

A linear account of session types in the pi calculus

Marco Giunti¹ and Vasco T. Vasconcelos²

¹ Faculty of Planning, University IUAV of Venice

² LaSIGE, Faculty of Sciences, University of Lisbon

Abstract. We present a reconstruction of session types in a conventional pi calculus where types are qualified as linear or unrestricted. Linearly typed communication channels are guaranteed to occur in exactly one thread, possibly multiple times. We equip types with a constructor that denotes the two ends of a same communication channel. In order to assess the flexibility of the new type system, we provide three distinct encodings (from the linear lambda calculus, from the linear pi calculus, and from the pi calculus with polarized variables) into our system. For each language we present operational and typing correspondences, showing that our system effectively subsumes the linear pi calculus as well as foregoing works on session types.

1 Introduction

Session types allow a concise description of protocols by detailing the sequence of messages involved in each particular run of the protocol. Introduced for a dialect of the pi calculus [6, 13], the concept has been transferred to different realms, including functional and object-oriented programming and operating systems; refer to [3] for a recent overview.

By way of motivation, consider a service allowing to create online petitions. Petition creators receive from the petition service a channel on which they provide the title of the petition, the petition text and the due date. After the initial setup, the exact same channel is ready to be distributed among the client's acquaintances to collect thousands of signatures, but not without the creator signing the petition first. The code for the creator can be written as follows,

$$petitionOnline(p).\bar{p} title.\bar{p} description.\bar{p} dueDate.\bar{p} signature.(\bar{a}_1 p \mid \dots \mid \bar{a}_n p)$$

where $x(y)$ denotes reading value y on channel x , $\bar{x}v$ denotes sending value v on channel x , and the vertical bar denotes parallel composition. Each of the acquaintances (not shown in the example), after reading p on channel a_i , can sign the petition and further distribute the channel at will.

The protocol for channel p can be concisely described by a type T of the form below, composed of an initial linear part that becomes shared (or unrestricted) in the later part.

$$\text{lin} !String.\text{lin} !String.\text{lin} !Date.S \quad \text{where} \quad S = \text{un} !String.S$$

The final part is unrestricted because it is desirable, but not absolutely necessary, that acquaintances (including the petition creator) sign the petition; conversely, the initial part is linear because petitions cannot be signed without first setting up the title, the description and the due date.

In the process above, each channel a_i forwards p at type S . Such a channel may be given the type $\text{lin}!S.\text{un end}$, if we require that acquaintances eventually receive the petition channel; the continuation is un end (the type of a channel on which no further interaction is possible) allowing channel a_i to be discarded thereafter. Concentrating on the type of *petitionOnline*, we see that petition creators need it a type $S_1 = \text{un}?T.S_1$ so that they may create as many petitions (of type T) as required. It should be easy to see that the service itself sees the same channel at the dual type $S_2 = \text{un}!T.S_2$. The whole system, composed by the service running in parallel with petition creators can be typed by reconciling the two *end point types* S_1 and S_2 in a single, unordered, *channel type* of the form (S_1, S_2) .

The language of the pi calculus, when considered in conjunction with a type system with session types, is known to require a means to distinguish the two ends of a session channel (S_1 and S_2 above) in order to preserve type soundness [4, 5, 17]. Alternative solutions not requiring such a distinction rely on the restriction of channel passing to bound output. Such systems include the original formulation of delegation in session types [6] as well as more recent works [2, 11].

Two approaches for distinguishing the ends of a channel are available in the literature: polarized channel variables [5], and form of channel double binder [15]. In the pi calculus with polarities the two ends of a channel x are distinguished by labelling each of its ends with a different label: x^+ and x^- denote the two ends of channel x . Given that from a given channel name one may find its two ends, one can restrict (the two ends of) a channel x with the usual pi calculus restriction operator $(\nu x)P$. Typing contexts, however accept two different entries for the same channel, one labelled with $+$, the other with $-$, as in the typing sequent below.

$$\Gamma, x^+ : S_1, x^- : S_2 \vdash \overline{x^+}v.P_1 \mid x^-(w).P_2$$

A variant of the above work, [15], uses distinct variables to describe the two ends of a same channel. In this case one cannot obtain the second end of a channel from the other end. It is restriction that puts together the two channel ends, by binding them together, as in $(\nu yz)P$. The assumptions in typing contexts are for simple variables, as in the example below where y and z denote the two ends of a same channel.

$$\Gamma, y : S_1, z : S_2 \vdash \overline{y}v.P_1 \mid z(w).P_2$$

The first work can be criticized for using non-conventional typing contexts, where typing information of a same channel x is split among two different entries, x^+ and x^- . The second work uses standard contexts but relies on a new scope restriction operator that binds two variables together. The goal of this work is to equip types with a constructor able to denote the two ends of a same channel. We then have the best of both worlds where we use the standard pi calculus (Milner et al. [10]) with standard typing sequents.

Syntax

$b ::=$	Booleans:	$P ::=$	Processes:
true	true	$\bar{x}v.P$	output
false	false	$x(x).P$	input
$v ::=$	Values:	$P \mid P$	composition
b	boolean value	if v then P else P	conditional
x	variable	$(\nu x)P$	restriction
		$*P$	replication
		$\mathbf{0}$	inaction

Rules for structural congruence

$$\begin{array}{l}
P \mid Q \equiv Q \mid P \quad (P \mid Q) \mid R \equiv P \mid (Q \mid R) \quad P \mid \mathbf{0} \equiv P \quad *P \equiv P \mid *P \\
(\nu x)P \mid Q \equiv (\nu x)(P \mid Q) \quad (\nu x)\mathbf{0} \equiv \mathbf{0} \quad (\nu x)(\nu y)P \equiv (\nu y)(\nu x)P
\end{array}$$

Rules for reduction

$$\begin{array}{l}
\bar{x}v.P \mid x(y).Q \rightarrow P \mid Q[v/y] \quad \text{[R-COM]} \\
\text{if true then } P \text{ else } Q \rightarrow P \quad \text{if false then } P \text{ else } Q \rightarrow Q \quad \text{[R-IFT] [R-IFF]} \\
\frac{P \rightarrow Q}{(\nu x)P \rightarrow (\nu x)Q} \quad \frac{P \rightarrow Q}{P \mid R \rightarrow Q \mid R} \quad \frac{P \equiv P' \quad P' \rightarrow Q' \quad Q' \equiv Q}{P \rightarrow Q} \\
\text{[R-RES] [R-PAR] [R-STRUCT]}
\end{array}$$

Fig. 1. Pi calculus: Syntax and operational semantics

We test the flexibility of our type system by embedding the pi calculus with polarities and session types [5] (hence the conventional pi calculus [10]). We do the same for the linear pi calculus [7], and for the linear (call by value) lambda calculus as in [16]. For each of these languages we prove an operational and a typing correspondence result. From the two first embeddings we learn that our type system is an extension of advanced type systems for pi calculi. The embedding of the linear lambda calculus crucially takes linearity into consideration generating code accordingly (replicated or non replicated) for shared and linear resources.

The outline of the paper is as follows. The next section recalls the pi calculus and introduces our type system. Then, the subsequent three sections present the embeddings of the three languages mentioned above: pi calculus with polarities, linear pi calculus and linear lambda calculus. The last section presents some related as well as future work.

2 Pi Calculus

This section introduces the pi-calculus, its syntax and semantics, as well as our type system. The syntax is in Figure 1. We rely on a set of variables, ranged over

$q ::=$	Qualifiers:	a	type variable
lin	linear	$\mu a.S$	recursive type
un	unrestricted	$T ::=$	Types:
$p ::=$	Pre Types:	bool	boolean
$?T.S$	receive	S	end point
$!T.S$	send	(S, S)	channel
end	termination	$\Gamma ::=$	Contexts:
$S ::=$	End Point Types:	\emptyset	empty context
$q p$	qualified channel	$\Gamma, x: T$	variable binding

Fig. 2. Pi calculus: Types and typing contexts

by x, y, z . Values include variables and the booleans `true` and `false`. For processes we have (synchronous, unary) output and input, in the forms $\bar{x}v.P$ and $x(y).P$, as well as a parallel composition, conditional, scope restriction, replication and the terminated process.

The binders for the language appear in parenthesis: x is bound in both $y(x).P$ and $(\nu x)P$. Free and bound variables in processes are defined accordingly, and so is alpha conversion, substitution of a variable x by a value v in a process P , denoted $P[v/x]$. We follow Barendregt's variable convention, requiring bound variables to be distinct from free variables in any mathematical context.

Structural congruence is the smallest relation on processes including the rules in the same figure. The first three rules say that parallel composition is commutative, associative and has $\mathbf{0}$ for neutral element. The last rule on the first line captures the essence of replication as an unbounded number of identical processes. The rules in the second line deal with scope restriction. The first, scope extrusion, allows the scope of x to encompass Q ; due to variable convention, x bound in $(\nu x)P$, cannot be free in Q . The other two rules state that restricting over a terminated process has no effect, and allow exchanging the order of restrictions.

The reduction relation is the smallest relation on processes including the rules in Figure 1. The [R-COM] rule communicates value v from an output prefixed one $\bar{x}v.P$ to an input prefixed process $x(y).Q$; the result is the parallel composition of the continuation processes, where the bound variable y is replaced by value v in the input process. The rules for the conditional are straightforward. The rules in the last line allow reduction to happen underneath scope restriction and parallel composition, and incorporate structural congruence into reduction.

The syntax of types is described in Figure 2. Types include the boolean type, end point types and channel types. The novelty with respect linear and session-based systems for the pi calculus is the introduction of a new type constructor to describe the two ends of a same channel, (S_1, S_2) , where S_1 details the behaviour of one end, whereas S_2 details that of the other end. An end point type S can

be a pre type qualified with `lin` or `un`, a recursive type or a type variable. Each qualifier in a type controls the number of times the channel can be used at that point: exactly once for `lin`; zero or more times for `un`. A pre type of the form `!T.S` describes a channel end able to send a value of type T and to proceed as prescribed by S . Similarly, pre type `?T.S` describes a channel end able to receive a value of type T and continue as S . Pre type `end` describes a channel end on which no further interaction is possible. For recursive (end point) types we rely on a set of type variables, ranged over by a . Recursive types are required to be contractive, that is, containing no subexpression of the form $\mu a_1 \dots \mu a_n. a_1$.

Type equality is not syntactic. Instead, we define it as the equality of regular infinite trees obtained by the infinite unfolding of recursive types, *modulo pair commutation*. The formal definition, which we omit, is co-inductive. In this way we use types $(\mu a. \text{lin}! \text{bool}. \text{lin}? \text{bool}. a, \text{un end})$ and $(\text{un end}, \text{lin}! \text{bool}. \mu b. \text{lin}? \text{bool}. \text{lin}! \text{bool}. b)$ interchangeably, in any mathematical context. This allows us never to consider a type $\mu a. S$ explicitly (or a for that matter). Instead, we pick another type in the same equivalence class, namely $S[\mu a. S/a]$. If the result of the process turns out to start with a μ , we repeat the procedure. Unfolding is bound to terminate due to contractiveness. In other words, we take an equi-recursive view of types [12].

Type duality plays a central role in the theory of session types, ensuring that communication between the two ends of a channel proceeds smoothly. Intuitively, the dual of output is input and the dual of input is output. In particular if S_2 is dual of S_1 , then $q?T.S_1$ is dual of $q!T.S_2$. Session type `end` is dual of itself. Rather than providing a co-inductive definition of duality, we start by defining a function from end-point channels into end-point channels as follows.

$$\overline{q?T.S} = q!T.\overline{S} \quad \overline{q!T.S} = q?T.\overline{S} \quad \overline{q \text{ end}} = q \text{ end} \quad \overline{\mu a. S} = \mu a. \overline{S} \quad \overline{a} = a$$

Then, to check that a given end point type S_1 is dual of another type S_2 , we first build the dual of S_1 and then check that the thus obtained type is equivalent to S_2 . For example, to show that type $\overline{\mu a. \text{lin}? \text{bool}. \text{lin}! \text{bool}. a}$ is a dual of type $\text{lin}! \text{bool}. \mu b. \text{lin}? \text{bool}. \text{lin}! \text{bool}. b$, we build $\overline{\mu a. \text{lin}? \text{bool}. \text{lin}! \text{bool}. a} = \mu a. \text{lin}! \text{bool}. \text{lin}? \text{bool}. a$, and then show that $\mu a. \text{lin}! \text{bool}. \text{lin}? \text{bool}. a = \text{lin}! \text{bool}. \mu b. \text{lin}? \text{bool}. \text{lin}! \text{bool}. b$. Qualifiers are important: S and \overline{S} must be equally qualified so that a linear output process may find a linear input process to embark in reduction.

Contexts, or type environments, are inductively defined in Figure 2. In a context $\Gamma, x: T$ we assume that x does not occur in Γ ; we also assume the various variable bindings in Γ to be unordered. We define predicate `un` to be true of a) the empty context, as well as of b) context $\Gamma, x: \text{bool}$, context $\Gamma, x: \text{un } p$, and context $\Gamma, x: (\text{un } p_1, \text{un } p_2)$, whenever `un`(Γ).

Typing relies on the context splitting operation described in Figure 3. It should be easy to understand: unrestricted types are copied into both contexts, linear types are placed in one of the two resulting contexts. The first four rules are standard [16], the last three rules are new to this work; the philosophy however remains the same. We omit three rules, duals to the last three, obtained by interchanging the end point types in the channel type (e.g., $(\text{un } p_2, \text{lin } p_1)$ in the last rule), for the effect can be obtained by a suitable choice of the type in its equivalence class (recall that pair types are unordered).

Context splitting rules

$$\begin{array}{c}
\emptyset = \emptyset \cdot \emptyset \qquad \frac{\Gamma = \Gamma_1 \cdot \Gamma_2 \quad T = \mathbf{un} p \text{ or } (\mathbf{un} p_1, \mathbf{un} p_2)}{\Gamma, x: T = (\Gamma_1, x: T) \cdot (\Gamma_2, x: T)} \\
\frac{\Gamma = \Gamma_1 \cdot \Gamma_2 \quad T = \mathbf{lin} p \text{ or } (\mathbf{lin} p_1, \mathbf{lin} p_2)}{\Gamma, x: T = (\Gamma_1, x: T) \cdot \Gamma_2} \quad \frac{\Gamma = \Gamma_1 \cdot \Gamma_2 \quad T = \mathbf{lin} p \text{ or } (\mathbf{lin} p_1, \mathbf{lin} p_2)}{\Gamma, x: T = \Gamma_1 \cdot (\Gamma_2, x: T)} \\
\frac{\Gamma = \Gamma_1 \cdot \Gamma_2}{\Gamma, x: (\mathbf{lin} p_1, \mathbf{lin} p_2) = (\Gamma_1, x: \mathbf{lin} p_1) \cdot (\Gamma_2, x: \mathbf{lin} p_2)} \\
\frac{\Gamma = \Gamma_1 \cdot \Gamma_2}{\Gamma, x: (\mathbf{lin} p_1, \mathbf{un} p_2) = (\Gamma_1, x: (\mathbf{lin} p_1, \mathbf{un} p_2)) \cdot (\Gamma_2, x: \mathbf{un} p_2)} \\
\frac{\Gamma = \Gamma_1 \cdot \Gamma_2}{\Gamma, x: (\mathbf{lin} p_1, \mathbf{un} p_2) = (\Gamma_1, x: \mathbf{un} p_2) \cdot (\Gamma_2, x: (\mathbf{lin} p_1, \mathbf{un} p_2))}
\end{array}$$

Typing rules for values

$$\frac{\mathbf{un}(\Gamma)}{\Gamma \vdash b: \mathbf{bool}} \quad \frac{\mathbf{un}(\Gamma)}{\Gamma, x: T \vdash x: T} \quad \frac{\Gamma \vdash v: (S, \mathbf{un} p)}{\Gamma \vdash v: S} \quad [\mathbf{T-BOOL}] \quad [\mathbf{T-VAR}] \quad [\mathbf{T-STRENGTH}]$$

Typing rules for processes

$$\begin{array}{c}
\frac{\mathbf{un}(\Gamma)}{\Gamma \vdash \mathbf{0}} \quad \frac{\Gamma_1 \vdash P_1 \quad \Gamma_2 \vdash P_2}{\Gamma_1 \cdot \Gamma_2 \vdash P_1 \mid P_2} \quad \frac{\Gamma \vdash P \quad \mathbf{un}(\Gamma)}{\Gamma \vdash *P} \quad [\mathbf{T-INACT}] \quad [\mathbf{T-PAR}] \quad [\mathbf{T-REPL}] \\
\frac{\Gamma_1 \vdash v: \mathbf{bool} \quad \Gamma_2 \vdash P_1 \quad \Gamma_2 \vdash P_2}{\Gamma_1 \cdot \Gamma_2 \vdash \text{if } v \text{ then } P_1 \text{ else } P_2} \quad \frac{\Gamma, x: (S, \bar{S}) \vdash P}{\Gamma \vdash (\nu x)P} \quad [\mathbf{T-IF}] \quad [\mathbf{T-RES}] \\
\frac{\Gamma, x: S, y: T \vdash P \quad (*)}{\Gamma, x: q?T.S \vdash x(y).P} \quad \frac{\Gamma_1 \vdash v: T \quad \Gamma_2, x: S \vdash P \quad (**)}{\Gamma_1 \cdot (\Gamma_2, x: q!T.S) \vdash \bar{x}v.P} \quad [\mathbf{T-IN}], [\mathbf{T-OUT}] \\
\frac{\Gamma, x: (S, S'), y: T \vdash P \quad (*)}{\Gamma, x: (q?T.S, S') \vdash x(y).P} \quad \frac{\Gamma_1 \vdash v: T \quad \Gamma_2, x: (S, S') \vdash P \quad (**)}{\Gamma_1 \cdot (\Gamma_2, x: (q!T.S, S')) \vdash \bar{x}v.P} \quad [\mathbf{T-INC}], [\mathbf{T-OUTC}] \\
(*) \quad q = \mathbf{un} \Rightarrow q?T.S = S \qquad (**) \quad q = \mathbf{un} \Rightarrow q!T.S = S
\end{array}$$

Fig. 3. Pi calculus: Typing

Equipped with the notions of type duality, unrestricted contexts, and context splitting we are ready to introduce the typing rules in Figure 3. The first two typing rules for values are standard. Rule [T-STRENGTH] is central to our system with channels described as pairs of types; we discuss it after introducing the remaining typing rules.

For processes, rule [T-INACT] says that the terminated process can only be typed in an unrestricted context, ensuring that linear channels are given a chance to be consumed. Rule [T-PAR] uses context splitting to partition linearly typed variables between the two processes: the incoming context is split into Γ_1 and Γ_2 , and we use the former to type check process P_1 and the latter to type check process P_2 . Rule [T-REPL] for replication requires the typing context not to contain linear values, for P may be used an unrestricted number of times. Rule

[T-IF] for the conditional process splits the incoming context in two parts: one used to check the condition, the other to check both branches. The same context for the two branches is justified by the fact that only one of P_1 or P_2 will be executed. Rule [T-RES] allows restricting channels whose end points are dual, making sure that communication on the channel happens according to the plan. Allowing to restrict a end point type S type would not break type preservation, Theorem 1, but we believe that such an alternative rule does not fit well in a linear system, where we expect linear channels to be given an opportunity to be consumed. Even though unrestricted end point types cannot be directly restricted, we can show that, for each derivation of $\Gamma, x: S \vdash P$, there is a derivation of $\Gamma, x: (S, \text{un } p) \vdash P$, thus allowing to apply scope restriction to an otherwise unrestricted channel end.

We have two rules for input, [T-IN] and [T-INC], depending on the type for channel x in the context. Rule [T-IN] deals with end point types. If x is typed with $q?T.S$, we know that the bound variable y is of type T , and we type check P under the extra assumption $y: T$. Equally important is the fact that the continuation uses channel x at continuation type S , that is, process $x(y).P$ uses channel x at type $q?T.S$ whereas P may use the *same* channel this time at type S . Finally, unrestricted channels, given that they may be shared, must retain their behavior throughout computation, hence the side condition. A solution to the equation in the side condition is $\mu a?T.a$ for a not in T , which we abbreviate to $*?T$ (and similarly for output). Rule [T-INC] follows the same pattern, consuming one end point and keeping the other unchanged. Similarly to input, we have two rules for output. Rule [T-OUT], splits the context in two parts, one to check v and the other to check continuation P . Notice that the context in the conclusion, $\Gamma_1 \cdot (\Gamma_2, x: q!T.S)$ allows to type process $\bar{x}x$ with a context $x: S$ with type S such that $S = \text{un}!S.S$.

Rule [T-STRENGTH] allows for a fine grained control of the channel ends of a given channel. A process holding the two ends of a given channel x , say $(*!\text{bool}, *?\text{bool})$, may pass the output capability only by using [T-VAR] followed by [T-STRENGTH] to obtain $\Gamma, x: (*!\text{bool}, *?\text{bool}) \vdash x: *!\text{bool}$ and then compose with rule [T-OUT] or [T-OUTC] in a process of the form $\bar{y}x.P$. The rule is also fundamental in establishing the main result of this section.

To lighten the syntax in examples, we omit all unrestricted qualifiers and only annotate linear types. We also omit the trailing un end in types, as well as the trailing $\mathbf{0}$ in processes. As an example, consider the type $?(lin!\text{bool}).S$ of an unrestricted channel that receives a linear channel capable of outputting a boolean value. The following sequent is easy to establish,

$$x: ?(lin!\text{bool}).S \vdash x(z).\bar{z} \text{ true} \mid x(w).\bar{w} \text{ false}$$

but only for an appropriate type S . Reading rule [T-IN], we realize that S must be equivalent to $?(lin!\text{bool}).S$, that is S must be (equivalent to) $\mu a.?(lin!\text{bool}).a$, abbreviated to $*?(lin!\text{bool})$. Continuing with the example, if P is the above process, then $P \mid (\nu y)\bar{x}y$ is not typable, for the linear input capability of channel y is never exercised. But $P \mid (\nu y)(\bar{x}y \mid y(u))$ is typable under context $x: (*?(lin!\text{bool}), *!(lin?\text{bool}))$.

Given that a type $(\text{un } p_1, \text{lin } p_2)$ cannot possibly be restricted in a process (cf. rule [T-RES]), the reader may wonder why we consider them at all. It turns out that free channel output may lead to situations where a thread holds the two ends of a same channel [17]. For instance, process $\bar{z}x. | z(w).w(y).\bar{x}\text{true}$, typable under context $z: (*?S, *!S), x: (S, \bar{S})$ with $S = \text{lin?bool}$, reduces to process $P = x(y).\bar{x}\text{true}$, which we want to type under the same context. By applying rule [T-INC] to P we obtain a judgment with a un-lin type, namely $z: (*?S, *!S), x: (\text{end}, \text{lin!bool}) \vdash \bar{x}\text{true}$. A further application of rule [T-OUTC] gets rid of the un-lin type, yielding $z: (*?S, *!S), x: (\text{end}, \text{end}) \vdash \mathbf{0}$ typable under rule [T-INACT].

We conclude the section with the main result of our system. Reduction preserves typability only for a certain kind of contexts. To understand why reduction does not preserve typability in the presence of arbitrary contexts, take for P the process $x(z).\text{if } z \text{ then } \mathbf{0} \text{ else } \mathbf{0} \mid (\nu y)\bar{x}y$. We can easily see that P is typable under the (non balanced) context $x: (\text{lin!end}, \text{end}, \text{lin?bool}, \text{end})$. But P reduces to process $(\nu y)\text{if } y \text{ then } \mathbf{0} \text{ else } \mathbf{0}$ which is not typable. The whole purpose of balancing is to make sure that the type of y in the output is that of z in the input.

We define predicate *balanced* to be true of a) the empty context, and b) context $\Gamma, x: \text{bool}$ and context $\Gamma, x: (S, \bar{S})$ whenever Γ is balanced.

Theorem 1 (Type Preservation). *If $\Gamma_1 \vdash P_1$ with Γ_1 balanced and $P_1 \rightarrow P_2$, then $\Gamma_2 \vdash P_2$ with Γ_2 balanced.*

Proof (Sketch). Albeit standard, the proof is quite long due to the combinatorics introduced by (the various rules in) context splitting, the lin-un qualifiers and the four rules for input and for output. We rely on several standard auxiliary results, including unrestricted context weakening, context strengthening, a substitution lemma (stating that if $\Gamma_1 \vdash v: T$ and $\Gamma_2, x: T \vdash P$ and $\Gamma_1 \cdot \Gamma_2$ is defined then $\Gamma_1 \cdot \Gamma_2 \vdash P[v/x]$), and balanced context preservation for structural congruence. In order to proceed by induction on the derivation of the reduction step when [T-PAR] is the last rule, we need a stronger statement that details the relation between Γ_1 and Γ_2 , namely $\Gamma_2 = \Gamma_1$ or $\Gamma_1 = \Gamma, x: (q?T.S, q!T.\bar{S})$ and $\Gamma_2 = \Gamma, x: (S, \bar{S})$.

3 Embedding the Pi Calculus with Polarities

This section shows that our type system embeds the polarity system introduced by Gay and Hole [5]. Since Gay and Hole show that the pi calculus with polarities embeds the simply typed pi calculus; by transitivity our language embeds the simply typed pi calculus as well.

In Figure 4 we present the branch-select free fragment of the pi calculus with polarities. Variables may be *polarized*, occurring in processes as well as in typing contexts as x^+ or x^- or simply as x . We write x^p for a general polarized name, where p represents an optional polarity. Duality on polarities, written \bar{p} exchanges $+$ and $-$. The new constructors of the language, input and output,

New syntactic forms

$P ::= \dots$	Processes:	$\mu a.S$	recursive type
$\bar{x}^p . P$	output	$S ::=$	Session types:
$x^p(x).P$	input	end	termination
$T ::=$	Types:	$?T.S$	receive
\hat{T}	standard channel	$!T.S$	send
S	session channel	a	type variable
a	type variable	$\mu a.S$	recursive type

New reduction rules

$$\bar{x}^p z^q . P \mid x^{\bar{p}}(y).Q \rightarrow_p P \mid Q[z^q/y] \quad [\text{R-COM}]$$

Context updating

$$\begin{aligned} \Gamma + x^p : S = \Gamma, x^p : S & \quad \text{if } x^p, x \notin \text{dom}(\Gamma) \text{ and } S = ?T.S, !T.S, \text{end} \\ \Gamma + x : T = \Gamma, x : T & \quad \text{if } x, x^+, x^- \notin \text{dom}(\Gamma) \\ \Gamma, x : T + x : T = \Gamma, x : T & \quad \text{if } T = \hat{T}, \text{bool} \end{aligned}$$

Typing rules

$$\begin{array}{c} \frac{\Gamma \text{ completed}}{\Gamma \vdash_p \mathbf{0}} \quad \frac{\Gamma \vdash_p P \quad \Gamma \text{ unlimited}}{\Gamma \vdash_p *P} \quad \frac{\Gamma, x : \hat{T} \vdash_p P}{\Gamma \vdash_p (\nu x)P} \quad \frac{\Gamma, x^+ : S, x^- : \bar{S} \vdash_p P}{\Gamma \vdash_p (\nu x)P} \\ \text{[T-INACT] [T-REPL] [T-NEW] [T-NEWS]} \\ \frac{\Gamma_1 \vdash_p P \quad \Gamma_2 \vdash_p Q}{\Gamma_1 + \Gamma_2 \vdash_p P \mid Q} \quad \frac{\Gamma, x : \hat{T}, y : T \vdash_p P}{\Gamma, x : \hat{T} \vdash_p x(y).P} \quad \frac{\Gamma, x^p : S, y : T \vdash_p P}{\Gamma, x^p : ?T.S \vdash_p x^p(y).P} \\ \text{[T-PAR] [T-IN] [T-INS]} \\ \frac{\Gamma, x : \hat{T} \vdash_p P}{(\Gamma, x : \hat{T}) + y^q : T \vdash_p \bar{x} y^q . P} \quad \frac{\Gamma, x^p : S \vdash_p P}{(\Gamma, x^p : !T.S) + y^q : T \vdash_p \bar{x}^p y^q . P} \\ \text{[T-OUT] [T-OUTS]} \end{array}$$

Fig. 4. Pi calculus with polarities

are in Figure 4; the remaining are taken from Figure 1; the syntactic category for values in Figure 1 does not contribute to the language.

The reduction relation, denoted by \rightarrow_p , is defined inductively by the rules in Figure 1 with rule [R-COM] replaced by that in Figure 4. From the above description it should be obvious that the two languages differ in the (optional) polarity annotation on (non-bound occurrences of) variables. We define an erase function that removes from a polarized processes all occurrences of $+$ and $-$, to yield a process generated by the grammar in Figure 1. There is an obvious operational correspondence between the two languages, stated in Theorem 2. The converse is clearly not true. Take for P the polarized process $x^+ \mid x^+(\cdot)$. Then $\text{erase}(P) = \bar{x} \mid x(\cdot)$ reduces while P does not.

The language of types includes a distinct category S for (linear) session types. Since we restrict our language to the branch-select free fragment of [5], we ignore subtyping. Duality is defined as in Section 2, with the appropriate changes which amount erasing the qualifiers. Typing contexts now gather assumptions on polarized variables, in addition to simple variables as before. There is however one restriction on the variables occurring in a context: x and x^+ (or x^-) cannot occur simultaneously in a given context Γ , even though x^+ and x^- may. New assumptions are added to contexts by means of an update operation $+$, defined in Figure 4. Context updating is different from splitting (in Figure 3) on what concerns unrestricted types: $\emptyset + (x : \text{un } T)$ is defined, whereas $\emptyset \cdot (x : \text{un } p)$ is not. We say that a context is *unlimited* if it contains no session types, and is *completed* if every session in it is end.

The typing relation is inductively defined by the rules in Figure 4. Rule [T-NEWS] requires the types for the two channel end points to be of dual types; contrast with rule [T-RES] in Figure 3: our system merges the two end points in a single variable and requires the two components of the channel type to be of dual types. Rules [T-IN] and [T-INS] in Figure 4 have their counterpart in rules [T-IN] and [T-INC] in Figure 3. The choice here is not based on whether the type for the input channel is an end point or a channel type but rather on whether the qualifier is linear or unrestricted. The same can be said of rules [T-OUT] and [T-OUTS].

From the above description it should be obvious that the two systems are quite close to each other. In order to define the typing correspondence we need to translate types and contexts for the polarized language (as in Figure 4) to those in our language (Figure 1). The definition is as follows; recall from Section 2 that we use $*?T$ as an abbreviation for $\mu a. ?T.a$, for some a not in T . To translate typing contexts we assume that if both x^+ and x^- are in Γ then they occur in contiguous positions (and in this order). The translation of typing contexts is as follows, where the rules must be tried in the given order; the first rule for mapping non-empty contexts is for polarized pairs while the second rule is for single entries.

$$\begin{aligned}
\llbracket T \rrbracket &= (*?\llbracket T \rrbracket, *!\llbracket T \rrbracket) & \llbracket \text{end} \rrbracket &= \text{un end} & \llbracket \emptyset \rrbracket &= \emptyset \\
\llbracket ?T.S \rrbracket &= \text{lin}?\llbracket T \rrbracket.\llbracket S \rrbracket & \llbracket a \rrbracket &= a & \llbracket \Gamma, x^+ : S, x^- : S' \rrbracket &= \llbracket \Gamma \rrbracket, x : (\llbracket S \rrbracket, \llbracket S' \rrbracket) \\
\llbracket !T.S \rrbracket &= \text{lin}!\llbracket T \rrbracket.\llbracket S \rrbracket & \llbracket \mu a.S \rrbracket &= \mu a.\llbracket S \rrbracket & \llbracket \Gamma, x^p : T \rrbracket &= \llbracket \Gamma \rrbracket, x : \llbracket T \rrbracket \\
&& \llbracket \mu a.T \rrbracket &= \mu a.\llbracket T \rrbracket & &
\end{aligned}$$

We are now in a position to state the main result of this section.

Theorem 2 (Polarity-Pi To Pi Correspondence).

1. If $\Gamma \vdash_p P$ then $\llbracket \Gamma \rrbracket \vdash \text{erase}(P)$.
2. If $P \rightarrow_p Q$, then $\text{erase}(P) \rightarrow \text{erase}(Q)$.

New syntactic forms

$c ::=$	Capabilities:	$T ::=$	Types:
i	input	$q c T$	channel
o	output	bool	boolean
io	input and output		

Combination of types

$$\text{bool} + \text{bool} = \text{bool} \quad \text{un } c_1 T + \text{un } c_2 T = \text{un } (c_1 \cup c_2) T \quad \text{lin } i T + \text{lin } o T = \text{lin } io T$$

Combination of contexts

$$(\Gamma_1 + \Gamma_2)(x) = \begin{cases} \Gamma_1(x) + \Gamma_2(x) & \text{if } x \in \text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) \\ \Gamma_1(x) & \text{if } x \in \text{dom}(\Gamma_1) \text{ and } x \notin \text{dom}(\Gamma_2) \\ \Gamma_2(x) & \text{if } x \in \text{dom}(\Gamma_2) \text{ and } x \notin \text{dom}(\Gamma_1) \end{cases}$$

Typing rules for processes

$$\frac{\Gamma \vdash_1 P_1 \quad \Gamma \vdash_1 P_2}{\Gamma + v : \text{bool} \vdash_1 \text{if } v \text{ then } P_1 \text{ else } P_2} \quad \frac{\Gamma, x : q \text{io } T \vdash_1 P}{\Gamma \vdash_1 (\nu x) P} \quad \begin{array}{l} [\text{T-IF}] \quad [\text{T-RES}] \\ \\ [\text{T-IN}] \quad [\text{T-OUT}] \end{array}$$

$$\frac{\Gamma, y : T \vdash_1 P}{\Gamma + x : q i T \vdash_1 x(y).P} \quad \frac{\Gamma \vdash_1 P}{\Gamma + x : q o T + v : T \vdash_1 \bar{x} v.P}$$

Fig. 5. Linear pi calculus

4 Embedding the Linear Pi Calculus

In this section we analyse (a synchronous variant of) the linear pi calculus [7] and provide a typing-preserving encoding into our system.

The syntax of linear pi processes and the reduction relation are described in Figure 1. Figure 5 defines the syntax of types and the typing rules for processes. Types have now the form $q c T$ where c is a capability formally defined as one of the following sets.

$$i = \{i\} \quad o = \{o\} \quad io = \{i, o\}$$

The linear discipline is imposed by way of a $+$ combination operation over types, defined in Figure 5. The operator is extended point-wise to typing contexts. Notice that context combination is different from the context splitting operation defined in Figure 3 when in the presence of unrestricted types: context splitting does not allow composing $(\Gamma_1, x : \text{un } c T)$ with Γ_2 whenever $x \notin \text{dom}(\Gamma_2)$ or when $\Gamma_2(x) \neq \text{un } c T$.

The typing system for the linear pi-calculus is defined by the rules in Figure 3 together with rules [T-INACT], [T-REPL] and [T-PAR] in Figure 4. Rule [T-OUT] is an adaptation of that in [7] to the synchronous setting: we let the continuation be typed with context Γ while in the original paper the premise to the rule is

$\text{un}(T)$ since the (absent) continuation behaves as $\mathbf{0}$. We also adapt rule [T-RES] to require that the restricted channel uses both capabilities; the original system allows processes of the form $(\nu x)\bar{x}\text{true}$ to be typed by assigning to channel x type $\text{lin } \circ \text{bool}$; cf. discussion around rule [T-RES] in Section 2.

The compositional encoding of linear types is defined below and is useful to understand the reconstruction of session types introduced in Section 2. A linear input (output) type is embedded as a linear input (output) type whose continuation is un end , meaning that the continuation process cannot further use the channel. Unrestricted input (output) types are mapped into unrestricted recursive input (output) types. For instance, the type $\text{lin } i(\text{lin } i \circ (\text{un } i \circ \text{bool}))$ is mapped into the type $\text{lin } ?(\text{lin } !T.\text{un end}, \text{lin } ?T.\text{un end}).\text{un end}$ where $T = (*! \text{bool}, *? \text{bool})$.

$$\begin{array}{ll} \llbracket \text{lin } iT \rrbracket = \text{lin } ?\llbracket T \rrbracket.\text{un end} & \llbracket \text{lin } oT \rrbracket = \text{lin } !\llbracket T \rrbracket.\text{un end} \\ \llbracket \text{un } iT \rrbracket = *?\llbracket T \rrbracket & \llbracket \text{un } oT \rrbracket = *!\llbracket T \rrbracket \\ \llbracket q \text{ io } T \rrbracket = (\llbracket q \text{ iT} \rrbracket, \llbracket q \text{ oT} \rrbracket) & \llbracket \text{bool} \rrbracket = \text{bool} \end{array}$$

The main result of this section establishes the correspondence between the two systems.

Theorem 3 (Linear-Pi To Pi Correspondence). *If $\Gamma \vdash_1 P$ then $\llbracket \Gamma \rrbracket \vdash P$.*

5 Embedding the Linear Lambda Calculus

This section shows that the call-by-value linear lambda calculus can be faithfully encoded in our language. We follow the presentation of Walker [16], except that we use an implicitly typed language.

The syntax of the language is in Figure 6; we rely on the set of variables introduced in Section 2 for the pi calculus; the missing non-terminal symbols, q , b and so on, are in Figure 1. Values are qualified, linear or unrestricted, and include boolean values and abstractions. Terms are variables, values, and applications and are evaluated in an abstract machine with an explicit store. The store is a sequence of variable-value pairs, treated as a map from variables into values. To simplify the presentation of the evaluation relation, we use an auxiliary function, $S \stackrel{q}{\sim} x$ that deallocates the value associated with variable x in S when the qualifier q is lin , and leaves S unchanged otherwise. The evaluation reduction copies values into the store, associating them with a fresh variable (rule [E-VAL]). For function application the value associated with the function is looked upon in the store; if linear it is then deallocated (rule [E-APP]). The remaining two rules implement the call-by-value strategy. We denote by \rightarrow_λ the reduction relation in Figure 6.

For typing, rule [T-VAR] is that of the pi-calculus (Figure 3). Rule [T-BOOL] contrasts with its homonymous in Figure 3 in that values in the linear lambda calculus are qualified, the type of a value inheriting the qualifier of the value. The remaining two rules, for abstraction and application, are standard in the linear lambda calculus; notice the $q(\Gamma)$ in rule [T-ABS] requiring an unrestricted

Syntax

$v ::=$	Values:	$p ::=$	Pretypes:
$q b$	boolean	\mathbf{bool}	boolean
$q \lambda x.M$	abstraction	$T \rightarrow T$	function
$M ::=$	Terms:	$T ::=$	Types:
x	variable	$q p$	qualified pretype
v	value	$S ::=$	Stores:
MM	application	\emptyset	empty store
		$S, x \mapsto v$	store binding

Store deallocation

$$(S_1, x \mapsto v, S_2) \stackrel{\text{lin}}{\sim} x = S_1, S_2 \qquad S \stackrel{\text{un}}{\sim} x = S$$

Evaluation

$$\begin{array}{c} (S; v) \rightarrow_\lambda (S, x \mapsto v; x) \qquad \frac{S(x_1) = q\lambda y.M}{(S; x_1x_2) \rightarrow_\lambda (S \stackrel{q}{\sim} x_1; M[x_2/y])} \quad [\text{E-VAL}] \quad [\text{E-APP}] \\ \frac{(S; M_1) \rightarrow_\lambda (S'; M'_1)}{(S; M_1M_2) \rightarrow_\lambda (S'; M'_1M_2)} \quad \frac{(S; M) \rightarrow_\lambda (S'; M')}{(S; xM) \rightarrow_\lambda (S'; xM')} \quad [\text{E-FUN}] \quad [\text{E-ARG}] \end{array}$$

Typing

$$\begin{array}{c} \frac{\text{un}(\Gamma)}{\Gamma, x: T \vdash_\lambda x: T} \qquad \frac{\text{un}(\Gamma)}{\Gamma \vdash_\lambda q b: q \mathbf{bool}} \quad [\text{T-VAR}] \quad [\text{T-BOOL}] \\ \frac{\Gamma, x: T_1 \vdash_\lambda M: T_2 \quad q(\Gamma)}{\Gamma \vdash_\lambda q\lambda x.M: qT_1 \rightarrow T_2} \quad \frac{\Gamma_1 \vdash_\lambda M_1: qT_1 \rightarrow T_2 \quad \Gamma_2 \vdash_\lambda M_2: T_1}{\Gamma_1 \cdot \Gamma_2 \vdash_\lambda M_1M_2: T_2} \quad [\text{T-ABS}] \quad [\text{T-APP}] \end{array}$$

Fig. 6. Linear lambda calculus

function to contain only unrestricted free variables ($\text{un}(\Gamma)$ is defined in Section 2; $\text{lin}(\Gamma)$ is true).

For the translation we rely on a polyadic variant of the pi language, allowing channels to carry an arbitrary (but fixed) number of values. The extension is straightforward to incorporate: for processes we need polyadic output and input, $x(\vec{x}).P$ and $\bar{x}\vec{v}.P$; for types (pre types, rather) we need $?\langle\vec{T}\rangle.S$ and $!\langle\vec{T}\rangle.S$. The extension can be incorporated in the base theory or added as an encoding [15].

On what concerns the translation of types below, the interesting cases are the two forms of arrow types. An unrestricted $T_1 \rightarrow T_2$ type is translated as a pair of types: the $*?X$ part caters for the resource (the function proper) and the $*!X$ for its clients. Channels describing functions carry a pair X of values: the first element in the pair is the argument to the function ($\llbracket T_1 \rrbracket$ in X); the second is a channel that will convey the result and that will be used exactly once (a linear channel of type $\text{lin}!\llbracket T_2 \rrbracket.\text{un end}$). The type for linear resources is similar, only that

they are linear, rather than unrestricted. The translation of terms follows that of Milner [9] with two exceptions. On the one hand, the value qualifiers are taken into consideration in the translation: a linear value is translated into a simple output (in the case of a boolean value) or a simple input (in the case of an abstraction); only for unrestricted values replication is used. On the other hand, applications of the form xM are partially evaluated, allowing for a simple operational correspondence (cf. [14]).

$$\begin{array}{ll}
\llbracket q \text{ bool} \rrbracket = q \text{ bool} & \llbracket x \rrbracket_p = \bar{p} x \\
\llbracket \text{un } T_1 \rightarrow T_2 \rrbracket = (*?X, *!X) & \llbracket v \rrbracket_p = (\nu x)(\llbracket x \mapsto v \rrbracket \mid \bar{p} x) \\
\llbracket \text{lin } T_1 \rightarrow T_2 \rrbracket = (\text{lin}?X.\text{un end}, \text{lin}!X.\text{un end}) & \llbracket xM \rrbracket_p = (\nu r)(\llbracket M \rrbracket_r \mid r(y).\bar{x} y p) \\
\text{where } X = \langle \llbracket T_1 \rrbracket, \text{lin}!\llbracket T_2 \rrbracket.\text{un end} \rangle & \llbracket MN \rrbracket_p = (\nu s)(\llbracket M \rrbracket_s \mid s(x).(\nu r)(\llbracket N \rrbracket_r \mid r(y).\bar{x} y p)) \\
\llbracket \emptyset \rrbracket = \mathbf{0} & \\
\llbracket S, x \mapsto v \rrbracket = \llbracket S \rrbracket \mid \llbracket x \mapsto v \rrbracket & \llbracket S; M \rrbracket = (\nu \text{dom}(S))(\llbracket S \rrbracket \mid \llbracket M \rrbracket) \\
\llbracket x \mapsto q b \rrbracket = \llbracket q \rrbracket \bar{x} b & \llbracket \text{un} \rrbracket = * \\
\llbracket x \mapsto q \lambda y M \rrbracket = \llbracket q \rrbracket x(y p).\llbracket M \rrbracket_p & \llbracket \text{lin} \rrbracket = \text{the empty string}
\end{array}$$

We are now in a position to state the main result of this section. Let \rightarrow^* be the reflexive and transitive closure of the reduction relation \rightarrow defined in Figure 1.

Theorem 4 (Linear-Lambda to Pi Correspondence).

1. If $\Gamma \vdash_\lambda M : T$, then $\llbracket \Gamma \rrbracket, p : \text{lin}!\llbracket T \rrbracket.\text{un end} \vdash \llbracket M \rrbracket_p$.
2. If $(S; M) \rightarrow_\lambda (S'; M')$, then $\llbracket S; M \rrbracket \rightarrow^* \llbracket S'; M' \rrbracket$.

6 Conclusion

As mentioned in the introduction, the pi calculus equipped with a polarity-based typing system [5] and the (double binder) pi calculus equipped with a conventional typing system [15] are the works closest to ours. Here we try to obtain the same results, relying on the traditional pi calculus equipped with a conventional type system. Towards this end we introduce an unordered pair constructor denoting, at type level, the two ends of a same channel. In order to distinguish linear from unrestricted variables we use type qualifiers applied to pre types, inspired by Walker’s presentation of substructural type systems [16]. Caires and Pfenning take a different approach, closely adhering to linear logic, treating all variables as linear and using exponential “!” to describe shared resources [1].

Algorithmic type checking is left for further work; the language with double binders [15] is equipped with such a system. Gay and Hole address the problem of polarity inference for closed processes under certain restrictions [5]. Particularly promising is the F° (“F-pop”) system by Maruzak et al. [8], where a kinding system, instead of type qualifiers, simplifies the use of linearity in functional programming languages, including a novel form of subtyping between linear and unrestricted kinds, which we would like to explore.

Acknowledgments. This work was supported by FCT/MCTES via projects PTDC/EIA-CCO/105359/2008 and CMU-PT/NGN44-2009-12. The authors would like thank Simon Gay and Nils Gesbert for insightful comments and the anonymous referees for constructive criticisms and detailed comments.

References

1. Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In this volume.
2. Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, Elena Giachino, and Luca Padovani. Foundations of session types. In *PPDP*, pages 219–230. ACM, 2009.
3. Mariangiola Dezani-Ciancaglini and Ugo de'Liguoro. Sessions and session types: an overview. In *WS-FM'09*, volume 6194 of *LNCS*, pages 1–28. Springer-Verlag, 2010.
4. Mariangiola Dezani-Ciancaglini, Sophia Drossopoulou, Dimitris Mostrous, and Nobuko Yoshida. Objects and session types. *Information and Computation*, 207(5):595 – 641, 2009.
5. Simon J. Gay and Malcolm J. Hole. Subtyping for session types in the pi calculus. *Acta Informatica*, 42(2/3):191–225, 2005.
6. Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP'98*, volume 1381 of *LNCS*, pages 22–138. Springer-Verlag, 1998.
7. Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. Linearity and the pi-calculus. *ACM Transactions on Programming Languages and Systems*, 21(5):914–947, 1999.
8. Karl Mazurak, Jianzhou Zhao, and Steve Zdancewic. Lightweight linear types in system F° . In *TLDI*, pages 77–88. ACM, 2010.
9. Robin Milner. Functions as processes. *Mathematical Structures in Computer Science*, 2(2):119–141, 1992.
10. Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, part I/II. *Information and Computation*, 100:1–77, September 1992.
11. Luca Padovani. Session types at the mirror. *EPTCS*, 12:71–86, 2009.
12. Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
13. Kaku Takeuchi, Kohei Honda, and Makoto Kubo. An Interaction-based Language and its Typing System. In *PARLE*, volume 817 of *LNCS*, pages 398–413. Springer-Verlag, 1994.
14. Vasco T. Vasconcelos. Lambda and pi calculi, CAM and SECD machines. *Journal of Functional Programming*, 15(1):101–127, 2005.
15. Vasco T. Vasconcelos. *SFM 2009*, volume 5569 of *LNCS*, chapter Fundamentals of Session Types, pages 158–186. Springer-Verlag, 2009.
16. David Walker. *Advanced Topics in Types and Programming Languages*, chapter Substructural Type Systems. MIT Press, 2005.
17. Nobuko Yoshida and Vasco T. Vasconcelos. Language primitives and type discipline for structured communication-based programming revisited: Two systems for higher-order session communication. In *SecReT'07*, volume 171(4) of *ENTCS*, pages 73–93. Elsevier Science Publishers, 2007.