A higher-order specification of the π-calculus

Joëlle Despeyroux

INRIA*

Abstract. We present a formalization of a typed π-calculus in the Calculus of Inductive Constructions. We give the rules for type-checking and for evaluation and formalize a proof of type preservation in the Coq system. The encoding of the π-calculus in Coq uses Coq fonctions to represent bindings of variables. This kind of encoding is called a higher-order specification. It provides a concise description of the calculus, leading to simple proofs. The specification we propose for the pi-calculus formalizes communication by means of function application.

1 Introduction

The π-calculus [MPW92,Mil91] is a model of concurrent computation based upon the notion of naming. Processes receive and emit names, which denote channels.

We propose here new formalizations of both evaluation and typing rules for a typed π-calculus, with the aim to enable machine-checked proofs of various properties of languages based on this calculus. As a simple, typical illustration of such a proof, we give a proof of preservation of types, also called the subject-reduction property, for the calculus.

For our experiment, we chose the Coq system [BBC+97,HKPM97], based on the Calculus of Inductive Constructions [CH88,PM92]. In this system, proofs are performed by applying tactics, in a goal-directed manner. We like this system because of its rich meta-language (a higher-order functional language, allowing full dependant inductive types), and because it is based on Type Theory.

Thanks to the use of the *higher-order abstract syntax* technique, giving rise to so-called *higher-order specifications*, our formalization of the π-calculus does not need definitions for free or bound variables in a term. Nor does it need definitions of notions of substitutions, which are implemented using the meta-level application, i.e. application available in the Logical Framework used to implement our calculus (which in our case is the Calculus of Inductive Constructions).

Robin Milner [Mil91] introduced notions of *abstractions* and *concretions* leading to a presentation of the calculus in which input processes evaluate to functions and output processes evaluate to concretions, i.e. (intuitively) pairs of a value and a process. Our formalization starts from this idea, replacing concretions by higher-order functions, thus leading to a nicer, more uniform presentation, in which communication is formalized by function application.

Related works include formalizations of various π-calculi in the same system. These works were conducted with different goals in mind. Some of them had the same goals as ours. The others aimed at the study of bisimilation techniques. We shall discuss both kinds of related works in detail at the end of the paper.

* INRIA, B.P. 93, F-06902 Sophia-Antipolis Cedex, France. Joelle.Despeyroux@sophia.inria.fr

Note on the Coq syntax. Terms in Coq are terms of a typed λ-calculus, where the abstraction $\lambda x : t.a$ is written [x:t]a, the dependant product type $\forall x : t.a$ is written (x:t)a, while the non-dependant product type $A \to B$ is simply written $A \to B$. The application of a function f to x is written $(f\ x)$. The expression $\exists x : t.a$ is written (Ex [x:t]a). The type of sets is denotated as Set while the type of propositions in denotated as Prop. As usual in Type Theory, the type $A_1 \to \cdots \to A_n \to B$ where B is a proposition represents (through the Curry-Howard isomorphism) the inference rule with premisses $A_1, \cdots, A_n$ and conclusion B. Inductive types are defined by a list of constructors together with their types, as for example for the naturals:

```
Inductive nat : Set := 0 : nat | S : nat -> nat.
```

An inductive definition automatically generates both an inductive principle for the type being defined and elimination principles which enable the definition of recursive functions on this type. In the following, Coq text will be given as above in a verbatim style.

2 Syntax

The terms of the π-calculus are processes, which receive and emit names denoting channels. For the sake of simplicity, we don't consider here the operators of choice and matching. It will be straightforward to add them. Also we have an operator for replication instead of recursion; both operators give the same expressive power to the calculus.

2.1 The monadic π-calculus

Processes are defined by the following rule:

$$\text{Processes}: \quad P, Q ::= \bar{x}y.P \mid x(y).P \mid \nu x : t.P \mid (P \mid Q) \mid !\ P \mid 0$$

in which x and y belong to an enumerable set of names and t denotes the type of a name as described in section 5.

The expression $\bar{x}y.P$ denotes a process which sends a name y along the channel x (*output* operation), then behaves as P. The term $x(y).P$ denotes a process which receives a name z along the channel x (*input* operation), then behaving like $P[z/y]$ (P in which z replaces y). The expression $(P \mid Q)$ denotes the *parallel* composition of P and Q where P and Q are concurrently active. 0 denotes the *empty* process. $!\ P$ denotes the infinite process $P \mid P \mid \cdots$; the ! operator (pronounce "bang") is called the *replication* operator. Finally, $\nu x : t.P$ declares x as a private channel with type t in P. The ν operator is called a *restriction* as it restricts the use of the name x to P.

In the term $\nu x : t.P$, the name x is bound in the term P. Similarly, in the term $x(y).P$, the name y is bound in the term P, waiting to be substituted by another name before P can be executed.

Processes evaluate to other processes while performing *actions*. In the standard approach, there are four kinds of actions: input actions xy (input a value for variable y on a channel x), output actions $\bar{x}y$ (output value y on a channel x), bounded output action $\bar{x}(y : t)$, and silent actions τ. Communications are always silent.

$$\text{Actions}: \quad A ::= xy \mid \bar{x}y \mid \bar{x}(y : t) \mid \tau$$

An original, and central notion of the π-calculus is the notion of *scope extrusion*. A name is said to be extruded when it is a private name sent to a process. In such a case, the scope of the name is extended to the new process which receives it. This phenomena, which may make the notion of binding in the π-calculus unclear at first glance, will be illustrated and discussed later on.

The informal account of terms needs the definition of the set of variables in a term, that we denote as $Var(t)$. It also needs the definition of free variables ($F(t)$) and bound variables ($B(t)$) in a term. In the informal presentation of the calculus, we will need a notion of substitution of a variable for another variable in a term, written as usual as $P\{y/x\}$.

2.2 Formalization of the syntax

We encode the terms of the calculus in Coq as follows. First we introduce a new parameter *name*, and suppose two usual properties on the set of names, that we state as axioms:

$$\forall x, y : name.\ x = y \vee x \neq y; \quad \forall x : name. \exists y : name.\ x \neq y$$

```
Parameter name: Set.
Axiom name_decidable : (x,y:name) {x=y} + {~x=y}.
Axiom name_different : (x:name)(Ex [y:name](~x=y)).
```

Then we introduce the set of types as a parameter too, stating some properties on this set later on.

```
Parameter typ : Set.
```

The processes are then defined as follows:

```
Inductive proc : Set :=
    Nil  : proc
  | In   : name -> (name -> proc) -> proc
  | Out  : name -> name -> proc -> proc
  | Par  : proc -> proc -> proc
  | Bang : proc -> proc
  | Res  : typ  -> (name -> proc) -> proc.
```

The representation of processes requires some explanations. There are two binding operators in the π-calculus: the input and the restriction operator. We use here Coq functions to represent bindings of variables; hence the type ($name \rightarrow proc$) in the types of input processes and restrictions. This kind of encoding is called a *higher-order specification*. This method of formalization, called the higher-order abstract syntax technique, provides an elegant way to formalize the binding nature of the constructors. In particular, it provides for free a definition of terms up to *α-equivalence*. This method also provides an elegant way to formalize *substitutions*, as we shall see later on, when formalizing the rules for evaluation.

Note that the type name is declared as a parameter, not as a variable. This prevents name to be instanciated by an inductive set.

There is no application related to restriction. Nevertheless, using functions (i.e. binding operations) to represent the restriction yields a simple and elegant formalization of the notion of private names.

In our description of the π-calculus, the only visible argument of an action will be the name of a channel. In other words, actions only need to carry a name. As a consequence, there is no need for bound output actions in our development:

```
Inductive action : Set :=
    InA : name -> action | OutA : name -> action | Tau : action.
```

3 Simple examples of evaluation

We give here some simple examples of communication. The simplest case of communication is the following one, where a name y is sent to Q, along a channel x:

$$\bar{x}y.P \mid x(z).Q \xrightarrow{\tau} P \mid Q\{y/z\}$$

Then let us illustrate the phenomena of scope extrusion. This occurs in a communication where a process sends a private name to an external process, as in:

$$\nu y : t.(\bar{x}y.P) \mid x(z).Q \quad \text{which reduces to} \quad \nu y : t.(P \mid Q[y/z]).$$

The private name y has been passed to the external process Q. We says that P has extruded the scope of the private channel y. We have a similar phenomena in imperative languages when a local name y is sent to a procedure Q taking his parameters by reference.

4 Evaluation rules

The evaluation rules define a judgement $P \xrightarrow{a} Q$ meaning that P evaluates to Q producing an action a. The chosen semantics is the early transition semantics.

4.1 Informal evaluation rules

We need here the definition of all variables in a term ($Var(t)$), together with the definitions of fresh and bound variables ($F(t)$ and $B(t)$). Moreover, usual informal presentations of the calculus use the open and close rules, which *open and close the scope of a name*, in order to describe the phenomena of scope extrusion that we illustrated above (section 3). The informal description of the evaluation rules is the following one.

$input$:
$$x(y).P \stackrel{xz}{\rightarrow} P\{z/y\}$$

$output$:
$$\bar{x}y.P \stackrel{\bar{x}y}{\rightarrow} P$$

com :
$$\frac{P \stackrel{xy}{\rightarrow} P' \qquad Q \stackrel{\bar{x}y}{\rightarrow} Q'}{P \mid Q \stackrel{\tau}{\rightarrow} P' \mid Q'} \qquad \frac{P \stackrel{xy}{\rightarrow} P' \qquad Q \stackrel{\bar{x}y}{\rightarrow} Q'}{Q \mid P \stackrel{\tau}{\rightarrow} Q' \mid P'}$$

$open$:
$$\frac{P \stackrel{\bar{x}y}{\rightarrow} P'}{\nu y : t.P \stackrel{\bar{x}(y:t)}{\rightarrow} P'} \quad x \neq y$$

$close$:
$$\frac{P \stackrel{xy}{\rightarrow} P' \qquad Q \stackrel{\bar{x}(y:t)}{\rightarrow} Q'}{P \mid Q \stackrel{\tau}{\rightarrow} \nu y : t.(P' \mid Q')} \quad y \notin F(P) \qquad \frac{P \stackrel{xy}{\rightarrow} P' \qquad Q \stackrel{\bar{x}(y:t)}{\rightarrow} Q'}{Q \mid P \stackrel{\tau}{\rightarrow} \nu y : t.(Q' \mid P')} \quad y \notin F(P)$$

res :
$$\frac{P \stackrel{a}{\rightarrow} P'}{\nu x : t.P \stackrel{a}{\rightarrow} \nu x : t.P'} \quad x \notin Var(a)$$

par :
$$\frac{P \stackrel{a}{\rightarrow} P'}{P \mid Q \stackrel{a}{\rightarrow} P' \mid Q} \quad B(a) \cap F(Q) = \emptyset \qquad \frac{P \stackrel{a}{\rightarrow} P'}{Q \mid P \stackrel{a}{\rightarrow} Q \mid P'} \quad B(a) \cap F(Q) = \emptyset$$

$bang$:
$$\frac{!P \mid P \stackrel{a}{\rightarrow} P'}{!\, P \stackrel{a}{\rightarrow} P'}$$

In the above set of rules, rules for *input* and *output* describe the basic cases of input and output processes. The *com* and *close* rules describe communication, while the rest of the rules describe the propagation of evaluation through a restriction, a parallel composition, or a replication.

The replication rule could be simpler. The advantage of the chosen rule is that it enables the extension of the rules with a choice operator.

4.2 Formalization of the evaluation rules

We shall take full advantage here from our choice to use functions in the formalization of the syntax.

First, as input processes are described as functions, we just say that input processes evaluate to themselves. This follows Robin Milner approach [Mil91] where he introduced a notion of *abstraction* to denote the result of the evaluation of input processes. We will define an inductive predicate *evali* for processes performing an input action:

```
Inductive evali : proc -> action -> (name -> proc) -> Prop
```

with as first and simplest case the case for input processes:

```
evali_in : (x:name)(p:name->proc) (evali (In x p) (InA x) p)
```

Robin Milner also introduced a notion of *concretion* [Mil91], which (intuitively) was a pair of a value and a processus, to denote the result of the evaluation of output processes. Then he (informaly) described the communication as the "application" of an abstraction to a concretion. We thought it would be nice, in order to formalize the communication by a real application, to describe the evaluation of an output process as yielding a *higher-order function* which will then be applied to the result of the evaluation of an input function in order to describe a communication, as we shall see later on.

Let us say that the evaluation of output processes yield higher-order functions taking functions of type $(name \rightarrow proc)$ as arguments. We define an inductive predicate *evalo* for processes performing an output action:

```
Inductive evalo : proc -> action -> ((name->proc) -> proc) -> Prop
```

with as first and simplest case the case for output processes:

```
evalo_out : (x,y:name)(p:proc)
   (evalo (Out x y p) (OutA x) [f:name->proc](Par (f y) p))
```

Example We are now in a position to see how we can describe communication on a basic example:

$$\frac{P \stackrel{\bar{x}y}{\rightarrow} P' \qquad Q \stackrel{xy}{\rightarrow} Q'}{\bar{a}b.P \mid a(y).Q \stackrel{\tau}{\rightarrow} P' \mid Q'\{b/y\}}$$

The evaluation of the input process $a(y).Q$ gives the function Q itself. Then the result of the evaluation of the output $\bar{a}b.P$ is a function $\lambda f : name \rightarrow proc.((f\ b) \mid P)$. Applying this function to Q yields the expected result: $(Q\ b) \mid P$.

```
(Out a b P)  --(OutA a)--> [f:name->proc]((f b) | P) = P'
   (In a Q)   --(In a)-->  Q
-------------------------------------------------------
   (Out a b P) | (In a Q) --Tau--> (P' Q) = (Q b) | P
```

Thus the evaluation of both input and output processes yields functions. Communication is formalized by applying the result of the evaluation of the output process to the result of the evaluation of the input process. Our formalization has replaced concretions by higher-order functions, thus leading to a nicer, more uniform presentation, in which communication is formalized by a real application of a function. On a technical level, the higher-order abstract syntax method enables us to describe *substitution* in evaluation rules by simply using the *application* of a Coq function to terms -names in this case.

It is staightforward to prove the consistency of this presentation of the semantics of the π-calculus with functions with respect to the standard approach.

The rule for example for left communication will be formalized as follows, in a third set of inductive rules describing the evaluation of processes that perform silent actions:

```
Inductive eval :  proc -> action -> proc -> Prop :=
  | eval_coml:
      (p:proc)(x:name)(p':(name->proc)->proc)
      (evalo p (OutA x) p') ->
      (q:proc)(q':name->proc) (evali q (InA x) q') ->
      (eval (Par p q) Tau (p' q'))
  [...].
```

Note that communication is always simply formalized by an application, no matter whether names are extruded or not during the communication. This is the reason why we do not need rules for opening and closing the scope of names in our formalization.

The other rules -for the parallel, replication and restriction operators- need to be duplicated, one for each type of action involved (input, output or silent). Note however that all proofs, either formal or informal, proceed by a case analysis on the type of the action involved anyway. Thus this apparent defect should not be considered as a drawback of our formalization.

Thanks to the use of the higher-order abstract syntax technique, when evaluating a process which declares a private name, we do not need to check that this private name is a fresh name. The freshness of the name naturally comes from the use of a universal quantification on names in the premiss of the rule.

```
eval_res :
      (u:typ)(p:name->proc)(a:action)(q:name->proc)
      ((x:name) (eval (p x) a (q x))) ->
      (eval (Res u p) a (Res u q)).
```

The rules for parallel composition do not need the definition for free or bound variables either. This is because the actions do not carry values here, but only the name of a channel.

```
eval_parl : (p:proc)(a:action)(p':proc)
      (eval p a p') -> (q:proc) (eval (Par p q) a (Par p' q))
```

The rule for replication (omitted here) is straightforward.

5 Type checking rules

The type system that we chosed use a notion of *directionality* of names [PS95], which is their ability to be used as emitters, receivers or both (or none). This name usage discipline is implemented by type-checking rules. This type system, while being simple, is at the same time widly used as witnessed by the experiment of Pict [PT97] and non-trivial because of the sub-typing that it captures. For our purpose, it is more interesting than the type system for the polyadic π-calculus [Mil91], whose typing rules are totally straightforward.

5.1 Types

A type t of a name of a channel is the union of its *capacity*, which represents how the channel can be used (read, write, both or none) and the *types* of the names that can transit in it:

$$\begin{array}{ll} \text{Capacities}: & c ::= \mathit{read} \mid \mathit{write} \mid \mathit{both} \mid \mathit{none} \\ \text{Types}: & t ::= (c, t) \end{array}$$

We formalize the capacities of names as follows, where none is called top (for top of the mini-lattice):

```
Inductive cap : Set := read : cap | write : cap | both : cap | top : cap.
```

Remember we only defined types as:

```
Parameter typ : Set.
```

This means that we do not impose a particular implementation of types. Instead we state some properties this implementation must enjoy. First we must be able to extract both the capacity $\langle t \rangle$ and the type $[t]$ parts of a type t:

```
Parameter get_cap : typ -> cap.
Parameter get_typ : typ -> typ.
```

Then we have a subtyping relation on capacities, denoted as $\leq$ and defined as follows:

$$c \leq c \mid both \leq c \mid c \leq none$$

The subtyping relation on capacities induces a subtyping relation on types, that we still denote as $\leq$ and define as follows:

$$\begin{gathered} t \leq t \\ \langle t' \rangle = none \Rightarrow t \leq t' \\ \langle t \rangle \leq read \ \wedge \langle t' \rangle = read \ \wedge [t] \leq [t'] \Rightarrow t \leq t' \\ \langle t \rangle \leq write \wedge \langle t' \rangle = write \wedge [t'] \leq [t] \Rightarrow t \leq t' \end{gathered}$$

The $\leq$ predicates are formalized by the following inductive types, the obvious rules of which we omit here:

```
Inductive stype_cap : cap -> cap -> Prop
Inductive stype : typ -> typ -> Prop
```

The covariance of the read capability and the contravariance of the write capability introduce a subtyping in depth, as the typing rules will show.

A process is said to be well *typed in a particular environment* which gives types to each of its free name. In contrary to the usual case in typed programming languages, the type-checking rules do not give types to terms in the π-calculus.

5.2 Examples

The process $\bar{x}y \mid x(u).\bar{u}v \mid \bar{x}z$ is well typed in an environment where x has the capabilities *both* (both read and write). $\bar{x}y$ alone is well typed in an environment where x has a capabilities $c \leq write$, which means that c is either *both* or *write*.

The following example show how certain interferences between processes can be eliminated. Consider the process $P = \nu y : (both, t).(\bar{x}y.y(z))$, where x has type $(write, (write, t))$. The process P gives the writing capacity to the outside, while keeping the reading capacity for himself on the name y. After the transmission of y, a process sending a value to P on channel x is sure that this value cannot be captured by another process.

The previous example also illustrates the use of subtyping in depth.

5.3 Informal typing rules

The informal typing rules, defining a judgement $env \vdash proc$, are self-explanatory. (We omit the obvious rules for par and bang.) Note the need for the Var predicate in the input and res rules.

$$typ_nil: \quad \Gamma \vdash 0$$

$$typ_input: \quad \frac{\langle \Gamma(x) \rangle \leq read \qquad \forall z \notin Var(P)\ \Gamma.(z : [\Gamma(x)]) \vdash P\{z/y\}}{\Gamma \vdash x(y).P}$$

$$typ_output: \quad \frac{\langle \Gamma(x) \rangle \leq write \qquad \Gamma(y) \leq [\Gamma(x)] \qquad \Gamma \vdash P}{\Gamma \vdash \bar{x}y.P}$$

$$typ_res: \quad \frac{\forall y \notin Var(P)\ \Gamma.(y : t) \vdash P\{y/x\}}{\Gamma \vdash \nu x : t.P}$$

5.4 Formalization of the typing rules

An environment is encoded as a function from names to types:

```
Definition env : Type := name -> typ.
```

Then the type-checking rules are a straightforward translation from informal rules, where the only interesting point to note is that we do not need to look for fresh names in the rules for input and restriction. The freshness condition is implemented by an universal quantification on variables in the premisses involved.

The complete set of rules is as follows:

```
Inductive type : env -> proc -> Prop :=
    type_nil : (g:env) (type g Nil)
  | type_in  :
      (g:env)(x:name)(p:name->proc)
      (stype_cap (get_cap (g x)) read) ->
      ((y:name) ((g y)=(get_typ (g x))) -> (type g (p y))) ->
      (type g (In x p))
  | type_out :
      (g:env)(x,y:name)(p:proc)
      (stype_cap (get_cap (g x)) write) ->
      (stype (g y) (get_typ (g x))) ->
      (type g p) ->
      (type g (Out x y p))
  | type_res :
      (g:env)(t:typ)(p:name->proc)
      ((y:name) ((g y)=t) -> (type g (p y))) ->
      (type g (Res t p)).
```

6 Preservation of types

We have to prove three theorems, one for each kind of actions. The theorem for the silent actions simply states the following:

```
Theorem sr_tau:
  (p:proc)(q:proc) (eval p Tau q) -> (g:env) (type g p) -> (type g q).
```

The proof of this theorem uses the following natural property for communications. The property says that a communication between two processes resulting of the evaluation of well-typed processes is well-typed:

```
Theorem sr_com:
  (p:proc)(x:name)(p':(name->proc)->proc)
  (evalo p (OutA x) p') -> (g:env) (type g p) ->
  (q:proc)(q':name->proc) (evali q (InA x) q') -> (type g q) ->
  (type g (p' q')).
```

The proof of the above theorem needs the theorems stating the preservation of types properties for input and output actions:

```
Theorem sr_in:
  (p:proc)(a:action)(q:name->proc) (evali p a q) -> (x:name)(a=(InA x)) ->
  (g:env) (type g p) ->
  (y:name) (stype (g y) (get_typ (g x))) -> (type g (q y)).

Theorem sr_out:
  (p:proc)(a:action)(q:(name->proc)->proc) (evalo p a q) ->
  (x:name)(a=(OutA x)) -> (g:env) (type g p) ->
  (f:name->proc) ((y:name) (stype (g y) (get_typ (g x))) -> (type g (f y))) ->
  (type g (q f)).
```

Note in the above theorems how we deal with the typing of functions. We say that a function f from name to processes is well-typed if the process $(f\ y)$ is well-typed for every name y of the appropriate type. Similarly, a function F from (name to processes) to processes will be said to be well-typed if the process $(F\ f)$ is well-typed for every well-typed function f.

Except for some simple lemmas on the subtyping relation, no further lemmas are needed in the proof. In particular there is no need for tedious proofs of lemmas about renaming, freshness of variables and the like, contrary to almost all the other developments we have seen so far.

7 Related Work

Related works include several recent formalizations of various π-calculi, all done in the Coq system. As we mentioned in the introduction of the paper, some of these works had the same goals as ours, while others aimed at the study of bisimilation techniques. To be fair with the experiments in the latter case, we must say that the use of functions in our specification of the π-calculus, while providing excellent results for our goals, might make the proofs of correctness of bisimilation techniques a bit harder than usual.

An extensive implementation of the π-calculus, including proofs of correctness of bisimilation techniques, has been realized by Daniel Hirschkoff [Hir97,Hir99] using the de Bruijn method. Our contribution is much more modest in the sense that it only provides a new basis for such a study. However, the de Bruijn technique quickly leads to very obscur semantic descriptions and makes proofs long and tedious. 75 percent of the development by Daniel Hirschkoff concern manipulation of de Bruijn codes, which is not the subject of interest.

A different method has recently been investigated by Guillaume Gillard [Gil99], on a concurrent object calculus recently proposed by Andrew Gordon and Paul Hankin. Guillaume Gillard provides a proof of preservation of types for this calculus. The method used is due to Andrew Gordon [Gor94]. It consists in describing terms modulo alpha-conversion, over a predefined set of λ-terms implemented with de Bruijn codes. The part of the development devoted to the manipulation of de Bruijn terms, while still sizeable, is definitly more reasonable than in the standard approach.

A recent development has been realized by Loic Henry-Greard [Hen98]. The proof done is the same as ours: the proof of preservation of types for the monadic π-calculus. The proof uses a (slight extension of a) method due to Randy Pollack and James Mc Kinna, extensively used to formalize Lego in Lego. The basic idea of the method is to distinguish between bound variables (represented by variables) and free variables (represented by parameters). This gives a nice solution to the problem of capture-avoiding substitution. However, there is still an overhead in both the descriptions and the proofs, to manipulate the variables and the parameters. It is interesting to note that our proof, while being considerably shorter, has exactly the same struture as Loic Henry-Greard's proof.

Finally, Furio Honsell and Marino Miculan have proposed a higher-order description of the π-calculus (with the match operator) [HMS99], different from ours, which they have used in the proofs of correctness of various bisimilarity techniques. Their formalisation uses meta-level (Coq) functions to encode the input and restriction operators like us; However, the rest of the presentation greatly differs from ours. They need to duplicate the evaluation rules in a non-natural way: some of the rules yield processes while the others yield functions from name to processes. Communication is formalized using open and close rules, which, despite the use of higher-order syntax, need a fresh predicate. Meanwhile, they provide a formal study of some bisimilarity principles, which, as it uses rules of evaluation (partly) described by means of functions, should a-priori inspire a similar study using our specification.

8 Conclusion and Future Work

We have presented in this paper a higher-order formalization of a monadic π-calculus in the Calculus of Inductive Constructions, and provided a proof of type preservation for this calculus in the Coq system. The specification we proposed formalizes communication by a function application. Our proof is very short: it is only three pages long. We believe that the same proof performed on alternative descriptions of the same calculus would give much longer and more tedious proofs. Extending our development to the polyadic π-calculus, while requiring longer proofs, should be straightforward.

It might be worth noting that the technique of higher-order abstract syntax has been very rarely used to describe imperative languages up to now. It seems that this method, while being of great benefit in the descriptions of functional programming languages, is a bit less suitable

for imperative languages. However, the π-calculus is considered as an imperative language, as communication is described as a side-effect in the meta-theoretical studies of the calculus. We may hope that the specification proposed here might give new ideas for the description of imperative languages.

Another interesting point to note is that only a limited form of higher-order abstract syntax is required to describe first-order, usual π-calculi. The point is that we only need functions from name to processes to describe binders, while functions from processes to processes would be needed in the description of the higher-order π-calculus, where processes themselves can be passed on a channel. The problem is that such types are not allowed as types of constructors of an inductive type for processes, at least in the standard way as the one provided by almost all the current systems providing induction. Consequently, we are not able to use the technique of higher-order abstract syntax to describe the higher-order π-calculus today. However, there are already propositions for various meta-logics in the literature [DPS97,DL98,MM97,Hof99] which could be used as a basis for the implementation of the system we need.

Concerning future work, our main goal is to provide the basis for a machine-checked study of languages based on the π-calculus. Particularly relevant here are proofs of properties such as the preservation of types for languages for which the rich type system make those proofs challenging, even in their hand-written presentations. We think here of languages involving polymorphism as the Pict [PT97] or the Join [FGL$^+$96] languages, or to type-checking rules to prevent dead-lock as proposed in [Kob98].

Other interesting languages will be those arising from the various nice propositions for concurrent and object calculi, which are outcoming of active research in the field [Bou97,GH98].

As we already mentioned, future work of a different nature is the study of the theory of the π-calculus: bisimilarity techniques. It would be very interesting to see whether the higher-order description of the π-calculus we propose here can be of real benefit too in such formal studies.

In the long term, we hope such formal studies to be used to formalize proof of correctness of programs written in languages based on the π-calculus.

Acknowledgements. Thanks go to Davide Sangiorgi who introduced us to the π-calculus. Our description of the calculus comes out of many discussions we add, trying to understand and formalize the binding mechanism in the π-calculus.

References

[BBC$^+$97] B. Barras, S. Boutin, C. Cornes, J. Courant, J.C. Filliâtre, E. Giménez, H. Herbelin, G. Huet, C. Muñoz, C. Murthy, C. Parent, C. Paulin, A. Saïbi, and B. Werner. The Coq Proof Assistant Reference Manual – Version V6.1. Technical Report 0203, INRIA, August 1997.

[Bou97] Gérard Boudol. The pi-calculus in direct style. In *proceedings of the POPL ACM Conference on Principles of Programming Languages*, pages 228–241, 1997.

[CH88] Thierry Coquand and Gérard Huet. The calculus of constructions. *Information and Computation*, 76:95–120, 1988.

[DL98] Joëlle Despeyroux and Pierre Leleu. A modal λ-calcul with iteration and case constructs. In *proceedings of the annual Types for Proofs and Programs seminar*, Springer-Verlag LNCS 1657, March 1998.

[DPS97] Joëlle Despeyroux, Frank Pfenning, and Carsten Schürmann. Primitive recursion for higher-order abstract syntax. In Philippe de Groote and J. Roger Hindley, editors, *proceedings of the TLCA 97 Int. Conference on Typed Lambda Calculi and Applications, Nancy, France, April 2–4*, pages 147–163. Springer-Verlag LNCS 1210, April 1997. An extended version is available as CMU Technical Report CMU-CS-96-172.

[FGL+96] Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A calculus of mobile agents. In *Proc. of the CONCUR conference*, 1996.

[GH98] A. Gordon and P. Hankin. A concurrent object calculus: reduction and typing. In *proceedings of the HLCL Conference*, 1998.

[Gil99] Guillaume Gillard. A formalization of a concurrent object calculus up to alpha-conversion. submitted for publication, January 1999.

[Gor94] A. Gordon. A mechanisation of name-carrying syntax up to alpha-conversion. In *proceedings of the int. workshop on Higher Order Logic Theorem Proving and its Applications, Vancouver*, Springer-Verlag LNCS 780, pages 414–426, 1994.

[Hen98] Loïc Henry_Greard. A proof of type preservation for the pi-calculus in Coq. Research Report RR-3698, Inria, December 1998. Also available in the Coq Contrib library.

[Hir97] Daniel Hirschkoff. A full formalization of pi-calculus theory in the Calculus of Constructions. In Elsa Gunter, editor, *proceedings of the International Conference on Theorem Proving in Higher Order Logics*, Murray Hill, New Jersey, August 1997.

[Hir99] Daniel Hirschkoff. *Mise en oeuvre de preuves de bisimulation.* Phd thesis, École Nationale des Ponts et Chaussées (ENPC), January 1999. In French.

[HKPM97] Gérard Huet, Gilles Kahn, and Christine Paulin-Mohring. The coq proof assistant: a tutorial, version 6.1. Technical Report RT 0204, Inria, Rocquencourt, France, August 1997. page html: http://www.inria.fr/RRRT/RT-0204.html.

[HMS99] Furio Honsell, Marino Miculan, and Ivan Scagnetto. Pi-calculus in (co)Inductive Type Theory. *to appear in TCS*, 1999.

[Hof99] Martin Hofmann. Semantical analysis of higher-order abstract syntax. In IEEE, editor, *proceedings of the International Conference on Logic In Computer Sciences, LICS*, 1999.

[Kob98] Naoki Kobayashi. A partially deadlock-free typed process calculus. *ACM Transactions on Programming Languages and Systems*, 20-2, December 1998. A preliminary summary appeared in LICS'97.

[Mil91] Robin Milner. The polyadic pi-calculus: a tutorial. Technical Report ECS-LFCS-91-180, LFCS, Edinburgh University, October 1991.

[MM97] Raymond McDowell and Dale Miller. A logic for reasoning with higher-order abstract syntax: An extended abstract. In Glynn Winskel, editor, *Proceedings of the Twelfth Annual Symposium on Logic in Computer Science*, Warsaw, Poland, June 1997. To appear.

[MPW92] R. Milner, R. Parrow, and J. Walker. A calculus of mobile processes, (part I and II). *Information and Computation*, 100:1–77, 1992.

[PM92] Christine Paulin-Mohring. Inductive definitions in the system coq. rules and properties. In J.F. Groote M. Bezem, editor, *Proceedings of the International Conference on Typed Lambda Calculi and Applications, TLCA'93*, Springer-Verlag LNCS 664, pages 328–345, 1992. also available as a Research Report RR-92-49, Dec. 1992, ENS Lyon, France.

[PS95] Benjamin Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. *Mathematical Structures in Computer Science*, 11, 1995.

[PT97] Benjamin Pierce and David Turner. Pict: A programming language based on the pi-calculus. Technical Report ., IU, 1997.