

Pict Correctness Revisited ^{*} ^{**}

Philippe Bidinger¹ and Adriana Compagnoni²

¹ VERIMAG, Grenoble, France

² Stevens Institute of Technology, Hoboken, NJ, USA

Abstract. The Pict programming language is an implementation of the π -calculus in which executions of π -calculus terms are specified via an abstract machine. An important property of any concurrent programming language implementation is the fair execution of threads. After defining fairness for the π -calculus, we show that Pict abstract machine executions implement fair π -calculus executions. We also give new proofs of soundness and liveness for the Pict abstract machine.

1 Introduction

The π -calculus [14, 17] is a minimal language designed to capture and model key concepts of communicating concurrent systems in a formal setting. It emphasizes channel-based communication, dynamic channel creation and the ability to communicate channels as data. Pict [19, 16] is a high-level programming language purely based on π -calculus primitives, as well as to explore the applicability of theoretical work on type systems. Pict's runtime environment is based on a formal abstract machine specification, but little emphasis has been placed on its correctness.

The correctness of a programming language runtime is critical since, in order to be able to reason about programs, we need the guarantee that programs are executed according to their semantics. Correctness results of implementations usually relate executions of terms in a high-level language to its implementation in a low-level language. The low-level language can be an existing process calculus, or like in Pict, an abstract machine specification.

In recent years, many process calculi based on the π -calculus have been introduced to study the dynamics of existing or new paradigms of computation, such as distributed computing, global computing, or component-based programming. Much work has been done on the distributed implementation of these calculi [5, 7, 22, 20, 10, 8, 11, 12, 1]. On the other hand, since the definition of Pict, there has been no new insight for the local implementation of these calculi. Therefore, Pict is still a reference implementation of the π -calculus, and we think that proving its correctness is a first step toward more general proofs of correctness of implementations of these calculi.

The π -calculus is a concurrent language where concurrency is modeled using a non-deterministic reduction relation. The Pict Abstract Machine (PAM)

^{*} This work was funded in part by the US Army under contract W15QKN-05-D-0011.

^{**} A short version of this paper appeared in FMOODS 2007.

implements a particular scheduling strategy that corresponds to a subset of the possible executions in the π -calculus. It is therefore impossible to state an exact correspondence between π -calculus executions and PAM executions. Instead, we will prove the correctness of the abstract machine with three properties:

- A soundness property that states that PAM executions correspond to valid π -calculus executions.
- A liveness property that ensures that the abstract machine is not stuck when its state corresponds to a π -calculus term that can reduce.
- A fairness property that characterizes PAM executions among possible π -calculus executions.

These properties are fairly standard but have not been proven for Pict yet (see section 5 for details). The main contribution of this paper lies in the statement and proof of a fairness property for Pict. To our knowledge, no implementation of a process calculus has been proven fair so far, although fairness is conjectured in [21, 19]. Moreover, the technique we propose is general enough to be adapted to similar settings.

Informally, we say that an execution is *weakly fair* if a prefixed process able to communicate *continuously* will eventually do so. An execution is *strongly fair* if a prefixed process able to communicate *infinitely often* will eventually do so. Consider for instance the π -calculus term

$$x!(a) \mid *x?(z).x!(z) \mid y!(0) \mid y?(z).\mathbf{0}$$

where $*x?(z).P$ represents replicated input. There are valid infinite executions in which the communication on y never takes place even though at any time this communication is possible. Similarly, in the term

$$x!(a) \mid *x?(z).y!(z) \mid *y?(z).x!(z) \mid *x?(z).x!(z)$$

there are infinite executions in which the communication on the last process never takes place, even though such a communication might happen infinitely often (but not continuously). The intuitive expectation of a programmer is that all processes running in parallel will be interleaved *fairly* and so such executions are considered unsatisfactory.

Stating a fairness property for the π -calculus is not immediate. The definition of the π -calculus makes it difficult to identify subprocesses within a process, and in particular, it is difficult to state properties about fair executions of these processes. When considering π -calculus processes, mainly two kinds of confusion can arise.

- Processes are identified up to renaming of bound names and lead to possible confusion of channels. For instance, we have

$$\nu x.\nu y.x!() \mid y!() \mid R \rightarrow \nu x.\nu y.x!() \mid R'$$

Because of possible renamings of x and y , we do not know which channel reacted.

- Confusion of structurally equivalent processes.

$$\begin{aligned}
P &= x!() \mid x?().x!() \mid *x?().x!() \\
P' &= x!() \mid x?().x!() \mid x?().x!() \mid *x?().x!()
\end{aligned}$$

We have $P \equiv P'$ and $P' \rightarrow P$, and we do not know which receivers react with $x!()$.

A possible solution is to define an auxiliary calculus in which prefixes are annotated with *labels* in such a way that labels uniquely denote prefixes and that this property is invariant throughout reduction [4, 3]. A *live action* of a term is then defined as a couple of labels corresponding to prefixed processes that can react. An infinite labeled execution is strongly fair if there are no labels appearing in an infinity of live actions.

This paper is organized as follows. In section 2, we define the π -calculus and its labeled variant that allows us to define strong fairness. In section 3, we define the Pict abstract machine, and a labeled version. We discuss the correctness of the abstract machine in section 4. In section 5, we discuss related work. We conclude with future research directions in section 6.

2 Fairness in the π -calculus

2.1 The π -calculus

We suppose given a set of names \mathbf{N} ranged over by x, y, \dots . We define the set of π -calculus processes \mathbf{P} as follows:

$$\begin{aligned}
P, Q, \dots &::= \mathbf{0} \mid \pi.P \mid \nu x.P \mid P \mid P \\
\pi &::= x!(y) \mid x?(y) \mid *x?(y)
\end{aligned}$$

The π -calculus evaluation contexts are given by:

$$\mathbf{E} ::= \cdot \mid \nu x.\mathbf{E} \mid P \mid \mathbf{E} \mid \mathbf{E} \mid P$$

The operational semantics is defined as the smallest relation such that rules in Figure 2 hold. It makes use of a structural equivalence relation defined as the smallest equivalence relation satisfying the rules in Figure 1. As usual, $\mathbf{fn}(P)$ denotes the set of free names of process P , and $=_\alpha$ equates two processes that differ only by their bound names. We write $\mathbf{E}[P]$ for the context \mathbf{E} in which the hole \cdot has been substituted with P .

Without loss of generality, we restrict the usual replication operator to input processes. Rule R.REP models communication with a replicated input process.

2.2 A Labeled π -calculus

Informally, a fair execution of a process is an execution in which no subprocess is ready to participate in a communication infinitely often. To formalize this

$$\begin{array}{c}
(P \mid Q) \mid R \equiv P \mid (Q \mid R) \text{ S.PAR.ASSOC} \qquad P \mid Q \equiv Q \mid P \text{ S.PAR.COM} \\
P \mid \mathbf{0} \equiv P \text{ S.PAR.NIL} \qquad \nu x.\mathbf{0} \equiv \mathbf{0} \text{ S.NU.NIL} \qquad \nu x.\nu y.P \equiv \nu y.\nu x.P \text{ S.NU.COM} \\
\frac{x \notin \mathbf{fn}(Q)}{(\nu x.P) \mid Q \equiv \nu x.P \mid Q} \text{ S.NU.PAR} \qquad \frac{P =_{\alpha} Q}{P \equiv Q} \text{ S.}\alpha \qquad \frac{P \equiv Q}{\mathbf{E}[P] \equiv \mathbf{E}[Q]} \text{ S.CTX}
\end{array}$$

Fig. 1. Structural Equivalence

$$\begin{array}{c}
\frac{}{x!(y).P \mid x?(z).Q \rightarrow P \mid Q\{y/z\}} \text{ R.RED} \\
\frac{}{x!(y).P \mid *x?(z).Q \rightarrow P \mid Q\{y/z\} \mid *x?(z).Q} \text{ R.REP} \qquad \frac{P \rightarrow Q}{\mathbf{E}[P] \rightarrow \mathbf{E}[Q]} \text{ R.CTX} \\
\frac{P \equiv P' \quad P' \rightarrow Q' \quad Q' \equiv Q}{P \rightarrow Q} \text{ R.STR}
\end{array}$$

Fig. 2. Reduction Relation

statement, we need to identify in a process the subprocesses that can participate in a communication, and keep track of their identities throughout reductions.

To do so, we follow [4, 3] and define a labeled version of the π -calculus in which prefixes are annotated with labels. A label has to identify a prefix uniquely in an entire execution of a process. In other words, not only do prefixes have distinct labels in a process, but when new prefixes are created, their labels are new with respect to all the labels in the past execution of the term. We then characterize the labels belonging to prefixes that can participate in a communication. Finally, we give the definition of fairness.

We denote by \mathbf{L} a set of labels such that $\mathbf{L} \cap \mathbf{N} = \emptyset$. We use $\mathcal{P}^f(\mathbf{L})$ to denote the finite subsets of \mathbf{L} . A labeled process is a pair made of a π -calculus process in which prefixes are labeled, and a finite set of labels. The set of labeled processes \mathbf{LP} is generated by the grammar given below.

$$\begin{array}{l}
C, D, \dots ::= \bar{P}, \mathcal{L} \\
\bar{P}, \bar{Q}, \dots ::= \mathbf{0} \mid \pi_l.\bar{P} \mid \nu x.\bar{P} \mid \bar{P} \mid \bar{P} \\
l \in \mathbf{L} \\
\mathcal{L} \in \mathcal{P}^f(\mathbf{L})
\end{array}$$

We also extend contexts with labels and we denote labeled contexts $\bar{\mathbf{E}}$.

We need several auxiliary functions. The function lab returns the set of all labels of a process or a context. The function unl erases all labeling information from a labeled process.

In order to ensure that labels occur uniquely in a process, we define a well-formation predicate wf as the smallest relation on LP such that rules in Figure 3 hold. We write $A \uplus B$ for $A \cup B$ when $A \cap B = \emptyset$. A labeled process C is said to be well-formed if we have $\text{wf}(C)$. We denote by WFP the set of well-formed labeled processes.

$$\begin{array}{c}
\frac{}{\text{wf}(\mathbf{0}, \mathcal{L})} \text{WF.NIL} \qquad \frac{\text{lab}(\overline{P}) \uplus \text{lab}(\overline{P}') \subseteq \mathcal{L}}{\text{wf}(\overline{P} \mid \overline{P}', \mathcal{L})} \text{WF.PAR} \qquad \frac{\text{wf}(\overline{P}, \mathcal{L})}{\text{wf}(\nu x.\overline{P}, \mathcal{L})} \text{WF.NEW} \\
\frac{\text{wf}(\overline{P}, \mathcal{L})}{\text{wf}(\pi_l.\overline{P}, \mathcal{L} \uplus \{l\})} \text{WF.PREFIX}
\end{array}$$

Fig. 3. Well-Formed Labeled Process

The operational semantics is defined in the same way as for the π -calculus via a structural equivalence relation \equiv and a reduction relation \rightarrow , both binary relations over LP . The structural equivalence is defined, as before, as the smallest equivalence relation that verifies rules in Figure 1 (where prefixes are labeled and equivalent processes have the same set of labels). The reduction relation is the smallest relation that verifies the rules in Figure 4. The main difference with the unlabeled reduction relation appears in the rule LR.REP for replicated input in which fresh labels are generated.

$$\begin{array}{c}
\frac{}{x!(y)_l.\overline{P} \mid x?(z)_{l'}.\overline{Q}, \mathcal{L} \rightarrow \overline{P} \mid \overline{Q}\{y/z\}, \mathcal{L}} \text{LR.RED} \\
\frac{\alpha \text{ injective and } \mathcal{L}' = \mathcal{L} \uplus \alpha(\mathcal{L})}{x!(y)_l.\overline{P} \mid *x?(z)_{l'}.\overline{Q}, \mathcal{L} \rightarrow \overline{P} \mid \overline{Q}\{y/z\} \mid \alpha(*x?(z)_{l'}.\overline{Q}), \mathcal{L}'} \text{LR.REP} \\
\frac{\overline{P}, \mathcal{L} \rightarrow \overline{P}', \mathcal{L}' \quad \text{lab}(\overline{\mathbf{E}}) \subseteq \mathcal{L}}{\overline{\mathbf{E}}[\overline{P}], \mathcal{L} \rightarrow \overline{\mathbf{E}}[\overline{P}'], \mathcal{L}'} \text{LR.CTX} \qquad \frac{D \equiv D' \quad D' \rightarrow C' \quad C' \equiv C}{D \rightarrow C} \text{LR.STR}
\end{array}$$

Fig. 4. Labeled Reduction Relation

Labeling is stable under reduction and structural equivalence. Hence, in the following, we consider only well-formed processes.

Lemma 1 (Stability of Labeling).

- (i) If $C \equiv D$ and $C \in \text{WFP}$, then $D \in \text{WFP}$.
- (ii) If $C \rightarrow D$ and $C \in \text{WFP}$, then $D \in \text{WFP}$.

The following lemma shows that the labeling system has been designed so that no label can occur more than once in a labeled term, and once a label disappears, it does not reappear in the system.

Lemma 2 (Uniqueness of Labeling).

- (i) If $C \in \text{WFP}$ then no label l occurs more than once in C .
- (ii) If $C \in \text{WFP}$, $C \rightarrow^* C' \rightarrow^* C''$ and $l \in \text{lab}(C) \cap \text{lab}(C'')$, then $l \in \text{lab}(C')$.

The labeled π -calculus is a conservative extension of the π -calculus. A labeled process has exactly the same reductions as the corresponding unlabeled process. Moreover, we can label any process into a well-formed labeled process

Proposition 1 (Operational Correspondence). *Let $P \in \mathcal{P}$ and $C \in \text{WFP}$ such that $P = \text{unl}(C)$. We have*

- (i) $P \rightarrow P'$ implies $\exists C' \in \text{LP}$ such that $C \rightarrow C'$ and $\text{unl}(C') = P'$.
- (ii) $C \rightarrow C'$ implies $P \rightarrow \text{unl}(C')$.

Proposition 2 (Existence of a Labeling). *For all $P \in \mathcal{P}$ there exists $C \in \text{WFP}$ such that $\text{unl}(C) = P$.*

We now define the live actions of a labeled process. A live action is a pair of labels corresponding to prefixed processes that can immediately react.

Definition 1 (Live Actions). *The set of live actions of a labeled process $C = \overline{P}, \mathcal{L}$ is defined as*

$$LA(C) = \{ \{l, l'\} / C \equiv \nu \tilde{x}. y!(v)_l. \overline{P}_0 \mid y?(z)_{l'}. \overline{P}_1 \mid \overline{P}_2, \mathcal{L} \\ \text{or } C \equiv \nu \tilde{x}. y!(v)_l. \overline{P}_0 \mid *y?(z)_{l'}. \overline{P}_1 \mid \overline{P}_2, \mathcal{L} \}$$

We also define the set of labels belonging to a live action as

$$L(C) = \{l \in x / x \in LA(C)\}$$

The following lemma states a correspondence between live actions and reductions.

Lemma 3. $C \rightarrow C'$ for some C' if and only if $LA(C) \neq \emptyset$.

Definition 2 (Execution). *For an arbitrary relation \rightarrow , an execution is a sequence of terms T_0, T_1, \dots , possibly infinite, such that $T_0 \rightarrow \dots \rightarrow T_n \rightarrow \dots$*

We can now define a strong fairness property for the labeled calculus. An execution is fair if a prefix cannot potentially participate in a reduction infinitely often. According to this definition, we only need to consider infinite executions.

Definition 3 (Strong Fairness in the Labeled π -calculus). *An infinite execution $C_0 \rightarrow \dots \rightarrow C_n \rightarrow \dots$ is fair if for any strictly increasing sequence of integers $(u_n)_{n \in \mathbb{N}}$, we have $\bigcap_{n \in \mathbb{N}} L(C_{u_n}) = \emptyset$.*

An execution in the π -calculus is fair if any corresponding labeled execution is fair.

Definition 4 (Strong Fairness in the π -calculus). *An infinite execution $P_0 \rightarrow \dots \rightarrow P_n \rightarrow \dots$ is fair if any execution $C_0 \rightarrow \dots \rightarrow C_n \rightarrow \dots$ such that $\forall i \in \mathbb{N}. \text{unl}(C_i) = P_i$ is fair.*

In fact, as a consequence of the next lemma, we can use a weaker definition. An execution is fair if there exists a corresponding fair labeled execution.

Lemma 4. *Let $P_0 \rightarrow \dots \rightarrow P_n \rightarrow \dots$ be an infinite execution. The two following propositions are equivalent:*

- (i) *There exists a fair labeled execution execution $C_0 \rightarrow \dots \rightarrow C_n \rightarrow \dots$ such that $\text{unl}(C_i) = P_i$.*
- (ii) *Any labeled execution execution $C_0 \rightarrow \dots \rightarrow C_n \rightarrow \dots$ such that $\text{unl}(C_i) = P_i$ is fair.*

3 Abstract Machine

3.1 Syntax and Operational Semantics

The syntax of the Pict abstract machine is given in Figure 5 and follows closely [19]³. A machine state, or PAM term, consists of a queue of π -calculus processes \mathcal{P} (the runqueue), a heap \mathcal{H} and a set of names \mathcal{N} . A heap is a function that maps channel names to process queues. We denote by \mathbf{M} the set of machine states. We often omit the set of names \mathcal{N} in PAM terms when it is not important, in particular in reduction rules where it remains unchanged. We also write $\mathcal{P} :: \mathcal{Q}$ for the appending of \mathcal{P} and \mathcal{Q} . The operational semantics is defined via two reduction

$$\begin{array}{ll}
 \mathcal{M} ::= \langle \mathcal{P}, \mathcal{H}, \mathcal{N} \rangle & \text{State} \\
 \mathcal{P}, \mathcal{Q}, \dots ::= [] \mid P :: \mathcal{P} & \text{Processes Queue} \\
 \mathcal{H} ::= \{x \rightarrow \mathcal{P}_x\}_{x \in \mathbb{N}} & \text{Heap}
 \end{array}$$

Fig. 5. Syntax of PAM Terms

relations, defined as the smallest binary relations over machine states that satisfy the inference rules given in Figure 6. Intuitively, the relation \rightsquigarrow corresponds to

³ In particular, this presentation makes use of synchronous communications.

the implementation of \equiv , whereas \mapsto implements the actual communication. An actual implementation of this abstract machine does not need to distinguish these relations and would implement $\rightarrow = \rightsquigarrow \uplus \mapsto$, but this distinction will help us to prove correctness properties. In rule AM.NEW, we suppose there is a function $\text{freshn} : \mathcal{P}^f(\mathbf{N}) \rightarrow \mathbf{N}$ such that $\text{freshn}(\mathcal{N}) \notin \mathcal{N}$. We also suppose that names generated by the freshn function never appear in the π -calculus processes in the PAM term (this could be enforced by defining a new syntactic category of names).

We refer the reader to [19, 16] for detailed explanations of these rules. We briefly summarize the main ideas here. An execution of the abstract machine starts with an empty heap (we denote it with $\mathcal{H}_{[]}^{\square}$) that maps all names to empty queues of processes, and a runqueue containing the π -calculus process to be executed. Depending on the form of the process at the top of the runqueue, and the state of the heap, exactly one rule can apply. The execution stops when the runqueue is empty.

A nil process is discarded from the runqueue (rule AM.NIL). Parallel composition of processes is split into two processes. The left hand process is kept at the head of the runqueue, while the right hand process is pushed to the end (rule AM.PAR). Rule AM.NEW implements name restriction by generating new fresh names. When the first term of the runqueue is a prefixed process willing to communicate on a name x , there are two possible cases. If there is no corresponding process in the heap, the process is pushed on the heap queue for x (rules AM.PUSHMESSAGE, AM.PUSHRECEIVER, AM.PUSHREPRECEIVER). If there is a corresponding process in the heap queue (the first element), the communication is performed and the continuation of the receiver and sender are placed in the runqueue (rules AM.COM1, AM.RCOM1, AM.COM2, AM.RCOM2).

Remark 1. In the following, we will consider only machine terms \mathcal{M} for which there exists a process P such that $\langle P :: [], \mathcal{H}_{[]}^{\square}, \emptyset \rangle \rightarrow^* \mathcal{M}$. In particular, we can show that processes appearing in an association $x \rightarrow \mathcal{P}$ are of the form $\pi.P$, where all prefixes are either output on x , or input (replicated or not) on x . We can also show that \mathcal{H} is finite. Moreover, we can notice that the relation \rightarrow is deterministic. In particular, generated fresh names are fully determined by the function freshn in rule AM.NEW.

3.2 Labeled Abstract Machine

We define a labeled version of the Pict abstract machine and essentially follow section 2. This auxiliary calculus is a technical tool, and it is only used for proving the correctness of the abstract machine. Its syntax is defined by adding labels to π -calculus processes appearing in PAM terms. We also extend PAM terms with a finite set of labels. We write LM for the set of labeled PAM terms, and we use

$$\begin{array}{c}
\frac{}{\langle \mathbf{0} :: \mathcal{Q}, \mathcal{H} \rangle \rightsquigarrow \langle \mathcal{Q}, \mathcal{H} \rangle} \text{AM.NIL} \quad \frac{}{\langle (P \mid Q) :: \mathcal{Q}, \mathcal{H} \rangle \rightsquigarrow \langle P :: \mathcal{Q} :: Q, \mathcal{H} \rangle} \text{AM.PAR} \\
\\
\frac{z = \text{freshn}(\mathcal{N})}{\langle \nu x.P :: \mathcal{Q}, \mathcal{H}, \mathcal{N} \rangle \rightsquigarrow \langle P\{z/x\} :: \mathcal{Q}, \mathcal{H}, \mathcal{N} \uplus \{z\} \rangle} \text{AM.NEW} \\
\\
\frac{\mathcal{P} = [] \vee \mathcal{P} = x?(z).Q :: \mathcal{P}' \vee \mathcal{P} = *x?(z).Q :: \mathcal{Q}}{\langle x?(y).P :: \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}\} \rangle \rightsquigarrow \langle \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P} :: x?(y).P \rangle} \text{AM.PUSHRECEIVER} \\
\\
\frac{\mathcal{P} = [] \vee \mathcal{P} = x!(z).Q :: \mathcal{P}'}{\langle x!(y).P :: \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}\} \rangle \rightsquigarrow \langle \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P} :: x!(y).P \rangle} \text{AM.PUSHMESSAGE} \\
\\
\frac{\mathcal{P} = [] \vee \mathcal{P} = x?(z).Q :: \mathcal{P}' \vee \mathcal{P} = *x?(z).Q :: \mathcal{P}'}{\langle *x?(y).P :: \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}\} \rangle \rightsquigarrow \langle \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P} :: *x?(y).P \rangle} \text{AM.PUSHREPReceiver} \\
\\
\frac{\mathcal{P} = x!(z).Q :: \mathcal{P}'}{\langle x?(y).P :: \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}\} \rangle \mapsto \langle P\{z/y\} :: \mathcal{Q} :: Q, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}' \rangle} \text{AM.COM1} \\
\\
\frac{\mathcal{P} = x!(z).Q :: \mathcal{P}'}{\langle *x?(y).P :: \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}\} \rangle \mapsto \langle *x?(y).P :: \mathcal{Q} :: P\{z/y\} :: Q, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}' \rangle} \text{AM.RCOM1} \\
\\
\frac{\mathcal{P} = x?(z).Q :: \mathcal{P}'}{\langle x!(y).P :: \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}\} \rangle \mapsto \langle P :: \mathcal{Q} :: Q\{y/z\}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}' \rangle} \text{AM.COM2} \\
\\
\frac{\mathcal{P} = *x?(z).Q :: \mathcal{P}'}{\langle x!(y).P :: \mathcal{Q}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}\} \rangle \mapsto \langle P :: \mathcal{Q} :: Q\{y/z\}, \mathcal{H} \oplus \{x \rightarrow \mathcal{P}' :: *x?(z).Q \rangle} \text{AM.RCOM2}
\end{array}$$

Fig. 6. PAM Reduction Rules

$\overline{\mathcal{M}}$ and its variants to range over them.

$$\begin{array}{ll}
\overline{\mathcal{M}} ::= \langle \overline{\mathcal{P}}, \overline{\mathcal{H}}, \mathcal{N}, \mathcal{L} \rangle & \text{State} \\
\overline{\mathcal{P}}, \overline{\mathcal{Q}}, \dots ::= [] \mid \overline{P} :: \overline{\mathcal{P}} & \text{Processes Queue} \\
\overline{\mathcal{H}} ::= \{x \rightarrow \overline{\mathcal{P}}_x\}_{x \in \mathbb{N}} & \text{Heap} \\
\mathcal{L} \in \mathcal{P}^f(\mathbb{L}) &
\end{array}$$

We define the set of well-formed PAM terms in Figure 7 and call it WFM. Reduction of labeled PAM terms is defined almost exactly as in the unlabeled calculus, apart from the rules AM.RCOM1 and AM.RCOM2. The functions lab and unl extend as expected on process queues, heaps and machine states.

$$\begin{array}{c}
\frac{\text{wf}(\overline{\mathcal{P}}) \quad \text{wf}(\overline{\mathcal{H}}) \quad \text{lab}(\overline{\mathcal{P}}) \uplus \text{lab}(\overline{\mathcal{H}}) \subseteq \mathcal{L}}{\text{wf}(\langle \overline{\mathcal{P}}, \overline{\mathcal{H}}, \mathcal{N}, \mathcal{L} \rangle)} \text{WF.STATE} \\
\frac{\text{wf}(\overline{\mathcal{P}}) \quad \text{wf}(\overline{P}) \quad \text{lab}(\overline{P}) \cap \text{lab}(\overline{\mathcal{P}}) = \emptyset}{\text{wf}(\overline{P} :: \overline{\mathcal{P}})} \text{WF.PROCQUEUE} \\
\frac{\forall x, y \in \mathbb{N}. x \neq y \implies \text{lab}(\overline{\mathcal{H}}(x)) \cap \text{lab}(\overline{\mathcal{H}}(y)) = \emptyset \quad \forall x \in \mathbb{N}. \text{wf}(\overline{\mathcal{H}}(x))}{\text{wf}(\overline{\mathcal{H}})} \text{WF.HEAP}
\end{array}$$

Fig. 7. Well-Formed PAM Term

$$\begin{array}{c}
\frac{\mathcal{P} = x!(z)_{i'}. \overline{Q} :: \overline{\mathcal{P}}' \quad \alpha \text{ injective and } \mathcal{L}' = \mathcal{L} \uplus \alpha(\mathcal{L})}{\langle *x?(y)_i. \overline{P} :: \overline{\mathcal{Q}}, \overline{\mathcal{H}} \oplus \{x \rightarrow \overline{\mathcal{P}} \}, \mathcal{L} \rangle \mapsto \langle \alpha(*x?(y)_i. \overline{P}) :: \overline{\mathcal{Q}} :: \overline{P}\{z/y\} :: \overline{\mathcal{Q}}, \overline{\mathcal{H}} \oplus \{x \rightarrow \overline{\mathcal{P}}' \}, \mathcal{L}' \rangle} \text{AM.RCOM1} \\
\frac{\mathcal{P} = *x?(z)_{i'}. \overline{Q} :: \overline{\mathcal{P}}' \quad \alpha \text{ injective and } \mathcal{L}' = \mathcal{L} \uplus \alpha(\mathcal{L})}{\langle x!(y)_i. \overline{P} :: \overline{\mathcal{Q}}, \overline{\mathcal{H}} \oplus \{x \rightarrow \overline{\mathcal{P}} \}, \mathcal{L} \rangle \mapsto \langle \overline{P} :: \overline{\mathcal{Q}} :: \overline{Q}\{y/z\}, \overline{\mathcal{H}} \oplus \{x \rightarrow \overline{\mathcal{P}}' :: \alpha(*x?(z)_{i'}. \overline{Q}) \}, \mathcal{L}' \rangle} \text{AM.RCOM2}
\end{array}$$

Fig. 8. Labeled PAM Reduction Rules

Lemma 5 (Stability of Labeling).

- (i) If $\overline{\mathcal{M}} \equiv \overline{\mathcal{M}}'$ and $\overline{\mathcal{M}} \in \text{WFM}$, then $\overline{\mathcal{M}}' \in \text{WFM}$.
- (ii) If $\overline{\mathcal{M}} \rightarrow \overline{\mathcal{M}}'$ and $\overline{\mathcal{M}} \in \text{WFM}$, then $\overline{\mathcal{M}}' \in \text{WFM}$.

Proposition 3 (Operational Correspondence). *Let $\mathcal{M} \in \mathbf{M}$ and $\overline{\mathcal{M}} \in \mathbf{WFM}$ such that $\mathcal{M} = \text{unl}(\overline{\mathcal{M}})$. If \Rightarrow denotes either \mapsto or \rightsquigarrow , we have*

- (i) $\mathcal{M} \Rightarrow \mathcal{M}'$ implies there is $\overline{\mathcal{M}}' \in \mathbf{LM}$ such that $\overline{\mathcal{M}} \Rightarrow \overline{\mathcal{M}}'$ and $\text{unl}(\overline{\mathcal{M}}') = \mathcal{M}'$.
- (ii) $\overline{\mathcal{M}} \Rightarrow \overline{\mathcal{M}}'$ implies $\mathcal{M} \Rightarrow \text{unl}(\overline{\mathcal{M}}')$.

We now define the live actions of a labeled PAM term as the live actions of a corresponding π -calculus term given by a translation function $\llbracket \cdot \rrbracket^r$. Intuitively, $\{l, l'\}$ is a live action whenever there are two matching prefixed processes somewhere in the PAM term that could *potentially* react.

Definition 5 (Translation from PAM to π -calculus). *We define a function $\llbracket \cdot \rrbracket^r$ from \mathbf{LM} to \mathbf{LP} . $\llbracket \cdot \rrbracket^r$ is defined inductively on the structure of a labeled PAM term $\overline{\mathcal{M}}$.*

$$\begin{aligned} \llbracket \langle \overline{\mathcal{P}}, \overline{\mathcal{H}}, \mathcal{N}, \mathcal{L} \rangle \rrbracket^r &= \nu \mathcal{N}. \llbracket \overline{\mathcal{P}} \rrbracket^r \mid \llbracket \overline{\mathcal{H}} \rrbracket^r, \mathcal{L} \\ \llbracket \square \rrbracket^r &= \mathbf{0} \\ \llbracket \overline{P} :: \overline{\mathcal{P}} \rrbracket^r &= \overline{P} \mid \llbracket \overline{\mathcal{P}} \rrbracket^r \\ \llbracket \{x \rightarrow \overline{\mathcal{P}}_x\}_{x \in \mathbf{N}} \rrbracket^r &= \prod_{x \in \mathbf{N}} \llbracket \overline{\mathcal{P}}_x \rrbracket^r \end{aligned}$$

We can define similarly $\llbracket \cdot \rrbracket^r$ for unlabeled terms.

For brevity, we write $\nu \mathcal{N}.P$, but to be rigorous the names in \mathcal{N} should be substituted by fresh π -calculus names.

Definition 6 (Live Actions). *The set of live actions of a labeled PAM term $\overline{\mathcal{M}}$ is defined as*

$$LA(\overline{\mathcal{M}}) = LA(\llbracket \overline{\mathcal{M}} \rrbracket^r)$$

We also define the set of labels belonging to a live action as $L(\overline{\mathcal{M}}) = \{l \in x \mid x \in LA(\overline{\mathcal{M}})\}$.

Lemma 6. *If $LA(\overline{\mathcal{M}}) \neq \emptyset$ then $\exists \overline{\mathcal{M}}'. \overline{\mathcal{M}} \rightsquigarrow^* \mapsto \overline{\mathcal{M}}'$.*

The following theorem can be seen as a fairness property for the labeled abstract machine.

Theorem 1. *If $\overline{\mathcal{M}}_0 \rightarrow \dots \rightarrow \overline{\mathcal{M}}_n \rightarrow \dots$ is an infinite execution then for any strictly increasing sequence $(u_n)_{n \in \mathbf{N}}$, we have $\bigcap_{n \in \mathbf{N}} L(\overline{\mathcal{M}}_{u_n}) = \emptyset$*

The proof is technical but it relies on intuitive ideas. Informally, it follows from two key properties of the abstract machine reduction system:

- If a process $\pi_l.\overline{P}$ appears in an evaluation context in the runqueue, it will eventually reach the top of the runqueue. This is proven in Lemma 13.
- The heap queues are organized following a FIFO policy.

4 Correctness

From an operational point of view, the correctness of an abstract machine can be stated by relating abstract machine executions of a process P with π -calculus executions of the same process P executed by the abstract machine. The initial state of an abstract machine running P is $\langle P, \mathcal{H}_\square \rangle$, hence we introduce the following translation function.

Definition 7 (Translation from π -calculus to PAM).

$$\llbracket P \rrbracket = \langle P :: \square, \mathcal{H}_\square, \emptyset \rangle \quad \llbracket \bar{P}, \mathcal{L} \rrbracket = \langle \bar{P} :: \square, \mathcal{H}_\square, \emptyset, \mathcal{L} \rangle$$

The first property we consider is the soundness of the abstract machine with respect to the calculus. Intuitively, this means that abstract machine executions correspond to valid π -calculus executions. If a machine state \mathcal{M} , corresponding to a process state P , reduces to a machine state \mathcal{M}' , then \mathcal{M}' must correspond to a process state P' where P reduces to P' . One reduction in the π -calculus may be implemented by several reductions of the abstract machine. In order to model a one-to-one correspondence, we identify two kinds of reductions. Administrative reductions are denoted by \rightsquigarrow model structural equivalence. Communication reductions are denoted by \mapsto . We will establish a correspondence between the relations $\rightsquigarrow^* \mapsto$ over PAM terms and \rightarrow over π -calculus terms. For that, we define a relation $\mathcal{M} \preceq P$ to mean that P corresponds to \mathcal{M} , read \mathcal{M} implements P .

We still need to define the relation \preceq . It has to be convincing enough that it effectively relates equivalent process states and machine states. It should at least enjoy the following two properties:

- $\llbracket P \rrbracket \preceq P$.
- If $\mathcal{M} \rightsquigarrow \mathcal{M}'$ and $\mathcal{M} \preceq P$ then $\mathcal{M}' \preceq P$.

The first property follows the idea that the initial state of an abstract machine executing P is $\llbracket P \rrbracket$. The second property follows the intuition that \rightsquigarrow is a structural, or administrative, reduction and that abstract machine states still implement the same π -calculus process after such reductions. We define \preceq as the smallest relation enjoying these two properties.

Definition 8. $\mathcal{M} \preceq P \iff \llbracket P \rrbracket \rightsquigarrow^* \mathcal{M}$.

The definition of \preceq extends naturally to labeled processes.

Alternatively, we could define $\mathcal{M} \preceq P$ as $\llbracket P \rrbracket \simeq \mathcal{M}$ where \simeq is a suitable equivalence over machine terms. Since any reasonable equivalence must verify $\mathcal{M} \rightsquigarrow^* \mathcal{M}' \implies \mathcal{M} \simeq \mathcal{M}'$, our definition of \preceq provides a stronger result.

Note that we do not have a notion of observables, although it would make the correspondence relation \preceq more convincing. However, it should be straightforward to define an observation predicate on π -calculus processes and PAM terms (such as those in [1, 7, 11]) and show that \preceq preserves the observables.

The following lemma relates the live actions of a labeled PAM term and a labeled process it implements.

Lemma 7. *If $\overline{\mathcal{M}} \preceq C$ then $LA(\overline{\mathcal{M}}) = LA(C)$.*

To prove the soundness property, we use a translation function from PAM terms to π -calculus processes. Its full definition is given in the appendix. The following lemma states the properties needed for the proof of soundness.

Lemma 8. *There exists a function $[\cdot]^{-1}$ from \mathbf{M} to \mathbf{P} such that*

- $\mathcal{M} \preceq [\mathcal{M}]^{-1}$
- $[[P]]^{-1} \equiv P$
- if $\mathcal{M} \rightsquigarrow \mathcal{M}'$ then $[\mathcal{M}]^{-1} \equiv [\mathcal{M}']^{-1}$
- if $\mathcal{M} \mapsto \mathcal{M}'$ then $[\mathcal{M}]^{-1} \mapsto [\mathcal{M}']^{-1}$.

This translation function allows us to build an execution of π -calculus processes starting from an execution of PAM terms such that every PAM term corresponds to an implementation of the corresponding π -calculus term. Note that the translation function $[\cdot]^r$ given in Definition 5 does not satisfy the first property of Lemma 8.

Theorem 2 (Soundness). *If $(\mathcal{M} \rightsquigarrow^* \mapsto \mathcal{M}' \wedge \mathcal{M} \preceq P)$ then $(\exists P'. P \rightarrow P' \wedge \mathcal{M}' \preceq P')$.*

Proof. The theorem follows from $[[P]] \rightsquigarrow^* \mapsto \mathcal{M} \implies (\exists P'. P \rightarrow P' \wedge \mathcal{M} \preceq P')$ which is a consequence of Lemma 8 with $P' = [[\mathcal{M}]]^{-1}$.

This property is not sufficient to prove the correctness of the abstract machine. Other properties are needed to characterize which executions of the π -calculus are actually implemented. First, a liveness property ensures that a PAM term is never blocked when it corresponds to a π -calculus term that can reduce.

Theorem 3 (Liveness). *If $P \rightarrow P' \wedge \mathcal{M} \preceq P$ then $\exists \mathcal{M}'. \mathcal{M} \rightsquigarrow^* \mapsto \mathcal{M}'$.*

Proof. We first prove: $P \rightarrow P' \implies \exists \mathcal{M}. [[P]] \rightsquigarrow^* \mapsto \mathcal{M}$. If $P \rightarrow P'$, we have $C \rightarrow C'$ with $\text{unl}(C) = P$ and $\text{unl}(C') = P'$, by propositions 2 and 1. Moreover, by Lemma 7, $LA([C]) = LA(C)$ with $LA(C) \neq \emptyset$, by Lemma 3. We deduce $[C] \rightsquigarrow^* \mapsto \overline{\mathcal{M}'}$ for some $\overline{\mathcal{M}'}$, by Lemma 6. We conclude, by Proposition 3, that $[[P]] = \text{unl}([C]) \rightsquigarrow^* \mapsto \text{unl}(\overline{\mathcal{M}'})$.

We know now that $[[P]] \rightsquigarrow^* \mapsto \mathcal{M}'' \mapsto \mathcal{M}'$ for some \mathcal{M}'' and \mathcal{M}' . Moreover we have $[[P]] \rightsquigarrow^* \mathcal{M}$, by definition of $\mathcal{M} \preceq P$. Because \rightarrow is deterministic, we conclude $\mathcal{M} \rightsquigarrow^* [[P]] \rightsquigarrow^* \mathcal{M}'' \mapsto \mathcal{M}'$.

Finally, our main result is a fairness theorem.

Theorem 4 (Fairness). *If $\mathcal{M}_0 \rightsquigarrow^* \mapsto \dots \rightsquigarrow^* \mapsto \mathcal{M}_n \rightsquigarrow^* \mapsto \dots$ is an infinite execution then there exists a fair execution $P_0 \rightarrow \dots \rightarrow P_n \rightarrow \dots$ such that $\mathcal{M}_i \preceq P_i$ for all i .*

Proof. Let $\mathcal{M}_0 \rightsquigarrow^* \mathcal{M}_1 \rightsquigarrow^* \mathcal{M}_2 \rightsquigarrow^* \dots$ be an infinite execution. We have an execution $\overline{\mathcal{M}}_0 \rightsquigarrow^* \overline{\mathcal{M}}_1 \rightsquigarrow^* \overline{\mathcal{M}}_2 \rightsquigarrow^* \dots$ such that for all i , $\text{unl}(\overline{\mathcal{M}}_i) = \mathcal{M}_i$, by Proposition 3.

The soundness theorem (Theorem 2) extends to the labeled calculus and gives us an execution $C_0 \rightarrow \dots \rightarrow C_n \rightarrow \dots$ such that

$$\begin{array}{ccccccc} \overline{\mathcal{M}}_0 & \xrightarrow{\rightsquigarrow^*} & \overline{\mathcal{M}}_1 & \xrightarrow{\rightsquigarrow^*} & \overline{\mathcal{M}}_2 & \xrightarrow{\rightsquigarrow^*} & \dots \\ \downarrow \simeq & & \downarrow \simeq & & \downarrow \simeq & & \\ C_0 & \longrightarrow & C_1 & \longrightarrow & C_2 & \longrightarrow & \dots \end{array}$$

From Lemma 7, we have $LA(C_i) = LA(\overline{\mathcal{M}}_i)$ for all i . Then we deduce from Theorem 1 that the execution $C_0 \rightarrow \dots \rightarrow C_n \rightarrow \dots$ is fair. By erasing the labels in both executions, we deduce the result.

$$\begin{array}{ccccccc} \text{unl}(\overline{\mathcal{M}}_0) & \xrightarrow{\rightsquigarrow^*} & \text{unl}(\overline{\mathcal{M}}_1) & \xrightarrow{\rightsquigarrow^*} & \text{unl}(\overline{\mathcal{M}}_2) & \xrightarrow{\rightsquigarrow^*} & \dots \\ \downarrow \simeq & & \downarrow \simeq & & \downarrow \simeq & & \\ \text{unl}(C_0) & \longrightarrow & \text{unl}(C_1) & \longrightarrow & \text{unl}(C_2) & \longrightarrow & \dots \end{array}$$

5 Related Work

Comparison with Pict Correctness results in [19] include a soundness and a liveness property based on the translation function $\llbracket \cdot \rrbracket^r$ from PAM terms to π -calculus terms given in Definition 6:

- (i) $\mathcal{M} \rightarrow \mathcal{M}' \implies \llbracket \mathcal{M} \rrbracket^r \equiv \llbracket \mathcal{M}' \rrbracket^r \vee \llbracket \mathcal{M} \rrbracket^r \rightarrow \llbracket \mathcal{M}' \rrbracket^r$
- (ii) $P \rightarrow P' \implies \exists \mathcal{M}. \llbracket P \rrbracket \rightarrow \mathcal{M}$

However, these properties are not sufficient for proving soundness or liveness. The first property means that we can build a π -calculus reduction from a PAM reduction, but does not prove that PAM reductions implement π -calculus reductions. A property relating \mathcal{M} and $\llbracket \llbracket \mathcal{M} \rrbracket^r \rrbracket$, such as our Lemma 8 is missing. The second property tells us that if P reduces to P' , there is a PAM reduction $\llbracket P \rrbracket \rightarrow \mathcal{M}$. However, the property cannot be applied on more than the first step of execution, as we do not know if there is P'' such that $P \rightarrow P''$ and $\llbracket P'' \rrbracket = \mathcal{M}$.

In [18], the Pict abstract machine is proven correct using a notion of testing, and a realistic model of the interactions between the abstract machine and its environment. However, they do not consider fairness issues.

Fairness Fairness has been defined using labels in CCS [3] and in the π -calculus [4, 2]. We essentially followed the same idea but our presentation is simpler as we annotate labeled terms with a set of labels that allow us to generate fresh labels in the replication rules, without relying on a structured labeling language.

In [13], fairness is defined for the π -calculus by considering *normal* reductions where α -equivalence is restricted and tags similar to labels are used to distinguish processes. Fresh tags are generated using the π -calculus name restriction operator.

Correctness of Abstract Machine There have been several recent papers devoted to the formal description of implementations of process calculi based on the π -calculus or the Ambient calculus. In addition to Pict, one can notably cite the Jocaml distributed implementation of the Join calculus [6, 5], the Join calculus implementation of Mobile Ambients [7], Nomadic Pict [22, 20], the abstract machine for the M-calculus [10], the Fusion Machine [8], the PAN and GCPAN abstract machines for Safe Ambients [11, 12], the CAM abstract machine for Channel Ambients [15] and the abstract machine for the Kell calculus [1]. Most of these works [7, 22, 20, 10, 8, 11, 12, 15, 1] deal with distributed implementations of calculi, rather than local implementation of concurrent processes like in Pict. They are defined by a translation to a low-level calculus or abstract machine. Their correctness is proven in terms of bisimilarity that does not apply to our setting, since Pict implementation makes deterministic choice and PAM reductions do not match all π -calculus reductions. Implementations that consider scheduling of processes are given in [15, 10]. In [15], a soundness property is given, similar to the one given for Pict in [19]. In [10], scheduling of processes is done as in Pict using FIFO lists, but no proof of correctness is given.

6 Conclusion

In this paper, we first defined strong fairness in the π -calculus. We then proved that Pict abstract machine executions are sound with respect to π -calculus executions and that they enjoy fairness and liveness properties. These correctness results for Pict are new and in particular, fairness has not been proven for any implementation of process calculi based on the π -calculus. We believe that these techniques are simple and general enough to be adapted to other calculi.

Very little work has been done on the scheduling of processes in the π -calculus or its variants. For future research, we will investigate alternative scheduling strategies. In particular, we would like to extend Pict and its implementation with priority constraints. Processes could be prioritized in order to allocate more processor time to more important processes. In Pict, even though executions are strongly fair, in a term $P \mid Q$, P can monopolize the processor usage by spawning new subprocesses much faster than Q . One can imagine annotated processes like in $P_h \mid Q_l$ where the annotations are taken into account by the scheduler. Such a scheme would fit naturally in a calculus with hierarchical localities such as [1]. For instance, a term of the form $a[b[P] \mid c[Q]]$ can be interpreted as two (possibly untrusted) agents b and c executed by a site a . The parent site a should be able to control the processor usage of the agents it is executing.

Most correctness results of the implementations of process calculi with localities concern their distributed implementation, but do not deal with the correctness of their local implementation, *i.e.* the scheduling of processes. On the other hand, Pict defines a local implementation. It would be interesting to consider correctness results combining these two approaches. We are currently investigating the proof of a refined abstract machine based on [1].

Acknowledgments We are grateful to Healdene Goguen, Benjamin Pierce, Alan Schmitt and Jean-Bernard Stefani, as well as the anonymous reviewers, for their comments on earlier drafts. We thank Pablo Garralda, whose PhD thesis work inspired us to use labeled processes to study fairness [9].

References

1. Philippe Bidinger, Alan Schmitt, and Jean-Bernard Stefani. An abstract machine for the Kell calculus. In *7th IFIP International Conference on Formal Methods for Object-Based Distributed Systems (FMOODS)*, pages 43–58, Athens, Greece, June 2005.
2. Diletta Cacciagrano, Flavio Corradini, and Catuscia Palamidessi. Fairpi. In *Proceedings of EXPRESS'06, ENTCS*, 2006.
3. Gerardo Costa and Colin Stirling. Weak and strong fairness in CCS. In Michal Chytil and Václav Koubek, editors, *MFCS*, volume 176 of *Lecture Notes in Computer Science*, pages 245–254. Springer, 1984.
4. F. Corradini D.R. Cacciagrano. Fairness in the pi-calculus. Technical report, Dipartimenti di Informatica, Università di L'Aquila, TR 005/2004, 2004.
5. Fabrice Le Fessant. *JoCaml: Conception et Implantation d'un Langage à Agents Mobiles*. PhD thesis, Ecole Polytechnique, 2001.
6. C. Fournet, G. Gonthier, J.J. Levy, L. Maranget, and D. Remy. A calculus of mobile agents. In *Proceedings 7th International Conference on Concurrency Theory (CONCUR '96), Lecture Notes in Computer Science 1119*, pages 406–421. Springer Verlag, 1996.
7. C. Fournet, J.J. Levy, and A. Schmitt. An asynchronous distributed implementation of mobile ambients. In *Proceedings of the International IFIP Conference TCS 2000, Sendai, Japan, Lecture Notes in Computer Science 1872*, pages 348–364. Springer, 2000.
8. Philippa Gardner, Cosimo Laneve, and Lucian Wischik. The fusion machine. In L. Brim, P. Jancar, and M. Kretinsky, editors, *CONCUR 2002*, volume 2421 of *LNCS*, pages 418–433. Springer-Verlag, 2002.
9. Pablo Garralda. *Boxed Ambients for Global Computing*. PhD thesis, Stevens Institute of Technology, New Jersey, USA, 2007.
10. F. Germain, M. Lacoste, and J.B. Stefani. An abstract machine for a higher-order distributed process calculus. In *Proceedings of the EACTS Workshop on Foundations of Wide Area Network Computing (F-WAN)*, July 2002.
11. Paola Giannini, Davide Sangiorgi, and Andrea Valente. Safe ambients: abstract machine and distributed implementation. *Sci. Comput. Program.*, 59(3):209–249, 2006.
12. Daniel Hirschhoff, Damien Pous, and Davide Sangiorgi. A correct abstract machine for safe ambients. In *COORDINATION*, pages 17–32, 2005.
13. Naoki Kobayashi. A type system for lock-free processes. *Inf. Comput.*, 177(2):122–159, 2002.
14. Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, parts I and II. *Inf. Comput.*, 100(1):1–78, 1992.
15. A. Phillips, N. Yoshida, and S. Eisenbach. A distributed abstract machine for boxed ambient calculi. In *Proceedings of ESOP 2004, LNCS*. Springer-Verlag, April 2004.

16. Benjamin C. Pierce and David N. Turner. Pict: A programming language based on the pi-calculus. In Gordon Plotkin, Colin Stirling, and Mads Tofte, editors, *Proof, Language and Interaction: Essays in Honour of Robin Milner*, pages 455–494. MIT Press, 2000.
17. D. Sangiorgi and D. Walker. *The π -calculus: A Theory of Mobile Processes*. Cambridge University Press, 2001.
18. Peter Sewell. On implementations and semantics of a concurrent programming language. In *Proceedings of CONCUR 97: Concurrency Theory (Warsaw)*. LNCS 1243, pages 391–405, 1997.
19. D. Turner. The polymorphic π -calculus: Theory and implementation. Technical report, University of Edinburgh, GB, 1996.
20. A. Unyapoth and P. Sewell. Nomadic Pict: Correct Communication Infrastructures for Mobile Computation. In *Proceedings ACM Int. Conf. on Principles of Programming Languages (POPL)*, pages 116–127, 2001.
21. Lucian Wischik. *Explicit Fusions: Theory and Implementation*. PhD thesis, Computer Laboratory, University of Cambridge, 2001.
22. P. Wojciechowski and P. Sewell. Nomadic Pict: Language and Infrastructure. *IEEE Concurrency*, vol. 8, no 2, pages 42–52, 2000.

A Proofs

A.1 Labeled π -calculus

Proof (Proof of Lemma 1). For the first part of the Lemma, we prove by induction on the derivation of $C \equiv D$ that:

- $C \equiv D$ and $\text{wf}(C)$ implies $\text{wf}(D)$
- $C \equiv D$ and $\text{wf}(D)$ implies $\text{wf}(C)$

All cases are immediate. We prove the second part of the lemma by induction on the derivation of $\bar{P}, \mathcal{L} \rightarrow \bar{P}', \mathcal{L}'$.

LR.RED We have $\text{wf}(x!(y)_l.\bar{P} \mid x?(z)_{l'}.\bar{P}', \mathcal{L})$. By rules WF.PAR and WF.PREFIX, we have $\text{lab}(\bar{P}) \uplus \text{lab}(\bar{P}') \subseteq \mathcal{L}$. We conclude by rule WF.PAR by remarking that lab is stable by renaming.

LR.REP As before, we can show that $\text{lab}(\bar{P}) \uplus \text{lab}(\bar{Q}\{z/y\}) \subseteq \mathcal{L}$. By the premise of the rule, we have $\text{lab}(\bar{P}) \uplus \text{lab}(\bar{Q}\{z/y\}) \uplus \text{lab}(\alpha(*x?(z)_{l'}.\bar{Q})) \subseteq \mathcal{L}'$. We conclude with rule WF.PAR.

LR.STR By induction and the first part of the lemma.

LR.CTX By induction on the structure of $\bar{\mathbf{E}}$.

Proof (Proof of Lemma 2).

- (i) Immediate by induction on the derivation of $\text{wf}(C)$.
- (ii) We suppose first that $\bar{P}, \mathcal{L} \rightarrow \bar{P}', \mathcal{L}' \rightarrow \bar{P}'', \mathcal{L}''$ with $l \in \text{lab}(\bar{P}'') \cap \text{lab}(\bar{P})$. Suppose $l \notin \text{lab}(\bar{P}')$, we can show that $l \in \beta(\mathcal{L}')$ for an injective function β such that $\beta(\mathcal{L}') \uplus \mathcal{L}' = \mathcal{L}''$. Besides, we have $l \in \mathcal{L} \subseteq \mathcal{L}'$ which contradicts the disjointness of $\beta(\mathcal{L}')$ and \mathcal{L}' . The same reasoning applies for the general case.

Proof (Proof of Proposition 1).

- (i) Immediate by induction on the derivation of $P \rightarrow P'$ since for all inference rules defining \rightarrow , we have a corresponding labeled rule. For rule L.STRUCT, we need the auxiliary lemma

$$(P \equiv P' \wedge P = \text{unl}(\bar{P})) \implies (\exists \bar{P}'. \bar{P} \equiv \bar{P}' \wedge P' = \text{unl}(\bar{P}'))$$

that we prove similarly by induction on the derivation of $P \equiv P'$ (we consider both directions of \equiv to handle the symmetry of \equiv).

- (ii) Immediate by induction on the derivation of $\bar{P} \rightarrow \bar{P}'$ since erasing labels in the labeled rule gives the corresponding unlabeled rule. For rule LR.STRUCT, we need to remark that $\bar{P} \equiv \bar{P}' \implies \text{unl}(\bar{P}) \equiv \text{unl}(\bar{P}')$ (by erasing the labels in the derivation of $\bar{P} \equiv \bar{P}'$).

Proof (Proof of Proposition 2). Immediate, any injective labeling of the prefixes works.

Proof (Proof of Lemma 3). We prove the direct implication by induction on the derivation of $C \rightarrow C'$. Cases LR.RED and LR.REP are immediate. Case LR.STR is a consequence of Lemma 1 and the induction hypothesis. We prove Case LR.CTX by induction on the structure of the context. The converse implication is immediate.

Proof (Proof of Lemma 4). The implication (ii) \implies (i) follows from the existence of a labeled execution (Propositions 1 and 2). Let us prove the converse implication. Suppose that $C_0 \rightarrow \dots \rightarrow C_n \rightarrow \dots$ and $C'_0 \rightarrow \dots \rightarrow C'_n \rightarrow \dots$ are such that for all i , $\text{unl}(C_i) = \text{unl}(C'_i) = P_i$, $C_i = \overline{P}_i, \mathcal{L}_i$ and $C'_i = \overline{P}'_i, \mathcal{L}'_i$. We can show that there is a substitution of labels σ such that for all i , $\overline{P}'_i \sigma = \overline{P}_i$. Suppose now that $C_0 \rightarrow \dots \rightarrow C_n \rightarrow \dots$ is fair. Let $(u_n)_{n \in \mathbb{N}}$ be an increasing sequence of integers. We have $\bigcap_{n \in \mathbb{N}} L(C'_{u_n}) = \bigcap_{n \in \mathbb{N}} L(C_{u_n}) \sigma = (\bigcap_{n \in \mathbb{N}} L(C_{u_n})) \sigma = \emptyset \sigma = \emptyset$. Hence, $C'_0 \rightarrow \dots \rightarrow C'_n \rightarrow \dots$ is fair.

A.2 Labeled PAM

Proof (Proof of Lemma 5). Immediate, similar to the proof of 1.

Lemma 9 (Determinism). *If $\overline{\mathcal{M}} \rightarrow^n \overline{\mathcal{M}}'$ and $\overline{\mathcal{M}} \rightarrow^n \overline{\mathcal{M}}''$ then there is a bijective substitution of labels σ such that $\overline{\mathcal{M}}' \sigma = \overline{\mathcal{M}}''$ and $\text{lab}(\overline{\mathcal{M}})$ is invariant by σ .*

Proof. We show it for $n = 1$ (it generalizes trivially by induction for an arbitrary n) and reason by cases on the rule used to infer both reductions (it is the same rule for both reductions as the rule is uniquely determined by the form of $\overline{\mathcal{M}}$). The interesting rules are AM.RCOM1 and AM.RCOM2 because the premises allow for non-determinism. Consider Case AM.RCOM1, we have

$$\begin{aligned} \overline{\mathcal{M}} &= \langle *x?(y)_l.\overline{P} :: \overline{\mathcal{Q}}, \overline{\mathcal{H}}, \mathcal{L} \rangle \\ \overline{\mathcal{M}}' &= \langle \alpha(*x?(y)_l.\overline{P}) :: \overline{\mathcal{Q}}, \overline{\mathcal{H}}', \mathcal{L}' \rangle \\ \overline{\mathcal{M}}'' &= \langle \beta(*x?(y)_l.\overline{P}) :: \overline{\mathcal{Q}}, \overline{\mathcal{H}}', \mathcal{L}'' \rangle \end{aligned}$$

where α and β are injective functions and $\mathcal{L}' = \mathcal{L} \uplus \alpha(\mathcal{L})$ and $\mathcal{L}'' = \mathcal{L} \uplus \beta(\mathcal{L})$. We define σ as follows:

$$\begin{aligned} x \in \alpha(\mathcal{L}) \setminus (\alpha(\mathcal{L}) \cap \beta(\mathcal{L})) &\implies \sigma(x) = \beta(\alpha^{-1}(x)) \\ x \in \beta(\mathcal{L}) \setminus (\alpha(\mathcal{L}) \cap \beta(\mathcal{L})) &\implies \sigma(x) = \alpha(\beta^{-1}(x)) \\ \text{otherwise } \sigma(x) &= x \end{aligned}$$

We see easily that σ is bijective and $\overline{\mathcal{M}}' \sigma = \overline{\mathcal{M}}''$.

Proof (Proof of Proposition 3). The first part is immediate by cases on the derivation of $\mathcal{M} \Rightarrow \mathcal{M}'$. We simply need to notice that for every unlabeled rule there is a corresponding labeled rule. In particular, in rules AM.RCOM1 and AM.RCOM2, we can find α such that the premises hold.

The second part is by cases on the derivation of $\overline{\mathcal{M}} \Rightarrow \overline{\mathcal{M}}'$. We simply need to erase the labels in the rule used to infer the corresponding unlabeled reduction.

Lemma 10 (Uniqueness of Labeling).

- (i) If $\overline{\mathcal{M}} \in \text{WFM}$ then no label l occurs more than once in $\overline{\mathcal{M}}$.
- (ii) If $\overline{\mathcal{M}} \in \text{WFM}$, $\overline{\mathcal{M}} \rightarrow^* \overline{\mathcal{M}}' \rightarrow^* \overline{\mathcal{M}}''$ and $l \in \text{lab}(\overline{\mathcal{M}}) \cap \text{lab}(\overline{\mathcal{M}}'')$, then $l \in \text{lab}(\overline{\mathcal{M}}')$.

Proof. The proof is similar to the corresponding proof for the labeled π -calculus.

Lemma 11. If $\overline{\mathcal{M}} \rightsquigarrow \overline{\mathcal{M}}'$ then $LA(\overline{\mathcal{M}}) = LA(\overline{\mathcal{M}}')$.

Proof. The two following properties are direct consequences of the definition of LA :

- $C \equiv D \implies LA(C) = LA(D)$
- $\mathcal{L} \subseteq \mathcal{L}' \implies LA(\overline{\mathcal{P}}, \mathcal{L}) = LA(\overline{\mathcal{P}}, \mathcal{L}')$

If $\llbracket \overline{\mathcal{M}} \rrbracket^r = \overline{\mathcal{P}}, \mathcal{L}$ and $\llbracket \overline{\mathcal{M}}' \rrbracket^r = \overline{\mathcal{P}}', \mathcal{L}'$, we can show by cases on the derivation of $\overline{\mathcal{M}} \rightsquigarrow \overline{\mathcal{M}}'$ that $\llbracket \overline{\mathcal{P}}, \mathcal{L} \rrbracket^r \equiv \llbracket \overline{\mathcal{P}}', \mathcal{L}' \rrbracket^r$ (notice that two structurally equivalent processes must have the same set of labels). We then have

$$LA(\overline{\mathcal{M}}) = LA(\overline{\mathcal{P}}, \mathcal{L}) = LA(\overline{\mathcal{P}}', \mathcal{L}') = LA(\overline{\mathcal{M}}')$$

Definition 9 (Size of a Process). We define the size of a process and of a process queue.

$$\begin{aligned} \text{size}(\mathbf{0}) &= 1 \\ \text{size}(\pi.P) &= 1 \\ \text{size}(\nu x.P) &= \text{size}(P) + 1 \\ \text{size}(P \mid P') &= \text{size}(P) + \text{size}(P') \\ \text{size}(\llbracket \rrbracket) &= 0 \\ \text{size}(P :: \mathcal{P}) &= \text{size}(P) + \text{size}(\mathcal{P}) \end{aligned}$$

This definition extends to labeled processes and labeled Pict terms.

Proof (Proof of Lemma 6). We will prove the contrapositive proposition. Suppose that for any n and $\overline{\mathcal{M}}'$ such that $\overline{\mathcal{M}} \rightsquigarrow^n \overline{\mathcal{M}}'$ there is no $\overline{\mathcal{M}}''$ such that $\overline{\mathcal{M}}' \mapsto \overline{\mathcal{M}}''$.

We show first that $\overline{\mathcal{M}} \rightsquigarrow^* \overline{\mathcal{M}}'$ with $\overline{\mathcal{M}}' = \langle \llbracket \rrbracket, \overline{\mathcal{H}}' \rangle$. Suppose $\overline{\mathcal{M}} = \langle \overline{\mathcal{P}}, \overline{\mathcal{H}} \rangle$, we prove by induction on $n = \text{size}(\overline{\mathcal{P}})$ that there exists $p \leq n$ such that $\overline{\mathcal{M}} \rightsquigarrow^p \overline{\mathcal{M}}'$ with $\overline{\mathcal{M}}' = \langle \llbracket \rrbracket, \overline{\mathcal{H}}' \rangle$. If $n = 0$, $\overline{\mathcal{P}} = \llbracket \rrbracket$ which proves the result. If $n > 0$, $\overline{\mathcal{P}} \neq \llbracket \rrbracket$ and we have either $\overline{\mathcal{M}} \rightsquigarrow \overline{\mathcal{M}}'$ or $\overline{\mathcal{M}} \mapsto \overline{\mathcal{M}}'$. By hypothesis, only the first case is possible. Moreover if $\overline{\mathcal{M}}' = \langle \overline{\mathcal{P}}', \overline{\mathcal{H}}' \rangle$, we remark that $\text{size}(\overline{\mathcal{P}}') < \text{size}(\overline{\mathcal{P}})$. We can conclude by applying the induction hypothesis.

By Lemma 11, we have $LA(\overline{\mathcal{M}}) = LA(\overline{\mathcal{M}}')$. We conclude by remarking that $LA(\overline{\mathcal{M}}') = \emptyset$ since in the heap, the process queue for a given name can not contain two complementary actions (cf. Remark 1).

Lemma 12. *If $l \in L(\langle \bar{\mathcal{P}}, \bar{\mathcal{H}}, \mathcal{N}, \mathcal{L} \rangle)$, we have either $\bar{\mathcal{P}} = \bar{\mathcal{P}}_0 :: \bar{\mathbf{E}}[\pi_l.\bar{\mathcal{P}}] :: \bar{\mathcal{P}}_1$ for some E , $\bar{\mathcal{P}}_0$ and $\bar{\mathcal{P}}_1$ or $\bar{\mathcal{H}}(x) = \bar{\mathcal{Q}} :: \pi_l.\bar{\mathcal{P}} :: \bar{\mathcal{Q}}'$ for some x , $\bar{\mathcal{Q}}$, $\bar{\mathcal{Q}}'$ and $\pi_l.\bar{\mathcal{P}}$.*

Proof. Let us write $\bar{\mathcal{M}} = \langle \bar{\mathcal{P}}, \bar{\mathcal{H}}, \mathcal{N}, \mathcal{L} \rangle$. We define $\bar{\mathcal{P}}_0 = \llbracket \bar{\mathcal{P}} \rrbracket^r \mid \llbracket \bar{\mathcal{H}} \rrbracket^r$. We have $\llbracket \bar{\mathcal{M}} \rrbracket^r = \nu \mathcal{N}.\llbracket \bar{\mathcal{P}} \rrbracket^r \mid \llbracket \bar{\mathcal{H}} \rrbracket^r, \mathcal{L}$. From the definition of LA , we see that $LA(\bar{\mathcal{M}}) = LA(\llbracket \bar{\mathcal{M}} \rrbracket^r) = LA(\bar{\mathcal{P}}_0, \mathcal{L})$. Moreover, for some l' , $\bar{\mathcal{P}}_0, \mathcal{L} \equiv \nu \tilde{x}.\pi_l.\bar{\mathcal{Q}} \mid \pi_{l'}.\bar{\mathcal{Q}}' \mid \bar{\mathcal{Q}}'', \mathcal{L}$ where π_l and $\pi_{l'}$ are matching prefixes. Since $\bar{\mathcal{M}}$ is well-formed, we know that either $l \in \text{lab}(\bar{\mathcal{P}})$ or $l \in \text{lab}(\bar{\mathcal{H}})$.

In the first case, $\bar{\mathcal{P}} = \bar{\mathcal{P}}_0 :: \bar{\mathcal{P}} :: \bar{\mathcal{P}}_1$ and we have $\bar{\mathcal{P}}_0 = \bar{\mathcal{P}} \mid \bar{\mathcal{Q}}$ with $l \in \text{lab}(\bar{\mathcal{P}})$, for some $\bar{\mathcal{Q}}$. We deduce $\bar{\mathcal{P}} = \bar{\mathbf{E}}[\pi_l.\bar{\mathcal{Q}}]$ for some $\bar{\mathbf{E}}$.

In the second case, $\bar{\mathcal{H}} = \bar{\mathcal{H}}' \oplus \{x \rightarrow \bar{\mathcal{P}}_0 :: \pi_{l''}.\bar{\mathcal{P}} :: \bar{\mathcal{P}}_1\}$ and we have $\bar{\mathcal{P}}_0 = \pi_{l''}.\bar{\mathcal{P}} \mid \bar{\mathcal{Q}}$ for some $\bar{\mathcal{Q}}$ and $l \in \text{lab}(\pi_{l''}.\bar{\mathcal{P}})$. We deduce $l'' = l$.

The following lemma states a key property of the abstract machine reduction system. If a process $\pi_l.\bar{\mathcal{P}}$ appears in an evaluation context in the runqueue, it will eventually reach the top of the runqueue.

Lemma 13. *If $\bar{\mathcal{M}} = \langle \bar{\mathcal{P}}_0 :: \bar{\mathbf{E}}[\pi_l.\bar{\mathcal{P}}] :: \bar{\mathcal{P}}_1, \mathcal{H} \rangle$ then $\mathcal{M} \rightarrow^* \langle \mathcal{P}', \mathcal{H}' \rangle$ with $\bar{\mathcal{P}}' = \pi_l.\bar{\mathcal{P}} :: \bar{\mathcal{P}}'_0$.*

Proof. We give first an informal proof. If the runqueue is of the form $\bar{\mathcal{P}}_0 :: \bar{\mathbf{E}}[\pi_l.\bar{\mathcal{P}}] :: \bar{\mathcal{P}}_1$, after a finite number (usually one) of reduction steps, the size of the term $\bar{\mathcal{P}}_0 :: \bar{\mathbf{E}}[.]$ strictly decreases (for an appropriate definition of size). Eventually, the size reaches \perp which means $\pi_l.\bar{\mathcal{P}}$ is at the top of the runqueue. In rules AM.NIL as well as the three “push” rules, the first process of the queue is dropped. In rules AM.NEW, AM.PAR and all the communication rules besides AM.RCOM1, the first process is replaced by a smaller process. In the rule AM.RCOM1, the first process of the queue (a replicated receiver) is dropped after a number of steps corresponding to the number of processes in the heap that can react with it.

Formally, we first define a *distance* function $d(., .)$. If $\bar{\mathcal{P}} = \bar{\mathcal{P}}_0 :: \bar{\mathbf{E}}[\pi_l.\bar{\mathcal{P}}] :: \bar{\mathcal{P}}_1$ then $d(l, \bar{\mathcal{P}}) = (\text{size}(\bar{\mathcal{P}}_0), \text{size}(\bar{\mathbf{E}}[\pi_l.\bar{\mathcal{P}}]))$. Otherwise, $d(l, \bar{\mathcal{P}}) = \perp$. The lemma follows from these three properties:

- (i) If $d(l, \bar{\mathcal{P}}) = (u, v)$ then $u \geq 0$ and $v \geq 1$.
- (ii) $d(l, \bar{\mathcal{P}}) = (0, 1)$ if and only if $\bar{\mathcal{P}} = \pi_l.\bar{\mathcal{P}} :: \bar{\mathcal{P}}_0$.
- (iii) If $d(l, \bar{\mathcal{M}}) = (u, v)$ with $(u, v) \neq (0, 1)$ then $\exists \bar{\mathcal{M}}'. \bar{\mathcal{M}} \rightarrow^* \bar{\mathcal{M}}'$ with $d(l, \bar{\mathcal{M}}') = (u', v')$ such that $u' < u$ and $v' < v$.

The first two properties are immediate. We prove the third one. First, notice that if $d(l, \bar{\mathcal{M}}) \neq \perp$, then $\exists \bar{\mathcal{M}}'. \bar{\mathcal{M}} \rightarrow \bar{\mathcal{M}}'$. We reason by cases on the reduction rule used.

AM.NIL We have $\bar{\mathcal{P}} = \mathbf{0} :: \bar{\mathcal{Q}} = \bar{\mathcal{P}}_0 :: \bar{\mathbf{E}}[\pi_l.\bar{\mathcal{P}}] :: \bar{\mathcal{P}}_1$. $\bar{\mathcal{M}}$ reduces to $\bar{\mathcal{M}}'$ with $\bar{\mathcal{P}}' = \bar{\mathcal{Q}} = \bar{\mathcal{P}}'_0 :: \bar{\mathbf{E}}[\pi_l.\bar{\mathcal{P}}] :: \bar{\mathcal{P}}_1$. We deduce that $u' = u - 1 < u$.

AM.PAR We have $\bar{P} = (\bar{P} \mid \bar{Q}) :: \bar{Q} = \bar{P}_0 :: \bar{\mathbf{E}}[\pi_l.\bar{P}] :: \bar{P}_1$. $\bar{\mathcal{M}}$ reduces to $\bar{\mathcal{M}}'$ with $\bar{P}' = \bar{P} :: \bar{Q} :: \bar{Q}$. If $\bar{P}_0 \neq []$ then $u' = u - \text{size}(\bar{Q}) < u'$. If $\bar{P}_0 = []$ then we have either $\bar{\mathbf{E}}[\pi_l.\bar{P}] = \bar{\mathbf{E}}'[\pi_l.\bar{P}] \mid \bar{P}'$ or $\bar{\mathbf{E}}[\pi_l.\bar{P}] = \bar{P}' \mid \bar{\mathbf{E}}'[\pi_l.\bar{P}]$. In the first case $v' = v - \text{size}(\bar{P}') < v$, in the second $u' = u - \text{size}(\bar{P}')$.

AM.NEW Similar to the previous case, using the fact that $|P| = |P\{y/x\}|$.

AM.PUSHRECEIVER, AM.PUSHMESSAGE, AM.PUSHREPRECEIVER These three cases are identical. We have $\bar{P} = \pi_l'.\bar{P}' :: \bar{Q} = \bar{P}_0 :: \bar{\mathbf{E}}[\pi_l.\bar{P}] :: \bar{P}_1$. Necessarily, $\bar{P}_0 \neq []$. $\bar{P}' = \bar{Q} = \bar{P}'_0 :: \bar{\mathbf{E}}[\pi_l.\bar{P}] :: \bar{P}_1$. We deduce that $u' = u - 1 < u$.

AM.COM1, AM.COM2, AM.RCOM2 These three cases are similar to the previous case.

AM.RCOM1 In this case, $\bar{P} = *x?(y).\bar{P}' :: \bar{Q}$. After a finite number of reductions (corresponding to the number of terms in $\bar{\mathcal{H}}(x)$), we have $\bar{\mathcal{M}} \rightarrow^* \bar{\mathcal{M}}'$ with $\bar{P}' = \bar{Q}$. We conclude as for the other communication rules.

Lemma 14. *Let us suppose $\bar{\mathcal{M}} = \langle \bar{P}, \bar{\mathcal{H}} \rangle$*

- (i) *If $l \in L(\bar{\mathcal{M}})$ then we have $\bar{\mathcal{M}} \rightarrow^* \bar{\mathcal{M}}' = \langle \bar{P}', \bar{\mathcal{H}}' \rangle$ for some $\bar{\mathcal{M}}'$ and one of these two conditions holds:*
 - $l \notin \text{lab}(\bar{\mathcal{M}}')$.
 - *There is x such that $\bar{\mathcal{H}}(x)$ is of the form $\bar{Q} :: \pi_l.\bar{P} :: \bar{Q}'$.*
- (ii) *If for some y , $\bar{\mathcal{H}}(y) = \bar{Q} :: \pi_l.\bar{P} :: \bar{Q}'$ then for all $\bar{\mathcal{M}}'$ such that $\bar{\mathcal{M}} \rightarrow^* \bar{\mathcal{M}}' = \langle \bar{P}', \bar{\mathcal{H}}' \rangle$ one of these two conditions holds:*
 - $l \notin \text{lab}(\bar{\mathcal{M}}')$.
 - $\bar{\mathcal{H}}'(y) = \bar{Q}_0 :: \pi_l.\bar{P} :: \bar{Q}'_0$ for some \bar{Q}_0, \bar{Q}'_0 with $|\bar{Q}_0| \leq |\bar{Q}|$.
- (iii) *If for some y , $\bar{\mathcal{H}}(y) = \bar{Q} :: \pi_l.\bar{P} :: \bar{Q}'$ and $l \in L(\bar{\mathcal{M}})$ then there is $\bar{\mathcal{M}}'$ such that $\bar{\mathcal{M}} \rightarrow^* \bar{\mathcal{M}}' = \langle \bar{P}', \bar{\mathcal{H}}' \rangle$ and one of these two conditions holds*
 - $l \notin \text{lab}(\bar{\mathcal{M}}')$.
 - $\bar{\mathcal{H}}'(y) = \bar{Q}_0 :: \pi_l.\bar{P} :: \bar{Q}'_0$ for some \bar{Q}_0, \bar{Q}'_0 with $|\bar{Q}_0| < |\bar{Q}|$.
- (iv) *If for some x , $\bar{\mathcal{H}}(x) = \bar{Q} :: \pi_l.\bar{P} :: \bar{Q}'$ then there is $\bar{\mathcal{M}}'$ such that $\bar{\mathcal{M}} \rightarrow^* \bar{\mathcal{M}}'$ and $\forall \bar{\mathcal{M}}''. \bar{\mathcal{M}}' \rightarrow^* \bar{\mathcal{M}}'' \implies l \notin L(\bar{\mathcal{M}}'')$.*

Proof. (i) If $l \in L(\bar{\mathcal{M}})$, by Lemma 12 we have either $\bar{P} = \bar{P}_0 :: \bar{\mathbf{E}}[\pi_l.\bar{P}] :: \bar{P}_1$ for some context $\bar{\mathbf{E}}$ or $\bar{\mathcal{H}}(x) = \bar{Q} :: \pi_l.\bar{P} :: \bar{Q}'$ for some x, \bar{Q} and \bar{Q}' . In the second case, the lemma is proved with $\bar{\mathcal{M}}' = \bar{\mathcal{M}}$. In the first case, we apply Lemma 13 and we have $\bar{\mathcal{M}} \rightarrow^* \bar{\mathcal{M}}_0$ with $\bar{\mathcal{M}}_0 = \langle \pi_l.\bar{P} :: \bar{P}'_0, \bar{\mathcal{H}}_0 \rangle$. We then have $\bar{\mathcal{M}}_0 \rightarrow \bar{\mathcal{M}}'$. If this reduction is inferred by a communication rule, then $l \notin \text{lab}(\bar{\mathcal{M}}')$, otherwise it is inferred by a pushing rule and $\bar{\mathcal{M}}' = \langle \bar{P}', \bar{\mathcal{H}}' \rangle$ with $\bar{\mathcal{H}}' = \bar{\mathcal{H}}_0 :: \pi_l.\bar{P}$.

(ii) We prove it by induction on n in $\bar{\mathcal{M}} \rightarrow^n \bar{\mathcal{M}}'$. If $n = 0$, the result is immediate. If $n \geq 0$, we have $\bar{\mathcal{M}} \rightarrow \bar{\mathcal{M}}' \rightarrow^n \bar{\mathcal{M}}''$. Consider now all the possible rules to derive $\bar{\mathcal{M}} \rightarrow \bar{\mathcal{M}}'$. In the rules AM.NIL, AM.PAR and AM.RES, we have $\bar{\mathcal{H}}(y) = \bar{\mathcal{H}}(x)$. In the rules AM.PUSHMESSAGE, AM.PUSHRECEIVER and AM.PUSHREPRECEIVER, if $x \neq y$, we have $\bar{\mathcal{H}}'(y) = \bar{\mathcal{H}}(y)$. Otherwise,

$\overline{\mathcal{H}}'(y) = \overline{\mathcal{H}}(y) :: P$. In the rules AM.COM1, AM.COM2, AM.RCOM1 and AM.RCOM2, if $x = y$ then $\overline{\mathcal{H}}'(y) = \overline{\mathcal{H}}(y)$. If $x \neq y$, either $l \notin \text{lab}(\overline{\mathcal{M}}')$ or $\overline{\mathcal{H}}'(y) = \overline{\mathcal{Q}}_0 :: \pi'_l.\overline{P} :: \overline{\mathcal{Q}}'_0$ with $|\overline{\mathcal{Q}}_0| \leq |\overline{\mathcal{Q}}|$. In the first case, the result comes from Lemma 10. In the second case, we conclude by the induction hypothesis.

- (iii) We have $\{l, l'\} \in LA(\overline{\mathcal{M}})$ for some l' . Then by Lemma 12 we have either $\overline{P} = \overline{P}_0 :: \overline{\mathbf{E}}[\pi'_{l'}.\overline{P}'] :: \overline{P}'_0$ for some context $\overline{\mathbf{E}}$ or $\overline{\mathcal{H}}(x) = \overline{\mathcal{Q}} :: \pi'_{l'}.\overline{P}' :: \overline{\mathcal{Q}}'$ for some x , $\overline{\mathcal{Q}}$ and $\overline{\mathcal{Q}}'$. We can show that the second case contradicts the hypothesis $\{l, l'\} \in LA(\overline{\mathcal{M}})$, so only the first case is possible. By Lemma 13, we have $\overline{\mathcal{M}} \rightarrow^* \overline{\mathcal{M}}_0$ with $\overline{\mathcal{M}}_0 = \langle \pi'_{l'}.\overline{P}' :: \overline{P}'_0, \overline{\mathcal{H}}_0 \rangle$. By (ii), we have $l \notin \text{lab}(\overline{\mathcal{M}}_0)$ or $\overline{\mathcal{H}}_0(x) = \overline{\mathcal{Q}}_0 :: \pi_l.\overline{P} :: \overline{\mathcal{Q}}'_0$ with $|\overline{\mathcal{Q}}_0| \leq |\overline{\mathcal{Q}}|$. In the first case, the result is proven with $\overline{\mathcal{M}}' = \overline{\mathcal{M}}_0$. In the second case, we can show that $\{l, l'\} \in LA(\overline{\mathcal{M}}_0)$ and that $\overline{\mathcal{M}}_0 \rightarrow \overline{\mathcal{M}}_1$ for some $\overline{\mathcal{M}}_1$, where \rightarrow is inferred by a communication rule on y . If $|\overline{\mathcal{Q}}_0| = 0$, then $l \notin \text{lab}(\overline{\mathcal{M}}_1)$. If $|\overline{\mathcal{Q}}_0| > 0$, then $\overline{\mathcal{M}}_1 = \langle \overline{P}_1, \overline{\mathcal{H}}_1 \rangle$ with $\overline{\mathcal{H}}_1 = \overline{\mathcal{Q}}_1 :: \pi_l.\overline{P} :: \overline{\mathcal{Q}}'_1$ with $|\overline{\mathcal{Q}}_1| < |\overline{\mathcal{Q}}_0|$. In these two cases, the result is proven with $\overline{\mathcal{M}}' = \overline{\mathcal{M}}_1$.
- (iv) We define $I(k)$:

$$\begin{aligned} I(k) &\equiv \overline{\mathcal{M}} = \langle \overline{P}, \overline{\mathcal{H}} \rangle \wedge \exists x. \overline{\mathcal{H}}(x) = \overline{\mathcal{Q}} :: \pi_l.\overline{P} :: \overline{\mathcal{Q}}' \wedge |\overline{\mathcal{Q}}| = k \\ &\implies \exists \overline{\mathcal{M}}'. \overline{\mathcal{M}} \rightarrow^* \overline{\mathcal{M}}' \wedge (\overline{\mathcal{M}}' \rightarrow^* \overline{\mathcal{M}}'' \implies l \notin L(\overline{\mathcal{M}}'')) \end{aligned}$$

We show $\forall k \geq 0$. $I(k)$ by strong induction on k . We distinguish two cases. If $\forall \overline{\mathcal{M}}'. \overline{\mathcal{M}} \rightarrow^* \overline{\mathcal{M}}' \implies l \notin L(\overline{\mathcal{M}}')$, then the theorem holds trivially. Otherwise, we have $\exists \overline{\mathcal{M}}' = \langle \overline{P}', \overline{\mathcal{H}}' \rangle$. $\overline{\mathcal{M}} \rightarrow^* \overline{\mathcal{M}}' \wedge l \in L(\overline{\mathcal{M}}')$. By (ii), we have $\overline{\mathcal{H}}'(x) = \overline{\mathcal{Q}}_0 :: \pi_l.\overline{P} :: \overline{\mathcal{Q}}'_0$ with $|\overline{\mathcal{Q}}_0| \leq |\overline{\mathcal{Q}}|$. We can apply (iii). In the first case, the property is proven. In the second case, we conclude by the induction hypothesis.

Proof (Proof of Theorem 1). Let us consider an infinite execution

$$\overline{\mathcal{M}}_1 \rightarrow \dots \rightarrow \overline{\mathcal{M}}_n \rightarrow \dots$$

where $\overline{\mathcal{M}}_k = \langle \overline{P}_k, \overline{\mathcal{H}}_k \rangle$ for $k \in \mathbb{N}^*$. Let $i \in \mathbb{N}^*$ and $l \in L(\overline{\mathcal{M}}_i)$. We want to show that there is a j such that $\forall i > j$. $l \notin L(\overline{\mathcal{M}}_i)$.

According to Lemma 14 (i), we have some $\overline{\mathcal{M}}$ such that $\overline{\mathcal{M}}_i \rightarrow^* \overline{\mathcal{M}}$. Moreover, one of these two conditions hold:

- $l \notin \text{lab}(\overline{\mathcal{M}})$.
- $\overline{\mathcal{H}}(x) = \overline{\mathcal{Q}} :: \pi_l.\overline{P} :: \overline{\mathcal{Q}}'$ for some x .

By lemma 9, there is $j \geq i$ and a bijective substitution of labels σ such that $\overline{\mathcal{M}}\sigma = \overline{\mathcal{M}}_j$, and in particular $\sigma(l) = l$. The same two conditions hold for $\overline{\mathcal{M}}_j$:

- $l \notin \text{lab}(\overline{\mathcal{M}}_j)$.
- $\overline{\mathcal{H}}_j(x) = \overline{\mathcal{Q}} :: \pi_l.\overline{P} :: \overline{\mathcal{Q}}'$ for some x .

In the first case, we know from Lemma 10 that $\forall i \geq j. l \notin \text{lab}(\overline{\mathcal{M}}_i)$. We deduce $\forall i \geq j. l \notin L(\overline{\mathcal{M}}_i)$. In the second case, we have by Lemma 14(iv):

$$\exists \overline{\mathcal{M}}'. \overline{\mathcal{M}} \rightarrow^* \overline{\mathcal{M}}' \text{ and } (\forall \overline{\mathcal{M}}''. \overline{\mathcal{M}}' \rightarrow^* \overline{\mathcal{M}}'' \implies l \notin L(\overline{\mathcal{M}}'')) \quad (1)$$

By Lemma 9, we have $j' \geq j$ and a bijective substitution of labels σ such that $\overline{\mathcal{M}}' \sigma = \overline{\mathcal{M}}_{j'}$ and $\sigma(l) = l$. Now suppose $k \geq j'$, we have $\overline{\mathcal{M}}_{j'} \rightarrow^* \overline{\mathcal{M}}_k$. We can show that $\overline{\mathcal{M}}_{j'} \sigma^{-1} \rightarrow^* \overline{\mathcal{M}}_k \sigma^{-1}$. By (1), we deduce that $l \notin L(\overline{\mathcal{M}}_k \sigma^{-1})$ and that $l \notin L(\overline{\mathcal{M}}_k)$. We showed that $\forall k \geq j'. l \notin L(\overline{\mathcal{M}}_k)$.

A.3 Correctness

Proof (Proof of Lemma 7). We have $\llbracket \overline{P} \rrbracket^r = \overline{P}$, and $LA(\llbracket \overline{P} \rrbracket) = LA(\llbracket \overline{P} \rrbracket^r) = LA(\overline{P})$. We conclude by Lemma 11.

The following definition is a function $\llbracket \cdot \rrbracket^{-1}$ from PAM terms to π -calculus processes. It will be used in the soundness theorem to build an execution of π -calculus processes starting from an execution of PAM terms such that every PAM term corresponds to an implementation of the corresponding π -calculus term.

The functions $\llbracket \cdot \rrbracket^{-1}$ and $\llbracket \cdot \rrbracket^r$ (Definition 5) are very close. In particular, we have $\llbracket \mathcal{M} \rrbracket^{-1} \equiv \llbracket \mathcal{M} \rrbracket^r$ for any machine term \mathcal{M} . However, $\llbracket \cdot \rrbracket^r$ cannot be used to prove the soundness property because it does not satisfy the first property of Lemma 8.

Intuitively, this comes from the fact that $\llbracket P_1 \mid P_2 \rrbracket$ and $\llbracket P_2 \mid P_1 \rrbracket$ may behave differently in general. In particular, to have $\llbracket \llbracket \mathcal{M} \rrbracket^{-1} \rrbracket \rightsquigarrow^* \mathcal{M}$, processes in \mathcal{M} should be arranged in a proper order (with respect to parallel composition).

Definition 10. We define a function $\llbracket \cdot \rrbracket^{-1}$ from \mathbf{M} to \mathbf{P} . We suppose $\mathcal{M} = \langle \mathcal{P}_p, \mathcal{H}, \mathcal{N} \rangle$ with $\mathcal{H} = \{x_i \rightarrow \mathcal{P}_i\}_{i \in 0..p-1}$. We define $(P_i)_{i \in \{0..n\}}$ such that

$$\mathcal{P}_0 :: \dots :: \mathcal{P}_p = P_0 :: \dots :: P_n$$

Then we define $(Q_i)_{i \in 0..n}$ as follows:

$$\begin{aligned} Q_0 &= P_0 \\ Q_p &= Q_{p-1} \mid P_{n-p+1}, p \in \{1..n\} \end{aligned}$$

Finally, if $\{x_1, \dots, x_k\} = \mathcal{N}$, we define:

$$\llbracket \mathcal{M} \rrbracket^{-1} = \nu a_1 \dots a_k. Q_n \{a_1/x_1, \dots, a_k/x_k\}$$

where the names a_i for $i \in \{1..k\}$ are all distinct and different from names in Q_n . We suppose also that $\forall i \in \{1..k-1\}. \text{freshn}(\{x_1, \dots, x_i\}) = x_{i+i}$. This assumption is justified by Remark 1. Names in \mathcal{N} have been generated by rule AM.NEW, hence, there is always an ordering x_1, \dots, x_k of \mathcal{N} such that $\text{freshn}(\{x_1, \dots, x_i\}) = x_{i+i}$. Alternatively, we could define $\mathcal{M} \preceq P$ as $\llbracket P \rrbracket \rightsquigarrow^* =_{\alpha} \mathcal{M}$ where $=_{\alpha}$ identifies machine terms that differ only in their generated names.

Proof (Proof of Lemma 8). We suppose $\mathcal{M} = \langle \mathcal{P}_p, \mathcal{H}, \mathcal{N} \rangle$ with $\mathcal{H} = \{x_i \rightarrow \mathcal{P}_i\}_{\{i \in 0..p-1\}}$. We reuse the notations of definition 10. We show that:

$$\begin{aligned} \llbracket \mathcal{M} \rrbracket^{-1} &\rightsquigarrow^* \langle Q_n, \mathcal{H}_{\square}, \mathcal{N} \rangle \\ &\rightsquigarrow^* \langle \mathcal{P}_0 :: \dots :: \mathcal{P}_p, \mathcal{H}_{\square}, \mathcal{N} \rangle \\ &\rightsquigarrow^* \langle \mathcal{P}_p, \mathcal{H}, \mathcal{N} \rangle = \mathcal{M} \end{aligned}$$

Hence, we have $\llbracket \mathcal{M} \rrbracket^{-1} \rightsquigarrow^* \mathcal{M}$ and $\mathcal{M} \preceq \llbracket \mathcal{M} \rrbracket^{-1}$. We detail these three steps:

– We have

$$\llbracket \mathcal{M} \rrbracket^{-1} = \langle \nu a_1 \dots a_k. Q_n \{a_1/x_1, \dots, a_k/x_k\}, \mathcal{H}_{\square}, \emptyset \rangle \rightsquigarrow^* \langle Q_n, \mathcal{H}_{\square}, \mathcal{N} \rangle$$

by successive applications of rule AM.NEW, and the condition over generated names $\text{fresh}(\{x_1, \dots, x_i\}) = x_{i+i}$.

– We have

$$\langle Q_n, \mathcal{H}_{\square}, \emptyset \rangle \rightsquigarrow^* \langle P_0 :: \dots :: P_n, \mathcal{H}_{\square}, \mathcal{N} \rangle$$

by successive applications of rule AM.PAR.

– We have

$$\langle P_0 :: \dots :: P_n, \mathcal{H}_{\square}, \mathcal{N} \rangle \rightsquigarrow^* \langle \mathcal{P}_p, \mathcal{H}, \mathcal{N} \rangle$$

by successive applications of pushing rules (AM.PUSHRECEIVER, AM.PUSHMESSAGE and AM.PUSHREPRECEIVER).

The second property $\llbracket P \rrbracket^{-1} = P$ is immediate by definition of $\llbracket \cdot \rrbracket^{-1}$ and $\llbracket \cdot \rrbracket$. For the two remaining properties, we notice that $\llbracket \mathcal{M} \rrbracket^{-1} \equiv \llbracket \mathcal{M} \rrbracket^r$ so it is enough to prove the result for $\llbracket \cdot \rrbracket^r$ instead of $\llbracket \cdot \rrbracket^{-1}$. The proof is by cases on the definitions of \rightsquigarrow and \mapsto and has been given in [19].