

PILIB: A Hosted Language for Pi-Calculus Style Concurrency

Vincent Cremet and Martin Odersky

École Polytechnique Fédérale de Lausanne
1015 Lausanne, Switzerland
{vincent.cremet,martin.odersky}@epfl.ch

Abstract. PILIB is a library written in SCALA that implements the concurrency constructs of the π -calculus. Some features of the programming language SCALA, examined in the paper, make it possible to use almost the same syntax as in the π -calculus. The advantages of this library with respect to a simple π -calculus interpreter are that we can transmit any value along names, we can get control over them using the type system, and we have access to the full power of SCALA in terms of expressiveness and libraries.

1 Introduction

Support for concurrency is now almost universal in general purpose programming languages. But the supported concepts vary widely from one language to the next. For example, Concurrent ML [1] is based on an event-based model for process synchronization. Occam [2] is a programming language based on CSP [3]. In a similar way, Pict [4] is based on the π -calculus [5, 6]. Facile [7] also uses π -calculus style channel-based communication. JoCaml [8], Funnel [9], and Polyphonic C# [10] use synchronization patterns from join calculus [11]. Erlang [12] is based on an actor model where processes interpret messages in their mailboxes by pattern matching. Oz [13] uses logic variables for process communication. Id [14] and Concurrent Haskell [15] use a form of mutable variables called M-structures [16].

In contrast, most mainstream languages are based on a shared memory thread model. Recent popular languages such as Java [17] or the .NET common language runtime [18] augment this model with thread synchronization based on the Monitor concept [19, 20].

In spite of this confusing variety, some trends can be observed. A large body of theoretical research in concurrency is based on CCS or its successor, the π -calculus. The same holds for many specifications of concurrent systems. A π -calculus specification can be an executable program. Nevertheless, most concurrent systems are still implemented with shared-memory threads, which are synchronized with semaphores or monitors. Such programs are often much harder to reason about and verify than π -calculus specifications. But they are also often more efficient to implement.

Given this situation, it seems desirable to have a wide spectrum of programming solutions in which one can move easily from a high-level specification and prototype to a (thread-based) low-level implementation. One way to achieve this is with a programming language that has both low-level and high-level concurrency constructs. However, such a language would tend to be quite large. Moreover, in light of the great variety of high-level solutions that have been proposed it seems difficult to pick a concrete high-level syntax with confidence that it is the “right one”.

A better solution is to express the high-level language as a library abstraction in terms of the low-level one. Then the high-level language becomes, in effect, a domain-specific language where the domain is process coordination. This approach is similar to the use of skeletons in parallel programming [21, 22].

Being implemented as a library, the high-level language is embedded (or: hosted) in the low-level language. This has the advantages that no separate compiler and run-time environment needs to be designed, and that all of the facilities of the host-language can be used from the high-level language. However, if the high-level language is to be convenient to use, the host language needs to have a fair degree of expressiveness.

In this paper, we describe our experience with such an embedding. We have developed a process coordination language that is closely modeled after the π -calculus and that is implemented as a library in the host language Scala [23]. Scala is a new functional/object-oriented language that interacts smoothly with Java and C#. Compared to these environments, Scala has several additional language features that make it suitable as a host language for domain-specific languages. Among others, it supports the following concepts.

- A rich type system, with generics as well as abstract and dependent types.
- Object composition using mixin-style multiple inheritance.
- Named as well as anonymous functions as first-class values that can be nested.
- Pure object orientation, in the sense that every value is conceptually an object and every operator is a method call.

Scala is designed to operate in a JVM or .NET environment, so it can be regarded as an extension language for Java or C#. Scala does not have any constructs dealing with concurrency in the language proper. Instead, it re-uses the concurrency constructs of the underlying environment. These constructs are almost the same for Java and .NET – they consist in each case of a class-based thread model with monitors for synchronization.

The need to define a high-level process model for Scala came up when we taught a class on concurrency theory and concurrent programming. We wanted to be able to have students program concurrent systems that closely follow specifications in CCS and the π -calculus. At the same time, we wanted to relate this to Java’s classical thread/monitor model of concurrent programming. Embedding the high-level concurrent language in Scala provided an elegant means of going from specification to (classical) implementation and at the same time made available the rich facilities of the Java environment.

Implementing a high-level concurrent language as a library abstraction poses some challenges for the syntactic and type system abilities of the host language. In this paper we show how a convenient syntax and type discipline for the high-level language can be achieved with the help of Scala’s constructs for abstraction and composition. Particularly important in this respect are Scala’s generic type system, its ability to have arbitrary operators as methods, the dynamic, receiver-based interpretation of such operators, and its light-weight syntax for nestable anonymous functions (i.e., closures).

The rest of this paper is organized as follows. Section 2 is a short introduction to the π -calculus. Section 3 allows the reader to appreciate the close link between PiLIB and the π -calculus by comparing how a small example can be expressed in both formalisms. In Section 4 we present the elements of the PiLIB language in detail. Section 5 shows how these elements are mapped to SCALA constructs. Section 6 recapitulates the languages features necessary for the embedding. Section 7 gives a brief description of PiLIB’s implementation, and Section 8 concludes.

2 The π -Calculus in Short

The π -calculus is a model of concurrent computation in which processes can connect dynamically by exchanging the names of the channels they use to communicate. A calculus can be seen as a programming language that is limited to essential concepts in order to permit formal reasoning. So, as for programming languages, the precise definition of the π -calculus includes a description of the syntax of the terms that compose the language as well as a definition of their meaning.

2.1 Syntax

There are various versions of the π -calculus. The variant we chose for PiLIB has very few constructs and is similar to Milner’s definition in [6]. The only differences in our definition are the absence of an unobservable action τ and the use of recursive agent definitions instead of the replication operator !.

Here is the inductive definition of π -calculus processes.

Processes

P, Q	$::= \sum_{i=0}^n G_i$	Sum	$(n \geq 0)$
	$\left \nu x.P$	Restriction	
	$\left P \mid Q$	Parallel	
	$\left A(y_1, \dots, y_n)$	Identifier	$(n \geq 0)$

Guarded processes

G	$::= x(y).P$	Input	
	$\left \bar{x}\langle y \rangle.P$	Output	

Definitions

$$A(x_1, \dots, x_n) \stackrel{\text{def}}{=} P \quad (x_i \text{'s distinct, } \text{fn}(P) \subset \bar{x})$$

A *process* in the π -calculus is either a choice $\sum G_i$ of several guarded processes G_i , or a parallel composition $P \mid Q$ of two processes P and Q , or a restriction $\nu x.P$ of a private fresh name x in some process P , or a process identifier $A(y_1, \dots, y_n)$ which refers to a (potentially recursive) process definition. In a sum, processes are guarded by an input action $x(y)$ or an output action $\bar{x}\langle y \rangle$, which receives (respectively sends) a channel y via the channel x before the execution of the two communicating processes continues. In the process definition, $\text{fn}(P)$ is the set of free names of the process P , i.e. the names that are not bound through the parameter of an input guarded process or through the restriction operator.

2.2 Structural Equivalence

A first way of assigning a meaning to processes is to define an equivalence relation between them. For instance identifying processes that differ only by the order of parallel sub-processes is a way of giving a meaning to the \mid operator. The equivalence that we define in this section is structural in the sense that it is based only on the syntax of processes. This is in contrast to behavioral equivalences which try to identify programs that have similar behaviors.

The structural equivalence is defined as the smallest congruence that contains a set of equalities. The more interesting equalities are the scope extrusion rules (EXTR-PAR, EXTR-SUM) that allow to pull a restriction to the top-level of a process, and the rule (UNFOLD) which replaces the occurrence of an agent name with the body of its definition:

$$\begin{aligned} \text{EXTR-PAR} \quad & \frac{}{P \mid \nu x.Q \equiv \nu x.(P \mid Q)} \quad x \notin \text{fn}(P) \\ \text{EXTR-SUM} \quad & \frac{}{P + \nu x.Q \equiv \nu x.(P + Q)} \quad x \notin \text{fn}(P) \\ \text{UNFOLD} \quad & \frac{A(\tilde{x}) \stackrel{\text{def}}{=} P}{A(\tilde{y}) \equiv P[\tilde{y}/\tilde{x}]} \end{aligned}$$

In the third equality, \tilde{x} represents a sequence x_1, \dots, x_n of names and $P[\tilde{y}/\tilde{x}]$ denotes the substitution of y_i for every occurrence of x_i in P .

Beside these three equalities, there are also rules that identify processes that differ only in the names of bound variables, that declare $+$ and \mid to be associative and commutative, and that allow permuting restriction binders νx .

2.3 Operational Semantics

The standard semantics of the π -calculus is given by a labeled reduction relation. Labels associated to reduction steps are needed to define bisimulation, an equivalence relation that identifies processes that can simulate each other

mutually. But to understand the computational content of a π -calculus term a reduction relation without labels is simpler. We therefore define \rightarrow to be the smallest relation that satisfies the following inference rules:

$$\begin{array}{c} \text{PAR} \frac{P \rightarrow P'}{P \mid Q \rightarrow P' \mid Q} \\ \\ \text{NU} \frac{P \rightarrow Q}{\nu x.P \rightarrow \nu x.Q} \\ \\ \text{STRUCT} \frac{P \equiv P' \quad P' \rightarrow Q' \quad Q' \equiv Q}{P \rightarrow Q} \\ \\ \text{COM} \frac{}{a(x).P + P' \mid \bar{a}(y).Q + Q' \rightarrow P[y/x] \mid Q} \end{array}$$

The rule PAR tells that a process can evolve independently of other processes in parallel. The rule NU is a contextual rule that makes it possible to reduce under a restriction. The rule STRUCT is the one that allows to identify equivalent processes when considering reduction. Finally, rule COM is the key to understand the synchronization mechanism of the calculus: two parallel processes that can perform complementary actions, i.e. send and receive the name of some channel along the same channel, can communicate. They proceed next with the guarded processes associated to the involved send and receive action. Note that the name x potentially free in the continuation P of the input guarded process $a(x).P$ is bound to the transmitted name y thanks to the substitution $[y/x]$ applied to P . The alternatives in each communicating sum that do not play a role in the communication, namely P' and Q' , are discarded, for this reason a sum is also called a *choice*.

3 Example: The Two-Place Buffer

3.1 Specification in the π -Calculus

The specification of a two-place buffer in the π -calculus consists of a set of mutually recursive process definitions:

$$\begin{aligned} \text{Buffer}(\text{put}, \text{get}) &= B_0(\text{put}, \text{get}) \\ B_0(\text{put}, \text{get}) &= \text{put}(x).B_1(\text{put}, \text{get}, x) \\ B_1(\text{put}, \text{get}, x) &= \overline{\text{get}}\langle x \rangle.B_0(\text{put}, \text{get}) + \text{put}(y).B_2(\text{put}, \text{get}, x, y) \\ B_2(\text{put}, \text{get}, x, y) &= \overline{\text{get}}\langle x \rangle.B_1(\text{put}, \text{get}, y) \end{aligned}$$

The definitions B_0 , B_1 and B_2 represent respectively the state of the buffer when it contains zero, one or two elements, the buffer is initially empty. All the definitions are parameterized by the names of the channels used to put (*put*) or

get (*get*) an element. Additionally definitions B_1 and B_2 are parameterized by the names of the elements stored in the buffer.

The logic of these definitions is then very intuitive: if the buffer is empty (state B_0) it can only accept the input of an item on channel *put* and move to state B_1 . If it contains exactly one item (state B_1), it can either output this item and go back to state B_0 or accept a new item and move to state B_2 . Finally when the buffer contains two items it can only output the first stored item and go to state B_1 .

$put(x).B_1(put, get, x)$ and $\overline{get}\langle x \rangle.B_0(put, get)$ are two instances of *guarded processes*. The first one is an input guarded process which binds an input item on channel *put* to the name x in the continuation process $B_1(put, get, x)$. The second one is an output guarded process which outputs the item x on channel *get* before continuing with process $B_0(put, get)$.

3.2 Implementation with PiLIB

Because PiLIB is modeled after π -calculus concurrency primitives, it is possible to reproduce the specification above almost as it is:

```
def Buffer[a](put: Chan[a], get: Chan[a]): unit = {
  def B0: unit = choice ( put * { x  $\Rightarrow$  B1(x) } );
  def B1(x: a): unit = choice ( get(x) * B0, put * { y  $\Rightarrow$  B2(x, y) } );
  def B2(x: a, y: a): unit = choice ( get(x) * B1(y) );
  B0
}
```

Process definitions are replaced by function definitions (introduced by **def** in SCALA) with result type `unit`. This type contains just one value and corresponds to `void` in C or JAVA. The buffer is parameterized by the type a of the stored items. $\text{Chan}[a]$ is the type of channels carrying objects of type a . A π -calculus input guarded process like $put(x).B_1(put, get, x)$ is now written `put * { x \Rightarrow B1(x) }`, and an output guarded process like $\overline{get}\langle x \rangle.B_0(put, get)$ is now written `get(x) * B0` (the character `*` replaces the dot symbol). A π -calculus summation $+$ between several alternatives is now a sequence of guarded processes separated by a comma and enclosed in `choice (...)`.

Compared to the specification in the π -calculus, the implementation using PiLIB has several advantages. First the channels have a type and the typechecker phase of the compiler ensures that only objects of the right type are passed to channels. Second, since processes in the π -calculus are modeled as functions, it is possible to hide the internal states B_0 , B_1 and B_2 by nesting the corresponding functions inside the function `Buffer`. The parameters *put* and *get* are visible from these local functions and do not need to be passed as extra parameters.

3.3 Using the Buffer

Here is an example of using the buffer in the π -calculus, where *Producer* is a process that repeatedly creates a new channel and put it in the buffer, and

Consumer is a process that tirelessly gets an item from the buffer and discards it.

$$\begin{aligned} \text{Producer}(\text{put}, \text{get}) &= \nu x. \overline{\text{put}}\langle x \rangle. \text{Producer}(\text{put}, \text{get}) \\ \text{Consumer}(\text{put}, \text{get}) &= \text{get}(x). \text{Consumer}(\text{put}, \text{get}) \end{aligned}$$

$$\nu \text{put}, \text{get}. \text{Producer}(\text{put}, \text{get}) \mid \text{Buffer}(\text{put}, \text{get}) \mid \text{Consumer}(\text{put}, \text{get})$$

The three processes are put in parallel using the operator \mid and are linked together through the sharing of the fresh channels *put* and *get* introduced by the restriction operator ν .

In the example above the values added to the buffer by the producer and extracted from it by the consumer are π -calculus channels because there is nothing else to transmit in the π -calculus. So both channels and the values they carry are taken from the same domain. A typical way of typing such recursive channels in PILIB consists in using a recursive type definition:

```
class Channel extends Chan[Channel];
```

Such a definition can be read: “A π -calculus channel is a channel that can carry other π -calculus channels”.

Using this type definition, the π -calculus code above has now an exact counterpart in PILIB:

```
def Producer(put: Channel, get: Channel): unit = {
  val x = new Channel;
  choice ( put(x) * Producer(put, get) )
}
def Consumer(put: Channel, get: Channel): unit =
  choice ( get * { x  $\Rightarrow$  Consumer(put, get) } );
```

```
val put = new Channel, get = new Channel;
spawn < Producer(put, get) | Buffer(put, get) | Consumer(put, get) >
```

New features of PILIB appear only in the way of creating new channels and executing several processes in parallel using `spawn`.

As we have seen, the implementation using PILIB is very close to the specification in π -calculus. It seems as if SCALA has special syntax for π -calculus primitives, however PILIB is nothing more than a library implemented using low-level concurrency constructs of SCALA inherited from JAVA. In the rest of the paper we will try to demystify the magic.

Of course we can also write a two-place buffer using monitors as in JAVA and the implementation would certainly be more efficient. But then relating the implementation to the specification would be non-trivial. With PILIB we can closely relate the specification language and the implementation language, and thereby gain immediate confidence in our program.

4 Description of the PiLIB Grammar

In this section we briefly describe the primitive concurrency interface of SCALA, directly inherited from JAVA, and then introduce the PiLIB interface.

4.1 Original Concurrency Interface

The original concurrency interface of SCALA consists mainly of a class Monitor and a function fork with the following signatures.

```
class Monitor {
  def synchronized[a](def p: a): a;
  def wait(): unit;
  def wait(timeout: long): unit;
  def notify(): unit;
  def notifyAll(): unit;
}

def fork(def p: unit): unit;
```

The method `synchronized` in class Monitor executes its argument computation `p` in mutual exclusive mode – at any one time, only one thread can execute a synchronized argument of a given monitor.

A thread can suspend inside a monitor by calling the monitor’s `wait` method. At this point, the thread is suspended until a `notify` method of the same monitor is called by some other thread. If no thread is waiting on this monitor, a call to `notify` is ignored.

The `fork` method creates a new thread for the execution of its argument and returns immediately.

4.2 PiLIB Interface

Now we present the grammar of PiLIB. How this grammar is actually implemented as a library in SCALA is the topic of the next section. Here we just present the syntax of PiLIB constructs and give an informal description of their associated semantics.

Channel creation. At the basis of the PiLIB interface is the concept of *channel* (also called “name” in the π -calculus). A channel represents a communication medium. To get an object that represents a fresh channel that can carry objects of type A , one simply writes `new Chan[A]`.

Guarded processes. Assuming an expression a that evaluates to a channel carrying objects of type A , the term $a * \{ x \Rightarrow c \}$ is an *input guarded process* with continuation c . References to x in c are bound to the value of the transmitted

object. The type of a guarded process whose continuation has result type B is $\text{GP}[B]$. Note that instead of $\{x \Rightarrow c\}$ we could have used any expression of type $A \Rightarrow B$ (the type of functions from A to B).

Similarly, an *output guarded process* is written $a(v) * c$ where v is the value to be sent and c is any continuation expression. The type of the guarded process is again $\text{GP}[B]$, where B is the type of the continuation expression c .

As in the π -calculus on which it is based, communications in PiLIB are synchronous: when a thread tries to output (resp. input) a value on a given channel it blocks as long as there is no other thread that tries to perform an input (resp. output) on the same channel. When two threads are finally able to communicate, through input and output guarded processes on the same channel, the thread that performs the input proceeds with the continuation of the input guarded process applied to the transmitted value, and the thread that performs the output proceeds with the continuation of the output guarded process.

Summation. The next specific ingredient of PiLIB is the function choice which takes an arbitrary number of guarded processes as arguments and tries to establish communication with another thread using one of these guarded processes. The choice function blocks until communication takes place, as explained above. Once the communication takes place, the guarded processes among the arguments to the function choice that do not take part in the communication are discarded.

Forking multiple threads. The construction `spawn` is used to fork several threads at the same time. For instance to execute concurrently some expressions p, q, r , each of type `unit`, one writes

$$\text{spawn } \langle p \mid q \mid r \rangle$$

The figure 1 summarizes the concurrency primitives provided by PiLIB with their typing rules. All of these typing rules are admissible in the host type system.

$$\begin{array}{l}
 \text{NU} \quad \frac{}{\text{new Chan}[A] : \text{Chan}[A]} \\
 \text{INPUT} \quad \frac{a : \text{Chan}[A] \quad \{x \Rightarrow c\} : A \Rightarrow B}{a * \{x \Rightarrow c\} : \text{GP}[B]} \\
 \text{OUTPUT} \quad \frac{a : \text{Chan}[A] \quad v : A \quad c : B}{a(v) * c : \text{GP}[B]} \\
 \text{CHOICE} \quad \frac{g_i : \text{GP}[B] \quad i \in \{1, \dots, n\} \quad n \geq 0}{\text{choice } (g_1, \dots, g_n) : B} \\
 \text{SPAWN} \quad \frac{p_i : \text{unit} \quad i \in \{1, \dots, n\} \quad n \geq 1}{\text{spawn } \langle p_1 \mid \dots \mid p_n \rangle : \text{unit}}
 \end{array}$$

Fig. 1. PiLIB constructs

4.3 Derived Constructs

The class `Chan[A]` also contains methods to perform synchronous read and synchronous write. These constructs can be derived from the more primitive constructs presented so far.

If a is an expression of type `Chan[A]`, an example of a synchronous read is:

```
val x = a.read;
```

It is equivalent to (and implemented as)

```
var x: A = null;
choice (a * { y ⇒ x = y });
```

Similarly, a synchronous write `a.write(v)` is equivalent to `choice (a(v) * {})`. Each of the derived constructs corresponds to a sum containing just one alternative, what is expressed in PiLiB by a call to the function `choice` with only one argument.

5 Desugaring

In this section we explain how it is possible to implement PiLiB as an hosted language, i.e. as a SCALA library. We consider each PiLiB construct in turn and see how it is interpreted by the language.

Channel Creation

In PiLiB, a fresh channel carrying objects of type A is created using the syntax `new Chan[A]`. Indeed, this is the normal SCALA syntax for creating instances of the parameterized class `Chan`.

Input Guarded Process Expression

SCALA allows a one-argument method to be used as an infix operator, i.e. the SCALA parser will replace any expression `e f g` by `e.f(g)`. Furthermore, `*` is a valid identifier. So an input guarded process

```
a * { x ⇒ c }
```

is recognized by the parser as

```
a.*( { x ⇒ c } ) .
```

In the code above, a is an expression of type `Chan[A]`, and `{ x ⇒ c }` is an anonymous function of type $A \Rightarrow B$. It works because the class `Chan[T]` contains a polymorphic method named `*` with the following signature:

```
def *[U](f: T ⇒ U): GP[U];
```

The anonymous function is itself desugared as the creation of an anonymous object of type `Function1[A, B]` (if A is the type inferred for the parameter x and B the type inferred for the body c) with a method `apply` whose unique parameter is x and whose body is c . That is, $\{x \Rightarrow c\}$ is translated to

```
new Function1[A,B] {
  def apply(x: A): B = c;
}
```

Output Guarded Process Expression

In SCALA an expression $a(v)$ where a has type `Function1[T,U]` is replaced by $a.apply(v)$. So an expression describing an output guarded process $a(v) * c$ is recognized as $a.apply(v).*c$. It implies that the class `Chan[A]` implements the interface `Function1[A,Product[A]]` where A is the type of the transmitted value and `Product[A]` is an helper class that contains the method `*`, which is necessary for simulating the dot symbol of the natural syntax.

Guarded Process

The type of an input or output guarded process expression is `GP[A]` where A is the result type of the continuation.

The class `GP[A]` encapsulates the elements that compose a guarded process: the name of the channel and the continuation function for an input guarded process, and the name of the channel, the transmitted value and the continuation function for an output guarded process.

Summation

A summation is written `choice(g1, ..., gn)` in Pilib. It is just a call to the polymorphic function `choice` with signature:

```
def choice[A](gs: GP[A]*): A:
```

As specified by the star symbol at the end of the argument type, the argument can be a sequence of guarded processes of arbitrary length.

Forking of Multiple Threads

The syntax to fork several threads is `spawn < p | q | r >`. As we have seen previously, this expression is recognized as `spawn.<(p).(q).(r).>`. Here is the implementation of this construct:

```
abstract class Spawn {
  def <(def p: unit): Spawn;
  def |(def p: unit): Spawn;
  def > : unit;
}
```

```

object spawn extends Spawn {
  def <(def p: unit): Spawn = { fork(p); this }
  def |(def p: unit): Spawn = { fork(p); this }
  def > : unit = ()
}

```

In the code above, `fork` is the low-level function to fork a sub-process.

The `def` modifier applied to parameters in the methods of the abstract class `Spawn` specifies that the argument must be passed without being evaluated. It will be evaluated each time it is used in the body of the function. This strategy of evaluation is called *call-by-name*. Without this indication, the different threads would be executed sequentially in the order of evaluation of the arguments of the method.

6 Scala Constructs Used for the Embedding

In this section we summarize the features that proved to be useful for hosting PiLIB in SCALA.

Polymorphism

Type genericity allowed us to parameterize a channel by the type of the objects it can carry. The SCALA type system makes sure that a channel is always used with an object of the right type.

In SCALA, methods can also be polymorphic. As π -calculus agents are represented by methods in PiLIB, such agents can be polymorphic, like the two-place buffer in the example of Section 3.

Type parameters may have bounds in SCALA, and these bounds can be recursive. This feature known under the name of *F-bounded polymorphism* can be used to express recursive channels. For instance to define a kind of channel that can carry pairs of integers and channels of the same kind, we would write

```

class C extends Chan[Pair[int, C]]

```

Higher-Order Functions and Anonymous Functions

We already saw that an input guarded process is expressed as a call of the method `*` with a continuation function as argument. This is an example of a higher-order function. Furthermore it is natural to express the continuation as an anonymous function because most of the time it is not recursive and does not need to be referred elsewhere in the program.

The type inference performed by the compiler allows some type annotations to be omitted, like the type of the formal parameter in the continuation of an input guarded process.

Syntactic Sugar

It is convenient to write `x(v) * c` instead of `x.apply(v).*(c)`. This is permitted by the SCALA parser which perform simple but powerful transformations.

The fact that all functions are members of a class (i.e. methods) allows us to overload already existing operators such as `*` easily without introducing ambiguity.

Call-by-Name

The parameter modifier `def` makes it possible in SCALA to explicitly specify that some parameters must be called by name. This feature of the language is used on two occasions in PILIB. In an output guarded process `x(v) * c`, the continuation `c` must not be evaluated at the creation of the guarded process but only after communication takes place with another matching guarded process, this is a perfect candidate to be passed by name. Also in the construct `spawn < p | q >`, the call-by-name strategy for the parameters of the construct `spawn` allows to avoid a sequential evaluation of the arguments `p` and `q`.

The alternative way to delay the evaluation of an expression `e` consists in manipulating the closure `() => e` but it is more cumbersome for the programmer.

Sequence Type

Another feature of SCALA, which is used in the definition of the choice function, is the possibility to pass an arbitrary number of arguments to a method.

7 Overview of the Implementation

To explain the implementation we first describe the data structures that are used and then present the means by which processes interact.

Representation of guarded processes. Two guarded processes are said *complementary* if one is an input and the other is an output on the same channel. For the sake of simplicity, the implementation unifies the concepts of input and output guarded processes: at the time where communication takes place between two complementary guarded processes, each of them provides a value to the other (the input guarded process provides the dummy value `()`) and applies its continuation to the value it receives. The data structure for a guarded process then consists of four components: a channel, a kind (input or output), the value it provides, and the continuation function to apply to the received value.

Representation of sums. A *sum* consists of an ordered list of guarded processes, together with the continuation to execute after communication takes place. At the creation of a sum, it is not known which of its guarded processes will react and what value this guarded process will receive, so the continuation of the sum is initially undefined. Two sums are said *complementary* if they contain complementary guarded processes. The implementation keeps a global list of pending sums, i.e. sums that are waiting to communicate. We call it the *pending list*. An invariant of this list is that the sums it contains are pairwise not complementary.

Choice resolution. Now we will explain what happens when a thread calls the function choice. The argument of the function choice is a sequence of guarded processes. This sequence is first turned into a sum with an undefined continuation. From now this sum is designated as the *arriving sum*. The pending list is scanned to find a complementary sum. If there is none, the arriving sum is appended to the end of the pending list, the choice function waits until the sum continuation gets a value and then execute this continuation. If there is a complementary sum, it is extracted from the pending list, both communicating sums get the value for their continuation and their respective threads are woken up.

The implementation consists of about 150 lines of SCALA code with a large part dedicated to the syntactic sugar. The complete source code is available on the web [24].

8 Conclusion

We have presented PiLIB, a language for concurrent programming in the style of the π -calculus. The language is hosted as a library in SCALA. The hosting provides “for free” a highly higher-order π -calculus, because all parts of a PiLIB construct are values that can be passed to methods and transmitted over channels. This is the case for channels themselves, output prefixes, continuations of input and output guarded processes, and guarded processes.

A PiLIB program corresponds closely to a π -calculus specification. Of course, any run of that program will only perform one of the possible traces of the specification. Moreover, this trace is not necessarily fair, because of PiLIB’s method of choosing a matching sum in the pending list. A fair implementation remains a topic for future work.

There is a result of Palamidessi [25, 26] showing that it is not possible to implement the mixed choice π -calculus (when a sum can contain at the same time input and output guarded processes) into an asynchronous language in a distributed (deterministic) way. Our current implementation is centralized and in this case the mixed choice is not problematic. A possible continuation of this work is to implement a distributed version of PiLIB, which would force us to abandon the mixed choice.

Our experience with PiLIB in class has been positive. Because of the close connections to the π -calculus, students quickly became familiar with the syntax and programming methodology. The high-level process abstractions were a big help in developing correct solutions to concurrency problems. Generally, it was far easier for the students to develop a correct system using PiLIB than using Java’s native thread and monitor-based concurrency constructs.

References

1. Reppy, J.: CML: A higher-order concurrent language. In: *Programming Language Design and Implementation*, SIGPLAN, ACM (1991) 293–259
2. INMOS Ltd.: *OCCAM Programming Manual*. Prentice-Hall International (1984)
3. Hoare, C.A.R.: Communicating sequential processes. *Communications of the ACM* **21** (1978) 666–677 Reprinted in “*Distributed Computing: Concepts and Implementations*” edited by McEntire, O’Reilly and Larson, IEEE, 1984.
4. Pierce, B.C., Turner, D.N.: Pict: A programming language based on the pi-calculus. In Plotkin, G., Stirling, C., Tofte, M., eds.: *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press (2000) 455–494
5. Milner, R., Parrow, J., Walker, D.: A calculus of mobile processes (Parts I and II). *Information and Computation* **100** (1992) 1–77
6. Milner, R.: *Communicating and Mobile Systems: the Pi-Calculus*. Cambridge University Press (1999)
7. Giacalone, A., Mishra, P., Prasad, S.: Facile: A symmetric integration of concurrent and functional programming. *International Journal of Parallel Programming* **18** (1989) 121–160
8. Conchon, S., Fessant, F.L.: Jocaml: Mobile agents for Objective-Caml. In: *First International Symposium on Agent Systems and Applications (ASA’99)/Third International Symposium on Mobile Agents (MA’99)*, Palm Springs, CA, USA (1999)
9. Odersky, M.: Functional nets. In: *Proc. European Symposium on Programming*. Number 1782 in LNCS, Springer Verlag (2000) 1–25
10. Benton, N., Cardelli, L., Fournet, C.: Modern concurrency abstractions for C[#]. In: *Proceedings of the 16th European Conference on Object-Oriented Programming*, Springer-Verlag (2002) 415–440
11. Fournet, C., Gonthier, G.: The reflexive chemical abstract machine and the join-calculus. In: *Principles of Programming Languages*. (1996)
12. Armstrong, J., Virding, R., Wikström, C., Williams, M.: *Concurrent Programming in Erlang*, Second Edition. Prentice-Hall (1996)
13. Smolka, G., Henz, M., Würtz, J.: Object-oriented concurrent constraint programming in Oz. In van Hentenryck, P., Saraswat, V., eds.: *Principles and Practice of Constraint Programming*. The MIT Press (1995) 29–48
14. Arvind, Gostelow, K., Plouffe, W.: The ID-Report: An Asynchronous Programming Language and Computing Machine. Technical Report 114, University of California, Irvine, California, USA (1978)
15. Peyton Jones, S., Gordon, A., Finne, S.: Concurrent Haskell. In ACM, ed.: *Conference record of POPL ’96, 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages: papers presented at the Symposium: St. Petersburg Beach, Florida, 21–24 January 1996*, New York, NY, USA, ACM Press (1996) 295–308
16. Barth, P.S., Nikhil, R.S., Arvind: M-structures: Extending a parallel, non-strict, functional language with state. In Hughes, J., ed.: *Proceedings Functional Programming Languages and Computer Architecture, 5th ACM Conference*, Cambridge, MA, USA, Springer-Verlag (1991) 538–568 *Lecture Notes in Computer Science* 523.
17. Gosling, J., Joy, B., Steele, G., Bracha, G.: *The Java Language Specification*, Second Edition. Java Series, Sun Microsystems (2000) ISBN 0-201-31008-2.
18. Box, D.: *Essential .NET, Volume I: The Common Language Runtime*. Addison Wesley (2002)

19. Hoare, C.A.R.: Monitors: an operating system structuring concept. *Communications of the ACM* **17** (1974) 549–557
20. Hansen, P.B.: The programming language Concurrent Pascal. *IEEE Transactions on Software Engineering* **1** (1975) 199–207
21. Bischof, al.: Generic parallel programming using c++ templates and skeletons (2004) In this volume.
22. Kuchen, H.: Optimizing sequences of skeleton calls (2004) In this volume.
23. Odersky, M.: Report on the Programming Language Scala. (2002) <http://lampwww.epfl.ch/scala/>.
24. Cremet, V.: Pilib. <http://lampwww.epfl.ch/~cremet> (2003)
25. Palamidessi, C.: Comparing the expressive power of the synchronous and the asynchronous pi-calculus. In: Proc. of the 24th ACM Symposium on Principles of Programming Languages (POPL), ACM (1997) 256–265
26. Palamidessi, C.: Comparing the expressive power of the synchronous and the asynchronous pi-calculus. *Mathematical Structures in Computer Science* **13** (2003) 685–719