

# Programming in the Pi-Calculus

A Tutorial Introduction to Tamed Pict

(Tamed Pict Version 20070802 )

**Benjamin C. Pierce**

Computer Science Department  
Indiana University  
Lindley Hall 215  
Bloomington, Indiana 47405-4101  
USA  
`pierce@cs.indiana.edu`

**February 18, 2009**

## Abstract

Pict is a programming language in the ML tradition, formed by adding high-level derived forms and a powerful static type system to a tiny core language. The core, Milner's pi-calculus, is becoming popular as a theoretical foundation for a broad class of concurrent computations. The goal in Pict is to identify and support idioms that arise naturally when these primitives are used to build working programs — idioms such as basic data structures, protocols for returning results, higher-order programming, selective communication, and concurrent objects. The type system integrates a number of features found in recent work on theoretical foundations for typed object-oriented languages: higher-order polymorphism, simple recursive types, subtyping, and a useful partial type inference algorithm.

This is a tutorial introduction to Pict, with examples and exercises.

## Copying

Pict is copyright ©1993–1997 by Benjamin C. Pierce and David N. Turner. This program and its documentation are free software; you can redistribute them and/or modify them under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version. Pict is distributed in the hope that it will be useful, but *without any warranty*; without even the implied warranty of *merchantability* or *fitness for a particular purpose*. See the GNU General Public License for more details. You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

This document was further edited by Matej Košík [kosik@fiit.stuba.sk](mailto:kosik@fiit.stuba.sk).

# Contents

<b>1</b>	<b>Processes and Channels</b>	<b>6</b>
1.1	Simple Processes . . . . .	7
1.2	Channels and Communication . . . . .	8
1.3	Values and Patterns . . . . .	9
1.4	Static Typing vs. Dynamic Typing . . . . .	11
1.5	Types . . . . .	11
1.6	Channel Creation . . . . .	12
1.7	Replication . . . . .	15
1.8	Primitive Booleans . . . . .	17
1.9	Running the Compiler . . . . .	18
<b>2</b>	<b>Core Language Semantics</b>	<b>19</b>
2.1	Overview . . . . .	19
2.2	Syntax . . . . .	21
2.2.1	Notational Conventions . . . . .	21
2.2.2	Concrete Syntax . . . . .	21
2.3	Scoping . . . . .	23
2.4	Type Reconstruction . . . . .	24
2.5	Operational Semantics . . . . .	25
2.5.1	Structural Congruence . . . . .	25
2.5.2	Substitution and Matching . . . . .	26
2.5.3	Reduction . . . . .	27
<b>3</b>	<b>Subtyping</b>	<b>29</b>
3.1	Input and Output Channels . . . . .	29
3.2	Subsumption . . . . .	31

3.3	Responsive Output Channels . . . . .	32
<b>4</b>	<b>Simple Derived Forms</b>	<b>34</b>
4.1	Primitive Values . . . . .	34
4.1.1	Symbolic Identifiers . . . . .	34
4.1.2	Numbers . . . . .	35
4.1.3	Characters and Strings . . . . .	37
4.2	Derived Forms for Declarations . . . . .	38
4.2.1	Declaration Sequences . . . . .	38
4.2.2	run Declarations . . . . .	39
4.3	Parallel Composition . . . . .	39
4.4	Larger Programs . . . . .	40
4.4.1	Importing Other Files . . . . .	40
4.4.2	Separate Compilation . . . . .	40
4.4.3	Library Modules . . . . .	41
<b>5</b>	<b>Toward a Programming Language</b>	<b>42</b>
5.1	Complex Values . . . . .	42
5.1.1	“Continuation-Passing” Translation . . . . .	42
5.1.2	Value Declarations . . . . .	45
5.1.3	Application . . . . .	45
5.2	Derived Forms for Abstractions . . . . .	46
5.3	Sequencing . . . . .	47
<b>6</b>	<b>Simple Concurrent Objects</b>	<b>49</b>
<b>7</b>	<b>Advanced Language Features</b>	<b>52</b>
7.1	Lists . . . . .	52
7.2	Polymorphism . . . . .	53
7.3	Abstract Types . . . . .	55
7.4	User-defined Type Constructors . . . . .	56
7.5	Recursive Types . . . . .	57
<b>8</b>	<b>Security Concerned Programming</b>	<b>60</b>
8.1	Introduction . . . . .	60
8.2	Related Work . . . . .	61

8.3	The Pict Programming Language . . . . .	61
8.4	Refactorization of the Original Pict Library . . . . .	63
8.5	Powerbox . . . . .	64
8.6	Experiments in the Kernel Space . . . . .	66
8.7	Conclusion and Future Work . . . . .	69
<b>A</b>	<b>Solutions to Selected Exercises</b>	<b>70</b>
	<b>Bibliography</b>	<b>70</b>

# Acknowledgments

The Pict language design and implementation are joint work with David N. Turner.

Robin Milner's past and present work on programming languages, concurrency, and the  $\pi$ -calculus in particular is strongly in the background of this project; conversations with Robin have contributed specific insights too numerous to list. The idea of basing a programming language design on the  $\pi$ -calculus was planted by Bob Harper and developed into a research project in the summer of 1992 in discussions of concurrent object-oriented programming languages with the Edinburgh ML Club. From Davide Sangiorgi, we learned about the higher-order  $\pi$ -calculus and the many ways of encoding  $\lambda$ -calculi in the  $\pi$ -calculus; we also did a lot of thinking together about static type systems for the  $\pi$ -calculus [PS93; PS97]. Didier Rémy worked with Pierce on the original PIC compiler [PRT93] (on which an early version of the present Pict compiler was based) and joined in many discussions about the integration of processes and functions. Uwe Nestmann's research on proof techniques for compilations between concurrent calculi [NP96] sharpened our ideas about the formal foundations of Pict. Martin Steffen helped study the formal foundations of the core subtyping algorithm [PS96]. Dilip Sequeira contributed both code and ideas to an early implementation of type inference and record type checking. Kevin Millikin and Philip Wadler gave us helpful comments on the formal definition.

Conversations with Luca Cardelli, Georges Gonthier, Cliff Jones, Oscar Nierstrasz, and John Reppy have deepened our understanding of the  $\pi$ -calculus and concurrent programming languages.

This document began as notes for a series of lectures given at the Laboratory for Foundations of Computer Science, University of Edinburgh, during May and June, 1993. The language and notes were refined for another course at the *1993 Fränkische OOrientierungstage* sponsored by the University of Erlangen-Nürnberg, and again for a winter postgraduate course at the LFCS in 1994. Early drafts were written at INRIA-Roquencourt, with partial support from Esprit Basic Research Actions TYPES and CONFER. Work has continued at the University of Edinburgh, at the Computer Laboratory, University of Cambridge, with support from CONFER and from the British Science and Engineering Research Council, and at Indiana University.

# Chapter 1

## Processes and Channels

The  $\pi$ -calculus of Milner, Parrow, and Walker [MPW92] bears many similarities to the  $\lambda$ -calculus developed by Church and his students in the 1920's and 30's [Chu41]. Though its origins predate computer science itself, the  $\lambda$ -calculus has come to be regarded as a *canonical* calculus capturing the notion of sequential computation in a clean, mathematically tractable way. Many of the fundamental issues of sequential programming languages can be studied by considering them in the more abstract setting of the  $\lambda$ -calculus. Conversely, the  $\lambda$ -calculus has influenced the design of numerous programming languages, from Landin's ISWIM [Lan66] and McCarthy's LISP [McC78] to modern languages such as ML, Scheme, and Haskell.

The  $\pi$ -calculus represents a synthesis and generalization of many years of work on process calculi such as CCS [Mil80; Mil89, etc.]. In the concurrency community, the  $\pi$ -calculus and similar calculi are widely studied, and a substantial body of theoretical work has accrued. More important for our present purposes, though more difficult to quantify, is the observation that the  $\pi$ -calculus is a more *computationally complete* model of real-world concurrent programs than previous formal theories of concurrency. For example, in pure CCS there is no notion of “value”: the entities passed along communication channels are just signals, carrying no additional information. This is fine for studying the basic concepts of concurrency, but as soon as we want to write a program, we find that we need various primitive structures such as integers and booleans that are not present in the pure calculus. These structures can be added, yielding a somewhat more complex system that nevertheless remains theoretically tractable [Mil89]. But value-passing CCS lacks another fundamental property: the ability to perform higher-order programming. For example the fundamental operation of constructing process networks by connecting processes and channels cannot be expressed in CCS, with or without values. Such considerations imply that, although there are several programming languages whose communication facilities are based on CCS, we cannot design a complete language using *only* CCS as its formal foundation.

The  $\pi$ -calculus, on the other hand, does directly support both higher-order programming and natural encodings of primitive datatypes. The ability to pass channels themselves as values between processes—the defining characteristic of the  $\pi$ -calculus—turns out to yield sufficient power to construct dynamically evolving communication topologies and to express a broad range of higher-level constructs. Basic algebraic datatypes like numbers, queues, and trees can be encoded as processes, using techniques reminiscent of Church's encodings in the  $\lambda$ -calculus. Indeed, the  $\lambda$ -calculus itself can be encoded fairly straightforwardly by considering  $\beta$ -reduction as a kind of communication [Mil90]. Thus, the step from the  $\pi$ -calculus to a high-level notation suitable for general-purpose

concurrent programming should be of the same order of magnitude as the step from  $\lambda$ -calculus to early dialects of LISP.

The  $\pi$ -calculus in its present form is not the final word on foundational calculi for concurrency. Milner himself is now considering much more refined systems [Mil92; Mil95], and discussion continues in the concurrency community as to what should constitute a general theory of concurrency. Nevertheless, we've reached a good point to begin experimenting. If Lisp (or, if you prefer, ML, Scheme, or Haskell) is a language based directly on the  $\lambda$ -calculus, then what could a language based directly on the  $\pi$ -calculus look like? The programming language Pict is our attempt to learn the answer.

This chapter offers an informal introduction to a fragment of the Pict language closely related to the pure  $\pi$ -calculus.<sup>1</sup> Chapter 2 develops a more precise treatment of the operational semantics of an even smaller fragment, called the *core language*. Chapter 3 introduces some refinements in the type system sketched in Chapter 1, principally the idea of subtyping. Chapter 4 reintroduces some convenient syntactic forms that were dropped between Chapters 1 and 2 and shows how they can be understood via simple translations into the core. Chapter 5 adds some more complex translation rules yielding convenient high-level syntactic constructs such as function application. Chapter 6 develops an extended example, showing how reference cell objects can be programmed in Pict.

The full Pict language offers a number of features not discussed in this brief tutorial. See the *Pict Language Definition* [PT97b] for a formal description of the entire language.

## 1.1 Simple Processes

The  $\pi$ -calculus is a notation for describing concurrent computations as systems of communicating agents. The basic unit of computation is a *process*.

The simplest process, written `()`, has no observable behavior<sup>2</sup>. To make this process expression into a complete Pict program—albeit not a very useful one—we prefix it with the keyword `run`<sup>3</sup>:

```
run ()
```

Two or more processes may be executed in parallel by separating them with bars and enclosing them in parentheses.

```
run ((() | () | ()))
```

---

<sup>1</sup>Readers familiar with the theoretical literature will notice that the language presented here is not precisely the original formulation of the  $\pi$ -calculus. The primary differences are: (1) like the systems of Honda and Tokoro [HT91] and Boudol [Bou92], output in this fragment is asynchronous: the sender cannot tell when it has actually occurred; (2) channels are typed; (3) the polyadic  $\pi$ -calculus is slightly generalized to allow the communication not only of tuples of channels, but of tuples of tuples, etc; and (4) for technical convenience, booleans and process definitions are included in the core language. There are also many differences in concrete syntax.

<sup>2</sup>The `()` process in Pict has analogous meaning as the `0` process in the  $\pi$ -calculus. In Pict, the `0` symbol has the usual meaning, i.e. it represents the zero integer.

<sup>3</sup>The necessity to use additional `run` keyword might seem at first arbitrary but it has a good reason. Consider the Pict compiler as if it operated in two modes. The first mode (and the default) enables you to conveniently write functional as well as procedural programs. The second mode enables you to describe behavior via the  $\pi$ -calculus-like constructs. The `run` keyword here was used here to switch from the default (functional/procedural) mode to the  $\pi$ -calculus-like mode. How to inline functional/procedural mode into the  $\pi$ -calculus-like mode is described later in this document—via “complex values”.

The simplest thing that a process can actually *do* is to cause an observable event in the outside world. For example, a process of the form `print!"abc"` causes the string `abc` to be printed on the standard output stream.

```
run ( print!"peering"  
      | print!"absorbing"  
      | print!"translating"  
      )
```

```
peering  
absorbing  
translating
```

(In this document, lines of output from the running program are left-justified to distinguish them from program text.)

Those who had previously learned basics of the  $\pi$ -calculus but also intuitively might be puzzled. How can we so boldly claim that parallel composition of those three actions will lead to that particular output? The answer is that this is caused by the way how Pict runtime works. The language designers did not feel that something is lost if in these cases the system will behave deterministically.

## 1.2 Channels and Communication

Besides processes, the other significant entities in the  $\pi$ -calculus are *channels* (also called *names* in the  $\pi$ -calculus literature). A channel is a port over which one process may send messages to another.

Suppose `x` is a channel. Then the expression `x! []` denotes a *sender* process that transmits the *signal* `[]` along the channel `x`. This transmission is completed when another process is ready to accept a value along the same channel `x`; such a *receiver* process has the form `x? [] = e`, where `e` is a process expression indicating what to do after the signal has been received. When the two are placed in parallel, the communication can take place:<sup>4</sup>

```
run ( x? [] = print!"Got it!"  
      | x! []  
      )
```

```
Got it!
```

Note that the nearly identical program

```
run x? [] = (print!"Got it!" | x! [])
```

prints nothing, since the output `x! []` is now inside the body of the input `x? [] = ...` and so cannot take place until after the input has succeeded. In general, the body `e` of an input expression `x?y = e` remains completely inert until after a communication along `x` has occurred.

---

<sup>4</sup>This is not quite a complete program: if you try to run it, the compiler will complain about `x` being an unbound name.

Sending a signal from one process to another is a useful way of achieving synchronization between concurrent threads of activity. For example, a signal sent from process  $e$  to process  $f$  might carry the information that  $e$  has finished modifying some data structure that it shares with  $f$ , or that it has completed some action that  $f$  requested. This sort of synchronization is a ubiquitous feature of Pict programs.

It is often useful for two processes to exchange some value when they synchronize. In particular, the a *channel* can be passed from the sender to the receiver as part of the act of communication.

```
run ( x?z = print!"Got it!"
     | x!y
     )
```

Got it!

As we shall see, this ability is a primary source of the  $\pi$ -calculus's (and Pict's) expressive power.

The name  $x$  plays the same role in the expressions  $x!y$  and  $x?z = e$ : it is the "location" where the two processes meet and exchange information. But the roles of  $y$  and  $z$  are different:  $y$  can be thought of as an *argument* of the message, whereas  $z$  is a *bound variable* that, at the moment of communication, is replaced by the received value  $y$ .

Of course, the receiving process may do other things with channels it obtains by communication. In particular, it may use them for communication, either as the channel on which to send or receive a value...

```
run ( x!y
     | x?z = z!u
     | y?w = print!"Got it!"
     )
```

Got it!

... or as the value that it sends along some other channel:

```
run ( x!y
     | x?z = a!z
     | a?w = print!"Got it!"
     )
```

Got it!

### 1.3 Values and Patterns

More generally, each communication step involves the atomic transmission of a single *value* from sender to receiver. In the fragment of Pict we are considering at the moment, values may be constructed in just two ways:

1. a channel is a value;

- if  $v_1$  through  $v_n$  are values, then the tuple  $[v_1 \dots v_n]$  is a value. Note that we write tuples with just whitespace between adjacent elements.

For example, if  $x$  and  $y$  are channels, then  $x$ ,  $[x\ y]$ , and  $[x\ [[y\ x]]\ y\ []\ x]$  are values. The signal  $[]$  is just the empty tuple of values. Character strings like "Got it" are also values, but for the moment they may only appear as arguments to `print`.

The general form of a sender process is  $x!v$ , where  $x$  is a channel and  $v$  is a value. Symmetrically, a receiver process has the form  $x?p = e$ , where  $p$  is a *pattern* built according to the following rules:

- a variable is a pattern;
- if  $p_1$  through  $p_n$  are patterns binding distinct sets of variables, then  $[p_1 \dots p_n]$  is a pattern.

For example,  $[]$ ,  $[x\ y\ z]$ , and  $[[]\ x\ y\ [[]]]$  are patterns.

When a sender  $x!v$  communicates with a receiver  $x?p = e$ , the value  $v$  is matched against the pattern  $p$  to yield a set of bindings for the variables in  $p$ . For example, matching the value  $[a\ b\ []]$  against the pattern  $[m\ n\ o]$  yields the bindings  $\{m \mapsto a, n \mapsto b, o \mapsto []\}$ . More precisely:

- any value  $v$  matches a variable pattern  $x$ , yielding the singleton binding  $\{x \mapsto v\}$ ;
- if  $p$  has the form  $[p_1 \dots p_n]$  and  $v$  has the form  $[v_1 \dots v_n]$  and, for each  $i$ , the value  $v_i$  matches the subpattern  $p_i$  yielding the binding  $\Delta_i$ , then  $v$  matches the whole pattern  $p$ , yielding the set of bindings  $\Delta_1 \cup \dots \cup \Delta_n$ .

Two additional forms of patterns are often useful in programming:

- A *wildcard pattern*, written `_` (underscore), matches any value but yields no bindings.
- A *layered pattern*  $x@p$  matches whatever  $p$  matches, yielding all the bindings that  $p$  yields, and also yields a binding of  $x$  to the whole value matched by  $p$ .

The following example is slightly artificial, but it demonstrates usage of the layered pattern.

```
run ( new ch: ^[String String]
      ( ch!["Hello" "world"]
        | ch?a@[b c] = ( print!b
                        | print!c
                      ) )
    ) )
```

```
Hello
world
```

In the above program:

- `a` becomes bound to `["Hello" "world"]`
- `b` becomes bound to `"Hello"`
- and `c` becomes bound to `"world"`

After the `@` character might be any kind of reasonable pattern. However, before this character must be a single variable name<sup>5</sup>.

---

<sup>5</sup>This silently follows from “metavariable conventions” introduced in [PT97b].

## 1.4 Static Typing vs. Dynamic Typing

Languages such as C, C++ and Java created a bad reputation to static type systems. There, in order to do something useful, you had to constantly “cast” values from one type to another. We recommend you to ignore those examples. You should rather look at good examples such as typing in ML or Ocaml. Typing belongs to computer science.

The question whether ideal language should be statically or dynamically typed has no definite answer. Each approach has its advantages and disadvantages.

Static typing restricts expressivity of the original language. How much, this depends on the character of the employed typing system. From the personal experience we can say that automatic type inference embedded into the Pict compiler increases comfort<sup>6</sup> of the programmer significantly.

Typical problem with dynamically typed programming languages is “bad match” (as in Erlang) or “does not understand” (as in Smalltalk). These are very simple errors and with good debugging tools (as it is the case of Smalltalk) they can be tracked down easily. No Smalltalker misses types to avoid these kind of problems.

Statical type system would not make much sense for Erlang<sup>7</sup> or Smalltalk (if it at all were possible to design one for this language). However, in Pict, it makes perfect sense. We have channels and the typing system tells us what kind of values can “flow” through each particular channel. The “selective receive” (in terms of Erlang community) is the primary operation in Pict and we have it for free. In those cases when we want to block certain process on multiple channels, we can use the `Choice`. It is provided within its standard library.

## 1.5 Types

Given an arbitrary pattern `p` and value `v`, it is perfectly possible that `p` does *not* match `v`. As in the  $\lambda$ -calculus, it is convenient to impose a *type discipline* on the use of channels that permits such situations to be detected. Like ML, Pict is *statically typed*, in the sense that this check occurs at compile time; program that passes the compiler cannot encounter a situation at run time where the value provided by a sender does not match the pattern of a possible receiver.

In Pict, there are several (finite number of) primitive types such as:

type	example value
Bool	true
Int	50
Char	'c'
String	"The Pict programming language"

In your programs, you can introduce variables bound to whatever legal value with the `val` declaration. For example:

---

<sup>6</sup>The present type inference algorithm can be further enhanced to correctly reveal types in more situations.

<sup>7</sup>This is not rejection of Erlang, we only make few notes concerning one (and not the only) aspect of concerning these languages. Erlang advocates could here point out many features of Erlang that are missing in Pict. Such as taking advantage of multiple cores; support for distributed programming and many other things. Unless these features are implemented for Pict, everyone has full right to doubt whether they can be implemented at all.

```

val i = 50
val b = true
val message = "Hello"

```

You can use also tuples. The following<sup>8</sup> rule tells how to infer type of the tuple when you know the types of the particular elements from which the tuple is composed:

*If the values  $v_1$  through  $v_n$  have the types  $T_1$  through  $T_n$ , then the tuple value  $[v_1 \dots v_n]$  has the type  $[T_1 \dots T_n]$ .*

Some examples:

example value	its type
[]	[]
[5]	[Int]
[5 10 15]	[Int Int Int]
[5 "foo"]	[Int String]
[[48 8 41] [17 6 46]]	[[Int Int Int] [Int Int Int]]

In particular, the value [] has the type []. Although the value [] and its type are written using the same sequence of characters, there is no danger of confusion: it will always be clear whether a given expression should be regarded as a value or a type.

It is often convenient to make up abbreviations for commonly used types. The declaration

```
type X = T
```

makes the symbol X stand for the type T in what follows.

The type system of Pict is quite rich, incorporating features such as higher-order polymorphism, subtyping, records, recursive types, and a type-inference mechanism. But there is no need to discuss all of these at once (some of them will be introduced one by one in later chapters; the rest can be found in the *Pict Language Definition*).

## 1.6 Channel Creation

New channel names are introduced by the **new** declaration:

```
new x: ^T
```

It creates a fresh channel, different from any other channel, and makes the name **x** refer to this channel in the scope of the declaration. The values sent and received on **x** must have the type **T**.

Concerning type of channels, the following simplified rule can be stated:

*Each channel may be used to send and receive values of exactly one type. If this type is **T**, then the type of the channel is  $\hat{T}$ , read “channel of **T**” or “channel carrying **T**.”*

This rule embodies an essential restriction: a Pict channel may not be used in one place to send an integer and in another to send a channel. That is, compilation of the following program:

<sup>8</sup>Take it as a preliminary information. Complete rules are given in the *Pict Language Definition* [PT97b].

```

new x:^Int
new y:^Int
run x!y

```

will fail with the following complaint:

```

Expected type does not match actual since ^Int is not
a subtype of Int since ^Int and Int do not match.

```

Values of type `Int` (such as 10) can be transmitted only through channels of type `^Int` (such as `x` and `y` as defined above). Values of type `^Int` (such as `x` and `y`) can only be communicated over channels of type `^^Int`. The following program is thus well typed:

```

new x:^Int
new y:^Int
new z:^^Int
run ( x!50
    | y!100
    | z!x
    | z!y
    )

```

Although (as with any type system simple enough to be tractable) this restriction excludes some reasonable and perhaps even useful programs, relaxing it would mean using a dynamic flow analysis for typechecking programs. If each channel is used, throughout its lifetime, to carry values of the same shape, then well-typedness becomes a static property and a wide range of well-understood techniques can be applied in the engineering of the type system.

The keywords `new`, `run`, and `type` all introduce *declaration clauses*; we will see a few other kinds of declaration clauses later on. A Pict program is simply a sequence of declaration clauses, where the scope of variables introduced in each clause is all of the following clauses.

```

new x:^Int
run ( x?a = () | x!5)

```

Above, the channel `x` has type `^Int`. Only values of type `Int` (such as 5) can be sent to it and received from it.

Similarly, if you want to transmit couples of integers, you can create a different channel:

```

new y:^[Int Int]

```

you can communicate over that channel for example as follows:

```

run ( y?[a b] = ()
    | y![5 6]
    )

```

The above program fragment is well typed because `[5 6]` has type `[Int Int]` and it is sent over channel whose type is `^[Int Int]`.

It is also possible to prefix any process expression with a sequence  $d$  of private declarations:  $(d\ e)$  is a process expression in which the scope of the variables introduced by  $d$  is just  $e$ . So the previous example could just as well have been written

```
run ( new x:^[]
      ( x! []
        | x? [] = ()
      ) )
```

or even:

```
run ( new x:^[]
      run x! []
      x? [] = ()
    )
```

Two **new** declarations binding the same channel name may be used in different parts of the same program:

```
run ( ( new x:^[]
        ( x! []
          | x? [] = ()
        ) )
      | ( new x:^[]
          ( x! []
            | x? [] = ()
          ) )
    ) ) )
```

There is no possibility of confusion between the two communications on  $x$ : the output in the first line can only synchronize with the input in the same line, and likewise for the second line. Two declarations of the same name may even have overlapping scopes. In this case, the inner binding hides the outer one. For example, this program does *not* print “Got it”:

```
run ( new x:^[]
      ( (new x:^[] x! [])
        | x? [] = print!"Got it?"
      ) )
```

It is often useful to communicate a channel outside of the scope of the **new** declaration where it was created. For example, the program fragment

```
run ( new x:^[]
      ( z!x
        | x? [] = print!"Continuing..."
      ) )
```

creates a channel  $x$ , sends it along  $z$  to some colleague in the outside world, and then waits for a response along  $x$  before continuing. In the  $\pi$ -calculus literature, the possibility of names escaping beyond the scope of their declarations is called *scope extrusion*.

## 1.7 Replication

Concurrent programs typically involve infinite behavior. For example, they often contain “server” processes that repeatedly perform some actions in response to requests from other processes. Such a server is never “finished”: when a request has been handled, it simply waits for another, deals with it, waits again, and so on.

Infinite behaviors are provided in the  $\pi$ -calculus by so called *replication*. In Pict we can introduce them with `def` construct, like:

```
def hailer [] =  
  print!"Hail!"
```

Above, the `def` construct causes two things:

- it defines a new channel named `hailer` of type `[]`
- it defines a behavior of a process that will be forked whenever you send (well-typed) value to the `hail` channel.

Such a code is equivalent to:

```
new hailer : ^[]  
  
run ( hailer?[] = print!"Hail!"  
    | hailer?[] = print!"Hail!"  
    | hailer?[] = print!"Hail!"  
    | ...  
    )
```

if we could write infinite number of processes. This equivalent meaning of `def` (expressed in more primitive terms) also shows us that the bound variable `hailer` introduced by the `def` construct is visible:

- in the body of replicated processes
- after the `def` construct

So a client can talk to `hailer` as follows:

```
def hailer [] =  
  print!"Hail!"  
  
run ( hailer![]  
    | hailer![]  
    | hailer![]  
    | hailer![]  
    | hailer![]  
    )
```

```
Hail!  
Hail!  
Hail!  
Hail!  
Hail!
```

Since the `hailer` variable is bound also in the body of the replicated processes, we can write recursive processes. This is very frequent case. With a little bit of syntactic sugar and some additional helper function we can write re-write the `hailer` process so that it will print a message as many times as we tell it:

```
def hailer n:Int =  
  if (== n 0) then  
    ()  
  else  
    ( print!"Hail!"  
      | hailer!(dec n)  
      )  
  
run hailer!5
```

```
Hail!  
Hail!  
Hail!  
Hail!  
Hail!
```

We can also define finite number of mutually recursive processes with `def ... and ... and ...` construct. Below we show two such mutually recursive processes:

```
def ping n:Int =  
  if (== n 0) then  
    ()  
  else  
    ( print!"ping"  
      | pong!(dec n)  
      )  
and pong n:Int =  
  if (== n 0) then  
    ()  
  else  
    ( print!"pong"  
      | ping!(dec n)  
      )  
  
run ping!6
```

```
ping  
pong  
ping  
pong  
ping  
pong
```

## 1.8 Primitive Booleans

There are multiple possibilities how booleans (values `true` and `false`) could be introduced to the language. Now they are part of Core Pict<sup>9</sup>. These two constants:

- `true`
- `false`

are visible in every scope of the program. We have a simple `if b then P else Q` construct. It describes a process that behaves as `P` if `b` is `true` and it behaves as `Q` if `b` is `false`.

The following two examples show how we can use the `if` construct in a process context.

```
run if true then
  print!"branch1"
else
  print!"branch2"

run if false then
  print!"branch3"
else
  print!"branch4"
```

```
branch1
branch4
```

Equivalently, we could compose two `if` processes in parallel:

```
run ( if true then
  print!"branch1"
  else
  print!"branch2"
| if false then
  print!"branch3"
  else
  print!"branch4"
)
```

```
branch1
branch4
```

A different flavor of `if` construct is provided by a syntactic sugar for functional contexts. That other kind of `if` behaves as an expression that returns value of one of its branches according to the value of the boolean condition.

---

<sup>9</sup>The other possibility would be to introduce them as part of the syntactic-sugar layer

## 1.9 Running the Compiler

To execute a Pict program, first place it in a file `prog.pi` and then compile and execute this file with the Pict compiler:

```
pict prog.pi
```

Running the compiler this way produces a file `a.out`, which is executed automatically the first time, and which you can thereafter execute directly, without recompiling the original program:

```
./a.out
```

Alternatively, you can give the Pict compiler an explicit output file

```
pict -o prog prog.pi
```

and run the resulting executable yourself:

```
./prog
```

Command line options are documented in `pict` manual page. You can find more details therein.

## Chapter 2

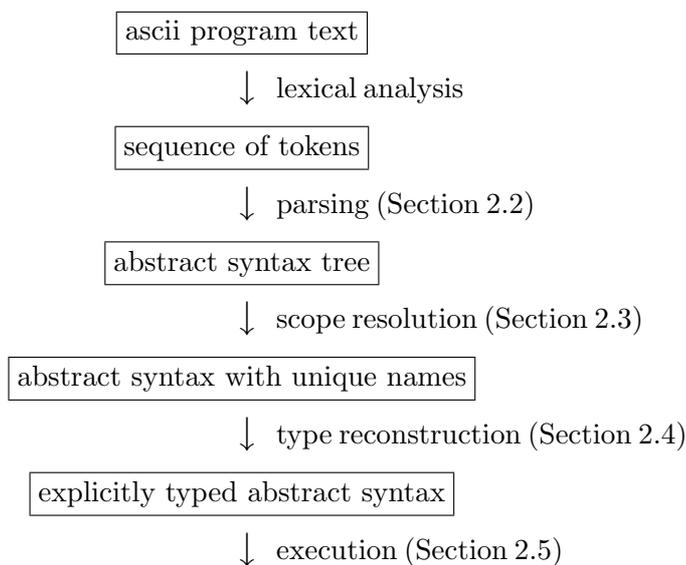
# Core Language Semantics

Defining a full-scale language rigorously requires some work. Our approach is to present first a very small, untyped sublanguage, called *Core Pict*, and explain the meaning of other language constructs by means of translations into this core. Chapter 1 introduced (a slight superset of) the core informally, relying on english descriptions, examples, and exercises to convey a practical understanding of the language. We now develop it precisely. The present chapter deals with the syntax and operational semantics of the core; a formal presentation of the type system can be found in the Pict Language Definition [PT97b].

Most of the material in this chapter is repeated, in less formal terms, elsewhere. It can be skimmed on a first reading.

### 2.1 Overview

The complete definition of core Pict semantics recapitulates the structure of a compiler:



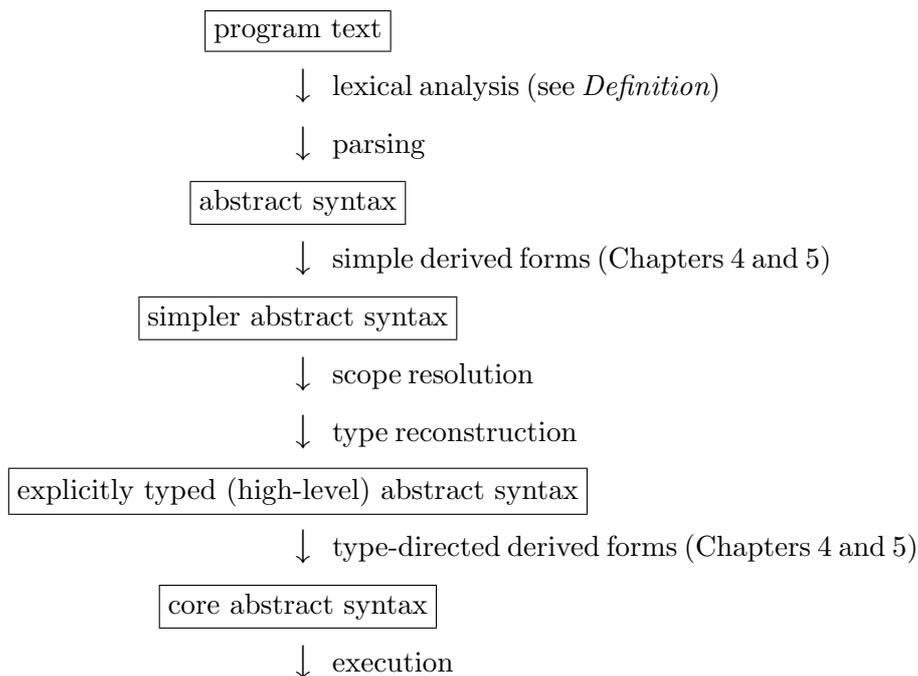
First, the source text of the program is transformed by a lexical analyzer into a sequence of *tokens*: identifiers, keywords, and so on. If the program is well formed, this sequence of tokens matches

the *concrete syntax* grammar of Pict in exactly one way, allowing the parser to transform it into an *abstract syntax* tree. The scope resolution phase identifies free occurrences of variables with the points at which they are bound, simplifying the description of later phases. The data structure generated during scope resolution is a new abstract syntax tree where each variable binder is named differently from any other binder elsewhere in the program. A type reconstruction pass walks over this tree, checking that it satisfies the typing rules of the language and filling in any type annotations that have been omitted by the programmer to yield a new abstract syntax tree with the types of all its subphrases explicitly given. This final abstract syntax tree is then executed<sup>1</sup>.

In an actual implementation of Pict, the execution phase has a number of subphases: optimisation, intermediate code generation, optimisation of the intermediate code, native code generation, linking, and finally execution on a real machine. Fortunately, we do not need to treat all of these formally: instead, we give a notion of *reduction of abstract syntax trees* that defines precisely, but at a fairly high-level, the set of allowable behaviors for every program. The compiler's job is to produce a machine-code program whose actual behavior exactly mimics some behavior in this set.

Though the execution phase operates on *typed* programs, the operational semantics is written in such a way that type annotations in a program do not affect its behavior. This justifies giving a formal treatment of the operational semantics while remaining informal (in this document) about the type system.

The definition of the full Pict language in Chapters 4 and 5 adds two more phases, which perform the translation from high-level to core abstract syntax.



The derived forms are divided into two separate phases. Some of them are so simple that they can be performed directly by the parser, thus eliminating the need to deal with them in the scope resolution or type reconstruction phases. Others (in particular, the CPS-conversion described in

---

<sup>1</sup>Internally, core Pict constructs are translated into the C programming language which in turn is compiled into a native executable.

Chapter 5) require that scopes be resolved and missing type annotations inserted before they can do their work.

## 2.2 Syntax

We now turn to defining the syntax of core Pict.

### 2.2.1 Notational Conventions

For describing syntax, we rely on a meta-syntactic notation similar to the Backus-Naur Form commonly used in language definitions. The possible forms of each production are listed on successive lines. Keywords are set in typewriter font. An expression of the form  $X \dots X$  denotes a list of zero or more occurrences of  $X$ . The expression  $\langle \textit{empty} \rangle$  denotes an empty production.

### 2.2.2 Concrete Syntax

We have said that the top level of a compilation unit (i.e., a file containing Pict source code) is a series of `import` statements followed by a series of declarations. For present purposes, though, it is convenient to consider programs in a much more restricted form: just a single `run` declaration with a process expression in its body.

*Program* = `run Proc` Program

A process expression can have several forms:

<i>Proc</i>	=	<code>Val ! Val</code>	Output atom
		<code>Val ? Abs</code>	Input prefix
		<code>( )</code>	Null process
		<code>( Proc   Proc )</code>	Parallel composition
		<code>( Dec Proc )</code>	Local declaration
		<code>if Val then Proc else Proc</code>	Conditional

Note that the syntax given here is a bit more restrictive than what we saw in Chapter 1: it allows only binary parallel composition and only one local declaration in front of a process. The more permissive forms will be recovered in Chapter 4 as derived forms.

On the other hand, there are also some ways in which this syntax is more permissive than the examples in Chapter 1 suggested. In particular, in the output expression  $v_1!v_2$ , both  $v_1$  and  $v_2$  are allowed to be arbitrary values, not just identifiers. This means that processes like `[]!x` are syntactically legal and must be weeded out during typechecking. (In principle, since the core abstract syntax also includes these forms, we might ask how such a nonsensical expression behaves. The answer will be “it doesn’t,” i.e. none of the evaluation rules allow any interaction between such a process and the rest of the program.)

Input processes are defined in terms of the syntactic class `Abs` of process expressions prefixed by patterns

*Abs* = *Pat* = *Proc* Process abstraction

where a pattern can be a variable, a record of (labeled or unlabeled) patterns, or a layered or wildcard pattern.

*Pat* = *Id RType* Variable pattern  
 [ *Label Pat* ... *Label Pat* ] Record pattern  
 \_ *RType* Wildcard pattern  
*Id RType* @ *Pat* Layered pattern

Every variable, wildcard, and layered pattern includes space for an optional type annotation.

*RType* =  $\langle \text{empty} \rangle$  Omitted type annotation  
 : *Type* Explicit type annotation

In the typechecking rules, operational semantics, and derived forms, we assume that all these annotations have been provided explicitly. It is the job of the *type reconstruction* phase to fill in omitted annotations, if possible, before these later phases need them.

The syntactic class of values—things that can be sent along channels—includes constants, variables, records, and field projections. For technical reasons, we syntactically restrict the “left hand side” of a field projection to be either a variable or another field projection by introducing a separate syntactic class of *paths*.

*Val* = *Const* Constant  
*Path* Path  
 [ *Label Val* ... *Label Val* ] Record

A path is a variable followed by zero or more field projections.

*Path* = *Id* *C* Variable  
*Path* . *Id* *C* Record field projection

Constant values include the booleans **true** and **false** and literal strings that we have already seen, plus integers and character constants (these last are discussed in Chapter 4):

*Const* = *String* String constant  
*Char* Character constant  
*Int* Integer constant  
**true** Boolean constant  
**false** Boolean constant

Type expressions correspond to the possible shapes of values: channels, constants

*Type* =  $\sim$  *Type* Input/output channel  
**Bool** Boolean type  
**String** String type  
**Int** Integer type  
**Char** Character type  
 [ *Label Type* ... *Label Type* ] Record type

Later in the text (Section 3.1) we also introduce further refinements of the input/output type.

Of course, user-defined typed (type aliases) are also permitted.

A declaration can be a channel definition, a process definition, or a type definition.

<i>Dec</i>	=	<code>new Id : Type</code>	Channel creation
		<code>def Id<sub>1</sub> Abs<sub>1</sub> and ... and Id<sub>n</sub> Abs<sub>n</sub></code>	Recursive definition ( $n \geq 1$ )
		<code>type Id = Type</code>	Type abbreviation

Note that the bodies of definitions are abstractions—the same syntactic form as the bodies of input processes. This means, in particular, that a process definition can actually use an arbitrary pattern, not just a tuple. Also, the type annotations on the pattern in a process definition are syntactically optional. However, because of the way the type reconstruction propagates information, it turns out that they can almost never be recovered during the type reconstruction phase, and must therefore be provided explicitly.

A group of definitions introduced by `def` and separated by `and` can be mutually recursive—i.e., each of the bodies can refer to any of the defined names.

The `run` declarations of Chapter 1 (except for the single outermost one) are missing from this syntax. They will be reintroduced in Chapter 4 as derived forms. Also, type declarations are treated informally in this document. In the *Language Definition* they are actually presented as derived forms, but their formalization depends on language features that fall outside the scope of this tutorial.

Finally, as we have seen, a field label (in a record pattern, record value, or record type) can be either an identifier followed by an `=` sign, or else empty.

<i>Label</i>	=	<code>&lt;empty&gt;</code>	Anonymous field
		<code>Id =</code>	Labeled field

## 2.3 Scoping

Since Pict has several syntactic categories, the scoping of variables is not quite so simple as in the  $\lambda$ -calculus (or the pure  $\pi$ -calculus, for that matter), where one can say, “The  $x$  in  $\lambda x. e$  is a binding occurrence whose scope is  $e$ ,” and be done with it. Instead, we need to identify the syntactic categories that can create new name bindings—*Dec* and *Pat*—and, wherever one of these categories is used in the abstract syntax, specify the scope of the names that it creates.

**2.3.1 Definition:** The rules for name creation and scoping are as follows:

- In a process expression of the form `(d e)`, the scope of names created by `d` is `e`.
- A declaration clause of the form `new x:T` creates the name `x`.
- A type declaration clause of the form `type X = T` creates the type name `X`.
- In an abstraction of the form `p = e`, the scope of names created by `p` is `e`.
- A variable pattern of the form `x:T` creates the name `x`.

- A record pattern of the form  $[l_1 p_1 \dots l_n p_n]$  creates all of the names created by the subpatterns  $p_1$  through  $p_n$ . The sets of names created by the subpatterns must be disjoint.
- A layered pattern of the form  $x:T@p$  creates all of the names created by the subpattern  $p$ , plus  $x$  (which must be distinct from the names created by  $p$ ).

where  $d$  represents *Dec* construct;  $e$  represents *Proc* construct;  $p$  represents *Pat* construct. These are non-terminals in the grammar introduced in the previous section. By  $x$  is meant name of ordinary variable. By  $X$  and  $T$  is meant name of a variable that holds type (as its value).

In the remaining phases of the core language definition (the operational semantics, typechecking, and translation rules), we assume that the names of bound variables are always different, silently renaming variables if necessary to achieve this form. A formal definition of this translation process can be found in the *Pict Language Definition*.

We use the notation  $FV(e)$  for the set of variables appearing free in an expression  $e$ . Formally,  $FV(e)$  can be thought of as the set of (binding occurrences of) variables appearing in  $e$  whose corresponding binding occurrences are not within  $e$ . Similarly,  $BV(d)$  stands for the set of variables bound (created) by the declaration  $d$ .

## 2.4 Type Reconstruction

In this tutorial, the typechecking rules and type reconstruction algorithm are treated informally (they appear in full glory in the *Definition*). Fortunately, the type system corresponding to the part of the language we are covering here is quite straightforward.

Type reconstruction, too, is based on a few simple ideas:

- During reconstruction, type information can be propagated either from subphrases to the larger phrases that contain them (“synthesis”) or from larger phrases to their constituents (“checking”). Synthesis mode is used to *calculate* the type of a phrase, while checking mode is used to *verify* that a phrase has some expected type.
- In synthesis mode, type annotations are generally required; in checking mode, an omitted annotation can often be filled in by using the expected type at that point.
- All bound variables must have their types determined at the point of binding, either by an explicit type annotation or by appearing in a checking context. This means that later occurrences always have a known type, whether they are encountered in a checking or a synthesis context.

In particular, the patterns in process definitions are analyzed in a synthesis context (since we do not know yet what shapes of values they are expected to match), while patterns in input expressions are analyzed in a checking context, since the type of the channel from which the input is taken determines the shape of the values that the pattern must match. The parameters to a process definition must therefore be explicitly typed, while the parameters of an input expression need not be.

## 2.5 Operational Semantics

The operational semantics of Pict programs is presented in two steps. First, we define a *structural congruence* relation  $e_1 \equiv e_2$  on process expressions; this relation captures the fact that, for example, the order of the branches in a parallel composition has no effect whatsoever on its behavior. Next, we define a *reduction relation*  $e_1 \rightarrow e_2$  on process expressions, specifying how processes evolve by means of communication.

For present purposes, a program is just the keyword `run` followed by a process expression  $e$ . Its behavior is the behavior of  $e$ .

### 2.5.1 Structural Congruence

Structural congruence plays an important technical role in simplifying the statement of the reduction relation. For example, we intend that the processes  $(x!v \mid x?y = e)$  and  $(x?y = e \mid x!v)$  both reduce to  $\{x \mapsto v\}e$ . By making these two structurally congruent, we can get away with writing the reduction rule just for the first case and adding a general stipulation that, when  $e$  contains some possibility of communication, any expression structurally congruent to  $e$  has the same possible behavior.

The first three structural congruence rules state that parallel composition is commutative

$$(e_1 \mid e_2) \equiv (e_2 \mid e_1) \quad (\text{STR-COMM})$$

and associative

$$((e_1 \mid e_2) \mid e_3) \equiv (e_1 \mid (e_2 \mid e_3)) \quad (\text{STR-ASSOC})$$

and that the null process  $()$  is an identity for parallel composition.

$$(e \mid ()) \equiv e \quad (\text{STR-NULL})$$

The next rule, often called the rule of *scope extrusion* in the  $\pi$ -calculus literature, plays a crucial role in the treatment of channels.

$$\frac{BV(d) \cap FV(e_2) = \emptyset}{((d \ e_1) \mid e_2) \equiv (d \ (e_1 \mid e_2))} \quad (\text{STR-EXTRUDE})$$

Informally, this rule says that a declaration can always be moved toward the root of the abstract syntax tree (“always,” because the precondition is always satisfied when the rule is read from left to right<sup>2</sup>). For example, the process expression

$$((\text{new } y: \text{ } [] \ x!y) \mid x?z = z! [])$$

may be transformed to:

---

<sup>2</sup>Indeed, since we have already performed scope resolution when the structural congruence rules are invoked, we are justified in assuming that the precondition *always* holds. We adopt this view formally, but retain the precondition as a reminder.

$(\text{new } y : \hat{\ } \ [] \ (x!y \mid x?z = z! \ []))$

It is precisely this rule that allows the new channel  $y$  to be communicated outside of its original scope.

Similarly, two adjacent **new** declarations can always be swapped (since, by the conventions introduced in Section 2.3, they must introduce channels with different names), and a **new** may be swapped with a **def** (or even a **def** with a **def**) as long as the body of the **def** does not use the name defined by the **new**. These cases are captured by a general rule for exchanging adjacent declarations:

$$\frac{BV(d_1) \cap FV(d_2) = \emptyset \quad BV(d_2) \cap FV(d_1) = \emptyset}{(d_1 \ (d_2 \ e)) \equiv (d_2 \ (d_1 \ e))} \quad (\text{STR-SWAPDEC})$$

Finally, two adjacent **def** clauses may be merged into one:

$$\frac{(\{x_1, \dots, x_m\} \cap \{x_{m+1}, \dots, x_n\} = \emptyset \quad (FV(a_1) \cup \dots \cup FV(a_m)) \cap \{x_{m+1}, \dots, x_n\} = \emptyset)}{(\text{def } x_1 a_1 \ \dots \ \text{and } x_m a_m \ (\text{def } x_{m+1} a_{m+1} \ \dots \ \text{and } x_n a_n \ e)) \equiv (\text{def } x_1 a_1 \ \dots \ \text{and } x_m a_m \ \text{and } x_{m+1} a_{m+1} \ \dots \ \text{and } x_n a_n \ e)} \quad (\text{STR-COALESCE})$$

Reading this rule in the other direction, it says that a single compound **def**...**and** clause may be split into two, so long as the bodies of definitions in the first part do not depend on names defined by the second part.

## 2.5.2 Substitution and Matching

To define precisely what happens when two processes communicate, we need some notation for matching values against patterns.

A *substitution* is a finite map from variables to values. The empty substitution is written  $\{\}$ . A substitution mapping just the variable  $x$  to the value  $v$  is written  $\{x \mapsto v\}$ . If  $\sigma_1$  and  $\sigma_2$  are substitutions with disjoint domains, then  $\sigma_1 \cup \sigma_2$  is a substitution that combines the effects of  $\sigma_1$  and  $\sigma_2$ . A substitution  $\sigma$  can be extended to a function from values to values by applying  $\sigma$  to variables that fall in its domain and leaving the rest of the value unchanged. For example, applying the substitution  $\sigma = \{x \mapsto a\} \cup \{y \mapsto []\}$  to the value  $[b \ [x] \ x \ y]$ , written  $\sigma([b \ [x] \ x \ y])$ , yields the value  $[b \ [a] \ a \ []]$ .

In order to support planned extensions of the language, we want to maintain the syntactic property that the head of every path is a variable, not an explicitly constructed record. To maintain this property during reduction, we need to make sure that, when we substitute a record value for a variable during a communication step, any projection expressions with this variable as their head are reduced at the same time by projecting out the appropriate field:

$$\{x \mapsto [\dots l_i v_i \dots]\} x.l_i = v_i$$

This refined version of substitution is the one used in the rules that follow.

Whenever some value  $v$  is matched against some pattern  $p$  it is usually interesting to know what kind of substitution should be applied to the process that follows (after the  $=$  terminal) the pattern. Abstractions can occur only in two context:

- input operation (the `?` construct)
- recursive definition (the `def` construct)

The function *getsubst* that computes this substitution as follows:

**2.5.2.1 Definition:** When a value  $v$  is successfully matched by a pattern  $p$ , the result is a substitution  $\{p \mapsto v\}$ , defined as follows

$$\begin{aligned}
\text{getsubst}(x:T, v) &\stackrel{\text{def}}{=} \{x \mapsto v\} \\
\text{getsubst}(\_ :T, v) &\stackrel{\text{def}}{=} \emptyset \\
\text{getsubst}(x:T@p, v) &\stackrel{\text{def}}{=} \{x \mapsto v\} \cup \text{getsubst}(p, v) \\
\text{getsubst}([l_1fp_1 \dots l_nfp_n], [l_1fv_1 \dots l_nfv_n]) &\stackrel{\text{def}}{=} \text{getsubst}(fp_1, fv_1) \cup \dots \cup \text{getsubst}(fp_n, fv_n)
\end{aligned}$$

If  $v$  and  $p$  do not have the same structure, then *getsubst*( $p, v$ ) is undefined.

The first clause is the base case of the definition: it states that a variable pattern matches any value and yields a substitution mapping that variable to the whole value. The remaining clauses traverse the structure of the pattern and the value in parallel, comparing their outermost constructors for consistency and then invoking *getsubst* recursively to match corresponding substructures.

### 2.5.3 Reduction

The reduction relation  $e \rightarrow e'$  may be read as “The process  $e$  can evolve to the process  $e'$ .” That is, the semantics is nondeterministic, specifying only what *can* happen as the evaluation of a program proceeds, not what *must* happen. Any particular execution of a Pict program will follow just one of the possible paths.

The most basic rule of reduction is the one specifying what happens when an input prefix meets an output atom:

$$\frac{\{p \mapsto v\} \text{ defined}}{(x!v \mid x?p = e) \rightarrow \{p \mapsto v\}(e)} \quad (\text{RED-COMM})$$

The rule for instantiating definitions is similar, except that the “input side” is some clause of a definition, not a simple input prefix.

$$\frac{\{p_i \mapsto v\} \text{ defined}}{(\text{def } x_1p_1 = e_1 \dots \text{ and } x_np_n = e_n \quad (x_i!v \mid e)) \rightarrow (\text{def } x_1p_1 = e_1 \dots \text{ and } x_np_n = e_n \quad (\{p_i \mapsto v\}(e_i) \mid e))} \quad (\text{RED-DEF})$$

A conditional expression reduces in one step to either its `then` part or its `else` part, depending on the value of the boolean guard:

$$\text{if true then } e_1 \text{ else } e_2 \rightarrow e_1 \quad (\text{RED-IF-T})$$

$$\text{if false then } e_1 \text{ else } e_2 \rightarrow e_2 \quad (\text{RED-IF-F})$$

The next two rules allow reduction to occur under declarations and parallel composition:

$$\frac{e_1 \rightarrow e_2}{(d \ e_1) \rightarrow (d \ e_2)} \quad (\text{RED-DEC})$$

$$\frac{e_1 \rightarrow e_3}{(e_1 \mid e_2) \rightarrow (e_3 \mid e_2)} \quad (\text{RED-PRL})$$

The body of an input expression, on the other hand, *cannot* participate in reductions until after the input has been discharged.

The structural congruence relation captures the distributed nature of reduction. Any two sub-processes at the “top level” of a process expression may be brought into proximity by structural manipulations and allowed to interact.

$$\frac{e_1 \equiv e_2 \rightarrow e_3 \equiv e_4}{e_1 \rightarrow e_4} \quad (\text{RED-STR})$$

In closing, it is worth mentioning that we have done here only a part of the work involved in giving a really complete definition of the semantics of Pict. For one thing, we have not talked about the fact that any reasonable implementation of this operational semantics must schedule processes for execution *fairly*. A short discussion of fairness in Pict appears in [PT97a].

For another thing, we have only specified the behavior of *closed programs*, with no connections to the outside world. Of course, real Pict programs do have external connections (such as the `print` channel and, using the libraries provided with the compiler, other external facilities such as file systems and X servers). Peter Sewell has shown how the simple semantics presented here can be extended to model the externally observable behavior of processes [Sew96].

## Chapter 3

# Subtyping

We have already introduced the essentials of Pict’s type system: values are assigned types describing their structure; in particular, the types of channels specify the types of the values that they carry. In this chapter, we discuss an important refinement of this basic type system.

### 3.1 Input and Output Channels

Channel types serve a useful role in ensuring that all parts of a program use a given channel in a consistent way, eliminating the possibility of pattern matching failure at run time. Of course, pattern matching failure is just one kind of bad behavior that programs may exhibit; especially in concurrent programs, the minefield of possible programming mistakes is vast: there may be unintended deadlocks, race conditions, and protocol violations of all kinds. Ultimately, one might wish for static analysis tools capable of detecting all of these errors—perhaps even capable of verifying that a program meets an arbitrary specification (expressed, for example, in some modal logic). But the technology required to do this is still a good distance away.

Fortunately, there are some simple ways in which our simple channel types can be enriched so as to capture useful properties of programs while remaining within the bounds established by current typechecker technology. One of the most important of these in Pict is based on the distinction between input and output capabilities for channels.

In practice, it is relatively rare for a channel to be used for both input and output in the same region of a program; the usual case is that some parts of a program use a given channel only for reading while in other regions it is used only for writing.

Let us consider a simple producer-consumer program. We create a channel `ch` and two processes:

- producer
- consumer

The producer generates some integers and sends them into channel `ch`. The consumer blocks itself receiving on channel `ch` and if some value arrives, it will print it on the screen. The whole program might look in the following way:

```

def consumer channel:^Int =
  channel?i = ( printi!i
                | consumer!channel
                )

def producer [channel:^Int i:Int] =
  if (== i 0)
  then ()
  else ( channel!i
         | producer![channel (dec i)]
         )

new ch:^Int
run consumer!ch
run producer![ch 4]

```

4  
3  
2  
1

We see that the producer as well as the consumer obtain value of `^Int` type. It means, that they can use this value both, for sending integers to the designated channel as well as receiving integers from the designated channel. This is violation of POLA. Therefore we rewrite the whole program as follows:

```

def consumer channel:?Int =
  channel?i = ( printi!i
                | consumer!channel
                )

def producer [channel:!Int i:Int] =
  if (== i 0)
  then ()
  else ( channel!i
         | producer![channel (dec i)]
         )

new ch:^Int
run consumer!ch
run producer![ch 4]

```

4  
3  
2  
1

The types of the `consumer` and the `producer` channels changed (from `^Int` to `?Int` or `!Int` respectively).

The consumer thus has the permission to receive integers from the designated channel.

The producer has the permission to send integers to the designated channel.

The final code fragment:

```

new ch: ^Int
run consumer!ch
run producer![ch 4]

```

might perhaps seem suspicious. It seems that we use types in inconsistent way:

- we have a value `ch` of type `^Int`
- we use this value as if it had type `?Int` (we pass it over the `consumer` channel)
- we use this value as if it had type `!Int` (we pass it over the `producer` channel)

Fortunately, these glitches have completely reasonable explanation. Pict programming language supports subtyping. In this case, the following holds:

- `^Int < ?Int` (`^Int` is subtype of `?Int`)
- `^Int < !Int` (`^Int` is subtype of `!Int`)

The compiler promoted value of type `^Int` to `?Int` in one case and it promoted the same value from type `^Int` to type `!Int` in the second case. You can see, that promotion in both cases led to dropping of some permissions. Therefore it does not introduce insecurity.

The following section introduces general rules related to subtyping.

## 3.2 Subsumption

The well-typedness of the programs in the previous section depends in several places on the observation that a value of type `S` can sometimes be passed along a channel of type `^T` (or `!T`) even though `S` and `T` are not identical types. For example, in the previous section we passed value of type `^Int` through channel of type `^!Int`<sup>1</sup>.

For any type `U`, the type `^U` is a subtype of the types `!U` and `?U`, but `!U` and `?U` themselves are unrelated: neither is a subtype of the other.

By the same reasoning, the tuple type `[^U]` is a subtype of `[!U]` and `[?U]`; in the example, we incur no risk of failure by passing the value `ch` of type `^Int` along the channel `consumer`, which nominally requires an argument of type `?Int`. More generally, if each `Si` is a subtype of the corresponding `Ti`, then the record type `[l1S1 . . . lnSn]` is a subtype of `[l1T1 . . . lnTn]`. That is, the labels in the two record types must match in each position (either both must be blank or both must be the same explicit label) and the types of corresponding fields must be in the subtype relation. More generally yet, we allow the smaller type to have some extra fields on the right, so that `[l1S1 . . . lmSm . . . lnSn]` is a subtype of `[l1T1 . . . lmTm]`. Note, though, that we do not allow fields to be reordered or extra fields to be added in the middle of the common ones. (In the absence of this restriction, separate compilation is much more difficult.)

---

<sup>1</sup>Technically, this is not true. But it is a very good approximation. The following section introduces so called *responsive types* where you will see the difference between *output type* (prefixed by `!`) and *responsive type* (prefixed by `/`).

Finally, when writing programs involving subtyping, it is occasionally convenient to have some type that is a supertype of every other type—a maximal element of the subtype relation, functioning as a kind of “don’t care” type. We call this type `Top` in Pict.

Pict also provides a type called `Bot` which is subtype of all existing types.

**3.2.1 Exercise [Recommended]:** *We have seen that the tuple constructor is monotone in the subtype relation: if each  $S_i < T_i$ , then  $[S_1, \dots, S_n] < [T_1, \dots, T_n]$ . What about channel types?*

1. *What relation should hold between  $S$  and  $T$  in order for  $!S$  to be a subtype of  $!T$ ? For example, would it be correct to allow  $!S < !T$  if  $S$  and  $T$  are not identical but  $S < T$ ? What about when  $T < S$ ?*
2. *What relation should hold between  $S$  and  $T$  in order for  $?S$  to be a subtype of  $?T$ ?*
3. *What relation should hold between  $S$  and  $T$  in order for  $\wedge S$  to be a subtype of  $\wedge T$ ?*

### 3.3 Responsive Output Channels

In fact, we take the refinement of channel types a step further in Pict, identifying one case of communication that is so common as to deserve special treatment.

Channels created by `def` clauses (as opposed to `new` clauses) have two important properties: (1) There is always a receiver (the body of the `def`) available, and (2) all communications are received by the *same* receiver—the same `def` body. These properties turn out to be extremely useful in compilation: sending on a channel that is known to have been created by a `def` can be implemented essentially as a branch, whereas general communication requires quite a bit more checking of channel status bits, manipulation of queues, etc. But, since channels can be passed along other channels before being used for communication, it is not possible to tell statically which outputs will actually be communicating with definitions when they are executed, just by looking at the program’s structure. Instead, we use the type system.

A channel created by a `def` has a type of a special form, written `/T` and pronounced “*responsive* channel carrying elements of type  $T$ .” When such channels are sent along other channels, the type system tracks this fact. For example, from the type of `x` in the following program, the output process `a!false` knows that it is communicating with a `def`.

```
new x :  $\wedge$ [/Bool]
def d b:Bool =
  if b then
    print!"True"
  else
    print!"False"
run x![d]
run x?[a] = a!false
```

False

The subtyping relation allows `/T < !T`, so that responsive channels can be used, if desired, as ordinary channels, forgetting their special properties.

```

new y : ^[!Bool]
run y![d]
run y?[a] = a!false

```

False

Allowing subtyping in the other direction, promoting ordinary channels to responsive ones, would clearly be unsound. But there are some cases where we may want to use an ordinary output channel where a responsive channel is required. For example, the Pict standard library provides a primitive called `pr` that (like `print`) prints a string on the standard output stream, but then sends a signal along another channel to indicate that it has completed its work. The type of `pr` is `/[String /[]]`—that is, `pr` is itself a responsive channel, and it expects its two arguments to be a string and another responsive channel (carrying data-less signals). We can invoke `pr` by sending it a channel `d` (that has type `/[]`) created by a definition...

```

def d [] = print!"done"
run pr!["pr... " d]

```

pr... done

...but if we try to send it an ordinary channel, from which we can later read to obtain the signal sent by `pr`, we get a typechecking failure:

```

new c : ^[]
run pr!["pr... " c]
run c?[] = print!"done"

```

```

example.pi:2.20:
Expected type does not match actual
since ^[] is not a subtype of /[] since ^[] and /[] do not match

```

In order to allow programs like the last one (which is arguably easier to read, since the “continuation” of the call to `pr` appears after the call itself), Pict provides an operator (`rchan ...`) than can be used to coerce an ordinary output channel into a responsive channel:

```

new c : ^[]
run pr!["pr... " (rchan c)]
run c?[] = print!"done"

```

pr... done

(In fact, `rchan` can be defined in Pict, using the high-level syntactic forms introduced in Chapter 5.)

## Chapter 4

# Simple Derived Forms

We now introduce several convenient higher-level programming constructs in the form of predefined channels and simple syntactic sugar. By the end of the chapter, we will have built up enough syntactic forms to do some more serious programming.

### 4.1 Primitive Values

Like most high-level programming languages, Pict provides special syntax and special treatment in the compiler for a few built-in types, including booleans, characters, strings, and numbers. We sketch some of the facilities available for manipulating these basic types and show how they can be understood in terms of encodings in the core language.

#### 4.1.1 Symbolic Identifiers

It is convenient to use standard symbolic names for operations like addition. Pict supports *symbolic identifiers* consisting of strings of symbols from the set `~`, `*`, `%`, `\`, `+`, `-`, `<`, `>`, `=`, `&`, `|`, `@`, `$`, `'`, and `,` (excluding those strings that are reserved words).

```
run ( new +++: ^ []
      ( +++! []
        | +++?*-- = print!"two together"
        | ( new +=&&%: ^ []
              ( +=&&% ? ** = +++!**
                | +=&&%! []
              )
            )
      ) ) )
```

two together

### 4.1.2 Numbers

Like the booleans, numbers and arithmetic operations can, in principle, be represented in core Pict via a “Church-like” encoding. Such encodings are nice theoretical exercises, illustrating the power of the core language, but they are too inefficient to be useful in practice. As usual in functional languages based on the  $\lambda$ -calculus, we need to add some primitive values to the  $\pi$ -calculus. However, we want to maintain the *illusion* that these primitive values are actually implemented as processes: a program that computes with numbers should not be able to tell whether they are represented as Church numerals or as machine integers. More to the point: a human being *reasoning about* a program that computes with numbers should not have to think about their concrete representation.

A number  $n$  can be thought of as a process “located at” a channel  $n$ ; it can be interrogated over  $n$  to gain information about its value. Higher-level arithmetic operations like  $+$  can be implemented as processes that interrogate their arguments and construct a new number as result. But if a program using numbers always manipulates them by means of processes like  $+$  rather than interrogating them directly<sup>1</sup>, then we can simply think of the channel  $n$  as *being* the number  $n$ . This done, we can introduce a special set of “numeric channels” and an efficient reimplementaion of  $+$ , and no one will be able to tell the difference.

Integers in Pict behave as any kind of legal value. They can be passed around through channels (with proper types) in a similar way as any other kind of value. Look at the following simple fragment of Pict code:

```
run ( new r:~Int
      ( r!30
        | r?a = ...
      ) )
```

There we define a new channel  $r$  and use it for passing an integer value 30 from one process to another.

Various useful processes that provide particular integer operations such as addition, subtraction, multiplication and so on are described in [Ko07] in section “Int: Integers”. In this section we avoid using syntactic sugar forms (they are described in Chapter 5) so dealing with integers will not seem elegant to you.

Let us look at addition. According to [Ko07] we see that we should use channel  $+$ . Its type is  $/[\text{Int Int } /\text{Int}]$ . It is a responsive channel. If you have two integers 10 and 20 and you want to compute their sum, you can do that as follows:

```
run +![10 20 printi]
```

30

The  $+$  process takes given two integers 10 and 20; computes their sum 30; and sends the result to the given channel (in this case `printi`). What happens with the result depends on the definition of the `printi` process.

In case when you do not want print the result of  $10 + 20$  immediately but you want to compute  $(10 + 20) * 30$  and print that result, you can do that as follows:

---

<sup>1</sup>That is, our programs does not try to send messages directly to channels leading to integers and does not try to receive messages directly from channels that lead to such integers.

```

run ( new r:^Int
      ( +![10 20 (rchan r)]
        | r?sum = *![sum 30 printi]
      )
    )

```

900

Note that we have used auxiliary channel `r` of type `^Int` to pass the result of the `+` operation to the input of the `*` operation. The result of the `*` is sent to the `printi` channel that prints it on the standard output.

Note also that we had to use the `rchan` channel to “cast”<sup>2</sup> the value `r` from `!Int` to the expected `/Int` type.

This way we (albeit not very conveniently) compute arbitrary complex arithmetic expressions. Chapter 5 introduced syntactic sugar that enables us to write such programs easier. The following example we print the value of this  $(10 + 20) * (30 + 40)$  arithmetic expression:

```

run ( new r1:^Int
      new r2:^Int
      ( +![10 20 (rchan r1)]
        | +![30 40 (rchan r2)]
        | r1?sum1 = r2?sum2 = *![sum1 sum2 printi]
      )
    )

```

2100

Above we

- we ask the `+` process to compute the sum of 10 and 20 and send the result to the `r1` channel
- we ask the `+` process to compute the sum of 30 and 40 and send the result to the `r2` channel
- we block until we receive a value from the `r1` channel and bind it to the `sum1` variable
- we block until we receive a value from the `r2` channel and bind it to the `sum2` variable
- we ask the `*` process to compute the product of `sum1` and `sum2` values and send the result to the `printi` channel (that will print it out)

We could equally well do the reception in the reversed order. In this case it would not affect the overall behavior of the program.

A more interesting example, which also illustrates the use of the `if` construct, is the factorial function. (Don’t be alarmed by the length of this example! The derived forms introduced in Chapter 5 make such programs much more concise.)

---

<sup>2</sup>More precisely, the `rchan` operation provides similar functionality as the `cast` (type coercion) operation albeit no type coercion actually happens. That would be against principles of the Pict programming language where programmers cannot arbitrarily cast one value from one type to another. The `rchan` operation only creates a “wrapper” of desired type `/Int` that acts as a forwarder to the capability of actual type `!Int`.

```

run ( def fact [n:Int r:!Int] =
  ( new br:^Bool
    ( {- calculate n=0 -}
      ==![n 0 (rchan br)]
    | {- is n=0? -}
      br?b =
        if b then
          {- yes: return 1 as result -}
          r!1
        else
          {- no... -}
          ( new nr:^Int
            ( {- subtract one from n -}
              -![n 1 (rchan nr)]
            | nr?nMinus1 =
              {- make a recursive call to compute fact(n-1) -}
              (new fr:^Int
                ( fact![nMinus1 fr]
                | fr?f =
                  {- multiply n by fact(n-1) and send the
                    result on the original result channel r -}
                  *![f n (rchan r)]
              ) ) ) ) )
    fact![5 printi]
  )

```

120

**4.1.2.1 Exercise [Recommended]:** *Use the numeric and boolean primitives to implement a simple algorithm for calculating the Fibonacci function:*

$$\begin{aligned}
 fib(0) &= 1 \\
 fib(1) &= 1 \\
 fib(n) &= fib(n-1) + fib(n-2) \quad \text{when } n > 1
 \end{aligned}$$

### 4.1.3 Characters and Strings

Besides booleans and integers, Pict provides the built-in types `Char` and `String`, with special syntax for values of these types. Character constants are written by enclosing a single character in single-quotes, as in `'a'`. Similarly, string constants are written by enclosing a sequence of zero or more characters in double-quotes. In both strings and character constants, special characters like double- and single-quote are written using the following *escape sequences*:

<code>\'</code>	single quote
<code>\"</code>	double quote
<code>\\</code>	backslash
<code>\n</code>	newline (ascii 13)
<code>\t</code>	tab (ascii 8)

The escape sequence `\ddd` (where `d` denotes a decimal digit) denotes the character with ascii code `ddd`.

The standard prelude provides a number of operations for manipulating characters and strings. For example, the operation `$$` converts an integer to its string representation, and the operation `+$` concatenates strings. Using these, the predefined operation `printi` can be expressed in terms of `print`:

```
def printi i:Int =
  ( new r1 : ^String
    ( $$![i (rchan r1)]
      | r1?s = print!s
    ) )

run printi!5
```

5

Indeed, `print` itself is defined in terms of the lower-level predefined channel `pr`: we just ignore the completion signal returned by `pr`.

```
def print s:String =
  ( def r[] = ()
    pr![s r]
  )
```

It is convenient to make the type `Char` a subtype of the type `Int`, so that any character value can implicitly be regarded as the integer representing its ASCII code. For example:

```
run printi!'a'
```

97

## 4.2 Derived Forms for Declarations

In this section, we extend the syntactic category of declarations with a number of handy constructs. Readers familiar with Standard ML [MTH90] will recognize our debt to its designers here.

### 4.2.1 Declaration Sequences

First, as in the examples in Chapter 1, we avoid proliferation of parentheses in a sequence of declarations like

```
(new x:A (new y:B (new z:C ...)))
```

by allowing a *Proc* to be preceded by a sequence of declaration clauses within a single set of parentheses:

```
(new x:A new y:B new z:C ...)
```

Processes taking advantage of this more liberal syntax are translated back into the core language during compilation, simply by reinserting the dropped parentheses. In the *Definition*, this is captured by the following rule:

$$(d_1 \dots d_n e) \Rightarrow (d_1 \dots (d_n e)) \quad (\text{TR-DECSEQ})$$

(We will not show all of the translation rules for the derived forms we discuss; see the *Definition* for full details.)

### 4.2.2 run Declarations

Again, as we saw in Chapter 1, in sequences of declarations it is often convenient to start some process running in parallel with the evaluation of the remainder of the declaration. For example, one often wants to define some server process and then start a single copy running. We introduce the declaration keyword `run` for this purpose. (The keyword `fork` might have been more intuitive, but we find programs easier to read if all declaration keywords are three characters long!) Once a declaration sequence has been translated into a nested collection of individual declarations, this `run` declaration may be translated into a simple parallel composition. For example, the process

```
( run print!"twittering"  
  run print!"rising"  
  print!"overhead passing"  
)
```

is equivalent to the following core-language process:

```
( print!"twittering"  
  | ( print!"rising"  
      | print!"overhead passing"  
      ) )
```

## 4.3 Parallel Composition

Strictly speaking, the core language syntax only allows two processes to be composed in parallel. We generalize this to arbitrary numbers ( $\geq 2$ ) of processes composed in parallel, translating “high level” process expressions like

```
(x![] | y![] | z![] | w![])
```

into:

```
(x![] | (y![] | (z![] | w![])))
```

## 4.4 Larger Programs

Section 1.9 already explained how to compile and run separate Pict programs whose source code is located within a single file.

The Pict compiler provides a simple facility for breaking up large programs into parts, storing the parts in separate files.

- If you change one module, you only have to recompile that particular module and to create executable you link it with the other already compiled modules.
- These modules usually solve one problem. It is often advantageous to study particular parts separately.
- This kind of modules enable us to re-use well written components and refactor common functionality into dedicated and shared modules.
- Even though it might at first sound strange, modules (as well as other Pict properties) enable us to solve some of interesting security problems related with execution of untrusted code. Chapter 8 is devoted to this topic.

### 4.4.1 Importing Other Files

A Pict program is organized as several files, each containing a declaration sequence preceded by a number of `import` clauses. Each `import` clause has the form

```
import "name"
```

where `name` is an absolute or relative path name (not including the suffix `.pi`). If a relative path name is used, both the current directory and a central directory of Pict library files are searched<sup>3</sup>. Imports may be nested; that is, imported files may themselves begin with `import` clauses.

Semantically, the first occurrence of an `import` clause for a given file name means exactly the same as if the whole imported file had been included at the point where the `import` appears. Subsequent occurrences of the same `import` clause have no effect.

If the Pict compiler finds a cycle in the `import` relationship, it will reject to compile such program.

### 4.4.2 Separate Compilation

Before a file `f.pi` can be imported (using `import "f"`) by other files, it must be processed by the Pict compiler, yielding a file `f.px`. The command-line flag `-set sep` tells the compiler that the file it is processing is a separately compiled module, not the main program.

For example, if the file `f.pi` contains

```
def x[] = print! "...endlessly rocking\n"
```

and the file `g.pi` contains

---

<sup>3</sup>It is located in the `/usr/lib/tamed-pict` directory.

```
import "f"  
run x! []
```

then we compile `g.pi` into an executable file `g` in two steps:

```
pict -set sep -o f.px f.pi  
pict -o g g.pi
```

### 4.4.3 Library Modules

The Pict distribution includes a library of precompiled modules implementing a variety of data structures, interfaces to operating system facilities, and other services. These are described in full in the *Tamed Pict Library* [Ko07]. A few of the most basic libraries, including all of the facilities described so far, are imported by default<sup>4</sup>; however, most must be imported explicitly if they are needed.

---

<sup>4</sup>See the Chapter *Standard Prelude* of [Ko07].

## Chapter 5

# Toward a Programming Language

Chapters 1 to 4 introduced the syntax and operational semantics of the Pict core language, as well as some of the simpler derived forms. We now proceed with (and, by the end of the chapter, essentially conclude) the task of defining a high-level programming notation based on these foundations.

### 5.1 Complex Values

So far, all the value expressions we have encountered have been built up in an extremely simple way, using just variables, channels (including built-in channels such as `pr` and `+`) and records of values. These *simple values* are important because they are the entities that can be passed along channels and participate in pattern matching.

#### 5.1.1 “Continuation-Passing” Translation

In programming, it is very common to write an expression that computes a simple value and immediately sends it along some channel. For example, the process `(new n:T x!n)` creates a fresh channel `n` and sends it off along `x`. More interestingly,

```
run (def f[x:Int res:/Int] = +![x x res]
    y!f)
```

creates a local definition `f` and sends its “request channel” along `y`.

An alternative syntax for such expressions, which can often make them easier to understand, puts the whole value-expression *inside* the output: `x!(new n:T n)`. In general, it is useful to allow such expressions in any position where a simple value is expected. Formally, we extend the syntactic category of values with declaration values of the form `(d v)`. We use the term *complex value* for an expression in the extended syntax that does not fall within the core language.

When we write `x!(new n:T n)`, we do not mean to send the *expression* `(new n:T n)` along `x`. A complex value is always evaluated “strictly” to yield a simple value, which is substituted for the complex expression.

In introducing complex values, we have taken a fairly serious step: we must now define the meaning of a complex value occurring in any position where simple values were formerly allowed. For

example, the nested expression  $x![23 \text{ (new } x:A \text{ } x) \text{ (new } y:B \text{ } y)]$  must be interpreted as a core language expression that creates two new channels, packages them into a simple tuple along with the integer 23 and sends the result along  $x$ .

To interpret arbitrary complex values, we introduce a general “continuation-passing” translation. Given a complex value  $v$  and a continuation channel  $c$ , the expression  $\llbracket v \rightarrow c \rrbracket$  will denote a process that evaluates  $v$  and sends the resulting simple value along  $c$ . We then introduce translation rules for process expressions containing complex values. For example, the rule

$$\frac{\Gamma \vdash v_1 \in T_1 \quad \Gamma \vdash v_2 \in T_2}{\llbracket v_1!v_2 \rrbracket = (\text{def } c_1 \text{ } x_1:T_1 = (\text{def } c_2 \text{ } x_2:T_2 = x_1!x_2 \llbracket v_2 \rightarrow c_2 \rrbracket) \llbracket v_1 \rightarrow c_1 \rrbracket)} \quad (\text{CPS-OUT})$$

translates an output  $v_1!v_2$  into a process expression that first allocates a fresh continuation channel  $c$ , next evaluates  $v_1$ , waits for its result to be sent along  $c$ , and then evaluates  $v_2$ , sending the result directly along the channel  $x$  that resulted from the evaluation of  $v_1$ . The notation  $\llbracket v \rightarrow c \rrbracket$  stands, intuitively, for a process that calculates  $v$  and sends the resulting simple value along  $c$ . The premises  $\Gamma \vdash v_1 \in T_1$  and  $\Gamma \vdash v_2 \in T_2$  are calls to the typechecker, which calculate the types  $T_1$  and  $T_2$  that should appear in the annotations of the two `defs`. (In general,  $\Gamma \vdash v \in T$  is read “Under the assumptions  $\Gamma$  (which give the types of all the free variables in  $v$ ), the value  $v$  has type  $T$ .”)

Input processes and conditionals containing complex values are translated similarly:

$$\frac{\Gamma \vdash v \in T}{\llbracket v?p=e \rrbracket = (\text{def } c \text{ } x:T = x?p=\llbracket e \rrbracket \llbracket v \rightarrow c \rrbracket)} \quad (\text{CPS-IN})$$

$$\llbracket \text{if } v \text{ then } e_1 \text{ else } e_2 \rrbracket = (\text{def } c \text{ } x:\text{Bool} = \text{if } x \text{ then } \llbracket e_1 \rrbracket \text{ else } \llbracket e_2 \rrbracket \llbracket v \rightarrow c \rrbracket) \quad (\text{CPS-IF})$$

The continuation passing translation itself is defined by induction on the syntax of value expressions. Variables and constants are easy to handle, since they already represent simple values:

$$\llbracket x \rightarrow c \rrbracket = c!x \quad (\text{CPS-VAR})$$

$$\llbracket k \rightarrow c \rrbracket = c!k \quad (\text{CPS-CONST})$$

The translation of values prefixed by declarations is also straightforward:

$$\llbracket (\text{new } x:T \text{ } v) \rightarrow c \rrbracket = (\text{new } x:T \llbracket v \rightarrow c \rrbracket) \quad (\text{CPS-NEWV})$$

$$\llbracket (\text{def } x_1a_1 \dots \text{ and } x_na_n \text{ } v) \rightarrow c \rrbracket = (\text{def } x_1\llbracket a_1 \rrbracket \dots \text{ and } x_n\llbracket a_n \rrbracket \llbracket v \rightarrow c \rrbracket) \quad (\text{CPS-DEFV})$$

$$\llbracket (\text{run } e \text{ } v) \rightarrow c \rrbracket = (\llbracket e \rrbracket \mid \llbracket v \rightarrow c \rrbracket) \quad (\text{CPS-RUNV})$$

The only slightly complex case is records, where each of the fields is CPS-converted separately, from left to right. To keep the size of the rule manageable, we rely on an auxiliary definition that recurses down the list of fields, transforming each one; we give it the list of fields and set a marker

(written  $|$ ) at the beginning of the list of fields to indicate that none of them have been translated yet.

$$\frac{\Gamma \vdash [l_1fv_1 \dots l_nfv_n] \in [l_1FT_1 \dots l_nFT_n]}{\llbracket [l_1fv_1 \dots l_nfv_n] \rightarrow c \rrbracket = \llbracket | l_1fv_1 \dots l_nfv_n \rightarrow c \rrbracket} \quad (\text{CPS-RECORD})$$

There are two rules for CPS-converting lists of field values, depending on whether the marker  $|$  is at the end of the list of fields, indicating that all fields have already been converted, or a value field, which must be CPS-converted. In the latter case the field following the marker is CPS-converted by creating a definition for a new channel  $c_i$  whose body is formed by recursively CPS-converting the remaining fields and starting a process to send the value for the  $i$ th field along  $c_i$ .

$$\llbracket l_1fv_1 \dots l_nfv_n | \rightarrow c \rrbracket = c! [l_1fv_1 \dots l_nfv_n] \quad (\text{CPSF-DONE})$$

$$\frac{\Gamma \vdash v_i \in T_i}{\llbracket l_1fv_1 \dots | l_iv_i \dots l_nfv_n \rightarrow c \rrbracket = (\text{def } c_i \ x_i:T_i = \llbracket l_1fv_1 \dots l_ix_i | l_{i+1}fv_{i+1} \dots l_nfv_n \rightarrow c \rrbracket \llbracket v_i \rightarrow c_i \rrbracket)} \quad (\text{CPSF-VALUE})$$

Using the mechanism of CPS-conversion, it is also easy to allow conditional value expressions and CPS-convert them to processes:

$$\llbracket \text{if:T } v \text{ then } v_1 \text{ else } v_2 \rightarrow c \rrbracket = (\text{def } d \ x:\text{Bool} = \text{if } x \text{ then } \llbracket v_1 \rightarrow c \rrbracket \text{ else } \llbracket v_2 \rightarrow c \rrbracket \llbracket v \rightarrow d \rrbracket) \quad (\text{CPS-IFV})$$

Conditional value expressions can be embedded to Pict programs wherever a value of an appropriate type is allowed. For example in the fragment:

```
(pr "AAA");
```

we could replace literal string "AAA" with an expression that evaluates to string. This means that instead of "AAA" we can put there conditional expression `if true the "AAA" else "BBB"`

```
(pr if true the "AAA" else "BBB")
```

which is equivalent to the original simpler statement because this conditional expression `if true the "AAA" else "BBB"` always returns value "AAA". All expressions evaluate to some value. All values have some type. The conditional expression is not exceptional in this case. Our `if true the "AAA" else "BBB"` expression has type `String`. The Pict programming language (see rule CPS-IFV) enables us to embed the type constraint to the conditional expression. This is not normally needed because type inference algorithm is able to infer type of the whole expression from types of its branches. But we does not stop us to specify it explicitly:

```
(pr if:String true the "AAA" else "BBB")
```

The typechecker only checks if the explicitly provided type is a supertype of the inferred type.

### 5.1.2 Value Declarations

Since complex value expressions may become long and may involve expensive computations, it is convenient to introduce a new declaration form that evaluates a complex value. For example, `(val x = (new n:T [n n]) e)` binds `x` to the result of evaluating `(new n:T [n n])` and then executes `e`. Formally, `val` declarations are translated into the core language using the same continuation-passing translation as above:

$$\llbracket (\text{val } p = v \ e) \rrbracket = (\text{def } c \ p = \llbracket e \rrbracket \ \llbracket v \rightarrow c \rrbracket) \quad (\text{CPS-VAL})$$

Since the expression on the left of the `=` can be an arbitrary pattern, a `val` declaration can be used to bind several variables at once. For example, `val [x y] = [23 [a]]` binds `x` to `23` and `y` to `[a]`. Note that, when a `val` declaration `(val p = v e)` is translated into the core language, the body `e` appears after an input prefix. This fact implies that `val` declarations are *strict* or *blocking*: the body cannot proceed until the bindings introduced by the `val` have actually been established.

Since declarations can also appear in value expressions, we also need to add a clause to the definition of CPS-conversion for values:

$$\llbracket (\text{val } p = v_1 \ v_2) \rightarrow c \rrbracket = (\text{def } d \ p = \llbracket v_2 \rightarrow c \rrbracket \ \llbracket v_1 \rightarrow d \rrbracket) \quad (\text{CPS-VALV})$$

### 5.1.3 Application

Of course, allowing declarations inside values represents only a minor convenience; the usefulness of this extension would not by itself justify all of the foregoing machinery—the distinction between complex and simple values, etc. But having established the basic pattern of simplifying complex value expressions by transformation rules, we can apply it to a much more useful extension.

In value expressions, we allow the *application* syntax `(v v1 ... vn)`. For example, if we define a `double` function by

```
def double [s:String r:/String] = +$![s s r]
```

(where `+$` is string concatenation), then, in the scope of the declaration, we can write `(double s)` as a value, dropping the explicit result channel `r`. For example,

```
run print!(double "soothe")
```

causes `"soothe``soothe"` to be sent along the built-in channel `print`.

We define the meaning of application values by adding a clause to the definition of the continuation-passing translation:

$$\llbracket (v \ l_1 f v_1 \dots l_n f v_n) \rightarrow c \rrbracket = \llbracket v! [l_1 f v_1 \dots l_n f v_n \ c] \rrbracket \quad (\text{CPS-APP})$$

Operationally, this rule encodes the intuition that the implicit final parameter in an application is the continuation of the function being invoked—the place where the function’s result should be sent in order for the rest of the computation to proceed.

**5.1.3.1 Exercise [Recommended]:** *What core-language program results from applying the translation rules to the following process expression?*

```
x![( * ( + 2 3 ) ( + 4 5 ) )]
```

**5.1.3.2 Exercise [Recommended]:** *Rewrite your solution to Exercise 4.1.2.1 using application syntax.*

We have now conveyed the most complex derived forms in Pict. It remains to discuss a few useful extensions of the syntax of processes, abstractions, and patterns.

## 5.2 Derived Forms for Abstractions

Although Pict’s core language and type system do not distinguish between “real functions” and “processes that act like functions,” it is often useful to write parts of Pict programs in a functional style. This is supported by a small extension to the syntactic class of abstractions, mirroring the ability to omit the names of result parameters in applications (Section 5.1.3). We replace a process definition of the form

```
def f[a1:A1 a2:A2 a3:A3 r:/T] = r!v
```

where the whole body of the definition consists of just an output of some (complex) value on the result channel, by a “function definition” that avoids explicitly giving a name to `r`:

```
def f (a1:A1 a2:A2 a3:A3):T = v
```

To avoid confusion with ordinary definitions (and for symmetry with application), we change the square brackets around the list of arguments to parentheses.

Since anonymous process declarations like `(def x () = e x)` are often useful, we provide a special form of value allowing the useless `x` to be omitted:

$$\backslash a \Rightarrow (\text{def } x \ a \ x) \qquad (\text{TR-ANONABS})$$

Anonymous abstractions are used quite heavily in Pict programs. For example, the standard library provides the operation `for`, which takes two integers `min` and `max`, a channel `f` of type `![Int /[]]`, and a completion channel `done`, and successively sends each integer between `min` and `max` to `f`, waiting each time for `f` to signal back before proceeding with the next. When `f` returns for the last time, `for` signals on `done`. The `for` function is already defined in the `Int` module as follows:

```
def for [min:Int max:Int f:![Int /[]] done:/[]] =
  ( def loop x:Int =
    if (<= x max) then
      ( new c : ^[]
        ( f![x (rchan c)]
          | c?[] = loop!(+ x 1)
        ) )
    ) )
```

```

    else
      done![]
  loop!min
)

```

The most common use of `for` is to pass it an anonymous abstraction for `f`, as in:

```

run ( new done : ^[]
      ( for![] 1 4
          \[x c] = (print!x | c![])
              (rchan done)
            ]
        | done?[] = print!"Done!") )

```

```

1
2
3
4
Done!

```

Another important use of anonymous process abstractions is in fields of records, where they can be thought of as the method bodies of an object:

```

val r = [
  one = \[] = print!"Low hangs the moon"
  two = \[] = print!"0 it is lagging"
]

run (r.one![] | r.two![])

```

```

0 it is lagging
Low hangs the moon

```

The connection with objects is continued in Chapter 6.

### 5.3 Sequencing

One very common form of result is a *continuation signal*, which carries no information but tells the calling process that its request has been satisfied and it is now safe to continue. For example, the standard library includes the output operation `pr`, which signals on its result channel when the output has been accomplished. So a sequence of outputs that are intended to appear in a particular order can be written:

```

run
( val [] = (pr "the ")
  val [] = (pr "musical ")
  val [] = (pr "shuttle")
  ()
)

```

the musical shuttle

(Note that `pr`, unlike `print`, does not append a carriage return to the string that is output.)

The type of `pr` is `![String /[]]`. The result channel has thus type `/[]`.

The idiom “invoke an operation, wait for a signal as a result, and continue” appears so frequently that it is worth providing some convenient syntax. Whenever `v` is a value expression whose result is an empty tuple, the expression `v;` is a declaration clause whose effect is to evaluate `v`, throw away the result, and then continue with its body. Like all declaration clauses, sequential declarations can appear in sequences, and can be mixed with other declaration clauses in arbitrary ways.

```
run ( (pr "Following ");
      val you = "you, "
      (pr you);
      (pr "my brother.\n");
      ()
    )
```

Following you, my brother.

Formally, we can use a `val` construct to wait for the evaluation of `v` to finish before proceeding with `e`.

$$v; \Rightarrow \text{val } [] = v \qquad \text{(TR-SEMI)}$$

In the Pict libraries, many basic operations (like `pr`) return a null value so that the caller can detect when they are finished. Even in situations where the caller does not care, the null result value must still be accepted and thrown away.

## Chapter 6

# Simple Concurrent Objects

As an example of many of the derived forms described in the previous chapters, let us see how a simple *reference cell* abstraction can be defined in Pict.

A reference cell can be modeled by a process with two channels connecting it to the outside world — one for receiving `set` requests and one for receiving `get` requests. For example, suppose that our cell holds an integer value and that it initially contains 0. Then its behavior can be defined like this:

```
new contents: ^Int          {- Create a local channel holding current contents -}
run contents!0              {- "Initialize" it by sending 0 -}

def set [v:Int c:/[]] =    {- Repeatedly read 'set' requests... -}
  contents?_ =             {- discard the current contents... -}
  (contents!v | c![])      {- install new contents and signal completion -}

def get [res:/Int] =       {- Repeatedly read 'get' requests... -}
  contents?v =             {- read the current contents... -}
  (contents!v | res!v)     {- restore contents and signal result -}
```

The current value of the cell is modeled by a waiting sender on the channel `contents`. The process definitions `set` and `get` must be careful to maintain the invariant that, at any given moment, there is at most one process waiting to send on `contents`; furthermore, when no instances of `get` or `set` are currently running, there should be exactly one sender on `contents`. We can test that our cell is behaving as we expect by sending a few requests and printing the results. Note the use of application and sequencing syntax.

```
run ((prNL ($$ (get)));
      (set 5);
      (prNL ($$ (get)));
      (set -3);
      (prNL ($$ (get)));
      ());
```

```
0
5
-3
```

(The operation `prNL` is a version of `pr` that prints its string argument followed by a newline character.)

This definition is fine if all we need is a single reference cell, but it would be awkward to have to repeat it over and over, choosing different names for the `set` and `get` channels of each whenever we needed a new reference cell. As we did for booleans, we can encapsulate it in a process definition that, each time it is invoked, generates a fresh reference cell and returns the `set` and `get` channels to the caller as a pair.

```
def refInt [res: /[[Int /[]] /[/Int]]] =
  (new contents:^Int
   run contents!0
   def set [v:Int c:/[]] = contents?_ = ( contents!v | c![] )
   def get [res:!Int]     = contents?v = ( contents!v | res!v )
   res![set get]
  )
```

Now we can build multiple reference cells and use them like this:

```
val [set1 get1] = (refInt)
val [set2 get2] = (refInt)

run ((set2 5);
     (prNL ($$ (get1)));
     (prNL ($$ (get2)));
     ())
```

0  
5

But it is not very convenient to have to bind two identifiers each time `refInt` is invoked. A cleaner solution is to bind a single identifier to the whole pair returned by `refInt`:

```
val ref1 = (refInt)
val ref2 = (refInt)
```

Moreover, if we modify `refInt` to return a two field record instead of a two-element tuple, then we can simply use record field-projection syntax to extract whichever request channels we need:

```
def refInt [res: /[set=/[Int /[]] get=/[/Int]]] =
  (new contents:^Int
   run contents!0
   def set [v:Int c:/[]] = contents?_ = ( contents!v | c![] )
   def get [res:/Int]     = contents?v = ( contents!v | res!v )
   res![set=set get=get] )

val ref1 = (refInt)
val ref2 = (refInt)

run ((ref2.set 5);
     (prNL ($$ (ref1.get)));
     (prNL ($$ (ref2.get)));
     ())
```

0  
5

The header of `refInt` will be easier to read if we move the long type of its result to a separate type definition:

```
type RefInt = [  
  set=/[Int /[]]  
  get=/[/Int]  
]
```

Finally, for a final touch of syntactic polish, we can move the definitions of `set` and `get` directly into the fields of the record that is being returned.

```
def refInt [res:/RefInt] =  
(new contents:^Int  
  run contents!0  
  res ! [  
    set = \[v:Int c:/[]] = contents?_ = ( contents!v | c![] )  
    get = \[res:!Int]     = contents?v = ( contents!v | res!v )  
  ])
```

and make the result `res` anonymous by making `refInt` a value abstraction instead of a process abstraction:

```
def refInt () : RefInt =  
(new contents:^Int  
  run contents!0  
  [  
    set = \[v:Int c:/[]] = contents?_ = ( contents!v | c![] )  
    get = \[res:!Int]     = contents?v = ( contents!v | res!v )  
  ])
```

What we have done, in effect, is to introduce a *function* (`refInt`) that creates reference cell *objects*, each consisting of

- a “server process” with some internal state that repeatedly services requests to query and manipulate the state, while carefully maintaining a state invariant, even in the presence of multiple requests, and
- two request channels used by clients to request services, packaged together in a record for convenience.

Active objects of this kind, reminiscent of (though lower-level than) the familiar idiom of *actors* [Hew77; Agh86] (also cf. [Nie92; Pap91; Vas94; PT95; SL96; Var96; NSL96, etc.]), seem to arise almost inevitably when programming in a process calculus. They are widely used in Pict’s libraries.

## Chapter 7

# Advanced Language Features

The full Pict language includes a number of features that cannot be treated fully in a short tutorial. This chapter surveys some of the most useful through a series of examples. See the *Pict Language Definition* for complete details.

### 7.1 Lists

The standard libraries of Pict include support for a variety of common data structures. Among the most important is the library `Untrusted/List`, which defines basic operations for constructing and manipulating lists. For example, the following program constructs a list of integers and then prints its second element:

```
import "Untrusted/List"

val l = (cons 6 (cons 7 (cons 8 nil)))
run print ! ($$ (car (cdr l)))
```

7

The first line makes the definitions in `Untrusted/List` available in the current compilation unit. The second uses the functions `cons` and `nil` to construct the list `l`. The third uses the function `cdr` to select the tail of `l`, then `car` to select its head, and prints the integer returned from `car`.

The next program defines a process abstraction that, when sent a list of integers, prints its second element:

```
def print2ndInt [l: (List Int)] =
  if (null l) then
    print!"Null list"
  else if (null (cdr l)) then
    print!"Null tail"
  else
    print!($$ (car (cdr l)))

run print2ndInt![(cons 6 (cons 7 (cons 8 nil)))]
```

(We elide the `import` clause from here on, to reduce clutter.)

To prevent the proliferation of parentheses in expressions like `(cons 6 (cons 7 (cons 8 (nil))))`, Pict provides a special syntax for repeated applications of the same function to a sequence of arguments (often called “folding”). The last line in the program above can be written

```
run print2ndInt![(cons > 6 7 8 nil)]
```

In general, `(f > a1 a2 ... an a)` means the same as `(f a1 (f a2 ... (f an a)))`. The expression `(f < a a1 a2 ... an)` stands for `(f (f (f a a1) a2) ... an)`, the analogous “right fold” of `f` over `a` and `a1` through `an`.

## 7.2 Polymorphism

Of course, we can build lists of values other than integers. For example, the program above can be rewritten to build and destruct a list of strings like this:

```
def print2ndString [l: (List String)] =
  if (null l) then
    print! "Null list"
  else if (null (cdr l)) then
    print! "Null tail"
  else
    print!(car (cdr l))

run print2ndString![(cons > "one" "two" "three" nil)]
```

two

Indeed, except for the type declaration in the pattern and the channel (`printi` or `print`) used for printing, the definitions of `print2ndInt` and `print2ndString` are identical. We may wish to combine them into a single definition by taking both the type of the elements and an appropriate printing function as parameters:

```
def print2nd [#X l:(List X) p:[X /String]] =
  if (null l) then
    print! "Null list"
  else if (null (cdr l)) then
    print! "Null tail"
  else
    print!(p (car (cdr l)))
```

The `#` before the parameter `X` indicates that it is a *type parameter*. When this abstraction receives a message, the occurrences of `X` in the types of the parameters `l` and `p` must be replaced, consistently,

with whatever type is passed as argument for `X` in the concrete types of whatever values are matched by `l` and `p`. For example, to use `print2nd` on a list of integers, we pass it `Int` for `X` and `$$` for `p`.

```
run print2nd! [#Int (cons > 6 7 8 nil) $$]
```

7

Similarly, to use `print2nd` on a list of strings, we pass `String` for `X` and the identity function for `p`.

```
run print2nd! [#String (cons > "one" "two" "three" nil) \(s:String)=s]
```

two

For the sake of brevity, the type reconstruction phase of the Pict compiler will attempt to fill in any missing types in an argument list that is passed to `print2nd`.

```
run print2nd! [(cons > 6 7 8 nil) $$]
run print2nd! [(cons > "one" "two" "three" nil) \(s:String):String=s]
```

two

7

When a type argument is omitted, the compiler attempts to determine its value by examining the types of the rest of the arguments and matching them against the type of the channel on which the arguments are being sent. If it succeeds in determining the omitted type uniquely, all is well. If not, a compile-time error occurs and the programmer must supply the missing type argument explicitly.

(One slightly subtle point illustrated by this example is that, when type arguments are omitted in this way, the typechecker then has less information to work with when synthesizing the types of the other arguments. Here, this shows up in the fact that we are obliged to add the explicit annotation `:String` showing the result type of the identity function, whereas this was inferred automatically before. Details of how all this works can be found in the *Definition*, but there is no need to understand exactly what information is required and what can be omitted: if the typechecker needs to know more, it will indicate exactly where.)

Indeed, we have been using this mechanism all along, whenever we invoked basic operations from the `List` library. Making all the type arguments explicit, list construction actually looks like this:

```
val l = (cons #Int 6 (cons #Int 7 (cons #Int 8 nil)))
```

The fully annotated version of `print2nd` is:

```
def print2nd [#X l:(List X) p:/[X /String]] =
  if (null #X l) then
    print!"Null list"
  else if (null #X (cdr #X l)) then
    print!"Null tail"
  else
    print!(p (car #X (cdr #X l)))
```

## 7.3 Abstract Types

The kind of polymorphic programming we saw in the previous section is not limited to polymorphic abstractions such as `print2nd`: more generally, Pict allows any tuple (including, of course, tuples with named fields) to include type fields marked by `#` signs; tuple patterns may, in general, include `#` fields as well, introducing dependencies in the types of later fields, as we have seen.

```
val [#X l:(List X) p:[X /String]]
  = [#String (cons > "one" "two" "three" nil) \ (s:String)=s]
val lcdr:(List X) = (cdr l)
run print ! (p (car lcdr))
```

two

Note that the pattern on the first line here binds not only the variables `l` and `p`, but also the type variable `X`, in the scope of further declarations. In other words (by the ordinary rules of variable binding), the `X` appearing in the types of the bound variables `l` and `p` is a different type from any type mentioned anywhere else in the program.

We can exploit this fact—that patterns with type bindings create fresh types—to build *abstract types* whose elements can be manipulated only by some given set of functions.<sup>1</sup> As, a simple (but useful) example, the following binding introduces a new *enumeration type* called `Weekday` and two operators, `sameday` and `tomorrow`:

```
val [
  #Weekday
  monday:Weekday tuesday:Weekday wednesday:Weekday thursday:Weekday
  friday:Weekday saturday:Weekday sunday:Weekday
  sameday:[Weekday Weekday /Bool]
  tomorrow:[Weekday /Weekday]
] = [
  #Int                                {- representation type -}
  0 1 2 3                             {- representations of days -}
  4 5 6
  \ (d1:Int d2:Int) = (== d1 d2)      {- same day? -}
  \ (d:Int) = (mod (+ d 1) 7)         {- tomorrow -}
]
```

The `Weekday` type defined this way is (in the scope where it is visible) completely unrelated to the implementation type `Int`; i.e. neither `Weekday = Int` nor `Weekday < Int` nor `Int < Weekday`. Actually, definition of the `sameday` operation could be done in a more terse way:

```
val [
  #Weekday
  monday:Weekday tuesday:Weekday wednesday:Weekday thursday:Weekday
  friday:Weekday saturday:Weekday sunday:Weekday
  sameday:[Weekday Weekday /Bool]
  tomorrow:[Weekday /Weekday]
```

---

<sup>1</sup>This programming technique directly follows Mitchell and Plotkin's explanation of abstract datatypes in terms of existential types in the lambda-calculus [MP88].

```

] = [
  #Int                {- representation type -}
  0 1 2 3            {- representations of days -}
  4 5 6
  ==                {- same day? -}
  \ (d: Int) = (mod (+ d 1) 7)  {- tomorrow -}
]

```

After this declaration, we can behave as though `Weekday` were a built-in type with constant elements `monday` through `sunday` and built-in operations `sameday` and `tomorrow`:

```

def weekend(d: Weekday): Bool =
  (|| (sameday d saturday) (sameday d sunday))

```

## 7.4 User-defined Type Constructors

In Chapter 6, we saw how to define a type `RefInt` of integer reference cell objects and an associated constructor `refInt` for creating new elements of this type.

```

type RefInt = [
  set=/[Int /[]]
  get=/[Int]
]

def refInt () : RefInt =
  ( new contents: ^Int
    run contents!0
    [ set = \[v: Int c:/[]] = contents?_ = ( contents!v | c![] )
      get = \[res: !Int]    = contents?v = ( contents!v | res!v )
    ]
  )

```

We now generalize these definitions to a *parametric type* `Ref` that, like `List`, describes a family of types—`(Ref T)` is the type of reference cells holding elements of `T`—and an associated *polymorphic constructor* `ref` for creating reference cell objects. First, the type constructor `Ref` is defined like this:

```

type (Ref X) = [
  set=/[X /[]]
  get=/[X]
]

```

The pattern `(Ref X)` on the left of the `=` declares the name `Ref` and gives the name `X` to the parameter type (this occurrence of `X` is a binder, whose scope is the type expression on the right-hand side of the `=`). Put differently, `Ref` is a function from types to types such that, for each type `T`, the instance `(Ref T)` means the same as `[set=/[X /[]] get=/[X]]`.

With this definition in hand, the polymorphic `ref` function is easy to write:

```

def ref (#X init:X) : (Ref X) =
  ( new contents:^X
    run contents!init
    [ set = \[v:X c:/[]] = contents?_ = ( contents!v | c![] )
      get = \[res:!X]    = contents?v = ( contents!v | res!v )
    ]
  )

```

The only real difference from `refInt` is that, here, we must take the initial value of the cell as a parameter to `ref`, since we do not have a uniform way to make up a default value of the parameter type `X`.

## 7.5 Recursive Types

We can combine what we know—parametric type definitions and polymorphic functions—plus one more feature—so-called *recursive types*<sup>2</sup>—to build an object-oriented version of the predefined `List` package. In this version, we will make each list value carry its own internal functions (methods) for the `null`, `car`, and `cdr` operations instead of writing these operations separately and passing the list as a parameter when they are called.

A list whose elements have type `X`, then, will be a record of three functions: (1) `null`, returning a `Bool`; (2) `car`, returning an element of `X`; and (3) `cdr`, returning a list of the same type. Informally, what we want to write is this:

```

type (OurList X) =
  [ null=/[Bool]
    car=/[X]
    cdr=/[OurList X]
  ]

```

But this definition is not well formed, since it mentions on the right of the `=` the very type constructor `OurList` that is being defined.

To handle such recursive type definitions, we introduce a type constructor `rec` that shows explicitly that `OurList` is recursively defined:

```

type (OurList X) =
  ( rec L =
    [ null=/[Bool]
      car=/[X]
      cdr=/[L]
    ]
  )

```

---

<sup>2</sup>The formulation of recursive types described here is an interim language feature, not a final solution. Although the required mechanisms are fairly simple to describe, we find them quite unwieldy in practice. (In particular, although it is possible in principle to use this mechanism to define *mutually recursive* types, it is too painful to do this in practice.)

The usual technique of making the typechecker automatically infer the required foldings and unfoldings of recursive types is not workable here, due to the complexity of the rest of the type system (in particular, subtyping and type operators). In our view, the best solution would be based on ML's `datatype` mechanism, where recursive types are combined with variants and explicit constructors are used to mark both variant tags and points where folding or unfolding is required.

```
    ]
  )
```

The bound variable `L` stands for the whole type (`rec L = ...`) in the type expression to the right of the `=`.

There are two ways of building lists: the empty list `nil` (we'll call it `ournil` to avoid confusion with the `List` library) and the constructor `cons` (we'll call it `ourcons`). As in the standard `List` library, these are both polymorphic functions, since we need to use them to create lists with elements of arbitrary types. Here is `ournil`:

```
def ournil (#X) : (OurList X) =
  ( rec
    [ null = \() = true
      car = \[r:/X] = ()
      cdr = \[r:/(OurList X)] = ()
    ]
  )
```

Note the value constructor `rec` here, which shows the typechecker explicitly that the type of the body (a record type) is to be “folded up” into the recursive type (`OurList X`). This operation must be performed explicitly: the unfolded record type and the folded recursive type are treated as completely distinct by the typechecker. Similarly, here's `ourcons`:

```
def ourcons (#X hd:X tl:(OurList X)) : (OurList X) =
  ( rec
    [ null = \()=false
      car = \()=hd
      cdr = \()=tl
    ]
  )
```

Using `ournil` and `ourcons`, we can build a list of two integers like this:

```
val l = (ourcons #Int 3 (ourcons #Int 4 (ournil #Int)))
```

Just as we folded up our record values above into recursively typed values, we must unfold our recursively typed value `l` before we can project its fields—if we try to project the `car` field of `l` itself, a typechecking error occurs:

```
run print!(l.car)
```

```
example.pi:22.14: expected a record type but found: (OurList Int)
```

The required unfolding is accomplished by the `rec` pattern constructor, which matches an element of a recursive type but whose body pattern matches an element of the unfolded body of the recursive type.

```
val (rec ll) = l
run print!(ll.car)
```

3

Just as before, we can drop the explicit type arguments to `ournil` and `ourcons` in the construction of `l`:

```
val l = (ourcons 3 (ourcons 4 ournil))
```

Finally, we can use our “function folding” syntactic sugar to reduce the number of parentheses:

```
val l = (ourcons > 3 4 5 6 7 8 9 ournil)
```

## Chapter 8

# Security Concerned Programming

This article presents additional necessary measures that enable us to use Pict as an object-capability programming language. It is desirable to be able to assess the worst possible threat that we—users—risk if we run a given program. If we know the threat, we are able to decide whether or not we are willing to risk running the program. The cost of a security audit that reveals such an assessment will be non-zero but it need not to be directly dependent on the size of the whole original program. It is possible to write programs in such a way that this analysis can be reliably performed on a fraction of the original program—on the trusted computing base. This technique does not always give the most accurate assessment but it gives sound and interesting assessment relatively cheaply. It does not prevent usage of other techniques that can further refine the initial assessment.

### 8.1 Introduction

There are two different points of view of a computer system. We can view it from an administrator's point of view and from a user's point of view. Users should be regarded as primary because the purpose of computers is not to be administered but to be used. The goal of the administrator is to ensure that none of the users is given excess authority. The goal of the user is (should be) to ensure that each of the processes runs with appropriate authority. Security mechanisms provided by operating system are practical for administrator but they do not help users with their security goals. Microsoft "Immutable" Law #1 states:

If a bad guy can persuade you to run his program on your computer, it's not your computer anymore.

The problem is, how to decide who is a good guy and who is a bad guy. More importantly, even good guys can make mistakes and their programs can cause damage. The purpose of the computer is that we—users—can run programs on it. This rule basically says that we are safe as long as we do not run any program on it. Let us stop here and think how ridiculous it is.

Noticeable progress has been made in the area of designing programming languages with respect to security. Outstanding example is the E programming language [Mil06]. From the security point of view<sup>1</sup>, it is interesting because it enables programmers to follow the principle of the least authority

---

<sup>1</sup>The E programming language addresses also other important problems.

(POLA). Multiple aspects of the language contribute to this fact:

- the authority to invoke methods of a particular object is an unforgeable capability
- when some subsystem decides to keep some capabilities as private, there are no language constructs that would enable other untrusted subsystems to “steal” them
- the reference graph can evolve only according to *rules of allowed reference graph dynamics* presented in Section 9.2 of [Mil06]

The contribution of this paper is that it shows how, through a refactorization of the libraries of the Pict programming language [PT00], the “ambient authority” is reduced to a minimum, and Pict can provide many of the benefits of existing object-capability languages. Provided examples illustrate the technique for determining authority of untrusted subsystems without the need to analyze their code.

## 8.2 Related Work

While this article is mostly concerned with taming of Pict—turning Pict into an object-capability programming language—this is not the first work of this kind. See for example: Oz-E [SR05], Emily [Sti07], Joe-E [MW06].

The Raw Metal occam Experiment (RMoX) [BJV03] can be regarded as a source of inspiration that languages, based on process calculi, can be used for defining of behavior of various operating system’s components and their mutual interaction. Using programming language constructs as a mechanism for isolation of various subsystems from each other instead of relying on awkward hardware support is also one of the points of the Singularity project [AFH<sup>+</sup>06].

## 8.3 The Pict Programming Language

The goal of the authors of the Pict programming language was to create a language that could play for the  $\pi$ -calculus a similar role as Haskell plays for the  $\lambda$ -calculus. It is defined in layers, see Figure 8.1. Syntax of the core language is formally described in the Pict Language Definition [PT97c] in Chapter 3; see rules tagged as *C* (as *Core*). Some of the syntactically correct programs can be further rejected by the typing rules at compile time. Semantics of the core language is defined in Chapter 13 of that document. It defines:

- structural congruence relation
- reduction relation

These together define behavior of all Pict programs.

Programs written in the core Pict cannot break *rules of allowed reference graph dynamics*. Derived forms make functional and sequential programming in Pict more convenient. By definition, they do not add expressivity to the core Pict language and thus can be used without concerns that the *rules of allowed reference graph dynamics* could be broken.

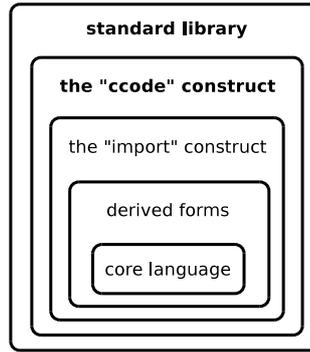


Figure 8.1: Layers of the Pict programming language. Programs that are composed solely from core constructs, derived forms and `import` directives are completely harmless because they have minimal authority.

The `import` directive is one of the two extralinguistic constructs of Pict. It enables us to split the whole program into multiple, separately compilable modules. These modules are related via `import` construct. This relation forms partial order with the biggest element—it is the main module of a complete program. There is no `export` directive via which the programmer could explicitly specify which bindings he wants to export from a given module. All the variables that are bound in the outer-most scope are automatically exported. The effect of the `import` directive is that all the variables exported by the imported module are visible in the importing module.

The `ccode` construct is the second of the two extralinguistic constructs of Pict. It enables the programmer to inline arbitrary C code into Pict programs. This is very useful and very dangerous at the same time. It is the sole mechanism that Pict programs can use to interact with their (non-Pict) environment such as the operating system. It is also used for implementation of certain operations in an efficient way. This is not absolutely essential<sup>2</sup> but it is pragmatic. Additional rules presented later in the text ensure that this construct cannot be directly or indirectly abused by untrusted modules to gain excess authority. These additional rules ensure that untrusted modules cannot break the *rules of allowed reference graph dynamics*.

The Standard Pict Library [PT97d] provides several reusable components. Figure 8.2 shows some modules that are part of this library. Some aspects of this original organization are logical and some are not logical. The `import` directive binds them into a partially ordered set. Minimal elements (`Misc`, `Prim`) are shown on the left. Maximal elements (`Random`, `Ref`, `Signals`, `Array2`, `Queue`, `Args`, `IO`) are shown on the right. The  $A \prec B$  means that module  $A$  is imported by module  $B$ .

Those names that are bound in the outer-most scope of some module are also exported by that module. Let  $en(A)$  denote a function that maps a given module  $A$  to the set of names that it exports. Then, by definition of the semantics of the `import` directive:

$$A \prec B \Rightarrow en(A) \subseteq en(B)$$

That is, all the names bound in the outer-most scope of module  $A$  are also bound in the outer-most scope of module  $B$  that imports  $A$ .

<sup>2</sup>The core language can model integers, booleans, strings and other values of basic data types together with operations with them.

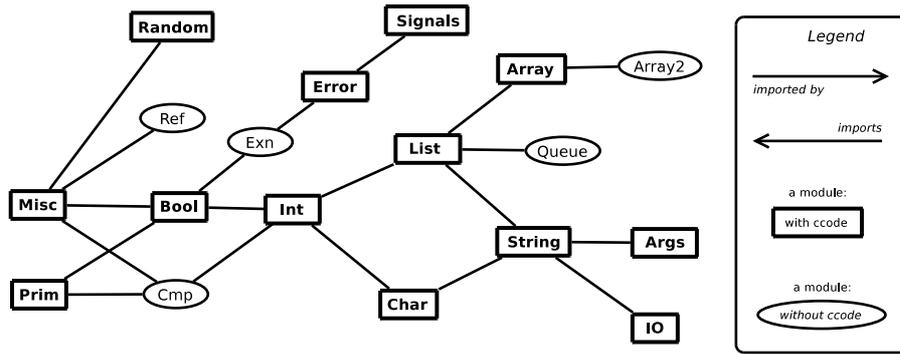


Figure 8.2: Partial ordering of modules with respect to the `import` relation. These modules are described in detail elsewhere [PT97d]. This figure is provided only for general impression concerning the structure of modules. Not all the aspects of this original version are completely logical.

When  $bn(A)$  denotes the set of all names bound in any scope within module  $A$  then for all modules  $A$ , by definition of the `import` construct, holds:

$$en(A) \subseteq bn(A)$$

All the names bound in  $A$  need not to be, and usually indeed are not, exported.

## 8.4 Refactorization of the Original Pict Library

The attempt to minimize the trusted computing base is inherently a good idea. In this light, the organization of the original Pict library is not optimal. Each module that employs the `ccode` construct must be considered as part of the trusted computing base. And, as you can see in Figure 8.2, there are many such modules. Additionally, the original set of primitives, expressed via `ccode` construct, is not orthogonal. Many of the existing primitives can be rewritten in terms of a pure Pict code. After we removed those superfluous primitives and we concentrated the originally scattered primitives in a few dedicated modules, the situation is different, see Figure 8.3. Now it has sense to discriminate among trusted and untrusted modules as follows:

- *trusted* modules can contain any (compilable) code
- *untrusted* modules:
  - cannot use the `ccode` construct
  - cannot import any trusted module except for the `Prim` module which provides harmless primitives

Due to the inherent properties of the Pict programming language, these measures are sufficient to ensure that rules of allowed reference graph dynamics hold for all our untrusted modules.

After all these measures, by *minimal authority* of untrusted modules written in Pict we mean:

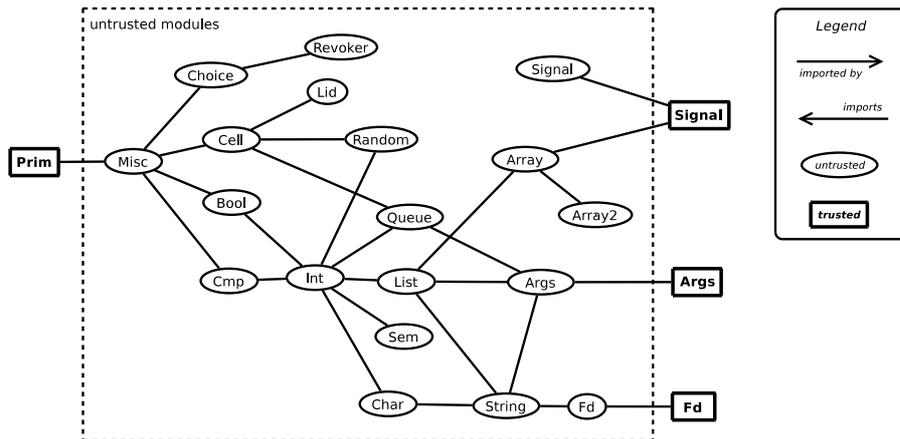


Figure 8.3: Refactored standard library of Pict with respect to security.

- Trusted modules have no straightforward way to influence policy how much memory can various sub-components allocate from the common heap of free memory which in Pict is bounded—it’s size is by default 1 MB. When untrusted components exhaust it, the Pict runtime prints out a relevant error message and terminates the whole system.
- Trusted modules have no straightforward way how influence the scheduling policy that would define rules for CPU utilization by untrusted modules. At present, there is a scheduler. It ensures fairness of the CPU utilization but this is not what we always want.
- Untrusted modules can make run-time errors (such as division by zero or they might attempt to access non-existent element of some array). These run-time errors are always detected and result in calling the `error` function that performs appropriate actions. At present, this means that some error message will be printed on the screen and the whole system will terminate. So indirectly, untrusted modules have the authority to terminate the whole system.

This situation is not at all fully satisfactory. But this is still far better than ambient authority of trusted modules. To address these remaining issues the whole Pict runtime must be redesigned with these problems in mind.

## 8.5 Powerbox

Powerbox is a fundamental security pattern. Its origin can be traced to the DarpaBrowser described in [WT02] and in [SM02]. It enables us to dynamically raise the level of authority of untrusted subsystems to a sufficient and acceptable level.

If we are concerned with some single purpose program then we have to identify the authority this program needs. There is nothing inherently wrong that various programs require some authority. As long as it is explicitly declared, users or security auditors can efficiently judge whether it is acceptable for us to grant such authority to the actual program.

To be able to follow POLA, the whole program must be split in at least two modules. The first of them will be trusted and the other one will be untrusted. The purpose of the trusted module is to

communicate part of its ambient authority to the untrusted module. The purpose of the untrusted module is to use the authority it is given and to do what we expect from it. It receives required capabilities “by introduction”. What kind of capabilities are communicated and through which channel depends on the contract between the trusted and the untrusted part.

Let us show a very simple example. We keep the untrusted module in the `Untrusted/Guest.pi` file and the trusted module in the `Trusted/Host.pi` file. The `Untrusted/Guest` module might look as follows:

```
new contract : ^!String

run contract?logger = ( logger!"0123456789"
                       | logger!"0123456789"
                       | logger!"0123456789"
                       | logger!"0123456789"
                       | logger!"0123456789"
                       )
```

It creates a fresh channel `contract` that can be used for passing values of the `!String` type<sup>3</sup>. The process in the untrusted module blocks until it receives a value from the `contract` channel. When such value arrives, it will be bound to the `logger` variable. The untrusted guest then has all the authority it needs to do its job. In this case it prints 50 characters on the screen. The above program is not very useful but it could very well perform various simulations and then print out the simulation report. The above code fragment is a mere illustration.

The `Trusted/Host` module is responsible for selecting parts of its ambient authority and communicating appropriate capabilities to the untrusted module. For example:

```
import "Untrusted/Guest"

run contract!print
```

It imports two modules. The first one is `Untrusted/Guest`. This means that it will see the `contract` capability exported by that module (because it is bound in the outer-most lexical scope). It also imports the `Trusted/Fd` module. It means that it will see the `print` capability exported by that module. It is up to this trusted module to select proper capabilities. In this case it selects the `print` capability and sends it over the `contract` channel. Of course, there may be situations where some guest needs more than one capability. In those cases the trusted host sends an n-tuple of capabilities.

Appropriate makefile for building executable out of these two modules can look as follows:

```
Host: Trusted/Host.pi Untrusted/Guest.px
     pict -o Host Trusted/Host.pi
```

---

<sup>3</sup>Pict is a strongly typed programming language. Each channel has a type. This type determines what kind of values can be communicated over a given channel. An attempt to send a wrong type of value over some channel is detected at compile time. The `contract` capability has a type `^!String`. Our process holds this capability *by initial conditions*. The initial `^` character means that this capability can be used for sending as well as for receiving values of the `!String` type. Our process uses this capability to receive a value from it and binds this value to the `logger` capability. This capability is of `!String` type. That means that the `logger` capability can be used for sending strings along it. It cannot be used for receiving strings from this channel.

```

Untrusted/Guest.px: Untrusted/Guest.pi
    isUntrusted $< && pict -reset lib -set sep -o Untrusted/Guest.px Untrusted/Guest.pi

clean:
    rm -f {Trusted,Untrusted}/*.{o,px} Host

```

Please notice two things:

- `Untrusted/Guest.pi` module is checked with the `isUntrusted` script whether it indeed can be regarded as untrusted<sup>4</sup>
- we compile the `Untrusted/Guest.pi` module with the `-reset lib` flag that inhibits inclusion of the standard prelude<sup>5</sup>.

These two actions give us enough confidence to believe that the `Untrusted/Guest` module has initially *minimal authority*. Its authority is later raised to be able to print characters on the standard output. It is not given any other authority. It cannot tamper with files that can be accessed by the user that runs this program. The untrusted module cannot communicate with other processes on your local system. Neither it can communicate over network. It can only print as many characters on the standard output as it wishes. For some programs this kind of authority might be completely sufficient and as you can see it can be trivially implemented.

The same scheme has many variants. The derived forms make certain useful things such as functional programming as well as sequential programming easier. If we express our trusted host and our untrusted guest in so called “continuation passing style” then it would appear that the trusted host gives the untrusted guest the capability to call certain functions. In this case, it is completely up to the trusted host to choose the right set of function-capabilities. The chosen set determines the authority of the untrusted guest.

The Powerbox pattern can also be used in situations when our system consists of multiple untrusted subsystems. In that case, each subsystem will be placed in a separate powerbox. This way, each untrusted component can be given different capabilities and thus we can determine the authority of particular untrusted modules independently. The complexity of the trusted part is determined by the complexity of our security policy. It is independent from the complexity of the untrusted part that does the real job.

## 8.6 Experiments in the Kernel Space

Capability-secure languages are useful not only in user-space but they can have interesting applications in the kernel space, too. They can be a precursor to making progress in monolithic (in the traditional sense) kernels. We have a flexible alternative to the classical microkernel-based operating system architecture. In our preliminary experiment we use the Pict programming language because it was easier to adapt to run on a bare metal. From the security point of view Pict is

---

<sup>4</sup>Untrusted modules cannot employ the `ccode` constructs. Untrusted modules cannot import trusted modules except for the `Trusted/Prim` module.

<sup>5</sup>Precise information concerning the “standard prelude” can be found in [Ko07]. Basically, it is a default sequence of `import` directives that is desirable in case of trusted modules but it is undesirable in case of untrusted modules.

in principle as good as E. From the concurrency point of view, the E programming language is much better. It provides more advanced synchronization mechanisms than Pict so E is a very good alternative for the future. Figure 8.4 shows the structure of modules with respect to the `import`

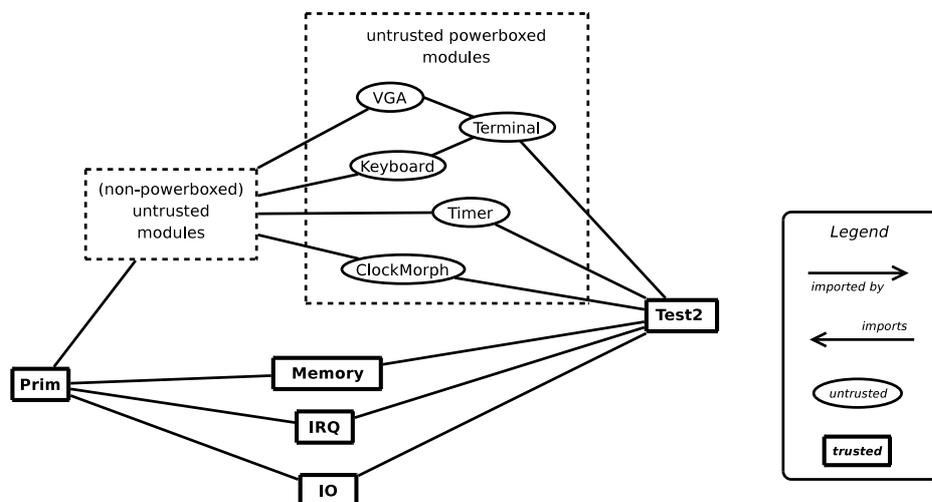


Figure 8.4: Structure (with respect to the `import` relation) of modules that form our experimental kernels. `Test2` is the main module.

relationship. This relationship determines the *connectivity by initial conditions* according to the semantics of the `import` directive. Capabilities that are exported from module *A* are also visible in module *B* if module *A* is imported by module *B*. Modules `Memory`, `IRQ` and `IO` provide various powerful primitives. For example the `Trusted/Memory` module exports a function

```
(memory.write.byte offset value)
```

that enables (those who see this function-capability) to write any value of byte size to any `offset` within current data segment (that spans through the whole physical memory). The `Trusted/IO` module exports two functions:

```
(io.write.byte port value)
(io.read.byte port)
```

Those who see the first function can write any value (byte) to any I/O port. Those who see the second function can read any I/O port of byte size.

The trusted `Test2` module has a single task, to disseminate proper capabilities to proper modules via the Powerbox pattern according to POLA. For example, one thing that the `VGA` module needs is the authority to write to the I/O port number 980 (0x3D4 in hexadecimal system). The `Test2` module sees the `io.write.byte` procedure so it could give the `VGA` module this capability. However, with respect to POLA, it gives it a different capability:

```
\(value) = (io.write.byte 980 value)
```

This abstraction (unnamed function, lambda-expression) is given to the `VGA` module. It gives it the authority to write any `byte` to the I/O port 980. The `Test2` module gives the `VGA` module few other similar capabilities. As a result, the `VGA` module has the authority:

- to write any byte to the I/O port 980
- to write any byte to the I/O port 981
- to read a byte from the I/O port 981
- to write any byte to the Video RAM (nowhere else)

This is enough for the VGA module to be able to provide expected services. Various abstractions

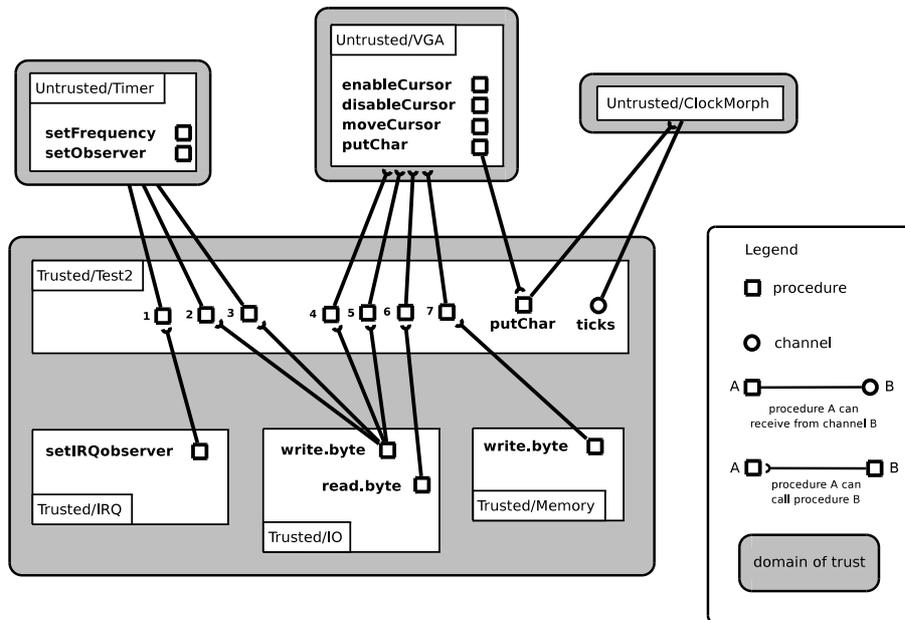


Figure 8.5: The actual reference graph in the `Test2` kernel.

created by the `Test2` module that act as proxies to more powerful capabilities are denoted as small numbered rectangles in the `Trusted/Test2` module in Figure 8.5.

If we have a complete and correct<sup>6</sup> knowledge concerning behavior of functions, procedures and processes in all the trusted modules (`Memory`, `IO`, `IRQ`, `Test2`), then we can safely assess the upper bound of authority of all untrusted powerboxed modules (`Timer`, `VGA`, `ClockMorph`) safely only with regard to the reference graph shown in Figure 8.5. It is possible because rules of allowed reference graph dynamics hold for our untrusted modules. In case of `VGA` and `Timer` modules the situation is trivial. We have complete knowledge about behavior of functions that we give to these two modules. Recall that we give the `VGA` module the following capability.

```
\(value) = (io.write.byte 980 value)
```

Since we have a complete information concerning the `io.write.byte` function, we also have a complete information concerning the behavior of the above abstraction. So the authority of the `VGA` and the `Timer` can be determined precisely; regardless of their actual implementation. These drivers provide various functions. One of them is:

<sup>6</sup>This is why we should try to keep the trusted computing base as small as possible.

```
(vga.putChar x y ch attribute)
```

When called, it puts any given character `ch` with any `attribute` anywhere on the screen. This is its assumed effect we believe it does.

In a similar way can we also give appropriate authority to the `ClockMorph` component. It is supposed to show the number of seconds from the boot-time in `HH:MM:SS` format. This kind of component obviously needs the authority to print eight consecutive characters somewhere on the screen. If we want to follow POLA also in this case, we have to implement a proxy function that will drop most of the `vga.putChar` authority and it will provide the ability to change eight consecutive characters on the screen, not more. We have defined the `putChar` function in the `Test2` module that does exactly this. It relies on the `vga.putChar` function, see Figure 8.5. Regardless how perfectly our trusted proxy function implements additional restrictions, unless we verify the correct behavior of the original untrusted `vga.putChar` function, we cannot claim anything stronger than: “The `ClockMorph` component has as much authority as the VGA driver plus it can receive messages from the `tick` channel.” But even with this simple technique, without studying the code of particular untrusted modules, we can see that the authority of the `ClockMorph` component is fairly limited.

## 8.7 Conclusion and Future Work

Our immediate goal is to address two immediate problems concerning minimal authority:

- there is no way how to give particular untrusted components only limited share of the CPU bandwidth
- there is no way how to give particular untrusted components only limited amount of memory

At present, when some of the untrusted components uses up the whole available memory, the runtime terminates the system. This is a show-stopper for using Pict for writing robust, from the traditional point of view, monolithic operating system kernels.

Sophisticated proof-techniques were developed for proving correctness of functional code, sequential (procedural) code. These can be very useful in analysis of procedures that are made visible to untrusted powerboxed modules. From the formally proved effects of these procedures (or processes) we can determine the authority of untrusted powerboxed modules that are part of the system.

### Acknowledgments.

This work was partially supported by the Slovak Research and Development Agency under the contract No. APVV-0391-06 and by the Scientific Grant Agency of Slovak Republic, grant No. VG1/3102/06. Thanks also go to our paper shepherd Frank Piessens whose insightful comments significantly helped to improve the paper.

# Appendix A

## Solutions to Selected Exercises

Solution to 5.1.3.2:

```
def fib[n:Int r:!Int] =  
  if (|| (== n 0) (== n 1)) then  
    r!1  
  else  
    r!(+ (fib (- n 1)) (fib (- n 2)))  
  
run print!(fib 7)
```

21

# Bibliography

- [AFH<sup>+</sup>06] Mark Aiken, Manuel Fähndrich, Chris Hawblitzel, Galen Hunt, and James Larus. Deconstructing process isolation. In *MSPC '06: Proceedings of the 2006 workshop on Memory system performance and correctness*, pages 1–10, New York, NY, USA, 2006. ACM.
- [Agh86] Gul A. Agha. *Actors: a Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, MA, 1986.
- [BJV03] Fred Barnes, Christian Jacobsen, and Brian Vinter. RMoX: A raw-metal occam experiment. In J.F. Broenink and G.H. Hilderink, editors, *Communicating Process Architectures 2003*, volume 61 of *Concurrent Systems Engineering Series*, pages 269–288, Amsterdam, The Netherlands, September 2003. IOS Press.
- [Bou92] Gérard Boudol. Asynchrony and the  $\pi$ -calculus (note). Rapport de Recherche 1702, INRIA Sofia-Antipolis, May 1992.
- [Chu41] Alonzo Church. *The Calculi of Lambda Conversion*. Princeton University Press, 1941.
- [Hew77] C. Hewitt. Viewing control structures as patterns of passing messages. *Artificial Intelligence*, 8:323–364, 1977.
- [HT91] Kohei Honda and Mario Tokoro. An object calculus for asynchronous communication. In Pierre America, editor, *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 512 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, Heidelberg, New York, Tokyo, 1991.
- [Ko07] Matej Kok. Tamed Pict Library, 2007.
- [Lan66] P. J. Landin. The next 700 programming languages. *Communications of the ACM*, 9(3):157–166, March 1966.
- [McC78] John McCarthy. History of Lisp. In *Proceedings of the first ACM conference on History of Programming Languages*, pages 217–223, 1978. ACM Sigplan Notices, Vol. 13, No 8, August 1978.
- [Mil80] Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer Verlag, 1980.
- [Mil89] Robin Milner. *Communication and Concurrency*. Prentice Hall, 1989.
- [Mil90] Robin Milner. Functions as processes. Research Report 1154, INRIA, Sofia Antipolis, 1990. Final version in *Journal of Mathematical Structures in Computer Science* 2(2):119–141, 1992.
- [Mil92] Robin Milner. Action structures. Technical Report ECS–LFCS–92–249, Laboratory for Foundations of Computer Science, University of Edinburgh, December 1992.
- [Mil95] Robin Milner. Calculi for interaction. *Acta Informatica*, 1995. To appear.

- [Mil06] Mark Samuel Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Johns Hopkins University, Baltimore, Maryland, USA, May 2006.
- [MP88] John Mitchell and Gordon Plotkin. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems*, 10(3), July 1988.
- [MPW92] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes (Parts I and II). *Information and Computation*, 100:1–77, 1992.
- [MTH90] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML*. The MIT Press, 1990.
- [MW06] Adrian Matthew Mettler and David Wagner. The Joe-E Language Specification (draft). Technical Report UCB/EECS-2006-26, EECS Department, University of California, Berkeley, 3 2006.
- [Nie92] Oscar Nierstrasz. Towards an object calculus. In M. Tokoro, O. Nierstrasz, and P. Wegner, editors, *Proceedings of the ECOOP '91 Workshop on Object-Based Concurrent Computing*, Lecture Notes in Computer Science number 612, pages 1–20. Springer-Verlag, 1992.
- [NP96] Uwe Nestmann and Benjamin C. Pierce. Decoding choice encodings. In *Proceedings of CONCUR '96*, August 1996.
- [NSL96] Oscar Nierstrasz, Jean-Guy Schneider, and Markus Lumpe. Formalizing composable software systems — a research agenda. In *Formal Methods in Open, Object-Based Distributed Systems (FMOODS '96)*, February 1996.
- [Pap91] M. Papathomas. A unifying framework for process calculus semantics of concurrent object-based languages and features. In Dennis Tsichritzis, editor, *Object composition Composition d'objets*, pages 205–224. Centre Universitaire d'Informatique, Universite de Geneve, [6] 1991.
- [PRT93] Benjamin C. Pierce, Didier Rémy, and David N. Turner. A typed higher-order programming language based on the pi-calculus. In *Workshop on Type Theory and its Application to Computer Systems, Kyoto University*, July 1993.
- [PS93] Benjamin Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. In *Logic in Computer Science*, 1993. Full version in *Mathematical Structures in Computer Science*, Vol. 6, No. 5, 1996.
- [PS96] Benjamin Pierce and Martin Steffen. Higher-order subtyping. *Theoretical Computer Science*, 1996. To appear. A preliminary version appeared in IFIP Working Conference on Programming Concepts, Methods and Calculi (PROCOMET), June 1994, and as University of Edinburgh technical report ECS-LFCS-94-280 and Universität Erlangen-Nürnberg Interner Bericht IMMD7-01/94, January 1994.
- [PS97] Benjamin Pierce and Davide Sangiorgi. Behavioral equivalence in the polymorphic pi-calculus. In *Principles of Programming Languages (POPL)*, 1997. Full version available as INRIA-Sophia Antipolis Rapport de Recherche No. 3042 and as Indiana University Computer Science Technical Report 468.
- [PT95] Benjamin C. Pierce and David N. Turner. Concurrent objects in a process calculus. In Takayasu Ito and Akinori Yonezawa, editors, *Theory and Practice of Parallel Programming (TPPP), Sendai, Japan (Nov. 1994)*, number 907 in Lecture Notes in Computer Science, pages 187–215. Springer-Verlag, April 1995.
- [PT97a] Benjamin C. Pierce and David N. Turner. Pict: A programming language based on the pi-calculus. Technical report, Computer Science Department, Indiana University, 1997. To appear in Milner *Festschrift*, MIT Press, 1997.

- [PT97b] Benjamin C. Pierce and David N. Turner. Pict language definition. Draft report; available electronically as part of the Pict distribution, 1997.
- [PT97c] Benjamin C. Pierce and David N. Turner. Pict language definition. <http://citeseer.ist.psu.edu/article/pierce96pict.html>, 1997.
- [PT97d] Benjamin C. Pierce and David N. Turner. Pict libraries manual. Available electronically, 1997.
- [PT00] Benjamin C. Pierce and David N. Turner. Pict: A programming language based on the pi-calculus. In G. Plotkin, C. Stirling, and M. Tofte, editors, *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press, 2000.
- [Sew96] Peter Sewell. Observations on Pict, a nondeterministic programming language. Manuscript, 1996.
- [SL96] Jean-Guy Schneider and Markus Lumpe. Modelling objects in Pict. Technical Report IAM-96-004, Universitaet Bern, Institut fuer Informatik und Angewandte Mathematik, January 1996.
- [SM02] Marc Stiegler and Mark S. Miller. A Capability Based Client: The DarpaBrowser. Technical Report Focused Research Topic 5 / BAA-00-06-SNK, Combex, Inc., 5 2002.
- [SR05] Fred Spiessens and Peter Van Roy. A Practical Formal Model for Safety Analysis in Capability-Based Systems, Revised Selected Papers. In *Trustworthy Global Computing*, pages 248–278. Springer Berlin, 2005. Lecture Notes in Computer Science.
- [Sti07] Marc Stiegler. Emily: A high performance language for enabling secure cooperation. *Creating, Connecting and Collaborating through Computing, International Conference on*, 0:163–169, 2007.
- [Var96] Patrick Varone. Implementation of “generic synchronization policies” in Pict. Technical Report IAM-96-005, Universitaet Bern, Institut fuer Informatik und Angewandte Mathematik, April 1996.
- [Vas94] Vasco T. Vasconcelos. Typed concurrent objects. In *Proceedings of the Eighth European Conference on Object-Oriented Programming (ECOOP)*, volume 821 of *Lecture Notes in Computer Science*, pages 100–117. Springer-Verlag, July 1994.
- [WT02] David Wagner and E. Dean Tribble. A Security Analysis of the Combex DarpaBrowser Architecture, 3 2002.