Types for Deadlock-Free Higher-Order Programs

Luca Padovani, Luca Novara

► To cite this version:

Luca Padovani, Luca Novara. Types for Deadlock-Free Higher-Order Programs. 35th International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE), Jun 2015, Grenoble, France. pp.3-18, 10.1007/978-3-319-19195-9_1 . hal-01767327

HAL Id: hal-01767327

<https://hal.inria.fr/hal-01767327>

Submitted on 16 Apr 2018

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

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

Distributed under a Creative Commons Attribution| 4.0 International License

Types for Deadlock-Free Higher-Order Programs

Luca Padovani and Luca Novara

Dipartimento di Informatica, Università di Torino, Italy

Abstract. Type systems for communicating processes are typically studied using abstract models – *e.g.*, *process algebras* – that distill the communication behavior of programs but overlook their structure in terms of functions, methods, objects, modules. It is not always obvious how to apply these type systems to structured programming languages. In this work we port a recently developed type system that ensures *deadlock freedom* in the π -calculus to a higher-order language.

1 Introduction

In this article we develop a type system that guarantees well-typed programs that communicate over channels to be free from deadlocks. Type systems ensuring this property already exist [7,8,10], but they all use the π -calculus as the reference language. This choice overlooks some aspects of concrete programming languages, like the fact that programs are structured into compartmentalized blocks (*e.g.*, functions) within which only the local structure of the program (the body of a function) is visible to the type system, and little if anything is known about the exterior of the block (the callers of the function). The structure of programs may hinder some kinds of analysis: for example, the type systems in [7,8,10] enforce an ordering of communication events and to do so they take advantage of the nature of π -calculus processes, where programs are flat sequences of communication actions. How do we reason on such ordering when the execution order is dictated by the reduction strategy of the language rather than by the syntax of programs, or when events occur within a function, and nothing is known about the events that are supposed to occur after the function terminates? We answer these questions by porting the type system in [10] to a higher-order functional language.

To illustrate the key ideas of the approach, let us consider the program

$$\langle \text{send } a \text{ (recv } b) \rangle \mid \langle \text{send } b \text{ (recv } a) \rangle \quad (1.1)$$

consisting of two parallel threads. The thread on the left is trying to send the message received from channel b on channel a ; the thread on the right is trying to do the opposite. The communications on a and b are mutually dependent, and the program is a deadlock. The basic idea used in [10] and derived from [7,8] for detecting deadlocks is to assign each channel a number – which we call *level* – and to verify that channels are used in order according to their levels. In (1.1) this mechanism requires b to have smaller level than a in the leftmost thread, and a to have a smaller level than b in the rightmost thread. No level assignment can simultaneously satisfy both constraints. In order to perform these checks with a type system, the first step is to attach levels to channel types. We therefore assign the types $![\text{int}]^m$ and $?[\text{int}]^n$ respectively to a and b

in the leftmost thread of (1.1), and $?[\text{int}]^m$ and $![\text{int}]^n$ to the same channels in the rightmost thread of (1.1). Crucially, distinct occurrences of the same channel have types with opposite polarities (input $?$ and output $!$) and equal level. We can also think of the assignments $\text{send} : \forall t. ![\text{int}]^t \rightarrow \text{int} \rightarrow \text{unit}$ and $\text{recv} : \forall t. ?[\text{int}]^t \rightarrow \text{int}$ for the communication primitives, where we allow polymorphism on channel levels. In this case, the application $\text{send } a \text{ (recv } b)$ consists of two subexpressions, the partial application $\text{send } a$ having type $\text{int} \rightarrow \text{unit}$ and its argument $\text{recv } b$ having type int . Neither of these types hints at the I/O operations performed in these expressions, let alone at the levels of the channels involved. To recover this information we pair types with *effects* [1]: the effect of an expression is an abstract description of the operations performed during its evaluation. In our case, we take as effect the level of channels used for I/O operations, or $\perp$ in the case of pure expressions that perform no I/O. So, the judgment

$$b : ?[\text{int}]^n \vdash \text{recv } b : \text{int} \& n$$

states that $\text{recv } b$ is an expression of type int whose evaluation performs an I/O operation on a channel with level n . As usual, function types are decorated with a *latent effect* saying what happens when the function is applied to its argument. So,

$$a : ![\text{int}]^m \vdash \text{send } a : \text{int} \rightarrow^m \text{unit} \& \perp$$

states that $\text{send } a$ is a function that, applied to an argument of type int , produces a result of type unit and, in doing so, performs an I/O operation on a channel with level m . By itself, $\text{send } a$ is a pure expression whose evaluation performs no I/O operations, hence the effect $\perp$. Effects help us detecting dangerous expressions: in a *call-by-value* language an application $e_1 e_2$ evaluates e_1 first, then e_2 , and finally the body of the function resulting from e_1 . Therefore, the channels used in e_1 must have smaller level than those occurring in e_2 and the channels used in e_2 must have smaller level than those occurring in the body of e_1 . In the specific case of $\text{send } a \text{ (recv } b)$ we have $\perp < n$ for the first condition, which is trivially satisfied, and $n < m$ for the second one. Since the same reasoning on $\text{send } b \text{ (recv } a)$ also requires the symmetric condition ($m < n$), we detect that the parallel composition of the two threads in (1.1) is ill typed, as desired.

It turns out that the information given by latent effects in function types is not sufficient for spotting some deadlocks. To see why, consider the function

$$f \stackrel{\text{def}}{=} \lambda x. (\text{send } a \ x; \text{send } b \ x)$$

which sends its argument x on both a and b and where $;$ denotes sequential composition. The level of a (say m) should be smaller than the level of b (say n), for a is used before b (we assume that communication is synchronous and that send is a potentially blocking operation). The question is, what is the latent effect that decorates the type of f , of the form $\text{int} \rightarrow^h \text{unit}$? Consider the two obvious possibilities: if we take $h = m$, then

$$\langle \text{recv } a \rangle \mid \langle f \ 3; \text{recv } b \rangle \tag{1.2}$$

is well typed because the effect m of $f \ 3$ is smaller than the level of b in $\text{recv } b$, which agrees with the fact that $f \ 3$ is evaluated *before* $\text{recv } b$; if we take $h = n$, then

$$\langle \text{recv } a; f \ 3 \rangle \mid \langle \text{recv } b \rangle \tag{1.3}$$

is well typed for similar reasons. This is unfortunate because both (1.3) and (1.2) reduce to a deadlock. To flag both of them as ill typed, we must refine the type of f to $\text{int} \rightarrow^{m,n} \text{unit}$ where we distinguish the smallest level of the channels that *occur* in the body of f (that is m) from the greatest level of the channels that *are used* by f when f is applied to an argument (that is n). The first annotation gives information on the channels in the function's closure, while the second annotation is the function's latent effect, as before. So (1.2) is ill typed because the effect of f 3 is the same as the level of b in $\text{recv } b$ and (1.3) is ill typed because the effect of $\text{recv } a$ is the same as the level of f in f 3.

In the following, we define a core multithreaded functional language with communication primitives (Section 2), we present a basic type and effect system, extend it to address recursive programs, and state its properties (Section 3). Finally, we briefly discuss closely related work and a few extensions (Section 4). *Proofs and additional material can be found in long version of the paper, on the first author's home page.*

2 Language syntax and semantics

In defining our language, we assume a synchronous communication model based on linear channels. This assumption limits the range of systems that we can model. However, asynchronous and structured communications can be encoded using linear channels: this has been shown to be the case for binary sessions [5] and for multiparty sessions to a large extent [10, technical report].

We use a countable set of *variables* $x, y, \dots$, a countable set of *channels* $a, b, \dots$, and a set of constants k . *Names* $u, \dots$ are either variables or channels. We consider a language of *expressions* and *processes* as defined below:

$$e ::= k \mid u \mid \lambda x.e \mid ee \qquad P, Q ::= \langle e \rangle \mid (\nu a)P \mid P \mid Q$$

Expressions comprise constants k , names u , abstractions $\lambda x.e$, and applications $e_1 e_2$. We write $_$ for unused/fresh variables. Constants include the unitary value $()$, the integer numbers $m, n, \dots$, as well as the primitives **fix**, **fork**, **new**, **send**, **recv** whose semantics will be explained shortly. Processes are either threads $\langle e \rangle$, or the restriction $(\nu a)P$ of a channel a with scope P , or the parallel composition $P \mid Q$ of processes.

The notions of free and bound names are as expected, given that the only binders are λ 's and ν 's. We identify terms modulo renaming of bound names and we write $\text{fn}(e)$ (respectively, $\text{fn}(P)$) for the set of names occurring free in e (respectively, in P).

The reduction semantics of the language is given by two relations, one for expressions, another for processes. We adopt a *call-by-value* reduction strategy, for which we need to define *reduction contexts* $\mathcal{E}, \dots$ and *values* $v, w, \dots$ respectively as:

$$\mathcal{E} ::= [] \mid \mathcal{E}e \mid \nu \mathcal{E} \qquad v, w ::= k \mid a \mid \lambda x.e \mid \text{send } v$$

The reduction relation $\longrightarrow$ for expressions is defined by standard rules

$$(\lambda x.e)v \longrightarrow e\{v/x\} \qquad \text{fix } \lambda x.e \longrightarrow e\{\text{fix } \lambda x.e/x\}$$

and closed under reduction contexts. As usual, $e\{e'/x\}$ denotes the capture-avoiding substitution of e' for the free occurrences of x in e .

Table 1. Reduction semantics of expressions and processes.

$\langle \mathcal{E}[\text{send } a \ v] \rangle \mid \langle \mathcal{E}'[\text{recv } a] \rangle \xrightarrow{a} \langle \mathcal{E}[] \rangle \mid \langle \mathcal{E}'[v] \rangle \quad \langle \mathcal{E}[\text{fork } v] \rangle \xrightarrow{\tau} \langle \mathcal{E}[] \rangle \mid \langle v() \rangle$	
$\frac{}{\langle \mathcal{E}[\text{new } ()] \rangle \xrightarrow{\tau} (va) \langle \mathcal{E}[a] \rangle} \quad a \notin \text{fn}(\mathcal{E}) \quad \frac{e \longrightarrow e'}{\langle e \rangle \xrightarrow{\tau} \langle e' \rangle}$	
$\frac{P \xrightarrow{\ell} P'}{P \mid Q \xrightarrow{\ell} P' \mid Q}$	$\frac{P \xrightarrow{\ell} Q}{(va)P \xrightarrow{\ell} (va)Q} \quad \ell \neq a$
$\frac{P \xrightarrow{a} Q}{(va)P \xrightarrow{\tau} Q}$	$\frac{P \equiv \xrightarrow{\ell} \equiv Q}{P \xrightarrow{\ell} Q}$

The reduction relation of processes (Table 1) has *labels* $\ell, \dots$ that are either a channel name a , signalling that a communication has occurred on a , or the special symbol τ denoting any other reduction. There are four base reductions for processes: a communication occurs between two threads when one is willing to send a message v on a channel a and the other is waiting for a message from the same channel; a thread that contains a subexpression `fork` v spawns a new thread that evaluates $v()$; a thread that contains a subexpression `new`() creates a new channel; the reduction of an expression causes a corresponding τ -labeled reduction of the thread in which it occurs. Reduction for processes is then closed under parallel compositions, restrictions, and structural congruence. The restriction of a disappears as soon as a communication on a occurs: in our model channels are *linear* and can be used for one communication only; structured forms of communication can be encoded on top of this simple model (see Example 2 and [5]). Structural congruence is defined by the standard rules rearranging parallel compositions and channel restrictions, where $\langle () \rangle$ plays the role of the inert process.

We conclude this section with two programs written using a slightly richer language equipped with `let` bindings, conditionals, and a few additional operators. All these constructs either have well-known encodings or can be easily accommodated.

Example 1 (parallel Fibonacci function). The `fib` function below computes the n -th number in the Fibonacci sequence and sends the result on a channel c :

```

1  fix λfib.λn.λc.if n ≤ 1 then send c n
2                      else let a = new() and b = new() in
3                          (fork λ_.fib (n - 1) a);
4                          (fork λ_.fib (n - 2) b);
5                          send c (recv a + recv b)

```

The fresh channels a and b are used to collect the results from the recursive, parallel invocations of `fib`. Note that expressions are intertwined with I/O operations. It is relevant to ask whether this version of `fib` is deadlock free, namely if it is able to reduce until a result is computed without blocking indefinitely on an I/O operation. ■

Example 2 (signal pipe). In this example we implement a function `pipe` that forwards signals received from an input stream x to an output stream y :

```

1  let cont = λx.let c = new() in (fork λ_.send x c); c in
2  let pipe = fix λpipe.λx.λy.pipe (recv x) (cont y)

```

Note that this pipe is only capable of forwarding handshaking signals. A more interesting pipe transmitting actual data can be realized by considering data types such as records and sums [5]. The simplified realization we consider here suffices to illustrate a relevant family of recursive functions that interleave actions on different channels.

Since linear channels are consumed after communication, each signal includes a *continuation channel* on which the subsequent signals in the stream will be sent/received. In particular, `cont x` sends a fresh continuation `c` on `x` and returns `c`, so that `c` can be used for subsequent communications, while `pipe x y` sends a fresh continuation on `y` after it has received a continuation from `x`, and then repeats this behavior on the continuations. The program below connects two pipes:

```
3  let a = new() and b = new() in
4    (fork λ_.pipe a b); (fork λ_.pipe b (cont a))
```

Even if the two pipes realize a cyclic network, we will see in Section 3 that this program is well typed and therefore deadlock free. Forgetting `cont` on line 4 or not forking the `send` on line 1, however, produces a deadlock. ■

3 Type and effect system

We present the features of the type system gradually, in three steps: we start with a monomorphic system (Section 3.1), then we introduce level polymorphism required by Examples 1 and 2 (Section 3.2), and finally recursive types required by Example 2 (Section 3.3). We end the section studying the properties of the type system (Section 3.4).

3.1 Core types

Let $\mathbb{L} \stackrel{\text{def}}{=} \mathbb{Z} \cup \{\perp, \top\}$ be the set of *channel levels* ordered in the obvious way ($\perp < n < \top$ for every $n \in \mathbb{Z}$); we use $\rho, \sigma, \dots$ to range over $\mathbb{L}$ and we write $\rho \sqcap \sigma$ (respectively, $\rho \sqcup \sigma$) for the *minimum* (respectively, the *maximum*) of ρ and σ . *Polarities* $p, q, \dots$ are non-empty subsets of $\{?, !\}$; we abbreviate $\{?\}$ and $\{!\}$ with $?$ and $!$ respectively, and $\{?, !\}$ with $\#$. *Types* $t, s, \dots$ are defined by

$$t, s ::= \mathbf{B} \mid p[t]^n \mid t \rightarrow^{\rho, \sigma} s$$

where *basic types* $\mathbf{B}, \dots$ include `unit` and `int`. The type $p[t]^n$ denotes a channel with polarity p and level n . The polarity describes the operations allowed on the channel: $?$ means input, $!$ means output, and $\#$ means both input and output. Channels are linear resources: they can be used once according to each element in their polarity. The type $t \rightarrow^{\rho, \sigma} s$ denotes a function with domain t and range s . The function has level ρ (its closure contains channels with level ρ or greater) and, when applied, it uses channels with level σ or smaller. If $\rho = \top$, the function has no channels in its closure; if $\sigma = \perp$, the function uses no channels when applied. We write $\rightarrow$ as an abbreviation for $\rightarrow^{\top, \perp}$, so $\rightarrow$ denotes pure functions not containing and not using any channel.

Recall from Section 1 that levels are meant to impose an order on the use of channels: roughly, the lower the level of a channel, the sooner the channel must be used. We

extend the notion of level from channel types to arbitrary types: basic types have level $\top$ because there is no need to use them as far as deadlock freedom is concerned; the level of functions is written in their type. Formally, the level of t , written $|t|$, is defined as:

$$|B| \stackrel{\text{def}}{=} \top \quad |p[t]^n| \stackrel{\text{def}}{=} n \quad |t \rightarrow^{\rho, \sigma} s| \stackrel{\text{def}}{=} \rho \quad (3.1)$$

Levels can be used to distinguish *linear types*, denoting values (such as channels) that *must* be used to guarantee deadlock freedom, from *unlimited types*, denoting values that have no effect on deadlock freedom and *may* be disregarded. We say that t is *linear* if $|t| \in \mathbb{Z}$; we say that t is *unlimited*, written $\text{un}(t)$, if $|t| = \top$.

Below are the type schemes of the constants that we consider. Some constants have many types (constraints are on the right); we write $\text{types}(\mathbf{k})$ for the *set* of types of $\mathbf{k}$.

$$\begin{array}{llll} () : \text{unit} & \text{fix} : (t \rightarrow t) \rightarrow t & \text{new} : \text{unit} \rightarrow \# [t]^n & n < |t| \\ n : \text{int} & \text{fork} : (\text{unit} \rightarrow^{\rho, \sigma} \text{unit}) \rightarrow \text{unit} & \text{recv} : ?[t]^n \rightarrow \top, n t & n < |t| \\ & & \text{send} : ![t]^n \rightarrow t \rightarrow^{n, n} \text{unit} & n < |t| \end{array}$$

The type of $()$, of the numbers, and of **fix** are ordinary. The primitive **new** creates a fresh channel with the full set $\#$ of polarities and arbitrary level n . The primitive **recv** takes a channel of type $?[t]^n$, blocks until a message is received, and returns the message. The primitive itself contains no free channels in its closure (hence the level $\top$) because the only channel it manipulates is its argument. The latent effect is the level of the channel, as expected. The primitive **send** takes a channel of type $![t]^n$, a message of type t , and sends the message on the channel. Note that the partial application **send** a is a function whose level and latent effect are both the level of a . Note also that in **new**, **recv**, and **send** the level of the message must be greater than the level of the channel: since levels are used to enforce an order on the use of channels, this condition follows from the observation that a message cannot be used until *after* it has been received, namely after the channel on which it travels has been used. Finally, **fork** accepts a thunk with arbitrary level ρ and latent effect σ and spawns the thunk into an independent thread (see Table 1). Note that **fork** is a pure function with no latent effect, regardless of the level and latent effect of the thunk. This phenomenon is called *effect masking* [1], whereby the effect of evaluating an expression becomes unobservable: in our case, **fork** discharges effects because the thunk runs in parallel with the code executing the **fork**.

We now turn to the typing rules. A *type environment* Γ is a finite map $u_1 : t_1, \dots, u_n : t_n$ from names to types. We write $\emptyset$ for the empty type environment, $\text{dom}(\Gamma)$ for the domain of Γ , and $\Gamma(u)$ for the type associated with u in Γ ; we write Γ_1, Γ_2 for the union of Γ_1 and Γ_2 when $\text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2) = \emptyset$. We also need a more flexible way of combining type environments. In particular, we make sure that every channel is used linearly by distributing different polarities of a channel to different parts of the program. To this aim, following [9], we define a partial *combination* operator $+$ between types:

$$\begin{array}{ll} t + t \stackrel{\text{def}}{=} t & \text{if } \text{un}(t) \\ p[t]^n + q[t]^n \stackrel{\text{def}}{=} (p \cup q)[t]^n & \text{if } p \cap q = \emptyset \end{array} \quad (3.2)$$

that we extend to type environments, thus:

$$\begin{array}{ll} \Gamma + \Gamma' \stackrel{\text{def}}{=} \Gamma, \Gamma' & \text{if } \text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset \\ (\Gamma, u : t) + (\Gamma', u : s) \stackrel{\text{def}}{=} (\Gamma + \Gamma'), u : t + s & \end{array} \quad (3.3)$$

For example, we have $(x : \text{int}, a : ![\text{int}]^n) + (a : ?[\text{int}]^n) = x : \text{int}, a : \#[\text{int}]^n$, so we might have some part of the program that (possibly) uses a variable x of type int along with channel a for sending an integer and another part of the program that uses the same channel a but this time for receiving an integer. The first part of the program would be typed in the environment $x : \text{int}, a : ![\text{int}]^n$ and the second one in the environment $a : ?[\text{int}]^n$. Overall, the two parts would be typed in the environment $x : \text{int}, a : \#[\text{int}]^n$ indicating that a is used for both sending *and* receiving an integer.

We extend the function $|\cdot|$ to type environments so that $|\Gamma| \stackrel{\text{def}}{=} \prod_{u \in \text{dom}(\Gamma)} |\Gamma(u)|$ with the convention that $|\emptyset| = \top$; we write $\text{un}(\Gamma)$ if $|\Gamma| = \top$.

Table 2. Core typing rules for expressions and processes.

Typing of expressions		
$\frac{}{\Gamma, u : t \vdash u : t \& \perp} \quad \text{un}(\Gamma)$	$\frac{}{\Gamma \vdash \mathbf{k} : t \& \perp} \quad \text{un}(\Gamma)$	$t \in \text{types}(\mathbf{k})$
$\frac{\Gamma, x : t \vdash e : s \& \rho}{\Gamma \vdash \lambda x. e : t \rightarrow^{ \Gamma , \rho} s \& \perp} \quad \text{[T-FUN]}$	$\frac{\Gamma_1 \vdash e_1 : t \rightarrow^{\rho, \sigma} s \& \tau_1 \quad \Gamma_2 \vdash e_2 : t \& \tau_2 \quad \tau_1 < \Gamma_2 }{\Gamma_1 + \Gamma_2 \vdash e_1 e_2 : s \& \sigma \sqcup \tau_1 \sqcup \tau_2} \quad \tau_2 < \rho \quad \text{[T-APP]}$	
Typing of processes		
$\frac{}{\Gamma \vdash e : \text{unit} \& \rho} \quad \text{[T-THREAD]}$	$\frac{\Gamma_1 \vdash P \quad \Gamma_2 \vdash Q}{\Gamma_1 + \Gamma_2 \vdash P \mid Q} \quad \text{[T-PAR]}$	$\frac{\Gamma, a : \#[t]^n \vdash P}{\Gamma \vdash (va)P} \quad \text{[T-NEW]}$

We are now ready to discuss the core typing rules, shown in Table 2. Judgments of the form $\Gamma \vdash e : t \& \rho$ denote that e is well typed in Γ , it has type t and effect ρ ; judgments of the form $\Gamma \vdash P$ simply denote that P is well typed in Γ .

Axioms [T-NAME] and [T-CONST] are unremarkable: as in all substructural type systems the unused part of the type environment must be unlimited. Names and constants have no effect ($\perp$); they are evaluated expressions that do not use (but may contain) channels.

In rule [T-FUN], the effect ρ caused by evaluating the body of the function becomes the latent effect in the arrow type of the function and the function itself has no effect. The level of the function is determined by that of the environment Γ in which the function is typed. Intuitively, the names in Γ are stored in the *closure* of the function; if any of these names is a channel, then we must be sure that the function is eventually used (*i.e.*, applied) to guarantee deadlock freedom. In fact, $|\Gamma|$ gives a slightly more precise information, since it records the smallest level of all channels that occur in the body of the function. We have seen in Section 1 why this information is useful. A few examples:

- the identity function $\lambda x. x$ has type $\text{int} \rightarrow^{\top, \perp} \text{int}$ in any unlimited environment;
- the function $\lambda_. a$ has type $\text{unit} \rightarrow^{n, \perp} ![\text{int}]^n$ in the environment $a : ![\text{int}]^n$; it contains channel a with level n in its closure (whence the level n in the arrow), but it does

- not use a for input/output (whence the latent effect $\perp$); it is nonetheless well typed because a , which is a linear value, is returned as result;
- the function $\lambda x. \text{send } x \ 3$ has type $![\text{int}]^n \rightarrow^{\top, n} \text{unit}$; it has no channels in its closure but it performs an output on the channel it receives as argument;
- the function $\lambda x. (\text{recv } a + x)$ has type $\text{int} \rightarrow^{n, n} \text{int}$ in the environment $a : ?[\text{int}]^n$; note that neither the domain nor the codomain of the function mention any channel, so the fact that the function has a channel in its closure (and that it performs some I/O) can only be inferred from the annotations on the arrow;
- the function $\lambda x. \text{send } x \ (\text{recv } a)$ has type $![\text{int}]^{n+1} \rightarrow^{n, n+1} \text{unit}$ in the environment $a : ![\text{int}]^n$; it contains channel a with level n in its closure and performs input/output operations on channels with level $n + 1$ (or smaller) when applied.

Rule $[\text{T-APP}]$ deals with applications $e_1 e_2$. The first thing to notice is the type environments in the premises for e_1 and e_2 . Normally, these are exactly the same as the type environment used for the whole application. In our setting, however, we want to distribute polarities in such a way that each channel is used for exactly one communication. For this reason, the type environment $\Gamma_1 + \Gamma_2$ in the conclusion is the combination of the type environments in the premises. Regarding effects, τ_i is the effect caused by the evaluation of e_i . As expected, e_1 must result in a function of type $t \rightarrow^{\rho, \sigma} s$ and e_2 in a value of type t . The evaluation of e_1 and e_2 may however involve blocking I/O operations on channels, and the two side conditions make sure that no deadlock can arise. To better understand them, recall that reduction is *call-by-value* and applications $e_1 e_2$ are evaluated *sequentially from left to right*. Now, the condition $\tau_1 < |\Gamma_2|$ makes sure that any I/O operation performed during the evaluation of e_1 involves only channels whose level is smaller than that of the channels occurring free in e_2 (the free channels of e_2 must necessarily be in Γ_2). This is enough to guarantee that the functional part of the application can be fully evaluated without blocking on operations concerning channels that occur *later* in the program. In principle, this condition should be paired with the symmetric one $\tau_2 < |\Gamma_1|$ making sure that any I/O operation performed during the evaluation of the argument does not involve channels that occur in the functional part. However, when the argument is being evaluated, we know that the functional part has already been reduced a value (see the definition of reduction contexts in Section 2). Therefore, the only really critical condition to check is that no channels involved in I/O operations during the evaluation of e_2 occur in the *value* of e_1 . This is expressed by the condition $\tau_2 < \rho$, where ρ is the level of the functional part. Note that, when e_1 is an abstraction, by rule $[\text{T-FUN}]$ ρ coincides with $|\Gamma_1|$, but in general ρ may be greater than $|\Gamma_1|$, so the condition $\tau_2 < \rho$ gives better accuracy. The effect of the whole application $e_1 e_2$ is, as expected, the combination of the effects of evaluating e_1 , e_2 , and the latent effect of the function being applied. In our case the “combination” is the greatest level of any channel involved in the application. Below are some examples:

- $(\lambda x.x) \ a$ is well typed, because both $\lambda x.x$ and a are pure expressions whose effect is $\perp$, hence the two side conditions of $[\text{T-APP}]$ are trivially satisfied;
- $(\lambda x.x) \ (\text{recv } a)$ is well typed in the environment $a : ?[\text{int}]^n$: the effect of $\text{recv } a$ is n (the level of a) which is smaller than the level $\top$ of the function;
- $\text{send } a \ (\text{recv } a)$ is ill typed in the environment $a : \#[\text{int}]^n$ because the effect of evaluating $\text{recv } a$, namely n , is the same as the level of $\text{send } a$;

- $(\text{recv } a) (\text{recv } b)$ is well typed in the environment $a : ?[\text{int} \rightarrow \text{int}]^0, b : ?[\text{int}]^1$. The effect of the argument is 1, which is *not* smaller than the level of the environment $a : ?[\text{int} \rightarrow \text{int}]^0$ used for typing the function. However, 1 is smaller than $\top$, which is the level of the *result* of the evaluation of the functional part of the application. This application would be illegal had we used the side condition $\tau_2 < |\Gamma_1|$ in [T-APP].

The typing rules for processes are standard: [T-PAR] splits contexts for typing the processes in parallel, [T-NEW] introduces a new channel in the environment, and [T-THREAD] types threads. The effect of threads is ignored: effects are used to prevent circular dependencies between channels used within the *sequential* parts of the program (*i.e.*, within expressions); circular dependencies that arise between *parallel* threads are indirectly detected by the fact that each occurrence of a channel is typed with the same level (see the discussion of (1.1) in Section 1).

3.2 Level polymorphism

Looking back at Example 1, we notice that `fibonacci` may generate two recursive calls with two corresponding fresh channels `a` and `b`. Since the `send` operation on `c` is blocked by `recv` operations on `a` and `b` (line 5), the level of `a` and `b` must be smaller than that of `c`. Also, since expressions are evaluated left-to-right and `recv a + recv b` is syntactic sugar for the application $(+)$ $(\text{recv } a) (\text{recv } b)$, the level of `a` must be smaller than that of `b`. Thus, to declare `fibonacci` well typed, we must allow different occurrences of `fibonacci` to be applied to channels with different levels. Even more critically, this form of level polymorphism of `fibonacci` is necessary *within* the definition of `fibonacci` itself, so it is an instance of *polymorphic recursion* [1].

The core typing rules in Table 2 do not support level polymorphism. Following the previous discussion on `fibonacci`, the idea is to realize level polymorphism by *shifting* levels in types. We define level shifting as a type operator $\uparrow^n$, thus:

$$\uparrow^n \mathbf{B} \stackrel{\text{def}}{=} \mathbf{B} \quad \uparrow^n p[t]^m \stackrel{\text{def}}{=} p[\uparrow^n t]^{n+m} \quad \uparrow^n (t \rightarrow^{\rho, \sigma} s) \stackrel{\text{def}}{=} \uparrow^n t \rightarrow^{n+\rho, n+\sigma} \uparrow^n s \quad (3.4)$$

where $+$ is extended from integers to levels so that $n + \top = \top$ and $n + \perp = \perp$. The effect of $\uparrow^n t$ is to shift all the finite level annotations in t by n , leaving $\top$ and $\perp$ unchanged.

Now, we have to understand in which cases we can use a value of type $\uparrow^n t$ where one of type t is expected. More specifically, when a value of type $\uparrow^n t$ can be passed to a function expecting an argument of type t . This is possible if the function has level $\top$. We express this form of level polymorphism with an additional typing rule for applications:

$$\frac{\text{[T-APP-POLY]} \quad \Gamma_1 \vdash e_1 : t \rightarrow^{\top, \sigma} s \ \& \ \tau_1 \quad \Gamma_2 \vdash e_2 : \uparrow^n t \ \& \ \tau_2 \quad \tau_1 < |\Gamma_2|}{\Gamma_1 + \Gamma_2 \vdash e_1 e_2 : \uparrow^n s \ \& \ (n + \sigma) \sqcup \tau_1 \sqcup \tau_2 \quad \tau_2 < \top}$$

This rule admits an arbitrary mismatch n between the level the argument expected by the function and that of the argument supplied to the function. The type of the application and the latent effect are consequently shifted by the same amount n .

Soundness of [T-APP-POLY] can be intuitively explained as follows: a function with level $\top$ has no channels in its closure. Therefore, the only channels possibly manipulated by the function are those contained in the argument to which the function is

applied or channels created within the function itself. Then, the fact that the argument has level $n + k$ rather than level k is completely irrelevant. Conversely, if the function has channels in its closure, then the absolute level of the argument might have to satisfy specific ordering constraints with respect to these channels (recall the two side conditions in $[T\text{-APP}]$). Since level polymorphism is a key distinguishing feature of our type system, and one that accounts for much of its expressiveness, we elaborate more on this intuition using an example. Consider the term

$$\text{fwd} \stackrel{\text{def}}{=} \lambda x. \lambda y. \text{send } y \text{ (recv } x)$$

which forwards on y the message received from x . The derivation

$$\frac{\frac{\vdots}{y : ![int]^1 \vdash \text{send } y : int \rightarrow^{1,1} \text{unit} \& \perp} [T\text{-APP}]}{\frac{x : ?[int]^0, y : ![int]^1 \vdash \text{send } y \text{ (recv } x) : \text{unit} \& 1} [T\text{-APP}]} \quad \frac{\vdots}{x : ?[int]^0 \vdash \text{recv } x : int \& 0} [T\text{-APP}]}{\frac{x : ?[int]^0, y : ![int]^1 \vdash \text{send } y \text{ (recv } x) : \text{unit} \& 1} [T\text{-FUN}]} \quad \frac{x : ?[int]^0 \vdash \lambda y. \text{send } y \text{ (recv } x) : ![int]^1 \rightarrow^{0,1} \text{unit} \& \perp} [T\text{-FUN}]}{\vdash \text{fwd} : ?[int]^0 \rightarrow ![int]^1 \rightarrow^{0,1} \text{unit} \& \perp} [T\text{-APP}]$$

does *not* depend on the absolute values 0 and 1, but only on the level of x being smaller than that of y , as required by the fact that the **send** operation on y is blocked by the **recv** operation on x . Now, consider an application $\text{fwd } a$, where a has type $?[int]^2$. The mismatch between the level of x (0) and that of a (2) is not critical, because all the levels in the derivation above can be *uniformly shifted up* by 2, yielding a derivation for

$$\vdash \text{fwd} : ?[int]^2 \rightarrow ![int]^3 \rightarrow^{2,3} \text{unit} \& \perp$$

This shifting is possible because fwd has no free channels in its body (indeed, it is typed in the empty environment). Therefore, using $[T\text{-APP-POLY}]$, we can derive

$$a : ?[int]^2 \vdash \text{fwd } a : ![int]^3 \rightarrow^{2,3} \text{unit} \& \perp$$

Note that $(\text{fwd } a)$ is a function having level 2. This means that $(\text{fwd } a)$ is *not* level polymorphic and can only be applied, through $[T\text{-APP}]$, to channels with level 3. If we allowed $(\text{fwd } a)$ to be applied to a channel with level 2 using $[T\text{-APP-POLY}]$ we could derive

$$a : \#[int]^2 \vdash \text{fwd } a \ a : \text{unit} \& 2$$

which reduces to a deadlock.

Example 3. To show that the term in Example 1 is well typed, consider the environment

$$\Gamma \stackrel{\text{def}}{=} \text{fibo} : int \rightarrow ![int]^0 \rightarrow^{\top,0} \text{unit}, n : int, c : ![int]^0$$

In the proof derivation for the body of **fibo**, this environment is eventually enriched with the assignments $a : \#[int]^{-2}$ and $b : \#[int]^{-1}$. Now we can derive

$$\frac{\frac{\vdots}{\Gamma \vdash \text{fibo } (n - 2) : ![int]^0 \rightarrow^{\top,0} \text{unit} \& \perp} [T\text{-APP}]}{\Gamma, a : \#[int]^{-2} \vdash \text{fibo } (n - 2) \ a : \text{unit} \& -2} [T\text{-APP-POLY}] \quad \frac{\vdots}{a : \#[int]^{-2} \vdash a : \#[int]^{-2} \& \perp} [T\text{-NAME}]$$

where the application `fibonacci (n - 2) a` is well typed despite the fact that `fibonacci (n - 2)` expects an argument of type $!\text{[int]}^0$, while `a` has type $!\text{[int]}^{-2}$. A similar derivation can be obtained for `fibonacci (n - 1) b`, and the proof derivation can now be completed. ■

3.3 Recursive types

Looking back at Example 2, we see that in a call `pipe x y` the channel `recv x` is used in the same position as `x`. Therefore, according to $[\text{T-APP-POLY}]$, `recv x` must have the same type as `x`, up to some shifting of its level. Similarly, channel `c` is both sent on `y` and then used in the same position as `y`, suggesting that `c` must have the same type as `y`, again up to some shifting of its level. This means that we need recursive types in order to properly describe `x` and `y`.

Instead of adding explicit syntax for recursive types, we just consider the possibly infinite trees generated by the productions for t shown earlier. In light of this broader notion of types, the inductive definition of type level (3.1) is still well founded, but type shift (3.4) must be reinterpreted coinductively, because it has to operate on possibly infinite trees. The formalities, nonetheless, are well understood.

It is folklore that, whenever infinite types are *regular* (that is, when they are made of finitely many distinct subtrees), they admit finite representations either using type variables and the familiar μ notation, or using systems of type equations [4]. Unfortunately, a careful analysis of Example 2 suggests that – at least in principle – we also need *non-regular* types. To see why, let `a` and `c` be the channels to which `(recv x)` and `(cont y)` respectively evaluate on line 2 of the example. Now:

- `x` must have smaller level than `a` since `a` is received from `x` (*cf.* the types of `recv`).
- `y` must have smaller level than `c` since `c` is sent on `y` (*cf.* the types of `send`).
- `x` must have smaller level than `y` since `x` is used in the functional part of an application in which `y` occurs in the argument (*cf.* line 2 and $[\text{T-APP-POLY}]$).

Overall, in order to type `pipe` in Example 2 we should assign `x` and `y` the types t^n and s^n that respectively satisfy the equations

$$t^n = ?[t^{n+2}]^n \quad s^n = ![t^{n+3}]^{n+1} \quad (3.5)$$

Unfortunately, these equations do not admit regular types as solutions. We recover typeability of `pipe` with regular types by introducing a new type constructor

$$t ::= \dots \mid [t]^n$$

that wraps types with a pending shift: intuitively $[t]^n$ and $\uparrow^n t$ denote the same type, except that in $[t]^n$ the shift $\uparrow^n$ on t is pending. For example, $[?\text{[int]}^0]^1$ and $[?\text{[int]}^2]^{-1}$ are both possible wrappings of $?\text{[int]}^1$, while $\text{int} \rightarrow^{0,\perp} !\text{[int]}^0$ is the unwrapping of $[\text{int} \rightarrow^{1,\perp} !\text{[int]}^1]^{-1}$. To exclude meaningless infinite types such as $[[[\dots]^n]^n]^n$ we impose a *contractiveness condition* requiring every infinite branch of a type to contain infinite occurrences of channel or arrow constructors. To see why wraps help finding regular representations for otherwise non-regular types, observe that the equations

$$t^n = ?[t^{n+2}]^n \quad s^n = ![[t^{n+1}]^2]^{n+1} \quad (3.6)$$

denote – up to pending shifts – the same types as the ones in (3.5), with the key difference that (3.6) admit regular solutions and therefore finite representations. For example, t^n could be finitely represented as a familiar-looking $\mu\alpha.?\llbracket\alpha\rrbracket^2\rrbracket^n$ term.

We should remark that $\llbracket t \rrbracket^n$ and $\uparrow^n t$ are *different* types, even though the former is morally equivalent to the latter: wrapping is a type *constructor*, whereas shift is a type *operator*. Having introduced a new constructor, we must suitably extend the notions of type level (3.1) and type shift (3.4) we have defined earlier. We postulate

$$\llbracket t \rrbracket^n \stackrel{\text{def}}{=} n + |t| \qquad \uparrow^n \llbracket t \rrbracket^m \stackrel{\text{def}}{=} \llbracket \uparrow^n t \rrbracket^m$$

in accordance with the fact that $\llbracket \cdot \rrbracket^n$ denotes a pending shift by n (note that $|\cdot|$ extended to wrappings is well defined thanks to the contractiveness condition).

We also have to define introduction and elimination rules for wrappings. To this aim, we conceive two constants, **wrap** and **unwrap**, having the following type schemes:

$$\text{wrap} : \uparrow^n t \rightarrow \llbracket t \rrbracket^n \qquad \text{unwrap} : \llbracket t \rrbracket^n \rightarrow \uparrow^n t$$

We add **wrap** v to the value forms. Operationally, we want **wrap** and **unwrap** to annihilate each other. This is done by enriching reduction for expressions with the axiom

$$\text{unwrap} (\text{wrap } v) \longrightarrow v$$

Example 4. We suitably dress the code in Example 2 using **wrap** and **unwrap**:

```
1  let cont = λx.let c = new() in (fork λ_.send x (wrap c)); c in
2  let pipe = fix λpipe.λx.λy.pipe (unwrap (recv x)) (cont y)
```

and we are now able to find a typing derivation for it that uses regular types. In particular, we assign `cont` the type $s^n \rightarrow s^{n+2}$ and `pipe` the type $t^n \rightarrow s^n \rightarrow^{n, \top} \text{unit}$ where t^n and s^n are the types defined in (3.6). Note that `cont` is a pure function because its effects are masked by **fork** and that `pipe` has latent effect $\top$ since it loops performing **recv** operations on channels with increasing level. Because of the side conditions in $\llbracket T\text{-APP} \rrbracket$ and $\llbracket T\text{-APP-POLY} \rrbracket$, this means that `pipe` can only be used in tail position, which is precisely what happens above and in Example 2. ■

3.4 Properties

To formulate subject reduction, we must take into account that linear channels are *consumed* after communication (last but one reduction in Table 1). This means that when a process P communicates on some channel a , a must be removed from the type environment used for typing the residual of P . To this aim, we define a partial operation $\Gamma - \ell$ that removes ℓ from Γ , when ℓ is a channel. Formally:

Theorem 1 (subject reduction). *If $\Gamma \vdash P$ and $P \xrightarrow{\ell} Q$, then $\Gamma - \ell \vdash Q$ where $\Gamma - \tau \stackrel{\text{def}}{=} \Gamma$ and $(\Gamma, a : \# \llbracket t \rrbracket^n) - a \stackrel{\text{def}}{=} \Gamma$.*

Note that $\Gamma - a$ is undefined if $a \notin \text{dom}(\Gamma)$. This means that well-typed programs never attempt at using the same channel twice, namely that channels in well-typed programs are indeed *linear channels*. This property has important practical consequences, since it allows the efficient implementation (and deallocation) of channels [9].

Deadlock freedom means that *if* the program halts, then there must be no pending I/O operations. In our language, the only halted program without pending operations is (structurally equivalent to) $\langle () \rangle$. We can therefore define deadlock freedom thus:

Definition 1. We say that P is deadlock free if $P \xrightarrow{\tau}^* Q \dashv\dashv$ implies $Q \equiv \langle () \rangle$.

As usual, $\xrightarrow{\tau}^*$ is the reflexive, transitive closure of $\xrightarrow{\tau}$ and $Q \dashv\dashv$ means that Q is unable to reduce further. Now, every well-typed, closed process is free from deadlocks:

Theorem 2 (soundness). If $\emptyset \vdash P$, then P is deadlock free.

Theorem 2 may look weaker than desirable, considering that every process P (even an ill-typed one) can be “fixed” and become part of a deadlock-free system if composed in parallel with the diverging thread $\langle \text{fix } \lambda x.x \rangle$. It is not easy to state an interesting property of well-typed *partial programs* – programs that are well typed in un-even environments – or of *partial computations* – computations that have not reached a stable (*i.e.*, irreducible) state. One might think that well-typed programs eventually use all of their channels. This property is false in general, for two reasons. First, our type system does not ensure termination of well-typed expressions, so a thread like $\langle \text{send } a \ (\text{fix } \lambda x.x) \rangle$ never uses channel a , because the evaluation of the message diverges. Second, there are threads that continuously generate (or receive) new channels, so that the set of channels they own is never empty; this happens in Example 2. What we can prove is that, *assuming* that a well-typed program does not internally diverge, then *each* channel it owns is eventually used for a communication or is sent to the environment in a message. To formalize this property, we need a labeled transition system describing the interaction of programs with their environment. *Labels* $\pi, \dots$ of transitions are defined by

$$\pi ::= \ell \mid a?e \mid a!v$$

and the transition relation $\xrightarrow{\pi}$ extends reduction with the rules

$$\frac{a \notin \text{bn}(\mathcal{C})}{\mathcal{C}[\text{send } a \ v] \xrightarrow{a!v} \mathcal{C}[\langle () \rangle]} \quad \frac{a \notin \text{bn}(\mathcal{C}) \quad \text{fn}(e) \cap \text{bn}(\mathcal{C}) = \emptyset}{\mathcal{C}[\text{recv } a] \xrightarrow{a?e} \mathcal{C}[e]}$$

where $\mathcal{C}$ ranges over *process contexts* $\mathcal{C} ::= \langle \mathcal{C} \rangle \mid (\mathcal{C} \mid P) \mid (P \mid \mathcal{C}) \mid (\nu a)\mathcal{C}$. Messages of input transitions have the form $a?e$ where e is an arbitrary expression instead of a value. This is just to allow a technically convenient formulation of Definition 2 below. We formalize the assumption concerning the absence of internal divergences as a property that we call *interactivity*. Interactivity is a property of *typed processes*, which we write as pairs $\Gamma \S P$, since the messages exchanged between a process and the environment in which it executes are not arbitrary in general.

Definition 2 (interactivity). Interactivity is the largest predicate on well-typed processes such that $\Gamma \S P$ interactive implies $\Gamma \vdash P$ and:

1. P has no infinite reduction $P \xrightarrow{\ell_1} P_1 \xrightarrow{\ell_2} P_2 \xrightarrow{\ell_3} \dots$, and
2. if $P \xrightarrow{\ell} Q$, then $\Gamma - \ell \circ Q$ is interactive, and
3. if $P \xrightarrow{a!v} Q$ and $\Gamma = \Gamma', a : ![t]^n$, then $\Gamma'' \circ Q$ is interactive for some $\Gamma'' \subseteq \Gamma'$, and
4. if $P \xrightarrow{a?x} Q$ and $\Gamma = \Gamma', a : ?[t]^n$, then $\Gamma'' \circ Q\{v/x\}$ is interactive for some v and $\Gamma'' \supseteq \Gamma'$ such that $n < |\Gamma'' \setminus \Gamma'|$.

Clause (1) says that an interactive process does not internally diverge: it will eventually halt either because it terminates or because it needs interaction with the environment in which it executes. Clause (2) states that internal reductions preserve interactivity. Clause (3) states that a process with a pending output on a channel a *must* reduce to an interactive process after the output is performed. Finally, clause (4) states that a process with a pending input on a channel a *may* reduce to an interactive process after the input of a particular message v is performed. The definition looks demanding, but many conditions are direct consequences of Theorem 1. The really new requirements besides well typedness are *convergence* of P (1) and the *existence* of v (4). It is now possible to prove that well-typed, interactive processes eventually use their channels.

Theorem 3 (interactivity). *Let $\Gamma \circ P$ be an interactive process such that $a \in \text{fn}(P)$. Then $P \xrightarrow{\pi_1} P_1 \xrightarrow{\pi_2} \dots \xrightarrow{\pi_n} P_n$ for some $\pi_1, \dots, \pi_n$ such that $a \notin \text{fn}(P_n)$.*

4 Concluding remarks

We have demonstrated the portability of a type system for deadlock freedom of π -calculus processes [10] to a higher-order language using an *effect system* [1]. We have shown that *effect masking* and *polymorphic recursion* are key ingredients of the type system (Examples 1 and 2), and also that latent effects must be paired with one more annotation – the function level. The approach may seem to hinder program modularity, since it requires storing levels in types and levels have global scope. In this respect, level polymorphism (Section 3.2) alleviates this shortcoming of levels by granting them a relative – rather than absolute – meaning at least for non-linear functions.

Other type systems for higher-order languages with session-based communication primitives have been recently investigated [6,14,2]. In addition to safety, types are used for estimating bounds in the size of message queues [6] and for detecting memory leaks [2]. Since binary sessions can be encoded using linear channels [5], our type system can address the same family of programs considered in these works with the advantage that, in our case, well-typed programs are guaranteed to be deadlock free also in presence of session interleaving. For instance, the pipe function in Example 2 interleaves communications on two different channels. The type system described by Wadler [14] is interesting because it guarantees deadlock freedom without resorting to any type annotation dedicated to this purpose. In his case the syntax of (well-typed) programs prevents the modeling of cyclic network topologies, which is a necessary condition for deadlocks. However, this also means that some useful program patterns cannot be modeled. For instance, the program in Example 2 is ill typed in [14].

The type system discussed in this paper lacks compelling features. *Structured data types* (records, sums) have been omitted for lack of space; an extended technical report [13] and previous works [11,10] show that they can be added without issues. The

same goes for *non-linear channels* [10], possibly with the help of dedicated `accept` and `request` primitives as in [6]. *True polymorphism* (with level and type variables) has also been studied in the technical report [13]. Its impact on the overall type system is significant, especially because level and type constraints (those appearing as side conditions in the type schemes of constants, Section 3.1) must be promoted from the metatheory to the type system. The realization of level polymorphism as type shifting that we have adopted in this paper is an interesting compromise between impact and flexibility. Our type system can also be relaxed with *subtyping*: arrow types are contravariant in the level and covariant in the latent effect, whereas channel types are invariant in the level. Invariance of channel levels can be relaxed refining levels to *pairs* of numbers as done in [7,8]. This can also improve the accuracy of the type system in some cases, as discussed in [10] and [3]. It would be interesting to investigate which of these features are actually necessary for typing concrete functional programs using threads and communication/synchronization primitives.

Type reconstruction algorithms for similar type systems have been defined [11,12]. We are confident to say that they scale to type systems with arrow types and effects.

Acknowledgments. The authors are grateful to the reviewers for their detailed comments and useful suggestions. The first author has been supported by Ateneo/CSP project SALT, ICT COST Action IC1201 BETTY, and MIUR project CINA.

References

1. T. Amtoft, F. Nielson, and H. Nielson. *Type and Effect Systems: Behaviours for Concurrency*. Imperial College Press, 1999.
2. V. Bono, L. Padovani, and A. Tosatto. Polymorphic Types for Leak Detection in a Session-Oriented Functional Language. In *FORTE'13*, LNCS 7892, pages 83–98. Springer, 2013.
3. M. Carbone, O. Dardha, and F. Montesi. Progress as compositional lock-freedom. In *COORDINATION'14*, LNCS 8459, pages 49–64. Springer, 2014.
4. B. Courcelle. Fundamental properties of infinite trees. *Theor. Comp. Sci.*, 25:95–169, 1983.
5. O. Dardha, E. Giachino, and D. Sangiorgi. Session types revisited. In *PPDP'12*, pages 139–150. ACM, 2012.
6. S. J. Gay and V. T. Vasconcelos. Linear type theory for asynchronous session types. *J. Funct. Program.*, 20(1):19–50, 2010.
7. N. Kobayashi. A type system for lock-free processes. *Inf. and Comp.*, 177(2):122–159, 2002.
8. N. Kobayashi. A new type system for deadlock-free processes. In *CONCUR'06*, LNCS 4137, pages 233–247. Springer, 2006.
9. N. Kobayashi, B. C. Pierce, and D. N. Turner. Linearity and the pi-calculus. *ACM Trans. Program. Lang. Syst.*, 21(5):914–947, 1999.
10. L. Padovani. Deadlock and Lock Freedom in the Linear π -Calculus. In *CSL-LICS'14*, pages 72:1–72:10. ACM, 2014. <http://hal.archives-ouvertes.fr/hal-00932356v2/>.
11. L. Padovani. Type Reconstruction for the Linear π -Calculus with Composite and Equi-Recursive Types. In *FoSSaCS'14*, LNCS 8412, pages 88–102. Springer, 2014.
12. L. Padovani, T.-C. Chen, and A. Tosatto. Type Reconstruction Algorithms for Deadlock-Free and Lock-Free Linear π -Calculi. In *COORDINATION'15*, LNCS 9037. Springer, 2015. to appear.
13. L. Padovani and L. Novara. Types for Deadlock-Free Higher-Order Concurrent Programs. Technical report, Università di Torino, 2014. <http://hal.inria.fr/hal-00954364>.
14. P. Wadler. Propositions as sessions. In *ICFP'12*, pages 273–286. ACM, 2012.