. Thus call-by-need uses explicit environments and β -reduction does not require substituting a term for a variable, as in call-by-name (or call-by-value) — just substituting a reference to a term for a variable. In this sense call-by-need is closer to the π -calculus than call-by-name, as substitutions in the π -calculus

only involve names. Again, the modifications that take us from call-by-name to call-by-need can be easily represented in a π -calculus encoding [24].

The π -calculus, having a rich and well-developed theory, as well as a remarkable expressiveness, has been advocated as a foundational model for reasoning about higher-order languages, including equivalence between programs and correctness of compilers and compiler optimisations [25, 26]. Indeed, the π -calculus and related languages have been used, via appropriate encodings, as a target language of compilers, for a number of experimental programming languages, beginning with Pict [18] and Join [7].

The above raises the question about how, and at which extent, the π -calculus and its current theory can be used to prove the correctness of call-by-need as an optimised implementation strategy for call-by-name. The only work on the correctness of the π -calculus representation of call-by-need is by Brock and Ostheimer [5]. The paper considers operational correspondence, between reduction in a call-by-need system and in the encoding π -calculus terms. However there are fundamental semantic issues that remain unexplored. A major one is the validity of β -reduction, namely the property that the processes encoding β -convertible λ -terms are behaviourally indistinguishable. The property holds in call-by-name (and it is at the heart of its theory), as well as in the π -calculus encoding of call-by-name. One would therefore hope to find analogous results for call-by-need. The correctness of the process representation of call-by-need is the topic of the present paper, focusing on the validity of β -reduction.

In a well-known and influential paper [17] Palamidessi has shown that the expressive power of the asynchronous π -calculus is strictly less than that of the full (synchronous) π -calculus. This gap in expressiveness has a correspondence, however, in sharper semantic properties for the former calculus, notably concerning algebraic laws. This paper may be seen as a demonstration of this, since most the proofs are carried out using algebraic laws that are only valid in the asynchronous π -calculus — precisely in the Asynchronous Local π -calculus, $AL\pi$, [12], where only the output capability of names may be exported.

In Section 2 we present $AL\pi$ and some of its laws. In Section 3 we briefly recall the call-by-name and call-by-need λ -calculus. In Section 4 we consider two encodings of call-by-need. We show that limited forms of validity β -reduction hold, and that the general property fails. The questions that follow from this, discussed in Section 5, may contribute to open some interesting directions for future work, which may also shed further light on the theory of the π -calculus and similar name-passing calculi.

2 The Asynchronous Local π -calculus

2.1 Syntax

Small letters a, b, \dots, x, y, \dots range over the infinite set of names, and P, Q, R, \dots over the set of all processes. A tilde represents a tuple. The i -th elements of a tuple \tilde{E} is referred to as \tilde{E}_i . Our notations are extended to tuples componentwise.

The Asynchronous Local π -calculus (AL π) [12] is built from the operators of inaction, input prefix, output, parallel composition, restriction, and replication:

$$P := \mathbf{0} \mid a(\tilde{b}).P \mid \bar{a}(\tilde{b}) \mid P_1 \mid P_2 \mid \nu a P \mid !a(\tilde{b}).P.$$

with the *syntactic constraint* that in processes $a(\tilde{b}).P$ and $!a(\tilde{b}).P$ names \tilde{b} may not occur free in P in input position.

When the tilde is empty, the surrounding brackets $()$ and $\langle \rangle$ will be omitted. $\mathbf{0}$ is the inactive process. An input-prefixed process $a(\tilde{b}).P$, where \tilde{b} has pairwise distinct components, waits for a tuple of names \tilde{c} to be sent along a and then behaves like $P\{\tilde{c}/\tilde{b}\}$, where $\{\tilde{c}/\tilde{b}\}$ is the simultaneous substitution of names \tilde{b} with names \tilde{c} . An output particle $\bar{a}(\tilde{b})$ emits names \tilde{b} at a . Parallel composition is to run two processes in parallel. The restriction $\nu a P$ makes name a local, or private, to P . A replication $!P$ stands for a countable infinite number of copies of P in parallel. We assign parallel composition the lowest precedence among the operators.

2.2 Terminologies and notations

We write $\bar{a}(b).b(\tilde{c}).Q$ as an abbreviation for $\nu b(\bar{a}b \mid b(\tilde{c}).Q)$, and similarly for $\bar{a}(b).!b(\tilde{c}).Q$. The prefix ' $\bar{a}(b)$ ' is called a *bound output*. In prefixes $a(\tilde{b})$ and $\bar{a}(\tilde{b})$, we call a the *subject* and \tilde{b} the *object*. We use α to range over prefixes. We often abbreviate $\alpha.\mathbf{0}$ as α , and $\nu a \nu b P$ as $(\nu a, b) P$. An input prefix $a(\tilde{b}).P$ and a restriction $\nu b P$ are binders for names \tilde{b} and b , respectively, and give rise in the expected way to the definition of *free names* (fn), *bound names* (bn) and *names* (n) of a term or a prefix, and *alpha conversion*. We identify processes or actions that only differ on the choice of the bound names. The symbol $=$ will mean ‘‘syntactic identity modulo alpha conversion’’. Sometimes, we use $\stackrel{\text{def}}{=}$ as abbreviation mechanism, to assign a name to an expression to which we want to refer later. In a statement, a name declared *fresh* is supposed to be different from any other name appearing in the objects of the statement, like processes or substitutions. Substitutions are of the form $\{\tilde{b}/\tilde{c}\}$, and are finite assignments of names to names. A context is a process expression with a hole $[\cdot]$ in it. We use C to range over contexts; then $C[P]$ is the process obtained from C by filling its hole with P .

2.3 Sorting

Following Milner [13], we only admit *well-sorted agents*, that is agents obeying a predefined *sorting* discipline in their manipulation of names. The sorting prevents arity mismatching in communications, like in $\bar{a}(b, c) \mid a(x).Q$. A sorting is an assignment of *sorts* to names, which specifies the arity of each name and, recursively, of the names carried by that name. We do not present the formal system of sorting because it is not essential in the exposition of the topics in the present paper.

We will however allow sorting to identify *linear* names, that is, names that are supposed to be used only once. Linearity will be used a few times to replace input-replicated prefixes with ordinary input prefixes. Again, we omit the details of linearity in type systems (sorting is a form of type system), as they are by now standard [9, 24].

$$\begin{array}{l}
\text{INP: } a(\tilde{b}).P \xrightarrow{a(\tilde{b})} P \qquad \text{OUT: } \bar{a}(\tilde{b}) \rightarrow \mathbf{0} \\
\text{REP: } \frac{P \mid !P \xrightarrow{\mu} P'}{!P \xrightarrow{\mu} P'} \qquad \text{PAR: } \frac{P \xrightarrow{\mu} P'}{P \mid Q \xrightarrow{\mu} P' \mid Q} \text{ if } \text{bn}(\mu) \cap \text{fn}(Q) = \emptyset \\
\text{COM: } \frac{P \xrightarrow{a(\tilde{c})} P' \quad Q \xrightarrow{(\nu \tilde{d}) \bar{a}(\tilde{b})} Q'}{P \mid Q \xrightarrow{\tau} \nu \tilde{d} (P' \{ \tilde{b}/\tilde{c} \} \mid Q')} \text{ if } \tilde{d} \cap \text{fn}(P) = \emptyset \\
\text{RES: } \frac{P \xrightarrow{\mu} P'}{\nu a P \xrightarrow{\mu} \nu a P'} \quad a \notin \text{n}(\mu) \quad \text{OPEN: } \frac{P \xrightarrow{(\nu \tilde{d}) \bar{a}(\tilde{b})} P'}{\nu c P \xrightarrow{(\nu c, \tilde{d}) \bar{a}(\tilde{b})} P'} \quad c \in \tilde{b} - \tilde{d}, \quad a \neq c.
\end{array}$$

Fig. 1. The transition system for $\text{AL}\pi$

2.4 Relations

A process has three possible forms of action. A *silent action* $P \xrightarrow{\tau} P'$ represents an interaction, i.e. an internal activity in P . *Input* and *output actions* are, respectively, of the form $P \xrightarrow{a(\tilde{d})} P'$ and $P \xrightarrow{(\nu \tilde{d}) \bar{a}(\tilde{b})} P'$. In both cases, the action occurs at a — the *subject* of the action. In the output action, \tilde{b} is the tuple of names which are emitted, and $\tilde{d} \subseteq \tilde{b}$ are private names which are carried out from their current scope. We use μ to represent the label of a generic action (not to be confused with α , which represents prefixes). In an input action $a(\tilde{d})$ and in an output action $(\nu \tilde{d}) \bar{a}(\tilde{b})$, names \tilde{d} are *bound*, the remaining ones free. Bound and free names of an action μ , respectively written $\text{bn}(\mu)$ and $\text{fn}(\mu)$, are defined accordingly. The *names* of μ , briefly $\text{n}(\mu)$, are $\text{bn}(\mu) \cup \text{fn}(\mu)$. The transition system of the calculus is presented in Figure 1. We have omitted the symmetric versions of rules PAR and COM. Alpha convertible processes have deemed to have the same transitions. We often abbreviate $P \xrightarrow{\tau} Q$ with $P \longrightarrow Q$. The ‘weak’ arrow \Longrightarrow is the reflexive and transitive closure of \longrightarrow .

We use the symbol \equiv to denote *structural congruence*, a relation used to rearrange the structure of processes [13]. We shall also use it to represent garbage-collection of restrictions and of inert terms.

Definition 1 (Structural congruence). Structural congruence, \equiv , is the smallest congruence relation satisfying the axioms below:

- $P \mid \mathbf{0} \equiv P$, $P \mid Q \equiv Q \mid P$, $P \mid (Q \mid R) \equiv (P \mid Q) \mid R$
- $!a(x).P \equiv a(x).P \mid !a(x).P$;
- $\nu a \mathbf{0} \equiv \mathbf{0}$, $\nu a \nu b P \equiv \nu b \nu a P$, $\nu a (P \mid Q) \equiv P \mid \nu a Q$ if $a \notin \text{fn}(P)$;
- $\nu a (P \mid !a(\tilde{b}).Q) \equiv P$ and $\nu a (P \mid a(\tilde{b}).Q) \equiv P$, if $a \notin \text{fn}(P)$.

(A derivable law is $\nu a P \equiv P$, for a not free in P .) A standard behavioural equivalence for the π -calculus is *barbed congruence*. Barbed congruence can be defined in any calculus possessing: (i) an *interaction relation* (the τ -steps in the π -calculus), modelling the evolution of the system; and (ii) an *observability predicate* \downarrow_a for each name a , which detects the possibility of a process of accepting a communication with the environment at a . More precisely, we write $P \downarrow_a$ if P can make an output action whose subject is a , that is, if there are P' and an output action μ with subject a such that $P \xrightarrow{\mu} P'$. We write $P \Downarrow_a$ if $P \Longrightarrow P'$ and $P' \downarrow_a$. Unlike synchronous π -calculus, in asynchronous calculi it is natural to restrict the observation to output actions [1]. The reason is that in asynchronous calculi the observer has no direct way of knowing when a message emitted is received.

Definition 2 (Barbed congruence). *A symmetric relation \mathcal{S} on π -calculus processes is a barbed bisimulation if $P \mathcal{S} Q$ implies:*

1. If $P \xrightarrow{\tau} P'$ then there exists Q' such that $Q \Longrightarrow Q'$ and $P' \mathcal{S} Q'$.
2. If $P \downarrow_a$ then $Q \downarrow_a$.

Two π -calculus processes P and Q are barbed bisimilar if $P \mathcal{S} Q$ for some barbed bisimulation \mathcal{S} . We say that P and Q are barbed congruent, written $P \approx Q$, if for each π -calculus context C , processes $C[P]$ and $C[Q]$ are barbed bisimilar.

Strong barbed congruence, written \sim , is defined analogously but replacing the weak arrows \Longrightarrow and \downarrow_a with the strong arrows \longrightarrow and \downarrow_a . As expected, we have $\equiv \subseteq \sim \subseteq \approx$; each containment is strict.

2.5 Further algebraic laws

Most of the proofs in the paper are carried out using algebraic reasoning. We report here some important laws. First some simple laws that are valid in the full (synchronous) π -calculus (Lemma 1). Then laws that are specific to $\text{AL}\pi$.

- Lemma 1.**
1. $\nu a (\bar{a}(b).P \mid a(y).Q) \approx (\nu a, b) (P \mid Q\{b/y\})$ and $\nu a (\bar{a}(b).P \mid !a(y).Q) \approx (\nu a, b) (P \mid Q\{b/y\} \mid !a(y).Q)$;
 2. $\nu a (\alpha.Q \mid !a(\tilde{x}).P) \sim \alpha.\nu a (Q \mid !a(\tilde{x}).P)$, if $\text{bn}(\alpha) \cap \text{fn}(a(\tilde{x}).P) = \emptyset$ and $a \notin \text{n}(\alpha)$ (a similar law holds without the replication);
 3. $a(\tilde{x}).(P \mid !a(\tilde{x}).P) \sim !a(\tilde{x}).P$.

Important laws of $\text{AL}\pi$ are the following ones. Their validity hinges on the asynchronous and output-capability properties of $\text{AL}\pi$. For simplicity we present them on monadic prefixes.

Lemma 2. *We have $\bar{a}b \approx \nu c (\bar{a}c \mid !c(x).\bar{b}x)$. Moreover, if b is linear, then the replication can be removed thus: $\bar{a}b \approx \nu c (\bar{a}c \mid c(x).\bar{b}x)$.*

Next, we report some distributivity laws for private replications, i.e., for systems of the form

$$\nu y (P \mid !y(\tilde{q}).Q)$$

in which y may occur free in P and Q only in output position. One should think of Q as a private resource of P , for P is the only process who can access Q ; indeed P can activate as many copies of Q as needed. One such law has already been given as Lemma 1(2). (The laws can be generalised to the full π -calculus, but need stronger assumptions.)

Lemma 3. *Suppose a occurs free in P, R, Q only in output position. Then:*

1. $\nu a (P \mid R \mid !a(\tilde{b}).Q) \sim \nu a (P \mid !a(\tilde{b}).Q) \mid \nu a (R \mid !a(\tilde{b}).Q)$;
2. $\nu a ((!P) \mid !a(\tilde{b}).Q) \sim !\nu a (P \mid !a(\tilde{b}).Q)$;
3. $\nu a (\alpha.P \mid !a(\tilde{b}).Q) \sim \alpha.\nu a (P \mid !a(\tilde{b}).Q)$, if $\text{bn}(\alpha) \cap \text{fn}(a(\tilde{b}).Q) = \emptyset$ and $a \notin \text{n}(\alpha)$;
4. $\nu a ((\nu c P) \mid !a(\tilde{b}).Q) \sim \nu c \nu a (P \mid !a(\tilde{b}).Q)$ if $c \notin \text{fn}(a(\tilde{b}).Q)$.

AL π has also sharper properties concerning labelled characterisation of bisimilarity and associated congruence properties [12, 3].

3 The λ -calculus

We use M, N to range over the set Λ of λ -terms, and x, y, z to range over variables. The set Λ of λ -terms is given by the grammar:

$$M ::= x \mid \lambda x. M \mid MN$$

A *redex* is a term of the form $(\lambda x. M)N$, and then its *contractum* is $M\{N/x\}$. In *call-by-name evaluation* [19], the redex is always at the extreme left of a term. We omit the standard evaluation rules.

Call-by-need [28, 2] optimises call-by-name as follows, so to guarantee that in the contractum $M\{N/x\}$ the evaluation of N is not performed more than once. Roughly, N is placed in an environment, and the evaluation continues on M . When x is needed (i.e., x reaches the leftmost position), then N is evaluated and, if a value (i.e., an abstraction) is obtained, say V , then V replaces x (value V can replace all occurrences of x or, more commonly, only the leftmost occurrence, and then other occurrences of x when they reach the outermost position). Call-by-need is best presented in a graph; or in a system with a **let** construct to represent sharing. We refer to Ariola et al. [2] for details, as they are not essential for understanding the remainder of the paper; see also the references in Section 5.

We sometimes omit λ in nested abstractions, thus for example, $\lambda x_1 x_2. M$ stands for $\lambda x_1. \lambda x_2. M$. We assume the standard concepts of free and bound

variables and substitutions, and identify α -convertible terms. Thus, throughout the paper ‘=’ is syntactic equality modulo α -conversion.

Following the call-by-value terminology, the set of abstractions and variables are the *values*. (Indeed, call-by-need may also be thought of as a modified form of call-by-value, in which the evaluation of the argument of a function $\lambda x. M$ is made only when x is used for the first time, rather than before performing the reduction.)

4 The encoding and its properties

4.1 Background material

Figure 2 presents the call-by-name and call-by-need encodings [16, 24]. The call-by-name one is a variant of the original encoding by Milner [14], with the advantage that it can be written in $AL\pi$ and can be easily modified to follow call-by-need.

We explain the encodings. The important part is the treatment of application. Both in call-by-name and in call-by-need, a function located at q (its ‘location’) is a process that signals to be a function on q , and then receives a pointer x to the argument N together with the location p for the next interaction. Now the evaluation of M continues. The difference between call-by-name and call-by-need arises when the argument N is needed. This is signaled by an output at x that also provides the location for the evaluation of a copy of N . In call-by-name, every output at x triggers the evaluation of a new copy of N . In call-by-need, in contrast, the evaluation is made only the first time. Precisely, in call-by-need N is evaluated at the first request and, when it becomes a value, a pointer to this value is returned (instantiating w , in the table). This pointer is returned to the process that requested N . When further requests for N are made, the pointer is returned immediately. Thus, for instance, in the call-by-name encoding of $(\lambda. xx)(II)$ term II is evaluated twice, whereas in the call-by-need encoding only once. In all encodings, the location names (in the table, the names ranged over by p, q, r) are linear.

Correctness of call-by-name has been studied in depth. In particular, it has been shown that β -reduction is validated by the encoding, that the encoding gives rise to a term model for the λ -calculus, and that the equivalence on λ -terms induced by the encoding corresponds to the best tree-structures of the λ -calculus — which are also at the heart of its denotational semantics — namely Böhm Trees and Lévy-Longo Trees [14, 23]. Correctness of the call-by-need encoding has been studied only by Brock and Ostheimer [5], and only for operational correspondence with respect to Ariola et al.’s system [2]. (The encoding in Figure 2 is actually a minor improvement over that in [5] — avoiding one reduction step during a β -reduction — and maintains the results of operational correspondence in [5] recalled below.) Following Ariola et al.’s system [2] we write $M \Downarrow$ if the call-by-need computation of M terminates, and $M \Uparrow$ if the computation does not terminate.

$$\begin{aligned}
& \text{call-by-name encoding} \\
\mathcal{M}[\lambda x. M]p & \stackrel{\text{def}}{=} \bar{p}(v). !v(x, q). \mathcal{M}[M]q \\
\mathcal{M}[x]p & \stackrel{\text{def}}{=} \bar{x}p \\
\mathcal{M}[MN]p & \stackrel{\text{def}}{=} (\nu q) \left(\mathcal{M}[M]q \mid q(v). \nu x \bar{v}\langle x, p \rangle. !x(r). \mathcal{M}[N]r \right) \\
& \text{call-by-need encoding} \\
\mathcal{N}[\lambda x. M]p & \stackrel{\text{def}}{=} \bar{p}(v). !v(x, q). \mathcal{N}[M]q \\
\mathcal{N}[x]p & \stackrel{\text{def}}{=} \bar{x}p \\
\mathcal{N}[MN]p & \stackrel{\text{def}}{=} \\
(\nu q) \left(\mathcal{N}[M]q \mid q(v). \nu x \bar{v}\langle x, p \rangle. x(r). \nu q' \left(\mathcal{N}[N]q' \mid \right. \right. \\
& \qquad \qquad \qquad \left. \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w) \right) \right)
\end{aligned}$$

Fig. 2. The encoding of call-by-name and call-by-need

Theorem 1 (Brock and Ostheimer, [5]). *We have, for M closed:*

1. $M \Downarrow$ iff $\mathcal{N}[M]p \Downarrow_p$;
2. $M \Uparrow$ iff $\mathcal{N}[M]p \Uparrow$.

The proof in [5] considers an extended version of the call-by-need system in [2], one that yields a closer (nearly one-to-one) correspondence between reductions in the call-by-need system and reductions on the encoding π -calculus processes.

Note that, since M is closed, the only free name of $\mathcal{N}[M]p$ is p ; and since p is used in $\mathcal{N}[M]p$ in output, the first visible action of $\mathcal{N}[M]p$ (if there is one) is an output at p .

However, operational correspondence alone is not fully satisfactory as a criterion for correctness. It does not ensure fundamental semantic properties of the source language terms. In the following sections we focus on the validity of β -reduction.

4.2 β -validity

We consider in this section a few cases of validity of β -reduction; that is, the property that a β -redex $(\lambda x. M)N$ and its contractum $M\{N/x\}$ are barbed congruent when represented as $\text{AL}\pi$ processes.

A form of β -reduction that is straightforward to handle is one in which the argument is never used.

Theorem 2. *If $x \notin \text{fv}(M)$ then $\mathcal{N}[(\lambda x. M)N]p \approx \mathcal{N}[M\{N/x\}]p$.*

A more interesting form deals with β -reduction between closed values.

Theorem 3. $\mathcal{N}[(\lambda x. M)(\lambda y. N)]p \approx \mathcal{N}[M\{\lambda y. N/x\}]p$.

Proof. Using algebraic reasoning, we first derive:

$$\begin{aligned}
& \mathcal{N}[(\lambda x. M)(\lambda y. N)]p \\
&= (\nu q) \left(\mathcal{N}[\lambda x. M]q \mid \right. \\
&\quad \left. q(v). \nu x \bar{v}\langle x, p \rangle. x(r). \nu q' (\mathcal{N}[\lambda y. N]q' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w)) \right) \\
&= (\nu q) \left(\bar{q}(v). !v(x, q'). \mathcal{N}[M]q' \mid \right. \\
&\quad \left. q(v). \nu x \bar{v}\langle x, p \rangle. x(r). \nu q' (\mathcal{N}[\lambda y. N]q' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w)) \right) \\
&\approx (\nu x) \left(\mathcal{N}[M]p \mid \right. \\
&\quad \left. x(r). \nu q' (\mathcal{N}[\lambda y. N]q' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w)) \right) \\
&= (\nu x) \left(\mathcal{N}[M]p \mid \right. \\
&\quad \left. x(r). \nu q' (\bar{q}'(v). !v(y, q''). \mathcal{N}[N]q'' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w)) \right) \\
&\approx (\nu x) \left(\mathcal{N}[M]p \mid \right. \\
&\quad \left. x(r). \nu v (!v(y, q''). \mathcal{N}[N]q'' \mid \right. \\
&\quad \quad \left. (\bar{r}v \mid !x(r'). \bar{r}'v)) \right) \\
&\sim (\nu x, v) \left(\mathcal{N}[M]p \mid \right. \\
&\quad \left. !v(y, q''). \mathcal{N}[N]q'' \mid \right. \\
&\quad \left. x(r). (\bar{r}v \mid !x(r'). \bar{r}'v) \right) \\
&\sim (\nu x, v) \left(\mathcal{N}[M]p \mid \right. \\
&\quad \left. !v(y, q''). \mathcal{N}[N]q'' \mid \right. \\
&\quad \left. !x(r'). \bar{r}'v \right)
\end{aligned}$$

where the two occurrences of \approx represent applications of law (1) of Lemma 1, and the two occurrences of \sim are due to laws (2) and (3) of the same lemma, respectively.

Now we proceed by induction on the structure of M . If M is variable different from x then the two replications at v and x can be garbage-collected and we are done. If $M = x$, then

$$\begin{aligned}
& (\nu x, v) (\mathcal{N}[M]p \mid !v(y, q''). \mathcal{N}[N]q'' \mid !x(r'). \bar{r}'v) = \\
& (\nu x, v) (\bar{x}p \mid !v(y, q''). \mathcal{N}[N]q'' \mid !x(r'). \bar{r}'v) \approx \\
& (\nu x, v) (\bar{p}v \mid !v(y, q''). \mathcal{N}[N]q'' \mid !x(r'). \bar{r}'v) \equiv \\
& (\nu v) (\bar{p}v \mid !v(y, q''). \mathcal{N}[N]q'') = \\
& \quad \mathcal{N}[\lambda y. N]p
\end{aligned}$$

where \approx is obtained from law (1) of Lemma 1, and \equiv from the garbage-collection laws of Definition 1.

replicated), and some rearrangements of the scopes of some restrictions; the use of \sim is due to (2).

$$\begin{aligned}
& \mathcal{N}[(\lambda z. zz)y]p \\
&= (\nu q) \left(\mathcal{N}[\lambda z. zz]q \mid \right. \\
&\quad \left. q(v). \nu x \bar{v}(x, p). x(r). \nu q' (\mathcal{N}[y]q' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w)) \right) \\
&= (\nu q) \left(\bar{q}(v). !v(z, q'). \mathcal{N}[zz]q' \mid \right. \\
&\quad \left. q(v). \nu x \bar{v}(x, p). x(r). \nu q' (\bar{y}q' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w)) \right) \\
&\xrightarrow{\tau} \xrightarrow{\tau} \equiv \nu x \left(\mathcal{N}[xx]p \mid \right. \\
&\quad \left. x(r). \nu q' (\bar{y}q' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w)) \right) \\
&\sim \nu x \left(\right. \\
&\quad \left. \bar{x}(q). q(v). \nu x' \bar{v}(x', p). x'(r). \bar{x}(q'). q'(w). (\bar{r}w \mid !x'(r'). \bar{r}'w) \mid \right. \\
&\quad \left. x(r). \nu q' (\bar{y}q' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w)) \right) \\
&\xrightarrow{\tau} \equiv (\nu x, q, q') \\
&\quad \left(q(v). \nu x' \bar{v}(x', p). x'(r). \bar{x}(q'). q'(w). (\bar{r}w \mid !x'(r'). \bar{r}'w) \mid \right. \\
&\quad \quad \left. \bar{y}q' \mid \right. \\
&\quad \quad \left. q'(w). (\bar{q}w \mid !x(r'). \bar{r}'w) \right) \\
&\xrightarrow{\bar{y}(q')} \xrightarrow{q'(w)} (\nu x, q) \\
&\quad \left(q(v). \nu x' \bar{v}(x', p). x'(r). \bar{x}(q'). q'(w). (\bar{r}w \mid !x'(r'). \bar{r}'w) \mid \right. \\
&\quad \quad \left. (\bar{q}w \mid !x(r'). \bar{r}'w) \right) \\
&\xrightarrow{\tau} (\nu x) \left(\nu x' \bar{w}(x', p). x'(r). \bar{x}(q'). q'(w). (\bar{r}w \mid !x'(r'). \bar{r}'w) \mid \right. \\
&\quad \quad \left. !x(r'). \bar{r}'w \right) \\
&\xrightarrow{\nu x' \bar{w}(x', p)} \xrightarrow{x'(r)} (\nu x) \left(\bar{x}(q'). q'(w). (\bar{r}w \mid !x'(r'). \bar{r}'w) \mid \right. \\
&\quad \quad \left. !x(r'). \bar{r}'w \right) \\
&\xrightarrow{\tau} \xrightarrow{\tau} \equiv \bar{r}w \mid !x'(r'). \bar{r}'w
\end{aligned}$$

Again, the above is the only possible trace for the term. Up-to some renaming, the final derivative is the same as the final derivative of the trace emanating from $\mathcal{N}[yy]p$ examined earlier. However, the two traces are different — the first is longer. As a consequence, the two terms can be distinguished, for instance in the context

$$C \stackrel{\text{def}}{=} \nu y ([\cdot \mid y(q). \bar{q}(v). v(x, p). \nu r (\bar{x}r \mid y(q'). \bar{h})$$

The observable output at h becomes visible only when the context is filled with the first term, $\mathcal{N}[yy]p$ (the one that produces the longer trace).

In contrast, in the call-by-name encoding the validity of the full β -reduction holds [23]. Therefore the above counterexample not only shows that the call-by-need encoding is observably different from the call-by-name one; it also tells us that the properties that the two encoding satisfy are quite different.

4.4 A refined encoding

In this section we experiment with a refinement \mathcal{R} of the encoding so to improve on the problems described in Section 4.3.

Such encoding is shown in Figure 3. In the definition of application for call-by-need, the argument of a function is interrogated only once, the first time the argument is used. Future uses of the argument will directly use the answer so received, without repeating the interrogation — this is indeed the essence of the call-by-need optimisation over call-by-name. To mirror this policy we modify the encoding of an abstraction $\lambda x.M$, so that the body M will interrogate the parameter x only once. As a consequence, in this refined encoding when the head ‘ λx ’ of the function is consumed a local entry is left that takes care of the dialogue with x ; in particular the local entry makes sure that x is consulted only once. The refined encoding, while it exhibits more interactions than the original one, in a distributed setting may be thought of as an optimisation of the latter, as the interactions with the new local entry replace interactions with the possibly remote actual parameter for x . We write $\mathbf{LE}(x, y)$ for a local entry in which the internal resource is x and the external one is y ; it will be convenient to break its definition into two parts, using the auxiliary local entry $\mathbf{LE}'(s, r, x)$; see Figure 3.

The local entry is unnecessary if the internal resource is used only once.

$$\begin{aligned} \mathcal{R}[\lambda x.M]p &\stackrel{\text{def}}{=} \bar{p}(v).!v(x', q).\nu x \left(\mathcal{R}[M]q \right. \\ &\quad \left. | \mathbf{LE}(x, x') \right) \\ \mathcal{R}[x]p &\stackrel{\text{def}}{=} \bar{x}p \\ \mathcal{R}[MN]p &\stackrel{\text{def}}{=} \\ &\left(\nu q \right) \left(\mathcal{R}[M]q \mid q(v).\nu x \bar{v}(x, p).x(r).\nu q' \left(\mathcal{R}[N]q' \mid \right. \right. \\ &\quad \left. \left. \mathbf{LE}'(q', r, x) \right) \right) \end{aligned}$$

where

$$\begin{aligned} \mathbf{LE}(x, y) &\stackrel{\text{def}}{=} x(r).\bar{y}(s).\mathbf{LE}'(s, r, x) \\ \mathbf{LE}'(s, r, x) &\stackrel{\text{def}}{=} s(v).(\bar{r}v \mid !x(r').\bar{r}'v) \end{aligned}$$

Fig. 3. The refined call-by-need encoding

Lemma 4. *Suppose z' appears only once in M . Then*

$$\nu z' (\mathcal{R}\llbracket M \rrbracket p \mid \text{LE}(z', z)) \approx \mathcal{R}\llbracket M\{z/z'\} \rrbracket p$$

Proof. By induction on the structure of M . The most interesting case is when $M = z'$, in which case we exploit the (linear) law of $\text{AL}\pi$ in Lemma 2. When M is an application we exploit the hypothesis (z' occurring only once), and simple algebraic manipulation, so to be able to carry out the induction.

The next lemma shows that local entries compose.

Lemma 5. $\nu x (\text{LE}(z, x) \mid \text{LE}(x, y)) \approx \text{LE}(z, y)$.

Proof. We use laws (2) and (1) of Lemma 1, and the garbage-collection laws of structural congruence.

We revisit the counterexample of Section 4.3, that involves terms yy and $(\lambda z. zz)y$, under the refined encoding \mathcal{R} . All free variables should be protected under a local entry, except for the variables that occur only once (by Lemma 4). We begin by examining $\nu y' (\mathcal{R}\llbracket y'y' \rrbracket q \mid \text{LE}(y', y))$. We have:

$$\begin{aligned} & \nu y' (\mathcal{R}\llbracket y'y' \rrbracket q \mid \text{LE}(y', y)) \\ & \sim \nu y' (\bar{y}'(q).q(v). \nu x \bar{v}\langle x, p \rangle. x(r). \bar{y}'(q').q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w) \\ & \quad \mid y'(r). \bar{y}(r'). r'(v). (\bar{r}v \mid !y'(r). \bar{r}v)) \\ & \approx (\nu y', q) (\bar{q}(v). \nu x \bar{v}\langle x, p \rangle. x(r). \bar{y}'(q').q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w) \\ & \quad \mid \bar{y}(r'). r'(v). (\bar{q}v \mid !y'(q). \bar{q}v)) \\ & \xrightarrow{\bar{y}(r') \rightarrow r'(v)} (\nu y', q) (\bar{q}(v). \nu x \bar{v}\langle x, p \rangle. x(r). \bar{y}'(q').q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w) \\ & \quad \mid \bar{q}v \mid !y'(q). \bar{q}v) \\ & \xrightarrow{\tau} (\nu y') (\nu x \bar{v}\langle x, p \rangle. x(r). \bar{y}'(q').q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w) \\ & \quad \mid !y'(q). \bar{q}v) \\ & \xrightarrow{\nu x \bar{v}\langle x, p \rangle \rightarrow x(r)} (\nu y') (\bar{y}'(q').q'(w). (\bar{r}w \mid !x(r'). \bar{r}'w) \\ & \quad \mid !y'(q). \bar{q}v) \\ & \xrightarrow{\tau \rightarrow \tau} \bar{r}v \mid !x(r'). \bar{r}'v \end{aligned}$$

where the occurrence of \sim is justified by (2) (the two encodings coincide on terms that do not contain abstractions) and definition of $\text{LE}(y', y)$, and the occurrence of \approx comes from law (1) of Lemma 1.

We now consider the second term, $(\lambda z. zz)y$, under the refined encoding. First we note that, if the argument of a function L is a variable, then the refined encoding can be simplified thus:

$$\begin{aligned} \mathcal{R}\llbracket Ly \rrbracket p &= (\nu q) \left(\mathcal{R}\llbracket L \rrbracket q \mid \right. \\ & \quad \left. q(v). \nu x \bar{v}\langle x, p \rangle. x(r). \nu q' (\mathcal{N}\llbracket y \rrbracket q' \mid \right. \\ & \quad \left. \left. \text{LE}'(q', r, x) \right) \right) \\ &= (\nu q) \left(\mathcal{R}\llbracket L \rrbracket q \mid \right. \\ & \quad \left. q(v). \nu x \bar{v}\langle x, p \rangle. \text{LE}(x, y) \right) \end{aligned}$$

Using this property, we have:

$$\begin{aligned}
\mathcal{R}[(\lambda z. zz)y]p &= (\nu q) \left(\mathcal{R}[\lambda z. zz]q \mid \right. \\
&\quad \left. q(v). \nu x \bar{v}\langle x, p \rangle. \mathbf{LE}(x, y) \right) \\
&= (\nu q) (\bar{q}(v). !v(z', q). \nu z (\mathcal{R}[\lambda zz]q \mid \mathbf{LE}(z, z')) \\
&\quad \mid q(v). \nu x \bar{v}\langle x, p \rangle. \mathbf{LE}(x, y)) \\
&\xrightarrow{\tau} \xrightarrow{\tau} \nu z \nu x (\mathcal{R}[\lambda zz]p \mid \mathbf{LE}(z, x) \mid \mathbf{LE}(x, y)) \\
&\approx \nu z (\mathcal{R}[\lambda zz]p \mid \mathbf{LE}(z, y))
\end{aligned}$$

where \approx is obtained by composition of local entries (Lemma 5).

The above reasoning shows that the behaviour of the initial and refined encodings are the same on the term $(\lambda z. zz)y$:

$$\mathcal{R}[(\lambda z. zz)y]p \approx \mathcal{N}[(\lambda z. zz)y]p$$

And it is also the same as that of the refined encoding on the β -contractum yy , when the free variables y are protected under the appropriate local entry, i.e.,

$$\mathcal{R}[(\lambda zz)y]p \approx \nu y' (\mathcal{R}[y'y']p \mid \mathbf{LE}(y', y))$$

The result also holds by closing up the terms:

$$\mathcal{R}[\lambda y. (yy)]p \approx \mathcal{R}[\lambda y. ((\lambda z. zz)y)]p$$

In the refined encoding \mathcal{R} , the local entry (necessary for confining the behaviour of yy) is produced by the encoding of the abstraction.

More generally, proceeding as above one can show that, on the encoding \mathcal{R} , β -reduction is valid when the argument of a function is a variable. Moreover, β -reduction is also valid when the argument is itself a function, reasoning as in Theorem 3. We may therefore conclude that β -reduction is valid when the argument is a value (as it is the case for the encoding of the call-by-value strategy [24]).

We write $\mathbf{LE}(\tilde{x}, \tilde{y})$ for $\mathbf{LE}(x_1, y_1) \mid \dots \mid \mathbf{LE}(x_n, y_n)$ where n is the length of the tuples \tilde{x} and \tilde{y} . We use V to range over values.

Theorem 5. *Suppose $\text{fv}((\lambda z. MV)) \subseteq \tilde{x}$. Then, for any y and fresh \tilde{y} , we have*

$$\nu \tilde{x} (\mathcal{R}[(\lambda z. M)V]p \mid \mathbf{LE}(\tilde{x}, \tilde{y})) \approx \nu \tilde{x} (\mathcal{R}[M\{V/z\}]p \mid \mathbf{LE}(\tilde{x}, \tilde{y}))$$

Similarly to what done earlier, a local entry $\mathbf{LE}(x_i, y_i)$ may be removed when the variable x_i appears at most once.

However, even in the encoding \mathcal{R} , the full β -reduction is unvalid. As a counterexample we can use the terms $(xz)(xz)$ and $(\lambda y. yy)(xz)$. Indeed we have:

$$\nu xz (\mathcal{R}[(xz)(xz)]p \mid \mathbf{LE}(x, x') \mid \mathbf{LE}(z, z')) \not\approx \mathcal{R}[(\lambda y. yy)(xz)]p$$

We omit the calculations, which are rather involved. Intuitively, the difference appears because the full trace of the process computing the application xz is visible twice in the first process, whereas, in the second process, the second time, part of the trace is concealed.

5 Conclusions

Call-by-need was proposed by Wadsworth [28] as an implementation technique. Formalizations of call-by-need on a λ -calculus with a `let` construct or with environments include Ariola et al. [2], Launchbury [11], Purushothaman and Seaman [20], and Yoshida [29]. The uniform encoding in Section 4 is from [16]. A study of the correctness of the call-by-need encoding in Figure 2 is in [5]. Encodings of graph reductions, related to call-by-need, into π -calculus were given in [4, 8] but their correctness was not studied. Niehren [15] used encodings of call-by-name, call-by-value and call-by-need λ -calculi into π -calculus to compare the time complexity of the strategies.

In the paper we have used the theory of the Asynchronous Local π -calculus (AL π) [12] to reason about the encoding of the call-by-need λ -calculus strategy as processes. We have mainly focused on the validity of β -reduction. We have showed that various instances of the property on closed terms hold, though the general property fails. We have also considered a refined encoding in which β -reduction on arbitrary values (though not on arbitrary terms) holds. All this leaves us with some challenging questions, that we leave for future work:

1. In the refined encoding, we use special processes called *local entries* to protect the formal parameter of the function, thus improving the results about β -validity. Is it possible to further protect variables (or terms) so to recover the full β -validity?
2. Is there a different form of behavioural equivalence under which the full β -validity holds, in the initial or in the refined encoding?
3. What is an appropriate process preorder under which call-by-need can indeed be proved to be an optimisation of call-by-name?
4. What is the equivalence on λ -terms induced by the call-by-need encoding? Following the results for call-by-name and call-by-value, one expects to recover some kind of tree structure (Böhm Trees and Lévy-Longo Tree for call-by-name, Lassen’s trees [10] for call-by-value). We are not aware of similar tree structures for call-by-need. Hence investigating this question may also shed light on what should the appropriate forms of trees for call-by-need be.

In questions (2) and (3), ‘easy’ answers may be obtained by confining the testing contexts to be encodings of λ -calculus contexts. The challenge is to find more general and useful answers, with applications outside the realm of pure λ -calculi. One may consider forms of behavioural types.

In questions (1) and (2), perhaps requiring the validity of the full β -reduction, in the same way as for call-by-name, is too demanding. Indeed in this way probably the tree structure referred to in question (4) is likely to be the same as that for call-by-name. One may find it acceptable to limit β -validity to reductions between closed terms.

Acknowledgments Thanks to the reviewers for their careful reading of the paper and their suggestions. Research partly supported by the H2020-MSCA-RISE project ID 778233 “Behavioural Application Program Interfaces (BEHAPI)”.

References

1. Amadio, R., Castellani, I., Sangiorgi, D.: On bisimulations for the asynchronous π -calculus. *Theoretical Computer Science* 195, 291–324 (1998)
2. Ariola, Z., Felleisen, M., Maraist, J., Odersky, M., Wadler, P.: A call-by-need λ -calculus. In: *Proc. 22th POPL*. ACM Press (1995)
3. Boreale, M., Sangiorgi, D.: Some congruence properties for π -calculus bisimilarities. *Theoretical Computer Science* 198, 159–176 (1998)
4. Boudol, G.: Some chemical abstract machines. In: *Proc. REX Summer School/Symposium 1993*. Lecture Notes in Computer Science, vol. 803. Springer Verlag (1994)
5. Brock, S., Ostheimer, G.: Process semantics of graph reduction. In: Lee, I., Smolka, S. (eds.) *Proc. CONCUR '95*. Lecture Notes in Computer Science, vol. 962, pp. 471–485. Springer Verlag, Philadelphia, Pennsylvania (1995)
6. Durier, A., Hirschhoff, D., Sangiorgi, D.: Eager functions as processes. In: *33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018*. IEEE Computer Society (2018)
7. Fournet, C., Gonthier, G.: The join calculus: A language for distributed mobile programming. In: Barthe, G., Dybjer, P., Pinto, L., Saraiva, J. (eds.) *Applied Semantics, International Summer School, APPSEM 2000*. Lecture Notes in Computer Science, vol. 2395, pp. 268–332. Springer (2002)
8. Jeffrey, A.: A chemical abstract machine for graph reduction. In: *Proc. Ninth International Conference on the Mathematical Foundations of Programming Semantics (MFPS'93)*. Lecture Notes in Computer Science, vol. 802. Springer Verlag (1993)
9. Kobayashi, N., Pierce, B., Turner, D.: Linearity and the pi-calculus. *TOPLAS* 21(5), 914–947 (1999), preliminary summary appeared in *Proceedings of POPL'96*
10. Lassen, S.B.: Eager normal form bisimulation. In: *20th IEEE Symposium on Logic in Computer Science (LICS 2005)*, 26–29 June 2005, Chicago, IL, USA, *Proceedings*. pp. 345–354 (2005)
11. Launchbury, J.: A natural semantics for lazy evaluation. In: *Proc. 20th POPL*. ACM Press (1993)
12. Merro, M., Sangiorgi, D.: On asynchrony in name-passing calculi. *Journal of Mathematical Structures in Computer Science* 14(5), 715–767 (2004)
13. Milner, R.: The polyadic π -calculus: a tutorial. Tech. Rep. ECS-LFCS-91-180, LFCS, Dept. of Comp. Sci., Edinburgh Univ. (Oct 1991), Also in *Logic and Algebra of Specification*, ed. F.L. Bauer, W. Brauer and H. Schwichtenberg, Springer Verlag, 1993.
14. Milner, R.: Functions as processes. *Journal of Mathematical Structures in Computer Science* 2(2), 119–141 (1992)
15. Niehren, J.: Functional computation as concurrent computation. In: *Proc. 23th POPL*. ACM Press (1996)
16. Ostheimer, G., Davie, A.: π -calculus characterisations of some practical λ -calculus reductions strategies. Tech. Rep. CS/93/14, St. Andrews (1993)
17. Palamidessi, C.: Comparing the expressive power of the synchronous and asynchronous pi-calculi. *Journal of Mathematical Structures in Computer Science* 13(5), 685–719 (2003)

18. Pierce, B.C., Turner, D.N.: Pict: A programming language based on the pi-calculus. In: Plotkin, G., Stirling, C., Tofte, M. (eds.) *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press (2000)
19. Plotkin, G.: Call by name, call by value and the λ -calculus. *Theoretical Computer Science* 1, 125–159 (1975)
20. Purushothaman, S., Seaman, J.: An adequate operational semantics for sharing in lazy evaluation. In: Krieg-Brückner, B. (ed.) *ESOP'92. Lecture Notes in Computer Science*, vol. 582, pp. 435–450. Springer Verlag (1992)
21. Sangiorgi, D.: An investigation into functions as processes. In: *Proc. Ninth International Conference on the Mathematical Foundations of Programming Semantics (MFPS'93)*. *Lecture Notes in Computer Science*, vol. 802, pp. 143–159. Springer Verlag (1993)
22. Sangiorgi, D.: From λ to π , or: Rediscovering continuations. *Journal of Mathematical Structures in Computer Science* 9(4) (1999), special Issue on "Lambda-Calculus and Logic" in Honour of Roger Hindley
23. Sangiorgi, D.: Lazy functions and mobile processes. In: Plotkin, G., Stirling, C., Tofte, M. (eds.) *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press (2000)
24. Sangiorgi, D., Walker, D.: *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press (2001)
25. Sangiorgi, D.: Typed pi-calculus at work: A correctness proof of Jones's parallelisation transformation on concurrent objects. *TAPOS* 5(1), 25–33 (1999)
26. Turner, N.: *The polymorphic pi-calculus: Theory and Implementation*. Ph.D. thesis, Department of Computer Science, University of Edinburgh (1996)
27. Vasconcelos, V.T.: Lambda and pi calculi, CAM and SECD machines. *J. Funct. Program.* 15(1), 101–127 (2005)
28. Wadsworth, C.P.: *Semantics and pragmatics of the lambda calculus*. Ph.D. thesis, University of Oxford (1971)
29. Yoshida, N.: Optimal reduction in weak lambda-calculus with shared environments. In: *Proc. of FPCA'93, Functional Programming and Computer Architecture*. pp. 243–252 (1993)