

On the Relative Expressive Power of Asynchronous Communication Primitives*

Daniele Gorla

Dipartimento di Informatica
Università di Roma “La Sapienza”

Abstract. In this paper, we study eight asynchronous communication primitives, arising from the combination of three features: *arity* (monadic vs polyadic data), *communication medium* (message passing vs shared dataspace) and *pattern-matching*. Each primitive has been already used in at least one language appeared in literature; however, to uniformly reason on such primitives, we plugged them in a common framework inspired by the asynchronous π -calculus. By means of possibility/impossibility of ‘reasonable’ encodings, we compare every pair of primitives to obtain a hierarchy of languages based on their relative expressive power.

1 Introduction

In the last 25 years, several languages and formalisms for distributed and concurrent systems appeared in literature. Some of them (e.g., CCS [18] and the π -calculus [23]) are mostly mathematical models, mainly used to formally reason on concurrent systems; other ones (e.g., LINDA [15]) are closer to actual programming languages and are mainly focused on issues like usability and flexibility. As a consequence, the former ones are usually very essential, while the latter ones provide more sophisticated and powerful programming constructs.

Despite their differences, there are, however, some basic features that are somewhat implemented in all these languages. Roughly speaking, these features can be described as the possibility of having different *execution threads* (or *processes*) that *run concurrently* by interacting via some form of *communication*. At least at a first glance, the last feature (i.e., the inter-process communication) has yielded the highest variety of proposals. These arose from the possibility of having synchronous/asynchronous primitives, monadic/polyadic data, first-order/higher-order values, dataspace-based/channel-based communication media, local/remote exchanges (whenever processes are explicitly distributed, like in [8, 11]), built-in pattern-matching mechanisms, point-to-point/broadcasting primitives, and so on. The aim of this work is to formally study some of these proposals and to organise them in a clear hierarchy, based on their expressive power. Hopefully, our results should help to understand the peculiarities of every communication primitive and, as a consequence, they could be exploited to choose the ‘right’ primitive when designing new languages and formalisms.

We focus on *asynchronous* communication primitives, since they are the most basic ones. Among the remaining features mentioned above, we focus on *arity of data*, *communication medium* and possibility of *pattern-matching*. The expressiveness of the omitted features has been already dealt with elsewhere [25, 11, 13]; we leave as a future

* This work has been partially supported by Project SENSORIA, founded by EU-IST Programme, contract number 016004.

work the integration of these results in our framework. Notice that we studied pattern-matching because it is nowadays becoming more and more important, especially in languages that deal with complex data like XML [1, 5, 9]. However, for the sake of simplicity, we consider here a very basic form of pattern-matching, that only checks for name equality while retrieving a datum; the formal study of more flexible and powerful mechanisms (e.g., those in [12]) is left for future work.

By combining the three features chosen, we obtain eight communication primitives that have been all already employed elsewhere, e.g. in [17, 4, 15, 8, 11, 9]. However, to uniformly reason on such primitives, we plugged them in a common framework inspired by the asynchronous π -calculus; we choose the π -calculus because nowadays it is one of the best-established workbenches for theoretical reasoning on concurrent systems. By following [26, 10, 22], we shall compare the resulting languages by means of their *relative expressive power*, i.e. we shall try to encode one in the other and study the properties of the encoding function. More precisely, we shall exploit possibility/impossibility of ‘reasonable’ encodings (as introduced in [22]) to compare every pair of primitives, thus obtaining a hierarchy of languages based on their relative expressive power.

Our results show that the communication paradigm underlying LINDA [15] (polyadic, dataspace-based and with pattern-matching) is at the top of the hierarchy; not incidentally, LINDA’s paradigm has been used in actual programming languages [3, 14]. On the opposite extreme, we have the communication paradigm used in Ambient [8] (monadic, dataspace-based but without pattern-matching). Such a paradigm is very simple but also very poor; indeed, Ambient’s expressive power mostly arises from the mobility primitives. Strictly in the middle, we find the asynchronous π -calculus (channel-based and without pattern-matching), in its monadic and polyadic version. This result stresses the fact that the π -calculus is a good compromise between expressiveness and simplicity. As a further contribution, we also prove that the polyadic π -calculus is strictly more expressive than the monadic one. A posteriori, this fact justifies the use of type-systems [19, 27, 24] to obtain a fragment of the former calculus that can be reasonably translated in the latter one.

This paper is organised as follows. In Section 2, we present a family of eight π -based asynchronous calculi arising from the combination of the three features studied. In Section 3, we present the criteria an encoding should satisfy to be a reasonable means for language comparison; there, we also sum-up the results of the paper, that are proved in Sections 4 and 5. We start with the encodability results and then we present the impossibility results, that are the main contribution of our work. Finally, in Section 6, we conclude the paper by also touching upon related work.

2 A Family of π -based Calculi

As we said in the Introduction, we shall assess the expressiveness of the communication primitives studied by putting them in a common framework, inspired by the asynchronous π -calculus. We assume two disjoint and countable sets: *names*, \mathcal{N} , ranged over by a, b, x, y, n, m, \dots , and *process variables*, \mathcal{X} , ranged over by X, Y, \dots . Notationally, when a name is used as a channel, we shall prefer letters a, b, c, \dots ; when a

name is used as an input variable, we shall prefer letters x, y, z, \dots ; to denote a generic name, we shall use letters n, m, \dots . The (parametric) syntax of our calculi is

$$P, Q, R ::= \mathbf{0} \mid OUT \mid IN.P \mid (\nu n)P \mid P|Q \\ \mid \text{if } n = m \text{ then } P \text{ else } Q \mid \text{rec } X.P \mid X$$

The different calculi will be obtained by plugging into this basic syntax a proper definition for input (IN) and output (OUT) actions. As usual, $\mathbf{0}$ and $P|Q$ denote the terminated process and the parallel composition of two processes, while $(\nu n)P$ restricts to P the visibility of n ; finally, $\text{if } n = m \text{ then } P \text{ else } Q$, $\text{rec } X.P$ and X are the standard constructs for conditional evolution, process definition and process invocation.

In this paper, we study the possible combinations of three features for asynchronous communications: *arity* (monadic vs. polyadic data), *communication medium* (channels vs. shared dataspace) and *pattern-matching*. As a result, we have a family of eight calculi, denoted as $\Pi_{a,m,p}$, whose generic element is denoted as $\pi_{\beta_1\beta_2\beta_3}$, where $\beta_i \in \{0,1\}$. Intuitively, $\beta_1 = 1$ iff we have polyadic data, $\beta_2 = 1$ iff we have channel-based communications and $\beta_3 = 1$ iff we have pattern-matching. Thus, the full syntax of every calculus is obtained from the following productions:

$$\begin{array}{lll} \pi_{000} : & P, Q, R ::= \dots & IN ::= (x) \quad OUT ::= \langle b \rangle \\ \pi_{001} : & P, Q, R ::= \dots & IN ::= (T) \quad OUT ::= \langle b \rangle \\ \pi_{010} : & P, Q, R ::= \dots & IN ::= a(x) \quad OUT ::= \bar{a}\langle b \rangle \\ \pi_{011} : & P, Q, R ::= \dots & IN ::= a(T) \quad OUT ::= \bar{a}\langle b \rangle \\ \pi_{100} : & P, Q, R ::= \dots & IN ::= (\tilde{x}) \quad OUT ::= \langle \tilde{b} \rangle \\ \pi_{101} : & P, Q, R ::= \dots & IN ::= (\tilde{T}) \quad OUT ::= \langle \tilde{b} \rangle \\ \pi_{110} : & P, Q, R ::= \dots & IN ::= a(\tilde{x}) \quad OUT ::= \bar{a}\langle \tilde{b} \rangle \\ \pi_{111} : & P, Q, R ::= \dots & IN ::= a(\tilde{T}) \quad OUT ::= \bar{a}\langle \tilde{b} \rangle \end{array}$$

where

$$T ::= x \mid \ulcorner n \urcorner \quad (\text{Template})$$

and $\tilde{}$ denotes a (possibly empty) sequence of elements of kind (whenever useful, we shall write a tuple $\tilde{}$ as the sequence of its elements, separated by a comma). Template fields of kind x are called *formal* and can be replaced by every name upon withdrawal of a datum; template fields of kind $\ulcorner n \urcorner$ are called *actual* and impose that the datum withdrawn contains exactly name n .

$\Pi_{a,m,p}$ can be easily ordered by language containment; in particular, $\pi_{\beta_1\beta_2\beta_3}$ can be seen as a sub-language of $\pi_{\beta'_1\beta'_2\beta'_3}$ if and only if, for every $i \in \{1, 2, 3\}$, it holds that $\beta_i \leq \beta'_i$. As an extremal example, consider π_{000} and π_{111} : monadic data are a particular case of polyadic data (all of length one); a shared dataspace can be modelled by letting all communications happen on the same global channel, say *ether*; finally, absence of pattern-matching can be obtained by only considering templates without actual fields.

Notice that π_{010} and π_{110} are very similar to the (monadic/polyadic) asynchronous π -calculus [17, 4]; π_{101} relies on the communication paradigm adopted in LINDA [15];

π_{000} and π_{100} rely on the communication paradigm adopted in the (monadic/polyadic) Ambient Calculus [8]; π_{001} and π_{011} rely on the communication paradigm adopted in LCKLAIM and CKLAIM [11], respectively; finally, π_{111} relies on the communication paradigm adopted, e.g., in μ KLAIM [11] or in semantic- π [9].

As usual, $a(\cdot \dots, x, \cdot \dots).P$ and $(\nu x)P$ bind x in P , while $\text{rec } X.P$ binds X in P . The corresponding notions of free and bound names of a process, $\text{FN}(P)$ and $\text{BN}(P)$, and of alpha-conversion, $=_\alpha$, are assumed. We let $\text{N}(P)$ denote $\text{FN}(P) \cup \text{BN}(P)$.

Operational semantics The operational semantics of the calculi is given by means of a *labelled transition system* (LTS) describing the actions a process can perform to evolve. Judgements take the form $P \xrightarrow{\alpha} P'$, meaning that P can become P' upon execution of α . *Labels* take the form

$$\alpha ::= \tau \mid a? \tilde{b} \mid (\nu \tilde{c})a! \tilde{b} \mid ? \tilde{b} \mid (\nu \tilde{c})! \tilde{b}$$

Traditionally, τ denotes an internal computation; $a? \tilde{b}$ and $(\nu \tilde{c})a! \tilde{b}$ denote the reception/sending of a sequence of names \tilde{b} along channel a ; when channels are not present (namely, in π_{-0-}), $? \tilde{b}$ and $(\nu \tilde{c})! \tilde{b}$ denote the withdrawal/emission of \tilde{b} from/in the shared dataspace. In $(\nu \tilde{c})a! \tilde{b}$ and $(\nu \tilde{c})! \tilde{b}$, some of the sent names, viz. \tilde{c} ($\subseteq \tilde{b}$), are restricted. Notationally, $(\nu \tilde{c})-! \tilde{b}$ stands for either $(\nu \tilde{c})a! \tilde{b}$ or $(\nu \tilde{c})! \tilde{b}$. As usual, $\text{BN}((\nu \tilde{c})-! \tilde{b}) \triangleq \tilde{c}$; $\text{FN}(\alpha)$ and $\text{N}(\alpha)$ are defined accordingly.

The LTS provides some rules shared by all the calculi; the different semantics are obtained from the axioms for input/output actions. The common rules, reported below, are an easy adaptation of an early-style LTS for the π -calculus; thus, we do not comment them and refer the interested reader to [23].

$$\begin{array}{c} \frac{P \xrightarrow{? \tilde{b}} P' \quad Q \xrightarrow{! \tilde{b}} Q'}{P \mid Q \xrightarrow{\tau} P' \mid Q'} \qquad \frac{P \xrightarrow{a? \tilde{b}} P' \quad Q \xrightarrow{a! \tilde{b}} Q'}{P \mid Q \xrightarrow{\tau} P' \mid Q'} \\ \\ \frac{P \xrightarrow{\alpha} P' \quad n \notin \text{N}(\alpha)}{(\nu n)P \xrightarrow{\alpha} (\nu n)P'} \qquad \frac{P \xrightarrow{(\nu \tilde{c})-! \tilde{b}} P' \quad n \in \text{FN}(\tilde{b}) \setminus \{-, \tilde{c}\}}{(\nu n)P \xrightarrow{(\nu n, \tilde{c})-! \tilde{b}} P'} \\ \\ \frac{P \xrightarrow{\alpha} P' \quad \text{BN}(\alpha) \cap \text{FN}(Q) = \emptyset}{P \mid Q \xrightarrow{\alpha} P' \mid Q} \qquad \frac{P \equiv P_1 \xrightarrow{\alpha} P_2 \equiv P'}{P \xrightarrow{\alpha} P'} \end{array}$$

The structural equivalence, \equiv , rearranges a process to let it evolve according to the rules of the LTS. Its defining axioms are the standard π -calculus' ones [23]:

$$\begin{array}{l} P \mid \mathbf{0} \equiv P \qquad P \mid Q \equiv Q \mid P \qquad P \mid (Q \mid R) \equiv (P \mid Q) \mid R \\ \\ \text{if } n = n \text{ then } P \text{ else } Q \equiv P \qquad \text{if } n = m \text{ then } P \text{ else } Q \equiv Q \text{ if } n \neq m \\ \\ P \equiv P' \text{ if } P =_\alpha P' \qquad (\nu n)\mathbf{0} \equiv \mathbf{0} \qquad (\nu n)(\nu m)P \equiv (\nu m)(\nu n)P \\ \\ P \mid (\nu n)Q \equiv (\nu n)(P \mid Q) \text{ if } n \notin \text{FN}(P) \qquad \text{rec } X.P \equiv P\{\text{rec } X.P/X\} \end{array}$$

To define the semantics for the basic actions of the various calculi, we must specify when a template matches a datum. Intuitively, this happens whenever both have the same length and corresponding fields match (i.e., $\lceil n \rceil$ matches n and x matches every name). This can be formalised via a partial function, called *pattern-matching* and written MATCH , that also returns a substitution σ ; the latter will be applied to the process that performed the input to replace template formal fields with the corresponding names of the datum retrieved. These intuitions are formalised by the following rules:

$$\begin{aligned} \text{MATCH}(\ ; \) &= \epsilon & \text{MATCH}(\lceil n \rceil; n) &= \epsilon & \text{MATCH}(x; n) &= \{n/x\} \\ \frac{\text{MATCH}(T; b) = \sigma_1 \quad \text{MATCH}(\tilde{T}; \tilde{b}) = \sigma_2}{\text{MATCH}(T, \tilde{T}; b, \tilde{b}) = \sigma_1 \circ \sigma_2} \end{aligned}$$

where ‘ ϵ ’ denotes the empty substitution and ‘ \circ ’ denotes substitution composition. Now, the operational rules for output/input actions in calculi π_{-0-} are

$$\langle \tilde{b} \rangle \xrightarrow{! \tilde{b}} \mathbf{0} \quad (\tilde{T}).P \xrightarrow{? \tilde{b}} P\sigma \quad \text{if } \text{MATCH}(\tilde{T}; \tilde{b}) = \sigma$$

and, similarly, the rules for calculi π_{-1-} are

$$\bar{a} \langle \tilde{b} \rangle \xrightarrow{a ! \tilde{b}} \mathbf{0} \quad a(\tilde{T}).P \xrightarrow{a ? \tilde{b}} P\sigma \quad \text{if } \text{MATCH}(\tilde{T}; \tilde{b}) = \sigma$$

Notation A substitution σ is a finite partial mapping of names for names; $P\sigma$ denotes the (capture avoiding) application of σ to P . As usual, we let \Rightarrow stand for $\xrightarrow{\tau}^*$ (i.e., the reflexive and transitive closure of $\xrightarrow{\tau}$) and $\xRightarrow{\alpha}$ stand for $\Rightarrow \xrightarrow{\alpha} \Rightarrow$. We shall write $P \xrightarrow{\alpha}$ to mean that there exists a process P' such that $P \xrightarrow{\alpha} P'$; a similar notation is adopted for $P \Rightarrow$ and $P \xRightarrow{\alpha}$. Moreover, we let ϕ range over visible actions (i.e. labels different from τ) and ρ to range over (possibly empty) sequences of visible actions. Formally, $\rho ::= \epsilon \mid \phi \cdot \rho$, where ‘ ϵ ’ denotes the empty sequence of actions and ‘ \cdot ’ represents concatenation; then, $N \xRightarrow{\epsilon}$ is defined as $N \Rightarrow$ and $N \xRightarrow{\phi \rho}$ is defined as $N \xRightarrow{\phi} \xRightarrow{\rho}$.

3 Quality of an Encoding and Overview of our Results

We now study the relative expressive power of the calculi in $\Pi_{a,m,p}$ by trying to encode one in another. Formally, an *encoding* $\llbracket \cdot \rrbracket$ is a function mapping terms of the source language into terms of the target language. As pointed out elsewhere [26, 10, 22], the relative expressive power of our calculi can be established by defining some criteria to evaluate the quality of the encodings or to prove impossibility results.

The main requirement, that we call *faithfulness*, is that the encoding must not change the semantics of a source term, i.e. it must preserve the observable behaviour of the term without introducing new behaviours. As very clearly discussed in [21], there are several ways to formalise this idea; we shall define it in the simplest possible way, by means of *barbs* and *divergence*.

Definition 1 (Barbs and Divergence). P offers a barb, written $P \Downarrow$, iff $P \xrightarrow{(\nu \bar{c})^{-1} \bar{b}}$. P diverges, written $P \Uparrow$, iff $P \xrightarrow{\tau} \omega$.

The idea is to identify a basic observable behaviour (or *barb*) for the languages considered and require that the encoding preserves and reflects it (i.e., the encoding should maintain all the original barbs without introducing new ones). In the setting of an asynchronous language [2, 6], a barb is the possibility of emitting some datum.¹ Since barb preservation and reflection alone are too weak, it is also required that the computations of a process correspond to the computations of its encoding, and vice versa; this property is usually known as *operational correspondence*. Barb preservation and operational correspondence together yield (*weak*) *barbed bisimulation* [20, 2] that, however, is insensitive to divergence (i.e., it can equate a term with an infinite computation and a term with only finite computations). In our setting, it is clearly undesirable to have an encoding that turns a terminating term into a divergent one, since this would change the behaviour of the source term. So, we need a further requirement stating that also divergence must be preserved and reflected by the encoding.

Finally, a good encoding cannot depend on the particular names involved in the source process, since we are dealing with a family of name-passing calculi; we call this property *name invariance*. Furthermore, the encoding should not decrease the degree of parallelism in favour of centralised entities that control the behaviour of the encoded term: if we can find some process behaviour that cannot be implemented in the target language with the same degree of distribution as in the source one, then surely the former language will be “weaker” than the latter one. We express this last property as *homomorphism w.r.t. ‘|’*.

To sum up, we consider an encoding as a ‘reasonable’ means to compare the expressive power of two languages if it enjoys all the properties discussed so far.

Definition 2 (Reasonable Encoding). An encoding $\llbracket \cdot \rrbracket$ is reasonable if it enjoys the following properties:

1. (homomorphism w.r.t. ‘|’): $\llbracket P_1 | P_2 \rrbracket \triangleq \llbracket P_1 \rrbracket | \llbracket P_2 \rrbracket$
2. (name invariance): $\llbracket P \sigma \rrbracket \triangleq \llbracket P \rrbracket \sigma$, for every permutation of names σ
3. (faithfulness): $P \Downarrow$ iff $\llbracket P \rrbracket \Downarrow$; $P \Uparrow$ iff $\llbracket P \rrbracket \Uparrow$
4. (operational correspondence):
 - (a) if $P \xrightarrow{\tau} P'$ then $\llbracket P \rrbracket \xrightarrow{\tau} \llbracket P' \rrbracket$
 - (b) if $\llbracket P \rrbracket \xrightarrow{\tau} Q$ then there exists a P' such that $P \Rightarrow P'$ and $Q \Rightarrow \llbracket P' \rrbracket$

The results of our paper are summarised in Table 1. It is worth noting that all the languages are Turing complete: it is easy to show that π_{000} can encode (by introducing divergence²) π_{001} that, in turn, is Turing complete (actually, the fragment without template formal fields suffices, see [6]). Moreover, notice that Definition 2(4).b is a weak form of correspondence; this makes our impossibility results stronger. However, for the

¹ We choose here a very weak form of barbs. This fact strengthens our impossibility results; on the other hand, our possibility results are not undermined by this choice, since they would also enjoy properties expressed in terms of more significant barbs, such as those in [2, 6].

² To study Turing completeness, the fact that the encoding introduces divergence is irrelevant.

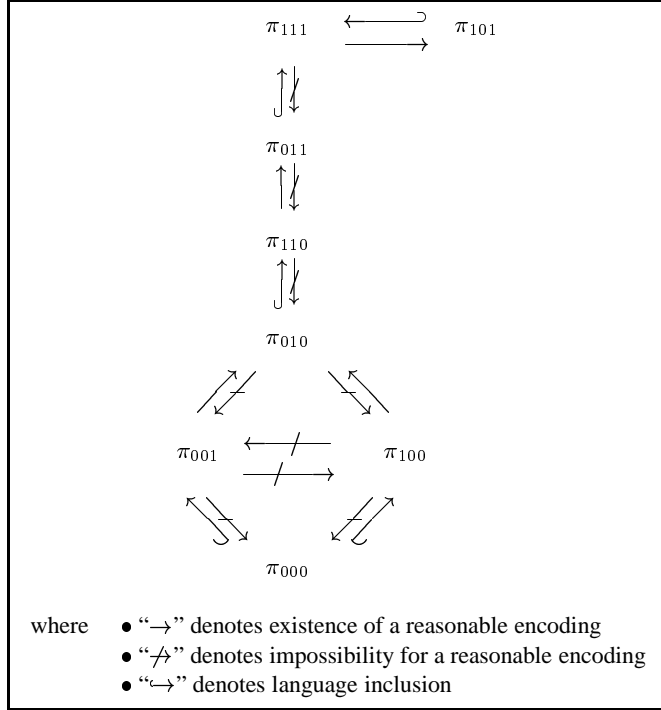


Table 1. Overview of the Results

encodability results, a better definition should keep into account every possible computation $\llbracket P \rrbracket \Rightarrow Q$. With this stronger property, the encodings provided in Sections 4.2 and 4.4 would not enjoy operational correspondence; we leave for future work the task of establishing whether encodings of π_{001} and π_{110} in π_{010} and π_{011} satisfying this stronger property exist or not.

4 Encodability Results

We start with the positive results, i.e. the “ \rightarrow ” arrows of Table 1. In all the cases, we shall describe only the translation of the input and output actions; the remaining operators will be translated homomorphically (this trivially satisfies Definition 2(1)). Moreover, in what follows we are going to prove only that the encodings do not introduce divergence; preservation of divergence is a trivial consequence of Definition 2(4).a; Definition 2(2) and barb preservation/reflection will hold by construction of the encodings; Definition 2(4) can be routinely proved.

4.1 An Encoding of π_{111} in π_{101}

The only feature of π_{111} not present in π_{101} is the possibility of specifying the name of a channel where the exchange happens. However, thanks to pattern-matching, this

feature can be very easily encoded in π_{101} : it suffices to impose that the first name of every datum represents the name of the channel where the interaction is scheduled and that every template starts with the corresponding actual field. This discipline is rendered by the following encoding:

$$\begin{aligned} \llbracket \bar{a}\langle \tilde{b} \rangle \rrbracket &\triangleq \langle a, \tilde{b} \rangle \\ \llbracket a(\tilde{T}).P \rrbracket &\triangleq (\tau a^\top, \tilde{T}).\llbracket P \rrbracket \end{aligned}$$

Proposition 1. *The encoding $\llbracket \cdot \rrbracket : \pi_{111} \longrightarrow \pi_{101}$ is reasonable.*

Proof. Definition 2(2) holds by construction. Definition 2(4).b can be proved as a stronger claim: if $\llbracket P \rrbracket \xrightarrow{\tau} Q$, then $Q \equiv \llbracket P' \rrbracket$ and $P \xrightarrow{\tau} P'$ (this result, like Definition 2(4).a, is proved by an easy induction over the shortest inference for $\xrightarrow{\tau}$). Definition 2(3) holds easily; just notice that the stronger formulation of operational correspondence mentioned above implies that the encoding cannot introduce divergence. \square

4.2 An Encoding of π_{001} in π_{010}

We now have to translate the monadic pattern-matching of π_{001} into the channel-based exchanges of π_{010} . This would have been an easy task, if only actual fields occurred in templates: indeed, $\langle b \rangle$ would have been translated in $\bar{b}\langle b \rangle$ and, correspondingly, $(\tau b^\top).P$ would have been translated in $b(y).\llbracket P \rrbracket$, for y fresh. This encoding, however, does not work well when trying to translate $(x).P$.

Thus, $\langle b \rangle$ in π_{001} should correspond to two outputs in π_{010} : one over b , to mimic name matching as described above, and one over a fresh and reserved channel `ether`, to enable inputs with formal fields. Symmetrically, an input action is translated in two successive inputs: the first one from `ether` and the second one from the received value, if we are translating an input with a formal field, and vice versa, otherwise. For example, $\langle b \rangle \mid \langle c \rangle \mid (\tau c^\top).P$ is translated to $\bar{b}\langle b \rangle \mid \overline{\text{ether}}\langle b \rangle \mid \bar{c}\langle c \rangle \mid \overline{\text{ether}}\langle c \rangle \mid c(y).\text{ether}(z).\llbracket P \rrbracket$. We believe that this encoding is reasonable, but proving that it does not introduce divergence is hard because of the possible interferences between parallel components (e.g., the above example could evolve in $\bar{b}\langle b \rangle \mid \overline{\text{ether}}\langle c \rangle \mid \llbracket P \rrbracket$, by performing a communication between $c(y)$ and $\bar{c}\langle c \rangle$ and between `ether`(z) and $\overline{\text{ether}}\langle b \rangle$).

The problem is that the two outputs in the translation of $\langle b \rangle$ are totally unrelated. This can be fixed by associating every output with a restricted name and by using such a name to reduce the effects of interferences. Formally,

$$\begin{aligned} \llbracket \langle b \rangle \rrbracket &\triangleq (\nu n)(\overline{\text{ether}}\langle n \rangle \mid \bar{b}\langle n \rangle \mid \bar{n}\langle b \rangle) \\ \llbracket (x).P \rrbracket &\triangleq \text{ether}(y).y(x).x(z).\text{if } y = z \text{ then } \llbracket P \rrbracket \\ \llbracket (\tau b^\top).P \rrbracket &\triangleq b(y).y(y').\text{ether}(z).\text{if } y = z \text{ then } \llbracket P \rrbracket \end{aligned}$$

for n, y, y' and z fresh names. Clearly, this solution does not rule out interferences; it simply blocks interfering processes. This suffices to make the proof of reasonableness easier; to this aim, the key result is the following Lemma.

Lemma 1. *Let κ be the number of top-level outputs in P . If $\llbracket P \rrbracket \xrightarrow{\tau}^{3\kappa+1} Q$, then there exists a P' such that $P \xrightarrow{\tau} P'$ and $\llbracket P \rrbracket \xrightarrow{\tau}^3 \llbracket P' \rrbracket \Rightarrow Q$.*

Proposition 2. *The encoding $\llbracket \cdot \rrbracket : \pi_{001} \longrightarrow \pi_{010}$ is reasonable.*

Proof. We only prove that $\llbracket \cdot \rrbracket$ does not introduce divergence; the other requirements are simple. Assume that $\llbracket P \rrbracket \uparrow$, i.e. there exists an infinite computation $\llbracket P \rrbracket \triangleq Q_0 \xrightarrow{\tau} Q_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} Q_{3\kappa+1} \xrightarrow{\tau} \dots$. By Lemma 1, there exists a P' such that $\llbracket P \rrbracket \xrightarrow{\tau}^3 \llbracket P' \rrbracket \Rightarrow Q_{3\kappa+1}$ and $P \xrightarrow{\tau} P'$. But $\llbracket P' \rrbracket$ is still divergent; indeed, $\llbracket P' \rrbracket \Rightarrow Q_{3\kappa+1} \xrightarrow{\tau} \dots$. By iterating this reasoning, we can build up a divergent computation for P , i.e. $P \xrightarrow{\tau} P' \xrightarrow{\tau} \dots$; hence, $\llbracket \cdot \rrbracket$ does not introduce divergence. \square

4.3 An Encoding of π_{100} in π_{010}

The only feature of π_{100} is that it can check the arity of a datum before retrieving it (see the definition of function MATCH). This, however, can be mimicked by the channel-based communication of π_{010} . Indeed, we assume a (reserved) channel for every possible arity: a datum of arity k will be represented as an output over channel k ; an input of arity k will be represented as an input from k ; a communication over k in π_{010} can happen if and only if pattern-matching succeeds in π_{100} ; finally, the exchanged datum is a restricted name that will be used in the actual data exchange.

The encoding assumes that $0, 1, \dots, k, \dots$ are fresh and reserved names; then

$$\begin{aligned} \llbracket \langle b_1, \dots, b_k \rangle \rrbracket &\triangleq (\nu n)(\bar{k}\langle n \mid n(y).(\bar{y}\langle b_1 \mid n(y).(\bar{y}\langle b_2 \mid \\ &\hspace{15em} n(y).(\dots n(y).\bar{y}\langle b_k \rangle) \dots) \rangle) \dots) \\ \llbracket (x_1, \dots, x_k).P \rrbracket &\triangleq k(x).(\nu m)(\bar{x}\langle m \mid m(x_1).(\bar{x}\langle m \mid \\ &\hspace{15em} m(x_2).(\dots (\bar{x}\langle m \mid m(x_k).\llbracket P \rrbracket) \dots) \rangle) \dots) \end{aligned}$$

for n, m, x and y fresh names. The datum emitted over k (viz. n) is used as a ‘‘synchroniser’’, to keep the order of the transmitted data and force the right name-to-variable association. The actual exchange takes place over a restricted channel created by the receiver (viz. m) and transmitted along n as an ack to the sender.

Like before, reasonableness of this encoding can be proved by using the following Lemma. Moreover, notice that for this encoding the stronger version of Definition 2(4).b holds: if $\llbracket P \rrbracket \Rightarrow Q$, then there is a P' such that $P \Rightarrow P'$ and $Q \Rightarrow \llbracket P' \rrbracket$.

Lemma 2. *Let κ and λ be the number and the maximum arity of top-level outputs in P , respectively. If $\llbracket P \rrbracket \xrightarrow{\tau}^{\kappa(2\lambda+1)} Q$, then there exists a P' such that $P \xrightarrow{\tau} P'$ and $\llbracket P \rrbracket \xrightarrow{\tau}^{2h+1} \llbracket P' \rrbracket \Rightarrow Q$, for $0 \leq h \leq \lambda$.*

Proposition 3. *The encoding $\llbracket \cdot \rrbracket : \pi_{100} \longrightarrow \pi_{010}$ is reasonable.*

4.4 An Encoding of π_{110} in π_{011}

In π_{110} , a communication succeeds if (and only if) a datum of a proper length is present over the channel specified by the inputting process. So, two kinds of information are

atomically verified: the length of the message and the channel where it should be transmitted. This can be mimicked in π_{011} by having one fresh and reserved name for every length (say, $0, 1, \dots, k, \dots$); a k -ary input from a is then translated into a process starting with $a(\ulcorner k \urcorner)$ and, correspondingly, a k -ary output on a is translated into a process offering k at a . Once this communication took place, we are sure that a k -ary datum is available on a ; we then proceed similarly to Section 4.3 for the actual data exchange: a new channel n is made available on a , to maintain the order of messages, while a new channel m is sent back on n to transmit the datum name by name.

However, we need to enhance the encoding of Section 4.3 to avoid interferences due to the fact that the existence of a k -ary output and the acquisition of the new name for the actual exchange are not atomic here. Indeed, the translation of $\bar{a}(b_1, b_2) \mid \bar{a}(c_1, c_2, c_3) \mid a(x_1, x_2).P \mid a(x_1, x_2, x_3).Q$ can originate interferences that can lead to divergence. Thus, like in Section 4.2, we shall verify at the end of the data exchange the consistency of the exchange, i.e. that a k -ary data has really been retrieved. To this aim, let end be another fresh and reserved name; then

$$\begin{aligned} \llbracket \bar{a}(b_1, \dots, b_k) \rrbracket &\triangleq \bar{a}(k) \mid (\nu n)(\bar{a}(n) \mid n(y).(\bar{y}(b_1) \mid \dots \mid n(y).(\bar{y}(b_k) \mid \\ &\quad n(y).\bar{y}(\text{end})))))) \\ \llbracket a(x_1, \dots, x_k).P \rrbracket &\triangleq a(\ulcorner k \urcorner).a(x).(\nu m)(\bar{x}(m) \mid m(x_1).(\dots \mid (\bar{x}(m) \mid m(x_k). \\ &\quad (\bar{x}(m) \mid m(\ulcorner \text{end} \urcorner).\llbracket P \rrbracket)))))) \end{aligned}$$

for n, m, x and y fresh names. Reasonableness of this encoding can be proved like in Proposition 2, as a consequence of the following Lemma.

Lemma 3. *Let κ and λ be the number and the maximum arity of top-level outputs in P , respectively. If $\llbracket P \rrbracket \xrightarrow{\tau, \kappa(2\lambda+3)+1} Q$, then there exists a P' such that $P \xrightarrow{\tau} P'$ and $\llbracket P \rrbracket \xrightarrow{\tau, 2h+4} \llbracket P' \rrbracket \Rightarrow Q$, for $0 \leq h \leq \lambda$.*

Proposition 4. *The encoding $\llbracket \cdot \rrbracket : \pi_{110} \longrightarrow \pi_{011}$ is reasonable.*

5 Impossibility Results

We now consider the impossibility results, i.e. the “ $\not\rightarrow$ ” arrows of Table 1, that are the main technical contribution of this paper. They are all proved by contradiction: we assume that a reasonable encoding exists and show that it introduces divergence. Often, the contradiction is obtained by exhibiting a process that cannot reduce but whose encoding reduces. This fact, together with operational correspondence, implies that the encoding introduces divergence, as stated by the following simple result.

Proposition 5. *Let $\llbracket \cdot \rrbracket$ be an operationally corresponding encoding. If there exists a process P such that $P \not\rightarrow$ but $\llbracket P \rrbracket \xrightarrow{\tau}$, then $\llbracket \cdot \rrbracket$ introduces divergence.*

Proof. The fact that $\llbracket P \rrbracket \xrightarrow{\tau} Q$ implies, by operational correspondence, that $P \Rightarrow P'$, for some P' such that $Q \Rightarrow \llbracket P' \rrbracket$. But the only P' such that $P \Rightarrow P'$ is P itself; thus, $\llbracket P \rrbracket \xrightarrow{\tau, +} \llbracket P \rrbracket$, i.e. $\llbracket P \rrbracket$ diverges. \square

Theorem 1. *There exists no reasonable encoding of π_{011} in π_{110} .*

Proof. Assume that $\llbracket \cdot \rrbracket$ is reasonable and consider the process $a(\ulcorner b \urcorner) \mid \bar{a}\langle b \rangle$, for $a \neq b$, that evolves in $\mathbf{0}$. By operational correspondence, $\llbracket a(\ulcorner b \urcorner) \mid \bar{a}\langle b \rangle \rrbracket \Rightarrow \llbracket \mathbf{0} \rrbracket$; moreover, by faithfulness, $\llbracket a(\ulcorner b \urcorner) \rrbracket \not\Downarrow$, $\llbracket \bar{a}\langle b \rangle \rrbracket \Downarrow$ and $\llbracket \mathbf{0} \rrbracket \not\Downarrow$. Thus, the barb of $\llbracket \bar{a}\langle b \rangle \rrbracket$ must be consumed in the computation leading $\llbracket a(\ulcorner b \urcorner) \mid \bar{a}\langle b \rangle \rrbracket$ to $\llbracket \mathbf{0} \rrbracket$.

Now, notice that $\llbracket \bar{a}\langle b \rangle \rrbracket$ cannot perform a τ -step otherwise, by Proposition 5, $\llbracket \cdot \rrbracket$ would introduce divergence. This fact, together with $\llbracket a(\ulcorner b \urcorner) \mid \bar{a}\langle b \rangle \rrbracket \stackrel{\Delta}{=} \llbracket a(\ulcorner b \urcorner) \rrbracket \mid \llbracket \bar{a}\langle b \rangle \rrbracket$, implies that $\llbracket a(\ulcorner b \urcorner) \rrbracket$ consumed the barb offered by $\llbracket \bar{a}\langle b \rangle \rrbracket$. Thus, it must be that $\llbracket a(\ulcorner b \urcorner) \rrbracket \xrightarrow{n\tilde{c}}$, for some n and \tilde{c} such that $\llbracket \bar{a}\langle b \rangle \rrbracket \xrightarrow{(\nu\tilde{c}')n!\tilde{c}}$; then $\llbracket a(\ulcorner b \urcorner) \rrbracket \xrightarrow{n\tilde{d}}$, for every \tilde{d} of the same arity as \tilde{c} , i.e. $|\tilde{d}| = |\tilde{c}|$.

If $n \neq b$, then pick up $e \notin \{a, b, n\}$ and the permutation of names swapping b and e ; by name invariance, it holds that $\llbracket \bar{a}\langle e \rangle \rrbracket \xrightarrow{(\nu\tilde{f}')n!\tilde{f}}$, where \tilde{f} and \tilde{f}' are the renaming of \tilde{c} and \tilde{c}' . In particular, $|\tilde{c}| = |\tilde{f}|$. Then, $\llbracket a(\ulcorner b \urcorner) \mid \bar{a}\langle e \rangle \rrbracket \xrightarrow{\tau}$, while $a(\ulcorner b \urcorner) \mid \bar{a}\langle e \rangle \not\xrightarrow{\tau}$. By Proposition 5, $\llbracket \cdot \rrbracket$ is not reasonable, as it introduces divergence; contradiction.

If $n = b$, then pick up $e \notin \{a, b\}$, the permutation of names swapping a and e , and work like before, with process $\llbracket a(\ulcorner b \urcorner) \mid \bar{e}\langle b \rangle \rrbracket$. \square

Corollary 1. *There exists no reasonable encoding of π_{001} in π_{100} .*

Proof. Trivial consequence of Theorem 1, since π_{001} and π_{100} can be seen as the subcalculi of π_{011} and π_{110} where all the communications happen on the same (unique and global) channel. \square

Theorem 2. *There exists no reasonable encoding of π_{010} in π_{100} .*

Proof. The proof is similar to that of Theorem 1. We start with process $\bar{a}\langle b \rangle \mid a(x)$, for $a \neq b$; it holds that $\llbracket \bar{a}\langle b \rangle \rrbracket \xrightarrow{(\nu\tilde{c}')!\tilde{c}}$ and $\llbracket a(x) \rrbracket \xrightarrow{?\tilde{c}}$, for some \tilde{c} . By name invariance, $\llbracket \bar{b}\langle a \rangle \rrbracket \xrightarrow{(\nu\tilde{d}')!\tilde{d}}$, where \tilde{d} and \tilde{d}' are obtained by swapping a and b in \tilde{c} and \tilde{c}' ; thus, $|\tilde{d}| = |\tilde{c}|$. Now, trivially, $\llbracket a(x) \rrbracket \xrightarrow{?\tilde{d}}$ and $\llbracket \bar{b}\langle a \rangle \mid a(x) \rrbracket \xrightarrow{\tau}$, while $\bar{b}\langle a \rangle \mid a(x) \not\xrightarrow{\tau}$. \square

Theorem 3. *There exists no reasonable encoding of π_{110} in π_{010} .*

Proof. Similarly to the proof of Theorem 1, consider the process $a(x, y) \mid \bar{a}\langle b, c \rangle$; again, $\llbracket \bar{a}\langle b, c \rangle \rrbracket \xrightarrow{(\nu\tilde{d}')n!\tilde{d}}$ and $\llbracket a(x, y) \rrbracket \xrightarrow{n\tilde{d}}$, for some d and \tilde{d}' . If $n \neq a$, choose $e \neq a$; by name invariance, $\llbracket \bar{e}\langle b, c \rangle \rrbracket \xrightarrow{(\nu\tilde{f}')n!\tilde{f}}$ and $\llbracket a(x, y) \mid \bar{e}\langle b, c \rangle \rrbracket \xrightarrow{\tau}$, while $a(x, y) \mid \bar{e}\langle b, c \rangle \not\xrightarrow{\tau}$. If $n = a$, consider $a(x, y, z) \mid \bar{a}\langle b, c, c \rangle$; like before, $\llbracket \bar{a}\langle b, c, c \rangle \rrbracket \xrightarrow{(\nu\tilde{e}')m!\tilde{e}}$ and $\llbracket a(x, y, z) \rrbracket \xrightarrow{m\tilde{e}}$. Then, if $m = a$, we have that $\llbracket a(x, y) \mid \bar{a}\langle b, c, c \rangle \rrbracket \xrightarrow{\tau}$; otherwise, choose $f \neq a$ and conclude that $\llbracket a(x, y, z) \mid \bar{f}\langle b, c, c \rangle \rrbracket \xrightarrow{\tau}$. \square

Corollary 2. *There exist no reasonable encodings of π_{001} and π_{100} in π_{000} .*

Proof. Easily derivable from Corollary 1 and Theorem 3, respectively. \square

Theorem 4. *There exists no reasonable encoding of π_{111} in π_{011} .*

Proof. Consider the process $\bar{a}\langle b, c \rangle \mid a(\ulcorner b \urcorner, \ulcorner c \urcorner)$, for a, b and c pairwise distinct. Like in Theorem 1, we have that $\llbracket a(\ulcorner b \urcorner, \ulcorner c \urcorner) \rrbracket \xrightarrow{n^2 m}$ and $\llbracket \bar{a}\langle b, c \rangle \rrbracket \xrightarrow{(\nu \bar{m})n!m}$. If the input of $\llbracket a(\ulcorner b \urcorner, \ulcorner c \urcorner) \rrbracket$ has been generated by relying on a template formal field, then $\llbracket a(\ulcorner b \urcorner, \ulcorner c \urcorner) \rrbracket \xrightarrow{n^2 l}$, for every l ; by Proposition 5, this would suffice to build up a divergent computation for $\llbracket a(\ulcorner b \urcorner, \ulcorner c \urcorner) \mid \bar{a}\langle b, d \rangle \rrbracket$, for every $d \neq c$. Otherwise, the input of $\llbracket a(\ulcorner b \urcorner, \ulcorner c \urcorner) \rrbracket$ relies on an actual field; we then consider the following possibilities for n and m :

1. $c \notin \{n, m\}$: let $d \neq c$ and consider the permutation that swaps c and d ; then, $\llbracket \bar{a}\langle b, d \rangle \rrbracket \xrightarrow{(\nu \bar{m})n!m}$ and $\llbracket \bar{a}\langle b, d \rangle \mid a(\ulcorner b \urcorner, \ulcorner c \urcorner) \rrbracket \xrightarrow{\tau}$.
2. $n = c, m \neq b$: let $d \neq b$ and consider the permutation that swaps b and d ; like before, $\llbracket \bar{a}\langle d, c \rangle \mid a(\ulcorner b \urcorner, \ulcorner c \urcorner) \rrbracket \xrightarrow{\tau}$.
3. $n = c, m = b$: let $d \neq a$ and consider the permutation that swaps a and d ; then, $\llbracket \bar{d}\langle b, c \rangle \mid a(\ulcorner b \urcorner, \ulcorner c \urcorner) \rrbracket \xrightarrow{\tau}$.
4. $m = c, n \neq b$: like case 2.
5. $m = c, n = b$: like case 3. □

Theorem 5. *There exists no reasonable encoding of π_{010} in π_{001} .*

Proof. By contradiction, assume that there exists a reasonable encoding $\llbracket \cdot \rrbracket$. Let a, b, c and d be pairwise distinct names; let Ω denote a divergent process and define $P \triangleq \text{if } x = d \text{ then } \Omega$. Faithfulness implies that $\llbracket \text{if } d = d \text{ then } \Omega \rrbracket$ diverges, $\llbracket a(x).P \rrbracket$ cannot offer data and $\llbracket \bar{a}\langle b \rangle \rrbracket$ must offer some datum. Moreover, because of Proposition 5, $\llbracket a(x).P \rrbracket$ and $\llbracket \bar{a}\langle b \rangle \rrbracket$ cannot perform τ -steps in isolation; however, because of operational correspondence, when put in parallel they must perform at least one τ -step to become $\llbracket P\{b/x\} \rrbracket$. If the input performed by $\llbracket a(x).P \rrbracket$ relies on a formal field, then we can obtain a divergent computation from $\llbracket a(x).P \mid \bar{c}\langle b \rangle \rrbracket$. So, it must be that $\llbracket a(x).P \rrbracket$ starts with an input relying on an actual template field $\ulcorner a \urcorner$ (it must be a otherwise, by name invariance, $\llbracket a(x).P \mid \bar{c}\langle b \rangle \rrbracket$ would diverge).

Now, it is easy to prove that $\llbracket a(x).P \rrbracket \mid \llbracket \bar{a}\langle b \rangle \rrbracket \Rightarrow \llbracket P\{b/x\} \rrbracket$ if and only if $\llbracket a(x).P \rrbracket \xrightarrow{\rho} R$, $\llbracket \bar{a}\langle b \rangle \rrbracket \xrightarrow{\bar{\rho}} R'$ and $R \mid R' \equiv \llbracket P\{b/x\} \rrbracket$. It must be that $\rho \triangleq ?a \cdot \rho_1 \cdot ?b \cdot \rho_2$, for b not occurring in ρ_1 and $?b$ generated by an input action with a formal template field; moreover, $\llbracket a(x).P \rrbracket \xrightarrow{?a \cdot \rho_1} R_1 \xrightarrow{?b} R_2 \xrightarrow{\rho_2} R$ and $\llbracket \bar{a}\langle b \rangle \rrbracket \xrightarrow{!a \cdot \bar{\rho}_1} R_3 \xrightarrow{!b} R_4 \xrightarrow{\bar{\rho}_2} R'$. In particular, $\rho_2 \triangleq \square_1 n_1 \cdot \dots \cdot \square_k n_k$, for $\square_i \in \{?, !\}$, and $\bar{\rho}_2 \triangleq \diamond_1 n_1 \cdot \dots \cdot \diamond_k n_k$, for $\diamond_i \in \{?, !\} - \{\square_i\}$.

Let σ be the permutation that swaps a and c and b and d ; by name invariance, $\llbracket c(x).P \rrbracket \xrightarrow{?c \cdot \rho'_1} R_1 \sigma \xrightarrow{?d} R_2 \sigma \xrightarrow{\rho'_2} R \sigma$ and $\llbracket \bar{c}\langle d \rangle \rrbracket \xrightarrow{!c \cdot \bar{\rho}'_1} R_3 \sigma \xrightarrow{!d} R_4 \sigma \xrightarrow{\bar{\rho}'_2} R' \sigma$, for $\rho'_1 = \rho_1 \sigma$ and $\rho'_2 = \rho_2 \sigma$. More precisely, $\rho'_2 \triangleq \square_1 n'_1 \cdot \dots \cdot \square_k n'_k$ and $\bar{\rho}'_2 \triangleq \diamond_1 n'_1 \cdot \dots \cdot \diamond_k n'_k$, for $n'_i \triangleq \sigma(n_i)$.

Now, consider $Q \triangleq a(x).P \mid \bar{a}\langle b \rangle \mid \bar{c}\langle d \rangle \mid c(x).P'$, where $P' \triangleq \text{if } x = b \text{ then } \Omega$; trivially, $Q \not\Downarrow$ while, as we shall see, $\llbracket Q \rrbracket \Uparrow$. This yields the desired contradiction. Consider

$$\llbracket Q \rrbracket \Rightarrow R_1 \mid R_3 \mid R_1 \sigma \mid R_3 \sigma \rightarrow \rightarrow R_2 \{d/b\} \mid R_4 \mid (R_2 \sigma) \{b/d\} \mid R_4 \sigma$$

where R_1 received d in place of b and $R_1\sigma$ received b in place of d (this is possible since these inputs do not rely on actual fields). Now, $R_2\{d/b\} \xrightarrow{\square_1 m_1 \dots \square_k m_k}$, where

$$m_k \triangleq \begin{cases} d & \text{if } n_i = b \\ n_i & \text{otherwise} \end{cases}$$

and $(R_2\sigma)\{b/d\} \xrightarrow{\square_1 m'_1 \dots \square_k m'_k}$, where

$$m'_k \triangleq \begin{cases} b & \text{if } n'_i = d \\ n'_i & \text{otherwise} \end{cases}$$

Finally, consider the computation

$$R_2\{d/b\} \mid R_4 \mid (R_2\sigma)\{b/d\} \mid R_4\sigma \Rightarrow R\{d/b\} \mid R'\{d/b\} \mid (R\sigma)\{b/d\} \mid (R'\sigma)\{b/d\}$$

obtained by performing a communication

- between $\square_i m_i$ and $\diamond_i n_i$ and between $\square_i m'_i$ and $\diamond_i n'_i$, if $n_i \neq b$, or
- between $\square_i m_i$ and $\diamond_i n'_i$ and between $\square_i m'_i$ and $\diamond_i n_i$, otherwise.

Now, $R\{d/b\} \mid R'\{d/b\} \triangleq (R \mid R')\{d/b\} \equiv \llbracket P\{b/x\} \rrbracket\{d/b\} \triangleq \llbracket \text{if } d = d \text{ then } \Omega \rrbracket$, that is a divergent process. \square

Theorem 6. *There exists no reasonable encoding of π_{100} in π_{001} .*

Proof. The proof is similar to that of Theorem 5. Assume that $\llbracket \cdot \rrbracket$ is reasonable; consider the process $P \triangleq (x, y).\text{if } x = a \text{ then if } y = d \text{ then } \Omega$; pick up $c \neq a$ and $d \neq b$; consider the permutation of names σ swapping a with c and b with d ; finally, show that $Q \triangleq P \mid \langle a, b \rangle \mid P\sigma \mid \langle c, d \rangle$ is not divergent, while $\llbracket Q \rrbracket \uparrow$. \square

6 Conclusion and Related Work

We have studied the expressive power of eight communication primitives, arising from the combination of three features: arity of data, communication medium and presence of pattern-matching. By relying on possibility/impossibility of ‘reasonable’ encodings, we obtained a clear hierarchy of communication primitives. Notably, LINDA’s communication paradigm [15] is at the top of this hierarchy, while the π -calculus is in the middle. A posteriori, this can justify the fact that the former one is usually exploited in actual programming languages [3, 14], where flexibility and expressive power are the driving issues, while the latter one is mostly used for theoretical reasoning.

One of the pioneering works in the study of communication primitives for distributed systems is [16]. There, the expressive power of several ‘classical’ primitives (like test-and-set, compare-and-swap, ...) is studied by associating to every primitive the highest number of parallel processes that can reach a distributed consensus with that primitive, under conditions similar to the ‘reasonableness’ of our Definition 2. It then follows that a primitive with number n is less expressive than every primitive with

number $m (> n)$: the latter one can solve a problem (i.e. the consensus among m processes) that the former one cannot reasonably solve. This idea is also exploited in [22] to assess the expressive power of the non-deterministic choice in the π -calculus.

In [10], the notion of relative expressive power is used to measure the expressiveness of programming languages. In particular, a simple class of three concurrent constraint languages is studied and organised in a strict hierarchy. The languages have guarded constructs and only differ in the features offered by the guards: a guard is always passed in the less expressive language; a guard is passed only if a given constraint is satisfied by the current knowledge; and, finally, a guard is passed only if a new constraint, that must be atomically added to the knowledge, is consistent with the current knowledge. Very roughly, the last kind of guards can be related to the pattern-matching construct of our calculi, for the possibility of atomically testing and modifying the environment; in both cases, this feature sensibly increases the expressive power of the language.

By the way, the form of pattern-matching considered here is very minimal: only the equality of names can be tested while retrieving a datum. However, many other forms of pattern-matching can be exploited [12], to yield more and more flexible formalisms; some proposals have been investigated from the expressiveness point of view in [28].

Finally, in [7] a form of atomic polyadic name matching is presented, but with a different approach w.r.t. ours. Indeed, while in our π_{111} the tuple of names to be matched is in the transmitted/received value (by using a standard π -calculus terminology, the tuple is in the ‘object’ part of an output/input), in [7] there are composite channel names that must be matched to enable a communication (thus, the tuple is in the ‘subject’ part of the output/input). This feature enables a nice modelling of distributed and cryptographic process calculi; nevertheless, our LINDA-like pattern-matching is stronger, since the possibility of using both formal and actual fields in a template yield a more flexible form of input actions.

To conclude, this paper is one of the first attempts to classify languages according to their communication primitive. A lot of work still remains to be done. For example, it would be interesting to study more concrete languages, maybe by encoding them in one of the calculi presented in this paper. Moreover, other common features (such as synchrony) could be added to the picture. Finally, it would also be interesting to prove stronger properties for the encodings of Section 4, whenever possible; indeed, since we were mostly interested in the impossibility results, we intentionally exploited quite a weak form of ‘reasonableness’.

Acknowledgements I would like to thank Rosario Pugliese for his suggestions that improved a first draft of this paper and Catuscia Palamidessi for her encouragements and discussions. Finally, I also thank the anonymous *FoSSaCS’06* referees for their positive attitude and fruitful comments.

References

1. L. Acciai and M. Boreale. X π : a typed process calculus for XML messaging. In *Proc. of FMOODS’05*, volume 3535 of *LNCS*, pages 47–66. Springer, 2005.
2. R. M. Amadio, I. Castellani and D. Sangiorgi. On bisimulations for the asynchronous π -calculus. *Theoretical Computer Science*, 195(2):291–324, 1998.

3. K. Arnold, E. Freeman and S. Hupfer. *JavaSpaces Principles, Patterns and Practice*. Addison-Wesley, 1999.
4. G. Boudol. Asynchrony and the π -calculus (note). Rapport de Recherche 1702, INRIA Sophia-Antipolis, May 1992.
5. A. Brown, C. Laneve and G. Meredith. π duce: a process calculus with native XML datatypes. In *Proc. of Services and Formal Methods*, volume 3670 of *LNCS*. Springer, 2005.
6. N. Busi, R. Gorrieri and G. Zavattaro. A process algebraic view of LINDA coordination primitives. *Theoretical Computer Science*, 192(2):167–199, 1998.
7. M. Carbone and S. Maffei. On the expressive power of polyadic synchronisation in π -calculus. *Nordic Journal of Computing*, 10(2):70–98, 2003.
8. L. Cardelli and A. D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, 2000.
9. G. Castagna, R. De Nicola and D. Varacca. Semantic subtyping for the π -calculus. In *Proc. of LICS'05*, pages 92–101. IEEE Computer Society, 2005.
10. F. de Boer and C. Palamidessi. Embedding as a tool for language comparison. *Information and Computation*, 108(1):128–157, 1994.
11. R. De Nicola, D. Gorla and R. Pugliese. On the expressive power of KLAIM-based calculi. Tech. Rep. 09/2004, Dip. Informatica, Univ. Roma “La Sapienza”. To appear in *TCS*.
12. R. De Nicola, D. Gorla and R. Pugliese. Pattern Matching over a Dynamic Network of Tuple Spaces. In *Proc. of FMOODS'05*, volume 3535 of *LNCS*, pages 1–14. Springer, 2005.
13. C. Ene and T. Muntean. Expressiveness of point-to-point versus broadcast communications. In *Proc. of FCT'99*, volume 1684 of *LNCS*, pages 258–268. Springer, 1999.
14. D. Ford, T. Lehman, S. McLaughry and P. Wyckoff. T Spaces. *IBM Systems Journal*, pages 454–474, August 1998.
15. D. Gelernter. Generative Communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, 1985.
16. M. Herlihy. Impossibility and universality results for wait-free synchronization. In *Proc. of PODC'88*, pages 276–290. ACM Press, 1988.
17. K. Honda and M. Tokoro. An object calculus for asynchronous communication. In *Proc. of ECOOP'91*, volume 512 of *LNCS*, pages 133–147. Springer, 1991.
18. R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
19. R. Milner. The polyadic π -calculus: A tutorial. In *Logic and Algebra of Specification*, volume 94 of *Series F*. NATO ASI, Springer, 1993.
20. R. Milner and D. Sangiorgi. Barbed bisimulation. In *Proc. of ICALP'92*, volume 623 of *LNCS*, pages 685–695. Springer, 1992.
21. U. Nestmann and B. Pierce. Decoding choice encodings. *Information and Computation*, 163:1–59, 2000.
22. C. Palamidessi. Comparing the expressive power of the synchronous and the asynchronous π -calculi. *Mathematical Structures in Computer Science*, 13(5):685–719, 2003.
23. J. Parrow. An introduction to the π -calculus. In *Handbook of Process Algebra*, pages 479–543. Elsevier Science, 2001.
24. P. Quaglia and D. Walker. On encoding $p\pi$ in $m\pi$. In *Proc. of FSTTCS'98*, volume 1530 of *LNCS*, pages 42–51. Springer, 1998.
25. D. Sangiorgi. Bisimulation in higher-order process calculi. *Information and Computation*, 131:141–178, 1996.
26. E. Shapiro. Separating concurrent languages with categories of language embeddings. In *Proc. of 23rd STOC*, pages 198–208. ACM Press, 1991.
27. N. Yoshida. Graph types for monadic mobile processes. In *Proc. of FSTTCS'96*, volume 1180 of *LNCS*, pages 371–386. Springer, 1996.
28. G. Zavattaro. Towards a hierarchy of negative test operators for generative communication. In *Proc. of EXPRESS'98*, volume 16 of *ENTCS*, 1998.