A Tutorial to the Clean Object I/O Library - version 1.1

Peter Achten
Martin Wierich
Department of Functional Programming
University of Nijmegen
The Netherlands

June 11, 1999

Contents

Chapter 1

Preface

The functional programming language Clean has an extensive library to build graphical user interface applications, the *object I/O library*. In this tutorial the basic concepts of the object I/O library are explained. Many of the concepts are illustrated by means of Clean program examples. Clean code will be typeset in `type writer` style. The example programs are also available as Clean sources in the corresponding 'Tutorial Examples' folder.

This tutorial is not a technical reference manual: you will not find an extensive and detailed description of every data structure and function of the library. It is also assumed that the reader is familiar with functional programming and Clean.

In Chapter 2 a very broad overview of the object I/O library is given, just to give the reader a taste of what the object I/O system is all about. Chapter 3 presents the global structure of the object I/O system.

The remaining part of this document explains the individual components of the library. But before we can explain the graphical user interface elements, we first talk about object identification (Chapter 4), and drawing (Chapter 5).

To users of graphical user interfaces *the* interface elements are of course the *windows* and *dialogues*. These are discussed in Chapter 6. Windows and dialogues can contain *controls*. Because there are many aspects about control handling their treatment deserves a separate chapter (7). In all graphical user interface systems, the set of available commands is presented by means of *menus*, see Chapter 8. To support timing features, *timers* can be used, see Chapter 9. Flexible communication of arbitrary expressions between components can be achieved by using *receivers* and *message passing*, see Chapter 10.

All of the above objects are elements of one interactive process. The object I/O library enables the programmer to split up a large interactive program into several interactive processes that can be created and closed dynamically. This is presented in Chapter 11.

Two final chapters remain that deal with issues that are also related to graphical user interface applications. In Chapter 12 we discuss clipboard handling, a simple user driven mechanism to transfer data between interactive applications. In Chapter 13, we demonstrate the means that exist to send output to a printer.

In Chapter 10 communication of arbitrary Clean expressions between interactive objects is discussed. In Chapter 14 it is demonstrated how Clean programs can communicate with arbitrary programs using TCP.

Appendix A contains the definition modules of the Clean object I/O library, version

1.1. in alphabetic order.

The Clean object I/O library is maintained by Peter Achten. Martin Wierich wrote the chapters on printing (Chapter 13) and TCP (Chapter 14) and also developed the corresponding library modules.

Chapter 2

Introduction

In this chapter we give a brief overview of the main features of the object I/O library. We first discuss what the basic components are and how they can be used to construct more complex components (Section 2.1). When these elements have been constructed, they must be opened to create an actual working image on the underlying platform. Elements can be opened and closed dynamically, but it is of course also possible to change them dynamically (Section 2.2). Once we know how to construct graphical user interfaces we can start an interactive program. This is explained in Section 2.3. Finally, to wrap things up Section 2.4 presents the first complete interactive Clean object I/O program of this tutorial, the ubiquitous "Hello world!".

2.1 What are interactive objects

One way of looking at the object I/O library is to regard it as a collection of building blocks, the *interactive objects*, that the programmer can use to construct graphical user interfaces. For instance, Figure 2.1 summarises the standard set of *control* objects that can be placed in a *window* object.

All interactive objects are defined by means of algebraic data types. As an example, to define the button control element in the table of Figure 2.1 one would write:

```
button = ButtonControl "Button" []
```

The constituents of this expression are the data constructor `ButtonControl`, applied to the string `"Button"`, and an empty list `[]` of control attributes. This is defined more concisely by the library type definition of a button control element:

```
:: ButtonControl lst pst
 = ButtonControl String [ControlAttribute *(lst,pst)]
```

Note that the names of the type constructor and the data constructor are identical (`ButtonControl`). This convention is used throughout the object I/O library.

The type definition of the button control is parameterised with two type variables: *lst* and *pst*. These correspond to another fundamental characteristic of interactive objects: an interactive object can have *local state*, and also have an effect on a *process state*. The type of the local state is identified by the *lst* type parameter, while the type of the process state is identified by the *pst* type parameter. The *effect*

Control object:	What does it look like:
RadioControl	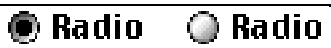
CheckControl	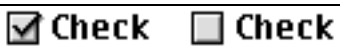
PopUpControl	
SliderControl	
TextControl	
EditControl	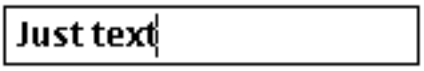
ButtonControl	
CustomButtonControl	Program defined button
CustomControl	Program defined button
CompoundControl	Program defined combination of controls

Figure 2.1: The standard set of controls.

of an interactive object that has a local state of type lst and a process state of type pst is defined by means of a function of type $(lst, pst) \rightarrow (lst, pst)$. Such a function is called a *callback function*. For most interactive objects, the callback function is an *attribute* of the object. Attributes are also defined by means of algebraic data types. For instance, among many other control attributes, one can find the callback function attribute of controls:

```
:: ControlAttribute state
=  ... | ControlFunction (state->state) | ...
```

Note that the pair of local state and process state constitute the *state* of an element. The *meaning* of attributing a control element with a callback function f is that when that element is selected by the user, and the current state is the value (l, p), then the new state will be $(f(l, p))$. In other words, a callback function defines a *state transition*.

Besides having a bag of interactive objects the object I/O library provides programmers *glue* to construct user interfaces. This glue serves two purposes: *(a)* from primitive objects one can construct new composite objects, and *(b)* it puts restrictions on what components are 'glue compatible'.

The object I/O library has one universal glue *:+:* that can be used to connect two interactive objects that operate on the same local state of type lst and process state of type pst. Its type definition is as follows:

```
:: :+:  t1 t2 lst pst
= (:+:) infixr 9 (t1 lst pst) (t2 lst pst)
```

In order to define what components are compatible to be glued *type constructor classes* are applied. The type constructor class `Controls` contains all control elements (see Figure 2.1) but also defines that only `Controls` members can be glued:

```
instance Controls RadioControl,
```

```
                CheckControl,
                PopUpControl,
                SliderControl,
                TextControl,
                EditControl,
                ButtonControl,
                CustomButtonControl,
                CustomControl,
                CompoundControl c | Controls c
                :+: c1 c2         | Controls c1 & Controls c2
```

The first nine instances are straightforward type constructor instances. The one but last instance states that if the type `c` is a `Controls` instance, then the type expression (`CompoundControl c`) is also a `Controls` instance. The last line of the list states that if expressions `e1` and `e2` have types `c1` and `c2` which are `Controls` instances, then the expression (`e1:+:e2`) has type (`(:+:) c1 c2`) and is also a `Controls` instance. Let `button` and `text` below define a button control and a text control respectively:

```
button = ButtonControl "Button"    []
text   = TextControl   "Just text" []
```

then the following expressions are all legal `Controls` instances: `button :+: button`, `button :+: text`, `text :+: button`, and `text :+: text`.

The collection of interactive objects that is supported by the object I/O library is ordered in four categories, called *abstract devices*.

Menus: Menus provide the set of commands that are available to the user of an interactive program. A program can have an arbitrary number of menus. Menus can be hierarchical, i.e. they can contain menus (*sub menus*) which can contain sub menus as well. Menu items correspond with the menu commands of the program.

Windows: Windows provide the primary interface element to the user of a program. Windows can either be *dialogues* or general purpose *windows*. Windows and dialogues can contain arbitrary collections of controls. Analogous to menus, these control collections can be hierarchical, i.e. they can contain collections of controls (*compound controls*), and so on.

Timers: Timers are used by a program to be able to *softly* synchronise actions. A timer basically triggers a callback function every passing of a given time interval.

Receivers: Receivers are the basic components that interactive objects can use to communicate messages in a flexible way.

2.2 How to manage running interactive objects

In the previous section we have had a glimpse of how to define (compositions of) interactive objects. In the object I/O library every interactive object can be *created* and *destroyed* dynamically, but we prefer to call this *opening* and *closing* which sounds more peacefully. Once an interactive object is running the program will need to *modify* it in several ways. Examples are to enable and disable menu elements,

and change the content of a window to reflect the state of the program, and so on. Closing an element is the ultimate modification of a running interactive object. So interactive objects have a *life-cycle* which consists of three consecutive phases: *opening*, *modification*, and *closing*. Below we discuss these phases.

2.2.1 Opening of interactive objects

For each abstract device a function is defined that will open an instance, given a definition. Again type constructor classes are used to control what elements are proper instances of each abstract device. As an example, dialogues are defined by the following type definition:

```
:: Dialog c lst pst
=  Dialog Title (c lst pst) [WindowAttribute *(lst,pst)]
```

Callback functions are state transition functions of type $(lst, pst) \rightarrow (lst, pst)$, so the window attribute type constructor is parameterised with (lst, pst).

The type constructor class `Dialogs` fixes the instances of dialogues. A (`Dialog c`) is a proper instance of this class, provided that `c` is a proper instance of the `Controls` type constructor class.

```
class Dialogs ddef where
  openDialog :: .ls (ddef .ls (PSt .l .p)) (PSt .l .p)
                -> (ErrorReport,PSt .l .p)

instance Dialogs (Dialog c) | Controls c
```

In Subsection 2.2.2 we will look more closely at the process state argument `PSt`.

Of each abstract device, the open function maps the definition of an abstract device instance (a value parameterised with callback functions) to a concrete 'physical' graphical user interface element that can be modified by the user (in case of windows, dialogues, and menus) or by the program (in case of all abstract devices). As an example we can open a very small and not very useful dialogue with the following expression:

```
(error,new_process_state)
   = openDialog
        my_local_state
        (Dialog "" (TextControl "Hello world!" []) [])
        the_process_state
```

which will have the following effect (in case no error occurs):

2.2.2 Modification of interactive objects

Once an interactive object has been opened and is in its running phase, it can be modified by the user and the program. For this purpose a running interactive object must be *stored* somewhere, and it must be *identified* by the program. Every running interactive object is stored in the *I/O state* of a program. In Section 2.1 we explained that every interactive object has access to a *local* state and a *process* state. The process state is a structured value defined by means of a record:

```
:: *PSt l p             // The process state record type
 = {  ls :: l           // The local  process state
   ,  ps :: p           // The public process state
   ,  io :: *IOSt l p   // The I/O state
   }
:: *IOSt l p            // IOSt is an abstract data type
```

A value of abstract type `IOSt` is created specifically for each interactive process by the object I/O system. In this special value the state of every running interactive object is stored. (The purpose of the local process state and public process state record fields `ls` and `ps` will be discussed in Chapter 11.)

The modication operations require the `IOSt` value (although some functions such as the abstract device open functions require the `PSt` record). But this is not sufficient because in general the `IOSt` will contain an arbitrary number of windows, dialogues, menus, timers, and receivers. So one has to specify which particular interactive object one intends to modify. For this purpose interactive objects (and their component structures) can be identified by means of *Ids*. An `Id` is an abstract type that is generated by the object I/O system. An interactive object is identified by means of a specific `Id` by adding this `Id` to the *attribute list* of the corresponding object definition (see Section 2.1). As an example, for controls the corresponding attribute is:

```
:: ControlAttribute state
=  ... | ControlId Id | ...
```

The major part of the object I/O library defines the modification functions that the programmer can use to modify a running interactive object.

As an example suppose we want to change the content of the text control of the "Hello world!" example to "Goodbye world!". To do this, we need to identify the text control. A control is identified uniquely by the `Id` of its parent window or dialogue and its personal `Id`. So the definition of both the "Hello world" dialogue as the text control need to be parameterised with an `Id` attribute. Let `dialogid` be the `Id` of the dialogue, and `textid` the `Id` of the text control. Then the adapted definition of opening the "Hello world!" dialogue is:

```
(error,new_process_state)
 = openDialog
    my_local_state
    (Dialog ""
        (TextControl "Hello world!" [ControlId textid])
        [WindowId dialogid]
    )
    the_process_state
```

```
changeText process_state=:{io}
 = {process_state & io=setWindow dialogid
                       (setControlTexts [(textid,"Goodbye world!")])
                       io
   }
```

2.2.3 Closing of interactive objects

As explained in the previous two subsections, once an interactive object has been opened it is in a running state and it will remain in that state until it is explicitly closed by the program. Note that although it may seem to the user that he is able to close a window, it is actually the program that responds to a *user request* to really close a window. For all interactive objects there are close operations. A close operation will remove the 'physical' graphical user interface element and free system resources that were required to operate the interactive object properly.

2.3 How to start an interactive program

The starting point of every Clean program (interactive or not) is the `Start` function. The essence of an interactive program is that it is a function that can change the world. So for interactive programs the `Start` function must have type `*World` → `*World`. The typical appearance of an interactive program looks something like this:

```
Start :: *World -> *World
Start world
   = ... world
```

In the previous sections we have seen how to define interactive objects and get them running. The abstract device open functions require a `PSt` value, containing an `IOSt` value. These are created by the object I/O system using the function `startIO`.

```
startIO :: !.l !.p !(ProcessInit (PSt .l .p))
              ![ProcessAttribute (PSt .l .p)] !*World -> *World

:: ProcessInit pst :== [IdFun pst]
:: IdFun       pst :== pst->pst
```

The details of this function will be explained in Chapter 11. Briefly, it creates an *initial* process state given the initial local and public process state components of type l and p. In particular this initial process state will contain a tailor-made `IOSt` value. In this way one switches from the world environment to the process state environment, and the interactive program can be initialised. This is done by the initialisation function (of type `ProcessInit`).

After initialisation the interactive process will be evaluated. Evaluation means that the object I/O system takes care that the proper process state transition functions will be applied that are specified by the program.

The evaluation of an interactive process terminates as soon as it becomes *closed*. The only way for an interactive process to be closed is by means of the process modification function `closeProcess`:

```
closeProcess :: (PSt .l .p) -> PSt .l .p
```

This function closes all currently running interactive objects of the interactive process and returns a process state that contains an *empty* `IOSt` value. All modification operations have no effect when applied to an empty `IOSt` value. Note that if an interactive program does not close itself it will run on forever.

2.4 My first Clean object I/O program

To complete the introduction we present the Clean object I/O version of the well known "Hello world!" program. Here it is:

```
module hello

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates a dialog that displays "Hello world!" text.
//  ********************************************************************************

import StdEnv, StdIO

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    = startIO NoState NoState [initialise] [] world
where
    initialise process
        # (error,process)   = openDialog NoState hello process
        | error<>NoError    = closeProcess process
        | otherwise         = process

    hello   = Dialog ""
                (   TextControl "Hello world!" []
                )
                [   WindowClose (noLS closeProcess)
                ]
```

This program creates an interactive process which opens the same dialogue as shown earlier in Section 2.2.1. The singleton type `NoState` is applied to state that this program has no interesting local and public process state components (the first two arguments of `startIO`), and also no interesting local dialogue state (the first argument of `openDialog`).

The dialogue has one attribute, the callback function that should be applied in case the user wants to close the dialogue. In this case closing the dialogue will close the "Hello world!" program. So the callback function can simply be `closeProcess`. However, the type of a callback function of an interactive object also operates on a local state (which is `NoState` in this case). To conveniently transform a function of type $(a \to b)$ into a function of type $(c, a) \to (c, b)$, the library function `noLS ::` $(a \to b)\ (c, a) \to (c, b)$ is used.

Chapter 3

Global structure of the object I/O library

The *application programmer's interface* of the Clean object I/O library currently consists of *fourty* modules. All corresponding definition modules can be found in Appendix A. These modules contain everything you need to create interactive Clean programs with. *No other modules and no other symbols should be imported from the object I/O library.* Violation of this rule can result in error-prone applications at worst and non portability at least.

In this chapter the module structure of the API is discussed. This will help you to find your way quickly in the object I/O library. As a global naming convention, all definition modules have the prefix `Std`. The module `StdIO` is a convenience module that collects all modules that you normally need for graphical user interface programs. The module `StdTCP` is a convenience module that collects all modules required for TCP programming. The fourty modules can be divided roughly into five major categories: *the abstract devices*, *interactive processes*, *drawing*, *channels* and *general*. Below they will be handled in the same order.

3.1 Abstract devices

Abstract devices have been introduced in Section 2.1 (page 13). These are the *menu*, *window*, *timer*, and *receiver* device. These devices occupy most of the modules and type definitions of the object I/O library. The following naming conventions have been employed for these modules:

- The names of the modules that contain type definitions to define abstract device instances have postfix `Def`. So *menu* definitions can be found in the module `Std`*Menu*`Def`, *receiver* definitions can be found in the module `Std`*Receiver*`Def`, and so on. Although *controls* are not an abstract device, a definition module also exists for *controls*, namely `Std`*Control*`Def`.

- Abstract device instances that consist of elements use type constructor classes to enumerate their elements. The corresponding type constructor classes and standard instances are defined in the modules which names end with `Class`. So *menu elements* can be found in the module `Std`*MenuElement*`Class`. The *controls* can be found in `Std`*Control*`Class`.

- Receivers are non standard elements of some abstract device instances. Given the name *Object* of the parent device instance, you can find the type constructor class instance declarations in the modules which names are formed like `Std`*Object*`Receiver`. So, receiver instances of timers can be found in `Std`*Timer*`Receiver`.
- The operations on an object *Object* can be found in the module named `Std`*Object*. So *window* operations can be found in the module `Std`*Window*, *menu element* operations can be found in the module `Std`*MenuElement*, and so on.

Below we discuss the module structure of each of the abstract devices.

3.1.1 Menus

The API for menus and menu elements consists of six modules that are related as schematised in Figure 3.1.

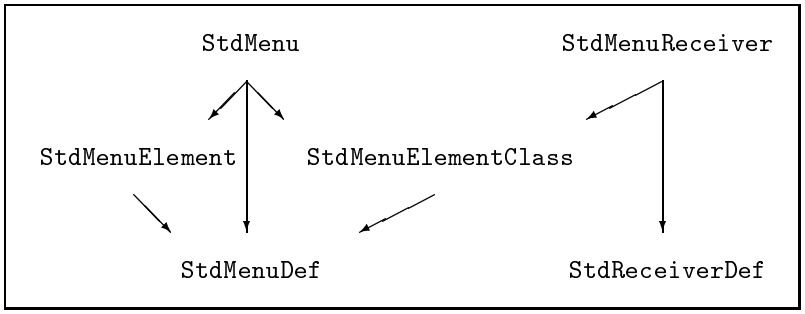

Figure 3.1: The module structure of menus.

3.1.2 Windows

The API for windows, dialogues, and controls consists of eight modules that are related as schematised in Figure 3.2.

3.1.3 Timers

The API for timers consists of six modules that are related as schematised in Figure 3.3.

3.1.4 Receivers

The API for receivers consists of two modules that are related as schematised in Figure 3.4.

3.2 Interactive processes

As stated in Section 2.3, every interactive program has to be opened as an interactive process. The API for interactive processes consists of two modules that are related as schematised in Figure 3.5.

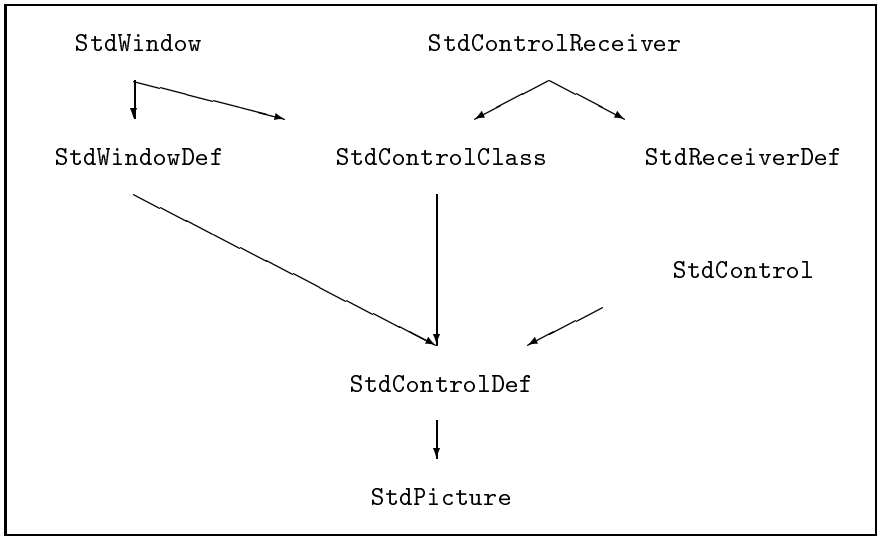

Figure 3.2: The module structure of windows.

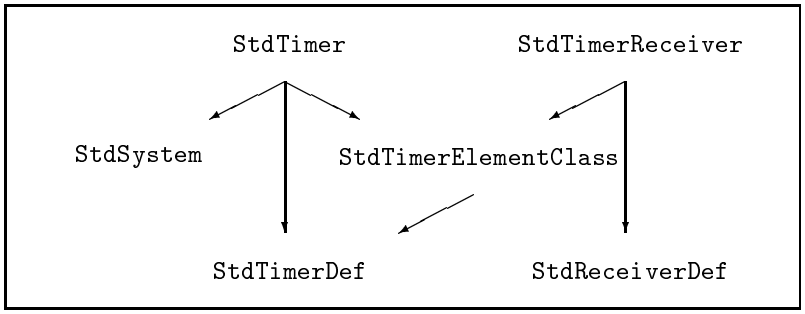

Figure 3.3: The module structure of timers.

3.3 Drawing

In graphical user interface applications graphics play an important role. Virtually every interface object has a visual representation that is drawn by the underlying platform. Drawing operations will be required by most applications to give the user visual feedback on the current documents that are being manipulated or the status of controls. The manipulation of text is also an issue in drawing information. Text can be presented in very different fonts, sizes, and variations. Drawing on screen and printing on paper is very analogous. With some care rendering can be abstracted from the output device.

The API for drawing consists of five modules that are related as schematised in Figure 3.6.

3.4 Channels

For communication via a network *channels* play an important role.

The API for using channels consists of five modules that are related as schematised

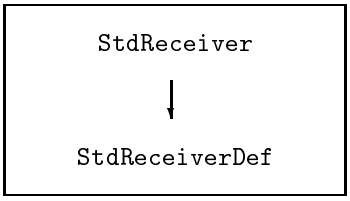

Figure 3.4: The module structure of receivers.

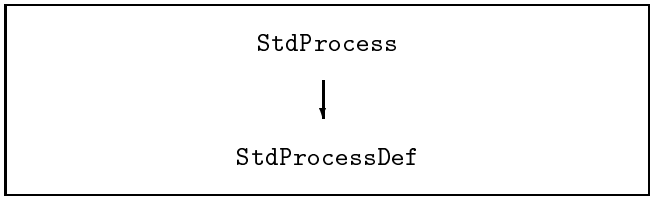

Figure 3.5: The module structure of interactive processes.

in Figure 3.7.

3.5 General

Finally, there are a number of modules that are less easily categorised. These are the following eight modules:

`StdClipboard`**:** In this module *clipboard* operations are defined. Clipboard operations are defined in more detail in Chapter 12.

`StdFileSelect`**:** In this module two functions are defined by which a user can open platform dependent directory browsing dialogues.

`StdId`**:** In this module the *identification* value generating functions are defined. This is discussed in more detail in Chapter 4.

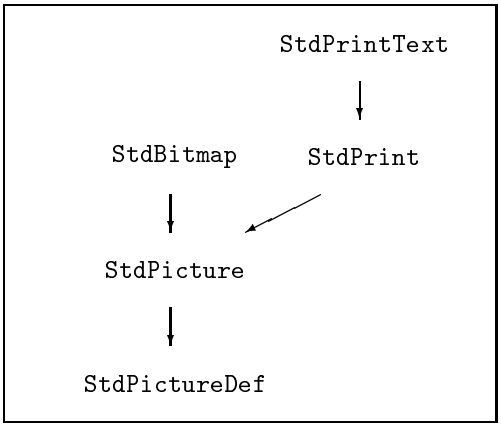

Figure 3.6: The module structure of drawing operations.

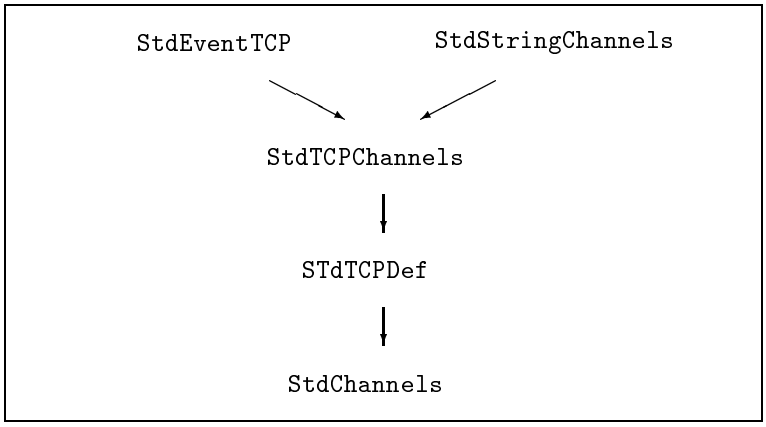

Figure 3.7: The module structure for channels.

`StdIOBasic` **and** `StdIOCommon`**:** These modules provide a lot of type definitions and functions that are needed by many of the abstract device modules introduced in Section 3.1.

`StdMaybe`**:** In this module a type is introduced that is very useful for providing optional results and optional arguments. It is used by many operations in the abstract device modules.

`StdPSt`**:** In this module operations are collected on the process state that are not related to any abstract device. It contains several type constructor class instance declarations for file and time access. Other frequently used functions are the 'lifting' functions defined on the process state. With these lifting functions one can easily transform for instance an `IOSt` transition function to a `PSt` transition function.

`StdTime`**:** In this module some time access operations can be found that can be used independently of the timer device.

Chapter 4

Object identification

Before we can really delve into the details of the object I/O library we first need to learn how to identify running interactive objects. As we have briefly discussed in Section 2.2.2, all objects can have an *identification attribute*. An identification attribute is a value of type `Id`. Because attributes are optional, the programmer is not forced to provide a value for this attribute. Objects without identification attribute can not be modified at run-time. If the program needs to modify an interactive object, it must have been provided with an identification attribute.

The type `Id` is an abstract data type, and you can import it via the module `StdId`. *All* `Id`s are generated by the system. The type constructor class `Ids` defines the creation functions. `Ids` can be created from the `World` environment and the `IOSt` environment. Every new `Id` taken from these environments is guaranteed to be fresh with respect to the other `Ids` generated by any of these functions.

```
::  Id

class Ids env where
    openId   ::      !*env -> (! Id,       !*env)
    openIds  :: !Int !*env -> (![Id],      !*env)

    openRId  ::      !*env -> (! RId m,    !*env)
    openRIds :: !Int !*env -> (![RId m],   !*env)

    openR2Id ::      !*env -> (! R2Id m r, !*env)
    openR2Ids:: !Int !*env -> (![R2Id m r],!*env)

instance Ids World
instance Ids (IOSt .l .p)
```

As the `Ids` class definition shows, `Id` values are not the only identification values that are used in the object I/O system. Two special kinds of identification values of type (`RId m`) and (`R2Id m r`) are used to identify receivers that respond to messages of type `m` in the first case, and in addition respond with a message of type `r` in the second case. Receivers are discussed in Chapter 10.

The purpose of having `Ids` is to unambiguously identify running interactive objects. When assigning `Ids` to interactive objects, a program must comply to the following `Id` *assignment rules*:

- Within one abstract device element instance, the `Ids` assigned to its elements

must be unique.

- Within one abstract device, the `Id`s assigned to the abstract device element instances must be unique.

For instance, the `Id`s assigned to windows must all be unique, and inside a window the `Id`s assigned to its controls must be unique. But for two different windows, the `Id`s assigned to the controls may overlap. Also the `Id`s assigned to other abstract device element instances such as menus and timers may overlap.

For `RId`s and `R2Id`s the assignment rules are more strict: at all times when a program is running, the `R(2)Id`s assigned to open receivers must be unique. The reason for this is explained in Chapter 10.

The abstract device open functions check whether the interactive object definition argument is valid with respect to the `Id` assignment rules. If this is not the case, the `ErrorReport` alternative `ErrorIdsInUse` is returned, and the interactive object will not be opened. If the `Id` assignment rules were not violated, and no other error occured, then the alternative `NoError` is returned.

Chapter 5

Drawing

In a graphical user interface most of the information that is presented to users is drawn, from the shape of windows, dialogues, controls, but also handling of text. This central issue is handled in this chapter. All drawing functions require an environment of type `*Picture`. A `Picture` is created for each interactive object that can be drawn in. The life-cycle of a `Picture` environment is equal to the life-cycle of its parent object. Also for printing a `Picture` environment is used (printing is discussed in detail in Chapter 13).

A `Picture` defines a coordinate system for drawing operations. Increasing x-axis coordinates run from left to right. Increasing y-axis coordinates run from top to bottom. The range for both axes is the full Clean `Int`eger range.

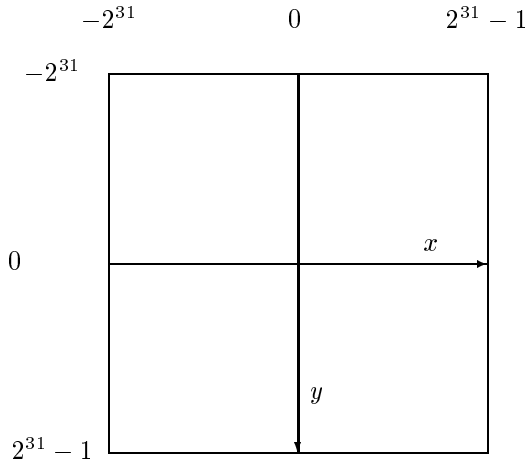

Figure 5.1: Picture coordinates.

Drawing operations on a `Picture` use the coordinate system to define where objects should be drawn. The objects themselves are made up of the pixels, lying between the coordinates. Figure 5.2 zooms in on the coordinate system from zero and increasing. The pixels on x-coordinates 0, 5, 10,... and y-coordinates 0, 5, 10,... are displayed.

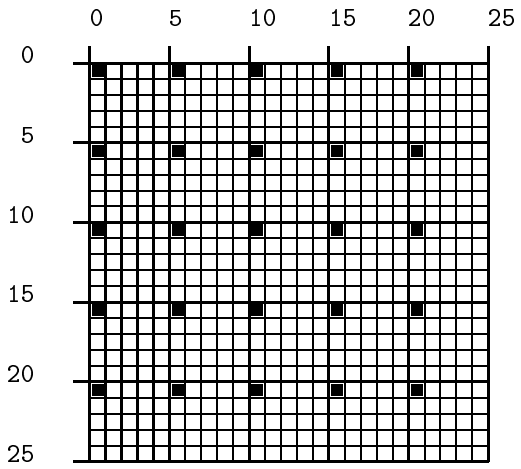

Figure 5.2: Picture coordinates and pixels

5.1 Picture attributes

Like most other interactive objects in the object I/O library, pictures have attributes. These are the following (they can be found in module `StdPicture`):

```
:: PictureAttribute
   =   PicturePenSize   Int
   |   PicturePenPos    Point
   |   PicturePenColour Colour
   |   PicturePenFont   Font
```

`PicturePenSize`**:** This attribute defines the width and height of the drawing pen. The default value is 1, which means that drawing a point will fill an area of 1 pixel wide and 1 pixel high. Negative or zero values are always set to 1.

`PicturePenPos`**:** This attribute determines the current position of the drawing pen. Its default value is `zero` (both fields are zero). The type definition of `Point` is:

```
:: Point = {x::!Int,y::!Int}
```

`PicturePenColour`**:** This attribute determines the colour that is used when drawing any element. The `Colour` data type is defined in module `StdPictureDef`:

```
:: Colour
   =   Black | DarkGrey | Grey | LightGrey | White
   |   Red   | Green    | Blue
   |   Cyan  | Magenta  | Yellow
   |   RGB RGBColour
:: RGBColour
   =   {r::!Int, g::!Int, b::!Int}
```

A colour can range between black and white (first five alternatives defining 100%, 75%, 50%, 25%, and 0% blackness), be one of red, green, blue, be one of

cyan, magenta, yellow, or some custom defined combination of red, green, blue components. Currently the library does not support colour tables or palette management operations, so the use of RGB colours tends to be speculative.

`PicturePenFont:` This attribute sets the current font that will be used when drawing text. Drawing text is not affected by the current width and height of the pen.

For each picture attribute there are functions to get and set them individually.

5.2 Drawing classes

The drawing operations are divided into three groups, ordered by means of type constructor classes:

Drawables draws (`draw` and `drawAt`) or erases (`undraw` and `undrawAt`) its instances. These are characters, strings, vectors, ovals, curves, boxes, rectangles, polygons, and bitmaps.

```
class Drawables figure where
    draw    ::        !figure !*Picture -> *Picture
    drawAt  :: !Point !figure !*Picture -> *Picture
    undraw  ::        !figure !*Picture -> *Picture
    undrawAt:: !Point !figure !*Picture -> *Picture
```

Fillables fills (`fill` and `fillAt`) or erases (`unfill` and `unfillAt`) its instances. These are ovals, curves, boxes, rectangles, and polygons.

```
class Fillables figure where
    fill    ::        !figure !*Picture -> *Picture
    fillAt  :: !Point !figure !*Picture -> *Picture
    unfill  ::        !figure !*Picture -> *Picture
    unfillAt:: !Point !figure !*Picture -> *Picture
```

Hilites fills the interior of its instances in such a way that the current picture content remains visible. Its instances are boxes and rectangles.

```
class Hilites figure where
    hilite  ::        !figure !*Picture -> *Picture
    hiliteAt:: !Point !figure !*Picture -> *Picture
```

Each of these type constructor classes allows its elements to be drawn at the *current* pen position or at an *absolute* pen position. Because of this reason the data type definition of most of these elements do not specify their location. Exceptions are rectangles, lines, and points.

5.3 Font and text handling

When working with text you frequently will want to know the dimensions of the text for layout purposes or simply to calculate the size of an element containing that text. The dimensions of a piece of text depends on two parameters:

- The *font* is an abstract value that describes the shape of a text. The usual way to identify a font is by its name, point size, and style variations.
- The *drawing environment* (`Picture`) determines the actual size in terms of a resolution dependent unit. The resolution of the screen is usually a lot smaller than the resolution of a laser writer.

Because there is a great variance of available fonts and drawing environments per machine writing a program that handles fonts properly requires some care. Font operations are ordered in three groups.

The first group of font operations return information about the currently available fonts. The function `getFontNames` returns a list of the names of all available fonts. Given an element of this list, the functions `getFontStyles` and `getFontSizes` return for that particular font the available style variations and sizes. Because in modern font management systems for many fonts no restriction exists on the size, the function `getFontSizes` is also parameterised with two bounding size arguments. Their type definitions are:

```
getFontNames ::          !*Picture ->(![FontName], !*Picture)
getFontStyles::          !FontName
                         !*Picture ->(![FontStyle],!*Picture)
getFontSizes ::!Int !Int !FontName
                         !*Picture ->(![FontSize], !*Picture)
```

The second group of font operations opens fonts. The function `openFont` creates a value of type `Font` given a font definition. A font definition is a record of type `FontDef` and is defined as follows:

```
:: FontDef
   =   { fName   :: !FontName
       , fStyles :: ![FontStyle]
       , fSize   :: !FontSize
       }
```

Because there are so many different font systems and style variations both font names and style variations are of type `String`. If you want to be sure that you are selecting an existing font use the functions of the first group. In any case, if the font definition argument of `openFont` does not correspond with an installed font, then it returns a `False` Boolean and a font value that is supposed to be the closest matching font available. If it succeeds to find a matching font it will return a `True` Boolean and that font. The type of `openFont` is:

```
openFont :: !FontDef !*Picture -> (!(!Bool,!Font),!*Picture)
```

There are two functions that open the font that is used by default in a document window (`openDefaultFont`) and the font that is used by the system for dialogs, controls, window titles and so on (`openDialogFont`). These fonts are of course always available. Their types are:

```
openDefaultFont :: !*Picture -> (!Font,!*Picture)
openDialogFont  :: !*Picture -> (!Font,!*Picture)
```

The last group of font operations is related to font and text metrics. Font metrics are retrieved by the functions `getPenFontMetrics` and `getFontMetrics`. The first retrieves the metrics of the current pen `Font`, while the second retrieves the metrics of the argument `Font`.

```
getPenFontMetrics::       !*Picture -> (!FontMetrics,!*Picture)
getFontMetrics   :: !Font !*Picture -> (!FontMetrics,!*Picture)
```

The metrics of a font consists of three height related values related (*leading*, *ascent*, and *descent*), illustrated in Figure 5.3.

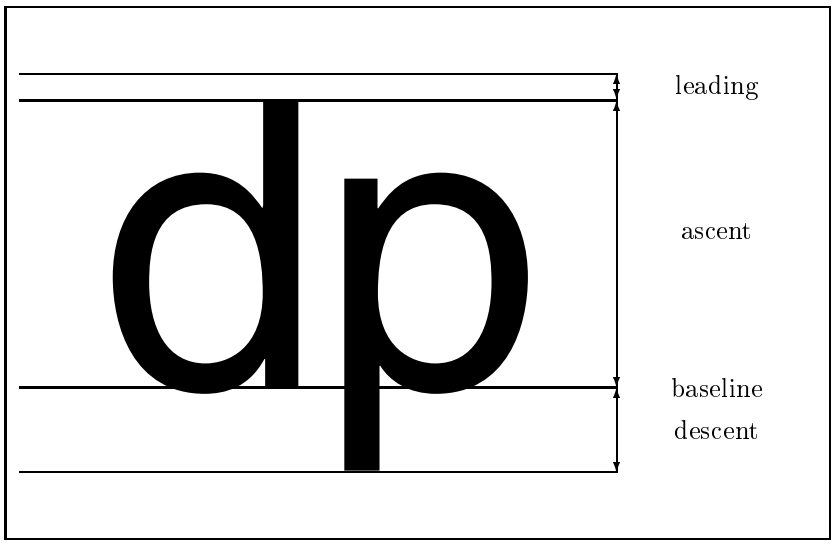

Figure 5.3: Font metrics.

The width related value (*max. width*) is the maximum width of all characters in that particular font. For *non-proportional* fonts such as `Courier` this implies that the width of all characters is identical, so "`iii`" is just as wide as "`mmm`". For *proportional* fonts such as this text the width of characters can vary a lot. Compare for instance the width of the text "iii" with "mmm". The metrics of a font are collected in a record of type `FontMetrics`:

```
:: FontMetrics
   =   { fAscent  :: !Int  // The ascent  of the font
       , fDescent :: !Int  // The descent of the font
       , fLeading :: !Int  // The leading of the font
       , fMaxWidth:: !Int  // The max. width of the font
       }
```

The final functions in the last group of font operations are used for calculating the width of a (list of) character(s), and a (list of) string(s). As with the font metrics functions, these functions also come in pairs, one using the current pen `Font`, and the other using the argument `Font`.

```
getPenFontCharWidth   ::    ! Char    !*Picture ->(! Int, !*Picture)
getPenFontCharWidths  ::    ![Char]   !*Picture ->(![Int],!*Picture)
getPenFontStringWidth ::    ! String  !*Picture ->(! Int, !*Picture)
```

```
getPenFontStringWidths::    ![String] !*Picture ->(![Int],!*Picture)

getFontCharWidth   :: !Font ! Char    !*Picture ->(! Int, !*Picture)
getFontCharWidths  :: !Font ![Char]   !*Picture ->(![Int],!*Picture)
getFontStringWidth :: !Font ! String  !*Picture ->(! Int, !*Picture)
getFontStringWidths:: !Font ![String] !*Picture ->(![Int],!*Picture)
```

There is a subtle difference in calculating the width of *one character* versus the width of *one string*. The width of a character is determined by the character only. The width of a string can depend on the order of the characters it contains. A font system can take advantage of the fact that some adjacent characters can be placed more closely together to obtain a better looking result when drawing the string. This is called *kerning*. In the object I/O system, the programmer can rely on the fact that if a piece of text is drawn character by character then the character width function returns the correct width of the drawn character. If a piece of text is drawn by using a string, then the string width function returns the correct width of the drawn string.

Because there is wide variety of fonts available the `StdPictureDef` module provides a number of macros that help you make a program less dependent on the set of available fonts. The following macros provide a number of font definitions that are guaranteed to be available on the platform:

Font macros:	**Example:**
`SerifFontDef`	Garamond, Times
`SansSerifFontDef`	Helvetica
`SmallFontDef`	"This is a small text"
`NonProportionalFontDef`	`Courier`
`SymbolFontDef`	$\forall\ \exists\ \alpha\ \beta \Leftarrow \Rightarrow$

The following macros provide a number of standard font variations that are guaranteed to be available on the platform:

Style macros:	**Example:**
`ItalicsStyle`	*Madam, I'm Adam*
`BoldStyle`	**Madam, I'm Adam**
`UnderlinedStyle`	Madam, I'm Adam

5.4 Examples

In this section we give small examples of all of the drawable elements. Each of the examples is defined by a drawing function `example` of type `*Picture -> *Picture`. To actually get something on the screen one can use the following drawing framework program `drawingframe`. Figure 5.4 gives a snapshot of the framework program.

```
module drawingframe

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program defines a framework in which one can test drawing functions.
//  The program relies on a function, example, of type *Picture -> *Picture.
//  ********************************************************************************
```

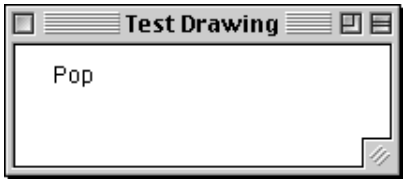

Figure 5.4: The framework window.

```
import StdEnv,StdIO

::  NoState = NoState

Start :: *World -> *World
Start world
    = startIO NoState
              NoState
              [snd o openWindow NoState testwindow]
              [ProcessClose closeProcess]
              world

testwindow
    = Window "Test Drawing" NilLS
        [   WindowSize       size
        ,   WindowClose      (noLS closeProcess)
        ,   WindowLook       (\_ _->example)
        ,   WindowResize
        ,   WindowViewDomain {corner1=origin,corner2=maxdomain}
        ]
where
    size        = {w=200,h=50}
    origin      = {x=(-20),y=(-20)}
    maxdomain   = {x=origin.x+size.w,y=origin.y+size.h}

// Here, example draws the string "Pop" at zero
example = drawAt zero "Pop"
```

5.4.1 Pen size and position

Given a pen position `{x,y}`, drawing a point, line, always occurs to the right and below the pen position. Figure 5.5 illustrates these cases: from left to right the following `example` functions are applied respectively:

```
example = drawPointAt zero o (setPenSize 1)
example = drawPointAt zero o (setPenSize 2)
example = drawPointAt zero o (setPenSize 3)
```

5.4.2 Drawing lines

The shape of a line is influenced by the `PicturePenSize` attribute in the same way as the shape of points are changed. Figure 5.6 shows the result of drawing lines with pen sizes 1, 2, and 3 respectively:

```
example = drawLine zero {x=5,y=5} o (setPenSize 1)
example = drawLine zero {x=5,y=5} o (setPenSize 2)
example = drawLine zero {x=5,y=5} o (setPenSize 3)
```

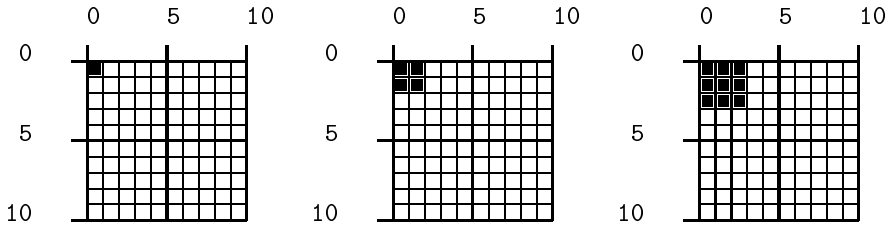

Figure 5.5: Drawing a point at zero with different pen sizes.

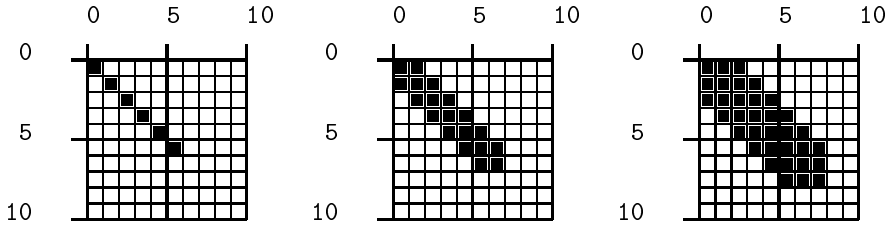

Figure 5.6: Drawing a line from zero to {x=5,y=5} with different pen sizes.

There are several ways to draw lines. The function `drawLineTo` draws a line from the current pen position to the argument point. If these happen to be equal, then the result is the same as `drawPointAt` with the same argument. The new pen position is the same as the target point. The function `drawLine` draws a line from the first argument point to the second argument point without changing the pen position.

Lines can also be drawn using the `Vector` instance functions from the `Drawables` type constructor class. The function `draw` applied to a vector `{vx,vy}` draws a line from the current pen position `{x,y}` to the point `{x=x+vx,y=y+vy}`. The function `drawAt` applied to a point `{x,y}` and a vector `{vx,vy}` draws a line from `{x,y}` to the point `{x=x+vx,y=y+vy}`.

5.4.3 Drawing text

Given a pen position `{x,y}`, drawing a piece of text (`Char` or `String`) will always draw the text using the current `PicturePenFont`. The shape of the drawn characters relies only on the font information, not on the current `PicturePenSize`. Text can be drawn in any of the available colours. The baseline of the particular font determines the position of the first character, which is drawn with its left baseline starting at `{x,y}`. Figure 5.7 shows the result of the following `example` function:

```
example = drawAt zero "Pop"
```

The new pen position is, in this case, `{x=24,y=0}`. One might expect the new pen position to be `{x=22,y=0}`, but usually the horizontal character space is also included. This facilitates drawing text character by character. But some caution should be taken. One might expect that the function

```
example = draw "p" o (draw "o") o (draw "P")
```

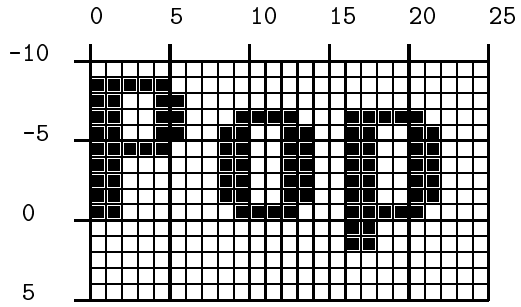

Figure 5.7: Drawing the text "Pop" at zero.

produces the same result, but this depends on the font (as explained in Section 5.3), so in general one should not assume that this is the case. The only certain way to know how much the pen position will change in case of text is by calculating the width of the *same* text, or by comparing the pen positions before and after drawing.

5.4.4 Drawing ovals

An `Oval` is a transformed unit circle defined by a horizontal radius, `oval_rx` and a vertical radius, `oval_ry`. For each point `{x,y}` on a unit circle, its corresponding point on the oval is given by `{x=x*oval_rx,y=y*oval_ry}`. The type definition of an `Oval` is:

```
:: Oval = {oval_rx::!Int, oval_ry::!Int}
```

Both radius values are always taken to be at least zero. If any of these values is negative, then zero is used instead. Ovals are drawn using the `Oval` instance functions from the `Drawables` type constructor class. The function `draw` uses the current pen position as the center of the oval. The function `drawAt` uses the argument `Point` as the center of the oval. Drawing an oval does not change the pen position. In case one of the radius values is taken to be zero drawing the oval displays nothing. Figure 5.8 shows the result of drawing three ovals at `zero` defined as follows:

```
example = drawAt zero {oval_rx=5,oval_ry=3}
example = drawAt zero {oval_rx=5,oval_ry=5}
example = drawAt zero {oval_rx=3,oval_ry=5}
```

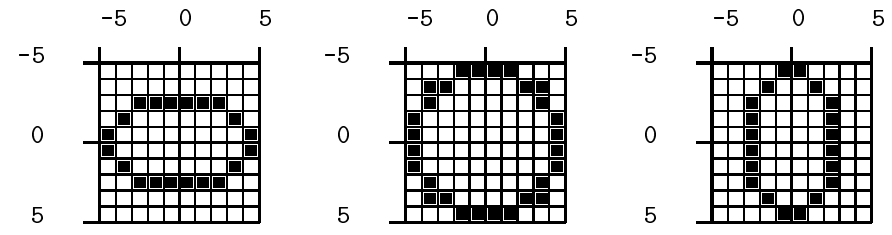

Figure 5.8: Three oval shapes drawn at zero.

The shape of an oval is also affected by the current `PicturePenSize` attribute. Increasing the pen size does not increase the outline of the oval. The only pixels that are affected are inside the oval. Figure 5.9 shows the center oval of Figure 5.8 when drawn with pen size of 1, 2, and 3.

```
example = drawAt zero {oval_rx=5,oval_ry=3} o (setPenSize 1)
example = drawAt zero {oval_rx=5,oval_ry=5} o (setPenSize 2)
example = drawAt zero {oval_rx=3,oval_ry=5} o (setPenSize 3)
```

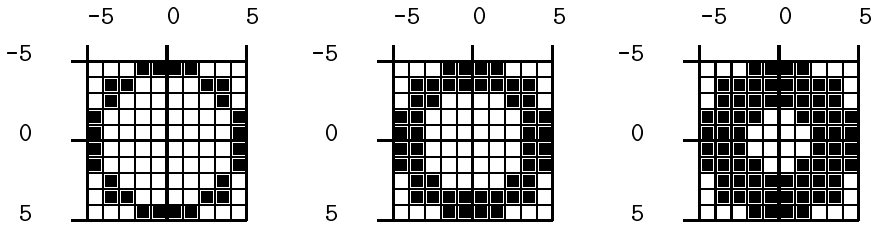

Figure 5.9: Three oval shapes drawn with increasing pen sizes.

`Ovals` are also an instance of the `Fillables` type constructor class. The function `fill` and `fillAt` *fill* rather than draw the oval. Filling an oval includes its outline and its interior. Figure 5.10 shows the same three ovals as given in Figure 5.8, but now filled.

```
example = fillAt zero {oval_rx=5,oval_ry=3}
example = fillAt zero {oval_rx=5,oval_ry=5}
example = fillAt zero {oval_rx=3,oval_ry=5}
```

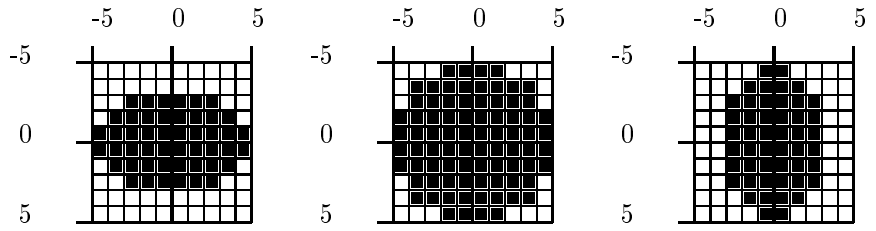

Figure 5.10: Three filled oval shapes.

5.4.5 Drawing curves

A `Curve` is a section of an `Oval`. A curve is defined by the source oval, `curve_oval`, a starting angle, `curve_from` and an ending angle, `curve_to`, both taken in radians, and the direction in which the section should be taken, `curve_clockwise` which is a Boolean value. The type definition of a `Curve` is:

```
:: Curve
   =   {   curve_oval      :: !Oval
       ,   curve_from      :: !Real
```

```
      ,   curve_to        :: !Real
      ,   curve_clockwise :: !Bool
      }
```

The start and end point of the section are again derived from the unit circle. Given an angle `alpha`, and a source oval defined by `{oval_rx, oval_ry}`, then the point on the curve (oval) corresponding with `alpha` is `{x=oval_rx*cos alpha, y=oval_ry*sin alpha}`. If `curve_clockwise` is `True` then the section is taken clockwise from the start point to the end point, otherwise it is taken counter clockwise. Figure 5.11 shows two sections of an `Oval`. In both cases the `curve_from` angle is $\frac{\pi}{6}$ and the `curve_to` angle is $\frac{3\pi}{2}$. The left section is taken counter clockwise, and the right section is taken clockwise. For your convenience, the value π is approximated in the module `StdPictureDef` by the macro `PI`.

```
example = drawAt zero { curve_oval      = oval
                      , curve_from      = PI/6.0
                      , curve_to        = 3.0*PI/2.0
                      , curve_clockwise = True
                      }
example = drawAt zero { curve_oval      = oval
                      , curve_from      = PI/6.0
                      , curve_to        = 3.0*PI/2.0
                      , curve_clockwise = False
                      }
oval    = {oval_rx=5,oval_ry=3}
```

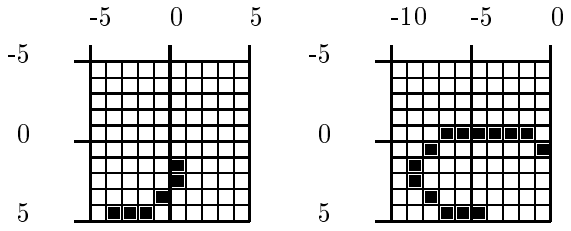

Figure 5.11: Two curves taken clockwise and counter clockwise.

Figure 5.11 not only shows the curve sections that are taken from an oval, but also what happens when these sections are drawn *at* a specific position. In both cases the curves are drawn at `zero`. The starting point, indicated by the starting angle, is determined by the current pen position in case of the `draw` function of the `Drawables` type constructor class, and is determined by the `Point` argument of the `drawAt` function of the `Drawables` type constructor class.

Drawing a `Curve` with varying `PicturePenSizes` is the same as taking the section of the corresponding `Oval` drawn with that pen size. Figure 5.12 shows three times the same curve taken but drawn with pen sizes 1,2, and 3 respectively. The source oval is the same as the one drawn in Figure 5.9. The section is taken counter clockwise from $\frac{\pi}{4}$ to $1\frac{3}{4}\pi$.

```
example = drawAt zero oval o (setPenSize 1)
example = drawAt zero oval o (setPenSize 2)
example = drawAt zero oval o (setPenSize 3)
```

```
oval    = { curve_oval      = {oval_rx=5,oval_ry=5}
          , curve_from      = PI/4.0
          , curve_to        = 1.75*PI
          , curve_clockwise = False
          }
```

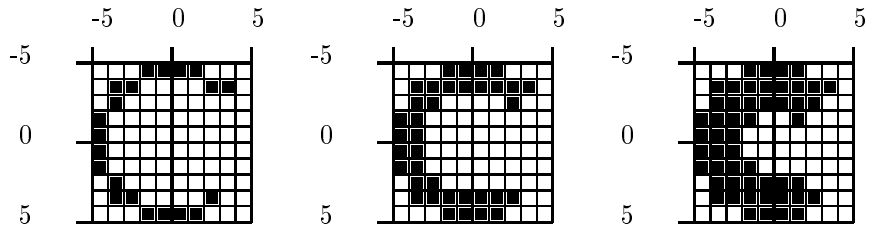

Figure 5.12: Three curves drawn with increasing pen sizes.

`Curves` are also an instance of the `Fillables` type constructor class. When filling a curve, the interior formed by the drawn curve and two lines connecting the center of the source oval and the end points of the curve is filled. Figure 5.13 shows the two curves of Figure 5.11, but now using `fill` rather than `draw`.

```
example = fillAt zero { curve_oval      = oval
                      , curve_from      = PI/6.0
                      , curve_to        = 3.0*PI/2.0
                      , curve_clockwise = True
                      }
example = fillAt zero { curve_oval      = oval
                      , curve_from      = PI/6.0
                      , curve_to        = 3.0*PI/2.0
                      , curve_clockwise = False
                      }
oval    = {oval_rx=5,oval_ry=3}
```

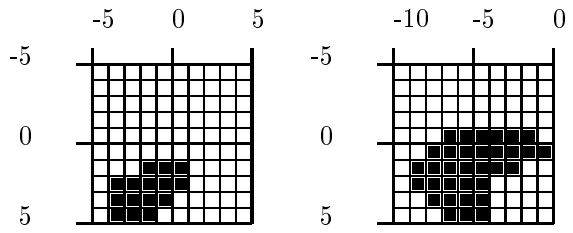

Figure 5.13: Two filled curves taken clockwise and counter clockwise.

5.4.6 Drawing rectangles

A `Rectangle` is a shape of four connected lines that is defined by two diagonally oriented corner `Points`, `corner1` and `corner2`. The type definition of a `Rectangle` is as follows:

```
:: Rectangle = {corner1::!Point, corner2::!Point}
```

The `Rectangle` type constructor is an instance of the `Drawables` type constructor class. The `drawAt` function is not very useful because it ignores its `Point` argument and proceeds as `draw`. Any two `Point`s are valid corner points of a `Rectangle`. In case a `Rectangle` has a zero width or zero height drawing that rectangle will show nothing. It does not matter in what order the two corner points are given. This is illustrated by the following `Rectangle` definitions, displayed in Figure 5.14.

```
example = draw {zero & corner2={x=10,y=6}}
example = draw {zero & corner2={x=6, y=6}}
example = draw {zero & corner1={x=10,y=6}}
```

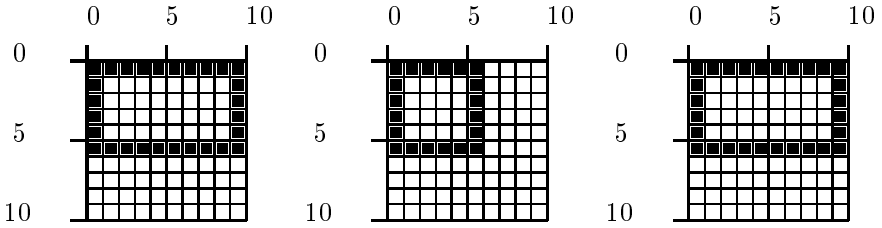

Figure 5.14: Three rectangle shapes.

The shape of a rectangle is also affected by the current `PicturePenSize` attribute. Increasing the pen size does not increase the outline of the rectangle. The only pixels that are affected are inside the rectangle. Figure 5.15 shows the `Rectangle {corner1=zero, corner2={x=10,y=10}}` when drawn with pen size 1, 2, and 3.

```
example = draw {zero & corner2={x=10,y=10}} o (setPenSize 1)
example = draw {zero & corner2={x=10,y=10}} o (setPenSize 2)
example = draw {zero & corner2={x=10,y=10}} o (setPenSize 3)
```

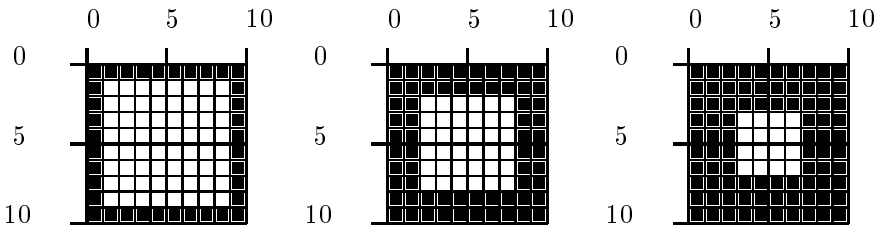

Figure 5.15: A rectangle drawn with increasing pen size.

`Rectangle`s are also an instance of the `Fillables` type constructor class. The function `fill` and `fillAt` *fill* rather than draw the rectangle. Filling a rectangle includes its outline and its interior. Figure 5.16 shows the same three rectangles as given in Figure 5.14, but now filled.

```
example = fill {zero & corner2={x=10,y=6}}
example = fill {zero & corner2={x=6, y=6}}
example = fill {zero & corner1={x=10,y=6}}
```

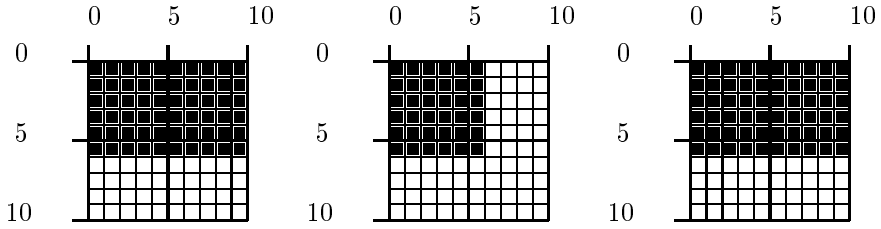

Figure 5.16: Three filled rectangles.

5.4.7 Drawing boxes

A `Box` is a `Rectangle` but without fixing a position. It is therefore only defined by a width, `box_w` and height, `box_h`. The type definition of a `Box` is:

```
:: Box = {box_w::!Int, box_h::!Int}
```

The position of a `Box` is determined by the drawing functions of the type constructor class `Drawables`. In case of `draw`, the current pen position is the base point. In case of `drawAt`, the `Point` argument is the base point. Given this base point `base={x,y}`, and a `Box {box_w,box_h}`, drawing the `Box` is the same as drawing the `Rectangle {corner1=base, corner2={x=x+box_w,y=y+box_h}}`. Any value for `box_w` or `box_h` is permitted (so also zero or negative values).

`Boxes` are drawn and filled in the same way as `Rectangles` are. So the effect of using different `PicturePenSizes` is the same as well as filling boxes.

5.4.8 Drawing polygons

A `Polygon` is an object which shape is formed by a number of `Vectors`, such as triangles, rectangles, but also more exotic shapes. The type definition of a `Polygon` is:

```
:: Polygon = {polygon_shape::![Vector]}
```

A `Polygon` is always a closed shape. A shape `polygon_shape` is closed if the following equation holds:

```
foldr (+) zero polygon_shape = zero
```

The object I/O library will always close the `polygon_shape` if this is not the case, so you don't have to worry about this. Drawing a polygon of shape `polygon_shape` is simply drawing the closed list of vectors in sequence:

```
seq (map draw polygon_shape)
```

Similar to `Boxes`, `Polygons` do not specify their location. Again, this is determined by the drawing functions of the type constructor class `Drawables`. In case of `draw`, the base point is defined by the current pen position. In case of `drawAt`, the base point is defined by the `Point` argument. Figure 5.17 shows three polygons defined by the following shapes:

```
example = drawAt zero {polygon_shape=[{vx=8,vy=0},{vx=(-4),vy=8}]}
example = drawAt zero {polygon_shape=[{vx=8,vy=0},{vx=0,   vy=8}
                                                  ,{vx=(-8),vy=0}]}
example = drawAt zero {polygon_shape=[{vx=8,vy=0},{vx=(-8),vy=8}
                                                  ,{vx=8,   vy=0}]}
```

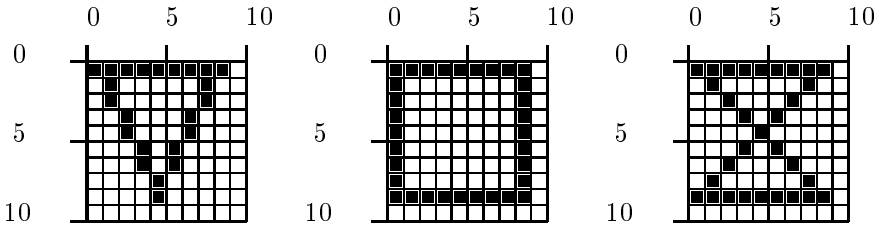

Figure 5.17: Three polygon shapes.

Because a polygon is a collection of vectors, its shape is affected by the current `PicturePenSize` attribute. Figure 5.18 shows the three polygons of Figure 5.17 drawn with pen size 2.

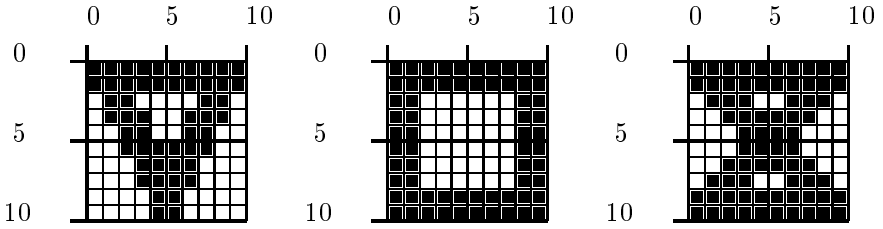

Figure 5.18: Three polygons drawn with pen size 2.

`Polygons` are also an instance of the `Fillables` type constructor class. The function `fill` and `fillAt` *fill* rather than draw the polygon. Filling a polygon includes its outline and its interior. Figure 5.19 shows the same three polygons as given in Figure 5.17, but now filled.

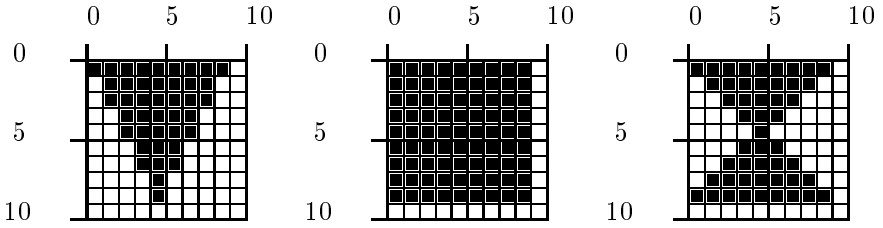

Figure 5.19: Three filled polygons.

5.4.9 Drawing bitmaps

Using the drawing operations discussed so far one can produce images that have an 'algorithmic' nature: they consist of text, lines, curves, and polygons. Not every

image can be expressed (easily) in this way (consider for instance the image in Figure 5.20). To produce more complex images *bitmaps* are very useful. A bitmap is a prefabricated image of a certain size (in pixels) that is stored in the file system.

Figure 5.20: A non algorithmic image.

You can use your favourite drawing package and create and store images as a bitmap. The file format depends on the platform. Currently the following formats are supported:

Platform:	**Format:**
Macintosh	PICT
Windows(95/NT)	BMP

The bitmap operations can be found in module `StdBitmap` (Appendix A.1). Bitmaps can be read from file using the function `openBitmap`:

```
openBitmap::!{#Char} !*env -> (!Maybe Bitmap,!*env) | FileSystem env
```

Given the full path name of a bitmap file, `openBitmap` reads the bitmap in memory. If this is successful, `(Just bitmap)` is returned. Reasons for failure are illegal file name arguments, wrong file formats, lack of heap space (in case of Windows(95/NT)), or lack of extra memory (in case of Macintosh[1]).

A `Bitmap` is an abstract data type. `Bitmaps` can be drawn in any size you like. For this purpose you can change its size by using `resizeBitmap`. The only information one can retrieve from a `Bitmap` value is its size.

```
resizeBitmap  :: !Bitmap !Size -> Bitmap
getBitmapSize :: !Bitmap       -> Size
```

`Bitmaps` are instances of the `Drawables` type constructor class. Given a current pen position `pos={x,y}` and a bitmap `bitmap` of size `{w,h}`, the functions `(draw bitmap)` and `(drawAt pos bitmap)` both place the bitmap exactly inside the rectangle `{corner1=pos, corner2={x=x+w,y=y+h}}`.

[1] In the current Macintosh implementation bitmaps are not garbage collected. This puts a restriction on the number of bitmaps that can be used inside one application. In a future version bitmaps will become garbage collected.

5.4.10 Drawing in XOR mode

For many programs it is sometimes useful to be able to temporarily draw figures over an existing drawing, and being able to remove them without affecting the source picture. Examples are flashing cursors, selection track boxes, and selection anchor points. For this purpose drawing in XOR mode is supported. The functions `appXorPicture` and `accXorPicture` handle this. Both apply their argument function to the argument picture in XOR mode. The latter function also allows its argument function to return a result.

```
appXorPicture:: !(IdFun *Picture) !*Picture -> *Picture
accXorPicture:: !(St *Picture .x) !*Picture -> (.x,!*Picture)
```

Drawing in XOR mode has the important property that drawing the same figure twice results in the same picture. Given a `picture`, and a drawing function `f`, the following equations hold:

```
          (appXorPicture f) o (appXorPicture f) = id
(snd o (accXorPicture f)) o (snd o (appXorPicture f)) = id
```

Let's explain what happens in the `Picture` when one uses XOR mode. Consider a source picture, `source`, shown left in Figure 5.21, which is a circle. Next to the source picture is the picture to be drawn in XOR mode, a fat rectangle, drawn by a drawing function `f`.

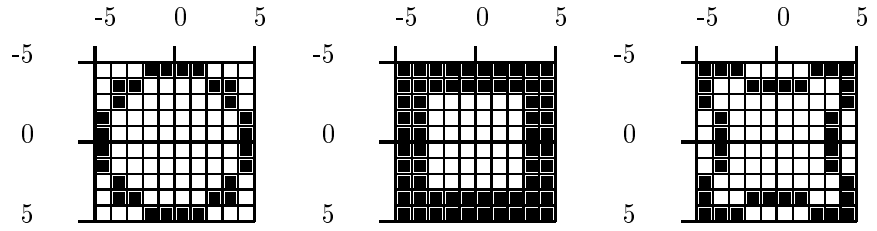

Figure 5.21: A source picture, the figure to be drawn in XOR mode, and the result.

If one interprets the white pixels of both pictures as `False`, and the black pixels of both pictures as `True` then the result of drawing `f` in XOR mode in `source` is the same as taking the Boolean exclusive or on all such interpreted pixels on the same coordinates. So all pixels that have the same colour become white (`False`), while all pixels of different colour become black (`True`). The result of this is shown in the right picture of Figure 5.21. Applying `f` once more in `xor` mode to this new picture yields a picture equal to `source`.

What happens when using more interesting colours than black and white is basically the same thing. In one way or another, the exclusive or is taken from the source picture and the drawing operations in such a way that repeating it gives the source picture again. What the colours of the 'xor-ed' picture are depends on the platform, and is not specified by the object I/O library.

5.4.11 Drawing in Hilite mode

Programs that want to indicate selections (for instance text segments in a word processor, or image components in a drawing program) can do this by drawing the

selected area in *hilite* mode. For this purpose the type constructor class `Hilites` is used. Its type definition is:

```
class Hilites figure where
   hilite   ::        !figure !*Picture -> *Picture
   hiliteAt :: !Point !figure !*Picture -> *Picture
```

The instances of `Hilites` are `Box` and `Rectangle`. The pixels that are affected by `hilite` and `hiliteAt` are the same as for `fill` and `fillAt`. Drawing in hilite mode has the same property as drawing in XOR mode that drawing the same figure twice on a source picture leaves the source picture unchanged. Given a `picture`, and a figure `figure`, the following equations hold:

```
hilite   figure (hilite   figure picture) = picture
hiliteAt figure (hiliteAt figure picture) = picture
```

The visual effect of hiliting these areas depends on the platform. On some platforms hiliting an area will change the colour of all pixels that have the background colour to a special hilite colour, ignoring all other pixels. Hiliting the area once more will revert the hilite colours back to the background colour. If a platform does not support hilite mode, the area will be drawn in XOR mode (shown in Figure 5.22). The source picture is the text "Pop" of Figure 5.7.

```
example = hilite {corner1={x=0,y=2},corner2={x=24,y=(-10)}}
        o (drawAt zero "Pop")
```

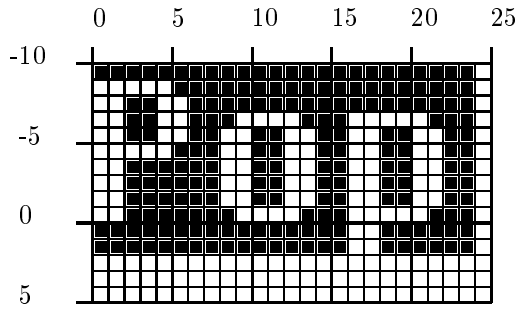

Figure 5.22: Hiliting a rectangular area using XOR mode.

5.4.12 Drawing in Clipping mode

Drawing in *clipping* mode is a powerful technique to create graphics that can not be drawn (or using much more complicated expressions) using only the drawing primitives discussed before. In clipping mode, the programmer specifies a *region* that works like a mask: drawing proceeds as described above, but only those pixels that are inside the clipping region are actually drawn.

A region is an abstract data type, `Region`. Regions are created by composing `Rectangle`s and `Polygon`s, using the type constructor class `toRegion`. When composing regions, using the list and `:^:` instance, the union of the argument regions is taken.

```
class toRegion area :: !area -> Region

::  PolygonAt
    =   {   polygon_pos :: !Point
        ,   polygon     :: !Polygon
        }

instance toRegion Rectangle
instance toRegion PolygonAt
instance toRegion [r]          | toRegion r
instance toRegion (:^: r1 r2) | toRegion r1 & toRegion r2
```

Clipping is done using the functions `appClipPicture` and `accClipPicture`:

```
appClipPicture::!Region !(IdFun *Picture) !*Picture-> *Picture
accClipPicture::!Region !(St *Picture .x) !*Picture-> (.x,!*Picture)
```

Suppose we have the following drawing function, `f`, which draws a number of horizontal lines with a result as shown in Figure 5.23:

```
f = seq [drawLine {x=0,y=y} {x=9,y=y} \\ y<-[0,2..8]]
```

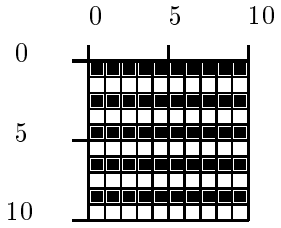

Figure 5.23: The source picture.

As clipping regions the polygons shown in Figure 5.17 are used. Figure 5.24 shows the result of the following clipping functions:

```
example = appClipPicture
            (toRegion { polygon_pos  =zero
                      , polygon_shape=[{vx=8,vy=0},{vx=(-4),vy=8}]
                      }
            ) f
example = appClipPicture
            (toRegion { polygon_pos  =zero
                      , polygon_shape=[{vx=8,vy=0},{vx=0,   vy=8}
                                                  ,{vx=(-8),vy=0}]
                      }
            ) f
example = appClipPicture
            (toRegion { polygon_pos  =zero
                      , polygon_shape=[{vx=8,vy=0},{vx=(-8),vy=8}
                                                  ,{vx=8,   vy=0}]
                      }
            ) f
```

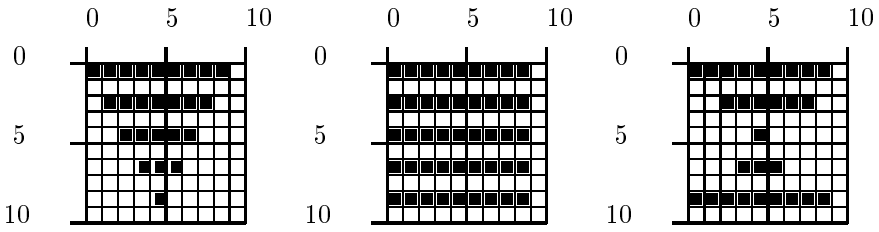

Figure 5.24: The clipped source picture.

Chapter 6

Windows and dialogues

Windows and *dialogues* are the major top level interactive objects of the object I/O library. In some aspects they are very similar. For instance, they can contain the same set of *controls*, and virtually all operations on windows and dialogues are similar. Dialogues differ from windows because they usually offer a special, enhanced, user interface to users. Another difference is that dialogues can be opened *modally*. In this mode the user can be forced by the program to handle the dialogue completely before continuing with the program. To emphasize the similarities of windows and dialogues, their algebraic type definitions are almost identical (these can be found in module `StdWindowDef`, Appendix A.39):

```
:: Window c lst pst
 = Window Title (c lst pst) [WindowAttribute *(lst,pst)]
:: Dialog c lst pst
 = Dialog Title (c lst pst) [WindowAttribute *(lst,pst)]
```

Before delving into details, we first introduce basic terminology for windows and dialogues in Section 6.1 and discuss the `WindowAttribute` alternatives in Section 6.2. Opening and closing of windows and dialogues is handled in Section 6.3. We then discuss the ways to handle the components windows and dialogues are made of in the Sections 6.4, 6.5, and 6.6. Handling keyboard and mouse input is presented in Section 6.7. Finally, this chapter is concluded with a treatment on modal dialogues in Section 6.8.

6.1 Basic terminology

The main purpose of a *window* is to present to the user a view on a document, graphically represented as an object of type `Picture`. `Picture`s have been discussed in Chapter 5. By using the mouse and keyboard, the user can manipulate the document. Controls in a window can add further manipulation functionality.

The main purpose of a *dialogue* is to present to the user a structured way of passing information to perform actions. This structured communication is realised by means of controls.

6.1.1 Anatomy of windows and dialogues

Although from an application user's perspective windows and dialogues appear to be 'solid' objects (Figure 6.1) it is illustrative to have a look at a window from a different perspective.

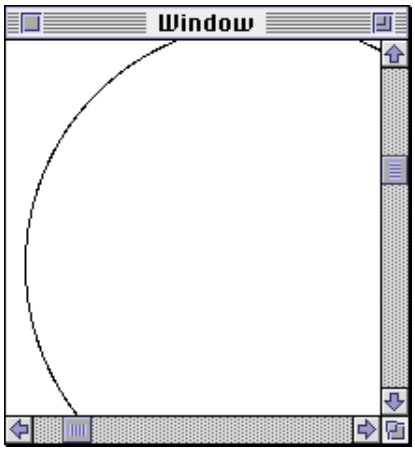

Figure 6.1: A window seen from the user perspective.

A window is composed of three layers, see Figure 6.2. The bottom layer, the *document layer*, is formed by the rendered document, the `Picture`. The middle layer, the *control layer*, contains all controls of the window. The top layer, the *window frame*, typically consists of a title bar, and window components to close and resize the window. The window frame can have any size and is restricted only by the screen size and a platform dependent minimum size. The window frame serves as a clipping area of the document layer. Windows usually contain scrollbars to help the user change the current view on the document layer. The default drawing domain of the document layer `Picture` ranges from 0 to $2^{31}-1$ in both axes. This is in general to large for rendering the document. A window can limit the displayable range of a `Picture` by setting a *picture domain*.

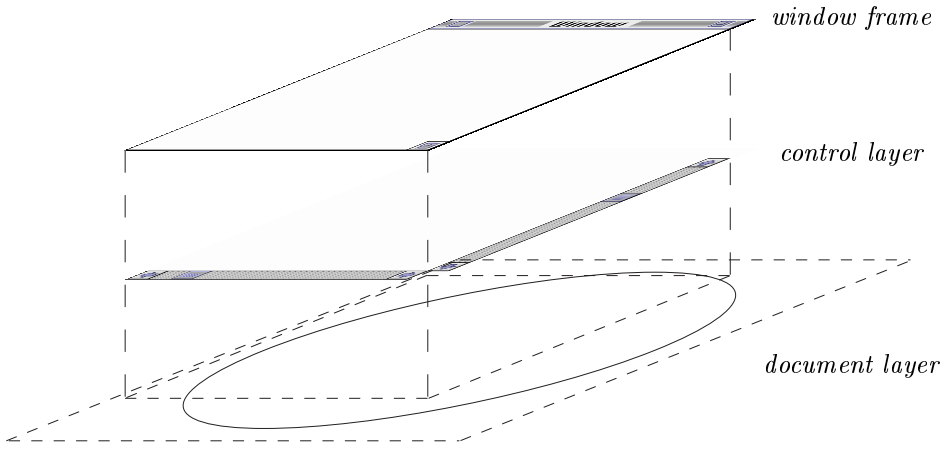

Figure 6.2: A different perspective at a window.

In contrast with windows, a dialogue is composed of only two layers, the control layer and the window frame, or also called the *dialogue frame*. Instead of a document layer, a dialogue has a platform dependent background. The program can not draw into nor navigate the background. The dialogue frame can not be resized by the user and is usually big enough to display the whole control layer.

For both windows and dialogues, the programmer has full control over the view frame, control layer, and document layer. Section 6.4 discusses how to handle the document layer. Handling the control layer is discussed in Section 6.5. The window and dialogue frame are controlled by the platform, see Section 6.6.

6.1.2 Stacking order

In general, a program can have an arbitrary number of windows and dialogues opened at the same time. These elements appear in a *stacking order*, seen by the application user in top to bottom order. For normal windows and dialogues the stacking order is not fixed. *Modal* dialogues however always appear topmost. Windows and modeless dialogues that are opened while (several) modal dialogues are open always appear below the bottom most modal dialogue.

6.1.3 Active window or dialogue

When a user works with a program, exactly one window or dialogue receives all keyboard and mouse input. This element is called the *active* window/dialogue and it has the *input focus*. With the exception of modal dialogues, the active window/dialogue does not necessarily occupy the top most stacking position.

6.2 Window and dialogue attributes

`WindowAttribute` is the type of window and dialogue attributes. The table below shows which attributes are valid for which element.

Window attributes	
For windows and dialogues:	**For windows only:**
`WindowCancel`	`WindowActivate`
`WindowClose`	`WindowCursor`
`WindowHide`	`WindowDeactivate`
`WindowHMargin`	`WindowHScroll`
`WindowId`	`WindowKeyboard`
`WindowIndex`	`WindowLook`
`WindowInit`	`WindowMinimumSize`
`WindowItemSpace`	`WindowMouse`
`WindowOk`	`WindowOrigin`
`WindowPos`	`WindowResize`
`WindowSize`	`WindowSelectState`
`WindowVMargin`	`WindowViewDomain`
	`WindowVScroll`

6.2.1 Window and Dialog attributes

`WindowCancel`, `WindowOk`: These attributes indicate which control should act conform the platform user interface 'confirm' control and 'cancel' control respec-

tively. If such an attribute is not provided, then no control is selected.

`WindowClose`: This attribute adds the platform dependent close control to the window/dialogue. The associated function will be evaluated in case this control is triggered. Actually closing the window/dialogue is the responsibility of this function. In case no `WindowClose` attribute is provided, the window/dialogue can not be closed in that way.

`WindowHide`: This attribute makes the given window/dialogue initially invisible. This attribute is ignored for modal dialogues. If the `WindowHide` attribute is not provided, then the window/dialogue will be opened visible.

`WindowHMargin`, `WindowVMargin`: These attributes determine the left-right, and top-bottom margin of a window/dialogue respectively. Their default value in case of windows is zero, and platform dependent for dialogues.

`WindowId`: This attribute identifies the window/dialogue to which it is associated. If you do not provide a `WindowId`, the object I/O system *open function* creates a fresh `Id` for the window/dialogue.

`WindowIndex`: This attribute determines the initial stacking position of the window/dialogue (see also 6.1.2). If no `WindowIndex` attribute is provided, then the window/dialogue will be opened frontmost. Modal dialogues are always opened frontmost.

`WindowInit`: This attribute defines an action that should be performed immediately after opening the window/dialogue. This is equivalent to the process initialisation action (see Section 2.3 and Chapter 11). If no `WindowInit` attribute is provided, no additional action is performed.

`WindowItemSpace`: This attribute determines the space between controls if no further offsets are provided in the layouts of controls. The default values are identical for windows and dialogues.

`WindowPos`: This attribute determines the initial position of the window/dialogue (see also Section 7.3). Relative `Id`s that occur in the `ItemPos` refer to other windows/dialogues. The object I/O system will always place a window/dialogue visibly on the current screen. If you do not provide a `WindowPos` to a window, then the window will be placed at the left top of the screen. If you do not provide a `WindowPos` to a dialogue, then the dialogue will be placed at a position conform the platform user interface (typically centered).

`WindowSize`: This attribute determines the initial size of the window/dialogue. If no `WindowSize` attribute is provided, then the system will derive a proper size. In case of dialogues the size is determined by the set of controls, margins, and item spaces. In case of windows the object I/O system exposes as much as possible of the view domain.

6.2.2 Window attributes

`WindowActivate`, `WindowDeactivate`: These attributes define the behaviour of the window in case the window becomes the active window (`WindowActivate`), and is no longer the active window (`WindowDeactivate`) respectively (see also 6.1.3). If no attribute is provided, this information will not be passed to the program.

`WindowCursor`: This attribute defines the shape of the cursor in case the mouse is over the window and not inside a control that may overrule this shape. In case no attribute is provided moving the mouse over the window will not change its shape.

`WindowHScroll`, `WindowVScroll`: These attributes add a horizontal scrollbar and a vertical scrollbar to the window. If no attribute is given, no scrollbars are added.

`WindowKeyboard`, `WindowMouse`: These attributes allow a window to respond to user actions with the mouse (`WindowMouse`) and keyboard (`WindowKeyboard`). If no attribute is provided, then this information will not be passed to the program. Both attributes can define an additional filter to ignore some input actions. If the `SelectState` of the window is `Unable` then neither function will obtain input.

`WindowLook`: This attribute defines a function that, given the current `SelectState` of the window and information about which part of the window should be displayed, defines what the window should look like. The default look fills the window with the platform dependent background colour. (The `Look` function also plays a role for `CustomButtonControls` (Section 7.1.9), `CustomControls` (Section 7.1.10), and `CompoundControls` (Section 7.1.11. It can also be used for printing (Chapter 13)).

`WindowMinimumSize`: This attribute determines the minimum size the window can obtain by resizing. If no attribute is given, a platform dependent minimum size is used.

`WindowOrigin`: This attribute determines the initial position of the `ViewFrame`, the rectangular part of the `ViewDomain` that is currently visible in the window. This position is always verified to be within the given `ViewDomain`. If no `WindowOrigin` attribute is provided, then the left-top coordinates of the `ViewDomain` are used.

`WindowResize`: This attribute allows the user to resize the window. If the attribute is not provided, then the window can not be resized. The size of a window may exceed the size of its `ViewDomain`. The area that is not part of the `ViewDomain` will be filled with a platform dependent background.

`WindowSelectState`: This attribute defines whether the window can be used by the user (`Able`) or not (`Unable`). The default value is `Able`. Note that although a window can be active, it can also be disabled. This only means that all input is ignored by the window.

`WindowViewDomain`: This attribute defines the drawing coordinate system of the document layer (see also 6.1.1). Drawing operations outside this area will be clipped. If no `ViewDomain` is provided, then the window will obtain the `ViewDomain` {`corner1=zero, corner2={x=maxint,y=maxint}`}.

6.3 Opening and closing of windows and dialogues

Windows and dialogues are opened using the appropriate instances of the respective type constructor classes `Windows` and `Dialogs` (module `StdWindow`, Appendix A.39).

```
class Windows wdef where
    openWindow :: .lst !(wdef .lst (PSt .l .p)) !(PSt .l .p)
                -> (!ErrorReport, !PSt .l .p)
    ...

class Dialogs ddef where
    openDialog :: .lst !(ddef .lst (PSt .l .p)) !(PSt .l .p)
                -> (!ErrorReport, !PSt .l .p)
    ...

instance Windows (Window c) | Controls c
instance Dialogs (Dialog c) | Controls c
```

To close windows and dialogues the function `closeWindow` has to be used.

```
closeWindow :: !Id !(PSt .l .p) -> PSt .l .p
```

6.4 Handling the document layer

The document layer of a window is used to present visual feedback to the user on the current status of the document that is being manipulated. Only windows have a document layer. Below we discuss two ways of drawing into the document layer. The first method, *indirect rendering* (Section 6.4.1), uses the look function, the second method, *direct rendering* (Section 6.4.2), draws directly into the document layer. Some programming pragmatics are discussed in Section 6.4.3. We conclude with an example in Section 6.4.4.

6.4.1 Indirect rendering

Two `WindowAttributes` play a paramount role with respect to the document layer: `WindowViewDomain` and `WindowLook`. Let's have a closer look at them.

The document layer is rendered using a `Picture`. As we have seen in Section 6.2, the default drawing range of a `Picture` is $(0, 2^{31} - 1)$ in both axes. This range can be changed by the `WindowViewDomain` attribute. It has a `ViewDomain` argument which is defined as a `Rectangle`. A `Rectangle` is a record:

```
::  Rectangle
    =   {   corner1 :: !Point
        ,   corner2 :: !Point
        }
::  Point
    =   {   x       :: !Int
        ,   y       :: !Int
        }
```

consisting of the two diagonally opposite corner points of the new drawing range. It is illegal to have identical x or y coordinates. This will cause a run-time error of the application. All drawing that occurs outside of the view domain of the document layer will be clipped.

The `Picture` of the document layer is rendered using the `WindowLook` attribute. This attribute has a `Look` function argument. It is defined as follows:

```
:: Look :== SelectState -> UpdateState -> *Picture -> *Picture
```

Whenever it is necessary to render (part of) the visible document layer, the object I/O system will apply the look function of the `WindowLook` attribute. It will be parameterised with the current `SelectState` of the window and detailed information about which part of the current view frame needs to be rendered. This information is presented by means of the `UpdateState` record:

```
:: UpdateState
 = { oldFrame  :: !ViewFrame
   , newFrame  :: !ViewFrame
   , updArea   :: !UpdateArea
   }
:: ViewFrame   :== Rectangle
:: UpdateArea  :== [ViewFrame]
```

This record consists of three fields:

`oldFrame`**:** If the size or orientation of the window was changed, then this field contains the view frame *before* that change. If this is not the cause, then this field contains the *current* view frame.

`newFrame`**:** If the size or orientation of the window was changed, then this field contains the view frame *after* that change (so, the current view frame). If this is not the cause, then this field contains the *current* view frame.

`updArea`**:** This field contains a list of `Rectangle`s (not necessarily disjoint) that define the parts of the visible *current* view frame that need to be drawn. This list always consists of at least one element. Each of its elements is always completely inside the current view frame.

Given the current `SelectState` of the window and the `UpdateState`, the look function is applied to the current `Picture` of the document layer. For this reason, the look function must be written in such a way that it works correctly for any argument `Picture`. Before the look function is applied, the object I/O system erases the rectangles of the `updArea` field.

The purpose of the `WindowLook` attribute function is to describe the look of the *current state* of the document layer. If the document does not change, then this function always correctly renders the document. However, it is very likely that the state of the document changes during the life-cycle of the window. The `WindowLook` attribute can be changed using the `StdWindow` function `setWindowLook` (Appendix A.39):

```
setWindowLook :: !Id !Bool !Look !(IOSt .l .p) -> IOSt .l .p
```

`setWindowLook` changes the current `WindowLook` attribute of the window indicated by the `Id` argument with the new `Look` function provided this window is present and does not refer to a dialogue. If the `Bool` argument is `False` then that's all. If the `Bool` argument is `True`, then the look function will be applied to the window in the way described above.

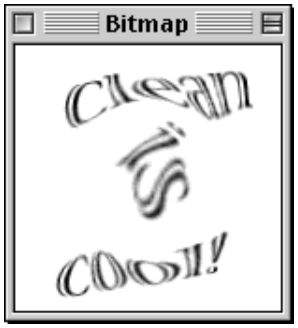

Figure 6.3: The bitmap program in action.

6.4.2 Direct rendering

Instead of using only the `Look` function of a window, one can also draw directly in the document layer `Picture`. This is done using the function `drawInWindow` (`StdWindow`, Appendix A.39):

```
drawInWindow :: !Id ![DrawFunction] !(IOSt .l .p) -> IOSt .l .p
```

`drawInWindow` applies a sequence of drawing functions to the `Picture` in the document layer of the window indicated by the `Id` argument if this window is present and does not refer to a dialogue. For some visual feedback such as drawing blinking cursors or track boxes this method is more suitable than the indirect way of using the `Look` function. The timer example in Section 9.1.1 shows a window in which the direct rendering method is applied.

6.4.3 Pragmatics

The `WindowLook` function is used by the object I/O system for all cases that the content of the window needs to rendered. Among others, causes are when the view frame, size, stacking order, or selectstate of a window changes. It is very annoying for the application user when these actions take to much time. Therefore it is worth your while to spend some effort in getting a good performance out of the list of drawing functions.

6.4.4 Example: displaying a bitmap

In this example we create a program that allows the user to select an arbitrary bitmap which will be displayed in a window (see Figure 6.3).

The first thing to do is to give the user the opportunity to select a bitmap. For this purpose the library function `selectInputFile` in module `StdFileSelect` (Appendix A.9) has been made. Applying this function opens the platform standard file selector dialogue. When a file has been selected by the user the return value contains the full pathname of the selected file. If no file has been selected, then `Nothing` is returned. It belongs to the `FileSelectEnv` type constructor class which contains one other function, `selectOutputFile`, that allows a user to select a (new) output file.

```
class FileSelectEnv env where
```

```
  selectInputFile ::                 !*env -> (!Maybe String,!*env)
  selectOutputFile:: !String !String !*env -> (!Maybe String,!*env)
```

The first part of the program looks as follows (if no file was selected the program simply terminates):

```
Start :: *World -> *World
Start world
  # (maybeFile,world) = selectInputFile world
  | isNothing maybeFile
    = world
```

Given a selected input file pathname, the function `openBitmap` (`StdBitmap`, Appendix A.1) can be used to read in the bitmap (see also Section 5.4.9).

```
openBitmap:: !{#Char} !*env-> (!Maybe Bitmap,!*env) | FileSystem env
```

`openBitmap` returns a bitmap if it could be read in successfully, otherwise it returns `Nothing` (in that case we also let the program terminate). This part of the program is as follows:

```
  # bitmapfile          = fromJust maybeFile
  # (maybeBitmap,world) = accFiles (openBitmap bitmapfile) world
  | isNothing maybeBitmap
    = world
```

Having successfully read in the bitmap, we can determine its size using the `StdBitmap` function `getBitmapSize`.

```
  # bitmap     = fromJust maybeBitmap
    bitmapsize = getBitmapSize bitmap
```

The window that displays the bitmap has the same initial size as the bitmap (setting the (`WindowSize bitmapsize`) attribute). In this example we use the indirect rendering method. This is done by setting the (`WindowLook (\_ _-> drawAt zero bitmap`) attribute. Finally, when the user requests closing of the window, the application simply terminates. This is done by setting the (`WindowClose (noLS closeProcess)`) attribute. So we obtain the following window definition:

```
  window = Window "Bitmap" NilLS
            [ WindowSize  bitmapsize
            , WindowLook  (\_ _->drawAt zero bitmap)
            , WindowClose (noLS closeProcess)
            ]
```

The final step is to create an interactive process that contains the window by applying the `startIO` function (`StdProcess`, Appendix A.24). The only initialisation action of the process is to open the window. In case the user requests the process to terminate, the application obeys. This is done by setting the (`ProcessClose closeProcess`) attribute.

```
  = startIO NoState
            NoState
            [snd o openWindow NoState window]
            [ProcessClose closeProcess]
            world
```

This completes the program. Here is the complete program code.

```
module showbitmap

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates a window that displays a user selected bitmap.
//  Make sure that this application has sufficient heap or extra memory.
//  ********************************************************************************

import StdIO, StdEnv

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    # (maybeFile,world)     = selectInputFile world
    | isNothing maybeFile
        = world
    # bitmapfile            = fromJust maybeFile
    # (maybeBitmap,world)   = accFiles (openBitmap bitmapfile) world
    | isNothing maybeBitmap
        = world
    # bitmap                = fromJust maybeBitmap
      bitmapsize            = getBitmapSize bitmap
      window                = Window "Bitmap" NilLS
                                [   WindowSize  bitmapsize
                                ,   WindowLook  (\_ _->drawAt zero bitmap)
                                ,   WindowClose (noLS closeProcess)
                                ]
    = startIO NoState NoState [snd o openWindow NoState window]
                             [ProcessClose closeProcess] world
```

6.5 Handling the control layer

The control layer contains the controls of a window or dialogue. Windows and dialogues can have the same set of controls. As we have seen in Section 6.3, this has been made explicit by the type constructor class instance declarations of the respective type constructor classes.

The standard set of `Controls` instances is defined in the module `StdControlClass` (Appendix A.5). Many operations with respect to the control layer are identical to operations on controls that are element of `CompoundControls`. Controls and their operations are discussed in detail in Chapter 7.

6.6 Handling the window and dialogue frame

The window and dialogue frame are the 'physical' borders of windows and dialogues. A user can grab them using the mouse or some keyboard interface and drag them around, change the size (in case of windows), or dispose of them. The program

has very limited influence on both the appearence and functionality of the frame. When *opening* a window or dialogue, the `WindowAttributes` are important. This is discussed in Section 6.6.1. User actions on the window or dialogue frame are handled by the program via the callback function mechanism. The program can also change frame properties. This is discussed in Section 6.6.2.

6.6.1 Opening a window or dialogue frame

The attributes of a window and dialogue definition that influence the window and dialogue frame are of course the `Title` and the following `WindowAttributes`:

`WindowSize`**:** This value gives the prefered size of the view frame of the window or dialogue. Be aware that this is not exterior size, which is in general larger and depends on the underlying platform.

`WindowClose`**:** If this attribute is present, then the window or dialogue frame is provided with a platform dependent interface element to allow the user to request the program to close that window or dialogue.

`WindowHScroll` **and** `WindowVScroll`**:** Although the horizontal and vertical scrollbar of a window are element of the control layer (Figure 6.2), their behaviour is intimately connected with the size of the view frame and therefore also with the size of the window and dialogue frame.

`WindowResize`**:** If this attribute is present, then the window is provided with a platform dependent interface element to allow the user to change the size of the window view frame.

6.6.2 Changing a window and dialogue frame

The two most apparant changes to the window or dialogue frame are changes of orientation and size. In case of windows, these can be caused by the application user, depending on the attributes as explained in Section 6.6.1.

The program can also change the size of the frame for both windows and dialogues. For dialogues this can be done only indirectly by opening or closing controls. For windows this can also be done directly. If the program changes the view frame (either its size or orientation) then the layout of controls is changed in the same way as if the user had caused this change.

6.7 Handling keyboard and mouse input

Of all windows and dialogues that are in control by an application, at most one receives the keyboard and mouse input. This window is the *active window*.

Except when modal dialogues are open, the application user can always select one of the visible windows or dialogues to become the new active window. Dialogues are not notified of these changes. Windows can be notified by adding two `WindowAttributes`: `WindowActivate` and `WindowDeactivate`. Both attributes are parameterised with a function that will be applied by the object I/O system as soon as that window becomes active or has become inactive respectively. It is guaranteed that the `WindowDeactivate` attribute function is applied before the `WindowActivate` function.

The underlying platform always gives visual clues to the application user about which window or dialogue is currently active. The program can retrieve this information using the `getActiveWindow` function (Appendix A.39). One should be aware that it is not correct to assume that the active window or dialogue has the top-most stack order position. As an example, one might try to get the `Id` of the active window by taking the `Id` from the first element of the result list of `getWindowStack`, but this only returns the top most window and not the active window.

The program can also activate windows and dialogues. This is done with the `activateWindow` function (Appendix A.39). Because modal dialogues are always front-most and the front most modal dialogue is active, one can not activate a window or modeless dialogue while modal dialogues are open. Instead, `activateWindow` restacks such a window or dialogue immediately behind the bottom most modal dialogue *without* making it the active window.

The active window receives all keyboard and mouse input. If this window contains controls it can be the case that the input is channelled to one of these controls. That particular control then has the *input focus*. If no control has the input focus, then all input is handled by the active window. In case the active window is a dialogue its response to input is defined entirely by the underlying platform. In case of windows the program can customise the behaviour by adding a `WindowKeyboard` attribute for keyboard input (Section 6.7.1) and by adding a `WindowMouse` attribute for mouse input (Section 6.7.2). Both attributes have a filter function (`KeyboardStateFilter` and `MouseStateFilter` respectively) which is applied before the actual callback function is evaluated. Only if the filter returns `True` then the callback function is evaluated.

6.7.1 Keyboard input

Every keyboard sensitive interface object has a `KeyboardFunction` which is a process state transition function that receives, as a first argument, a value of type `KeyboardState`. This value represents one keyboard event. Keyboard events are always generated in sequences that are characterised by a value of type `KeyState` in the following order:

`(KeyDown False)` $\{$`(KeyDown True)`$\}^*$ `KeyUp`

A keyboard event is either an ASCII character (`CharKey` alternative) or a special key (`SpecialKey` alternative).

```
:: KeyboardState
 = CharKey    Char       KeyState
 | SpecialKey SpecialKey KeyState Modifiers
```

The special keys are imported via the module `StdIOCommon` (Appendix A.13). Among others they define the function keys, arrow keys, page and line keys. The `Modifiers` type is also defined in `StdIOCommon`. It is a record that refers to the state of the meta keys of the keyboard. Because some ASCII characters are generated using these meta keys they are not provided at the `CharKey` alternative. So *shift 'a'* simply generates the `CharKey` alternative with `Char` value `'A'`.

The object I/O system guarantees that at all times only one keyboard alternative is being handled. Assume that a user is pressing the 'a' key on the keyboard. This generates a *character 'a' key down event* (`CharKey 'a' (KeyDown False)`), and then a sequence of *character 'a' repeat key events* (`CharKey 'a' (KeyDown True)`).

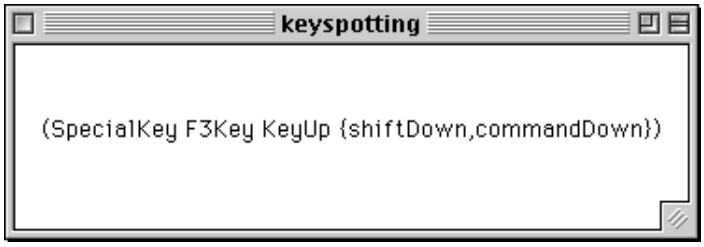

Figure 6.4: The keyspotting program in action.

If the user now also presses the 'b' key, the object I/O system inserts two virtual events that force the program to believe that the user first released the 'a' key with a *character 'a' key up event* (`CharKey 'a' KeyUp`), and then pressed the 'b' key with a *character 'b' key down event* (`CharKey 'b' (KeyDown False)`). These are followed by *character 'b' repeat key event*s.

Example: keyspotting

To illustrate the use of keyboard handling we create a program that has a window in which the last keyboard input is displayed. Figure 6.4 presents a snapshot of the program.

To allow the window track keyboard input, it must have the `WindowKeyboard` attribute. Because we intend to monitor every keyboard input the attribute's keyboard filter function argument must accept every `KeyboardState`. This can be defined conveniently by (`const True`), using the `StdFunc` library function `const`. Of course the keyboard function must handle keyboard input, so the `SelectState` attribute is `Able`. The keyboard function, `spotting` is not interested in the local state of the window which can be ignored by using the lifting function `noLS1`. The keyboard function is parameterised with the `Id` of the window, `wid`. This gives us the following definition of the `WindowKeyboard` attribute:

```
WindowKeyboard (const True) Able (noLS1 (spotting wid))
```

The keyboard function `spotting` uses the indirect rendering method (discussed in Section 6.4.1) to display the last keyboard input. It applies the `setWindowLook` function (`StdWindow`, Appendix A.39) to change the `Look` function of the window each time new keyboard input reaches the window. For this reason, `spotting` is parameterised with the `Id` of the window. We assume there exists a `toString` instance for the `KeyboardState` type. The type of `spotting` can be generalised so that it works for any second argument for which an instance of the overloaded function `toString` exists.

```
spotting :: Id x (PSt .l .p) -> PSt .l .p | toString x
spotting wid x pst
 = appPIO (setWindowLook wid True (look (toString x))) pst
```

The `Look` function of the window is parameterised with the string that should be displayed. It simply centers this string in the current view frame. The size of the view frame (a value of type `Rectangle`) can be determined using the `StdIOBasic` function `rectangleSize`. The size of the string, when drawn with the current pen, can be determined using the `StdPicture` function `getPenFontStringWidth` (see also Section 5.3).

```
look :: String SelectState UpdateState *Picture -> *Picture
look text _ {newFrame} picture
 # picture          = unfill newFrame picture
 # (width,picture) = getPenFontStringWidth text picture
 = drawAt {x=(w-width)/2,y=h/2} text picture
where
 {w,h}              = rectangleSize newFrame
```

Setting up `look` in this way has the advantage that the window's appearance adapts itself to its actual size. If we also make the window resizable (using the `WindowResize` attribute) this allows us to ignore to specify or calculate an initial size because the user can set the window size at wish.

The only things that need to be done to get a complete program is to create an `Id` value for the window, and start an interactive process that opens the window. This is shown below in the complete code of the keyspotting example.

```
module keyspotting

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program monitors keyboard input that is sent to a Window.
//  ********************************************************************************

import StdEnv,StdIO

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    # (wid,world)   = openId world
    # window        = Window "keyspotting" NilLS
                        [   WindowKeyboard  (const True) Able (noLS1 (spotting wid))
                        ,   WindowId        wid
                        ,   WindowClose     (noLS closeProcess)
                        ,   WindowResize
                        ]
    = startIO NoState NoState [snd o openWindow NoState window]
                              [ProcessClose closeProcess] world
where
    spotting :: Id x (PSt .l .p) -> PSt .l .p | toString x
    spotting wid x pst
        = appPIO (setWindowLook wid True (look (toString x))) pst

    look :: String SelectState UpdateState *Picture -> *Picture
    look text _ {newFrame} picture
        # picture           = unfill newFrame picture
        # (width,picture)   = getPenFontStringWidth text picture
        = drawAt {x=(w-width)/2,y=h/2} text picture
    where
        {w,h}               = rectangleSize newFrame
```

6.7.2 Mouse input

Every mouse sensitive interface object has a `MouseFunction` which is a process state transition function that receives, as a first argument, a value of type `MouseState`. This value represents one mouse event.

```
::  MouseState
```

```
    =   MouseMove   Point Modifiers
    |   MouseDown   Point Modifiers Int
    |   MouseDrag   Point Modifiers
    |   MouseUp     Point Modifiers
```

The `Point` and `Modifiers` types are defined in the module `StdIOCommon` (Appendix A.13). The `Point` type constructor states the position of the mouse at the mouse event in terms of the view frame coordinates of the interactive object that contains the specific `MouseFunction`. The `Modifiers` type constructor is a record that refers to the state of the meta keys of the keyboard that were pressed at the mouse event.

Mouse events are always generated in sequences that are characterised by the alternative constructor of the `MouseState` type constructor:

$$\{\texttt{MouseMove}\}^* \texttt{ [MouseDown } \{\texttt{MouseDrag}\}^* \texttt{ MouseUp]}$$

The `Int` argument of the `MouseDown` alternative gives the number of times the mouse was down within the *mouse double down time*. The mouse double down time is a platform dependent time interval that distinguishes two sequential mouse down events from a double click event. Although an integer is used for this count, its maximum value is usually three. If a mouse down event with count i has occured and a new mouse down event is generated within the mouse double down time, then the next mouse down event has count $i + 1$. If the next mouse down event is not generated within the mouse double down time, then the next mouse down event has count 1.

The object I/O system guarantees that every `MouseFunction` of a mouse sensitive interface object is applied to a sequence of mouse events as characterised above. Assume that a certain window is active and the user is pressing the mouse. This generates first a *mouse down event* (`MouseDown` alternative), followed by a sequence of *mouse drag events* (`MouseDrag` alternative). If for some reason another window is being activated, the object I/O system inserts a virtual event that forces the program to believe that the user has released the mouse button with a *mouse up event* (`MouseUp` alternative). If the new window is also mouse sensitive, then its `MouseFunction` is applied to a new virtual event that forces the program to believe that the user has pressed the mouse again with a *mouse down event* (`MouseDown` alternative). These are followed again by *mouse drag event*s.

Example: mousespotting

To illustrate the use of mouse handling we create a program that monitors the mouse input of a window (see Figure 6.5). This program is almost identical to the keyspotting example in the Section 6.7.1. The only differences are the title of the window and the replacement of the `WindowKeyboard` attribute by a `WindowMouse` attribute. Here we take advantage of the fact that the `spotting` function is overloaded. We assume there is an instance of `toString` for `MouseState` values. For completeness, the mousespotting example is shown here.

```
module mousespotting

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program monitors mouse input that is sent to a Window.
//  ********************************************************************************
```

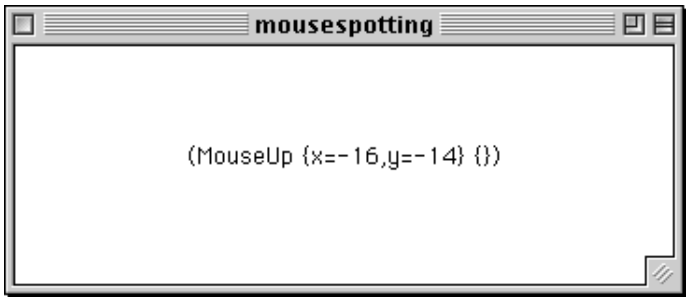

Figure 6.5: The mousespotting program in action.

```
import StdEnv,StdIO

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    # (wid,world)   = openId world
    # window        = Window "mousespotting" NilLS
                        [   WindowMouse     (const True) Able (noLS1 (spotting wid))
                        ,   WindowId        wid
                        ,   WindowClose     (noLS closeProcess)
                        ,   WindowResize
                        ]
    = startIO NoState NoState [snd o openWindow NoState window]
                             [ProcessClose closeProcess] world
where
    spotting :: Id x (PSt .l .p) -> PSt .l .p | toString x
    spotting wid x pst
        = appPIO (setWindowLook wid True (look (toString x))) pst

    look :: String SelectState UpdateState *Picture -> *Picture
    look text _ {newFrame} picture
        # picture           = unfill newFrame picture
        # (width,picture)   = getPenFontStringWidth text picture
        = drawAt {x=(w-width)/2,y=h/2} text picture
    where
        {w,h}               = rectangleSize newFrame
```

6.8 Modal dialogues

Windows and dialogues have many aspects in common. One important difference, which was also mentioned in the introduction of this chapter, is that dialogs can also be opened *modally*. When a program opens a modal dialog, the user is forced to completely handle the dialogue before any other operation can occur. When a modal dialogue is open, other modal dialogues can also be opened, but these must also be completely handled. In this way a stack of modal dialogues can be created. An example of such a situation is the behaviour of the platform dependent output file selector dialog (which can be created by the `StdFileSelect` function `selectOutputFile`), see Figure 6.6. If the user selects an existing file, another modal dialog is opened that asks the user if it is OK to overwrite that file. The user must answer this dialog before the output file selector dialog can be closed.

The type constructor class `Dialogs` contains an additional function to create dialogs

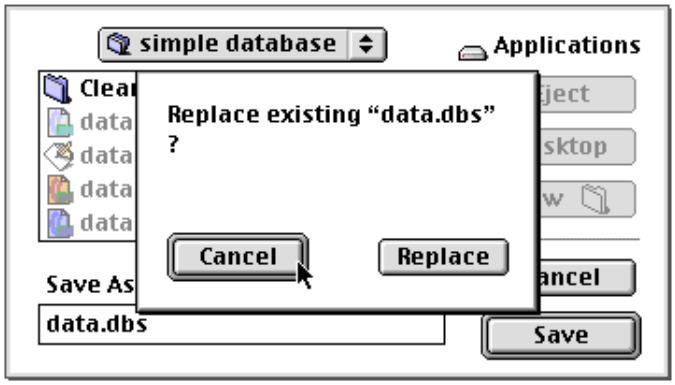

Figure 6.6: Two modal dialogues created by the `selectOutputFile` function.

in a modal way:

```
class Dialogs ddef where
    ...
    openModalDialog :: .lst !(ddef .lst (PSt .l .p)) !(PSt .l .p)
                    -> (!ErrorReport, !PSt .l .p)
    ...
```

Although the types of `openModalDialog` and `openDialog` are identical, their behaviour is different. The function `openDialog` opens its argument dialogue (if no errors occur) and terminates immediately. The function `openModalDialog` opens its argument dialogue (if no errors occur) and terminates only when its argument dialogue is closed. This situation can only be reached by applying `closeWindow` (`StdWindow`, Appendix A.39) or by terminating the parent interactive process using `closeProcess` (`StdProcess`, Appendix A.24).

6.8.1 Example: a notice extension

In this example we show the use of modal dialogues by making new `Dialogs` class instances for *notices*. A notice is a very simple dialog, containing only some line(s) of text, and a (number of) button(s) that close the notice. Notices can be used by programs to briefly inform the user about some situation. The confirmation dialogue in Figure 6.6 is an example of such a notice. We will first make the `notice` module which contains the new `Dialogs` instance in Section 6.8.1. Using this new interactive object we can rephrase the "Hello world" program (Section 2.4). This is done in Section 6.8.1.

The notice module

We want to make a special instance of the `Dialogs` type constructor class that specifies a notice. As said, a notice is a dialog that contains some line(s) of text, and a (number of) button(s). As done for every object I/O element, a notice will be defined by means of an algebraic data type:

```
:: Notice lst pst
 = Notice [TextLine] (NoticeButton *(lst,pst))
```

```
                    [NoticeButton *(lst,pst)]
:: NoticeButton st
 = NoticeButton TextLine (IdFun st)
```

The first argument of a `Notice` is a list of text lines. These will be shown below each other. By using one additional `NoticeButton` argument we can make sure that a notice always consists of at least one notice button. This mandatory notice button is placed to the right of the notice. The notice buttons will appear from right to left. A notice button is defined by `NoticeButton`. It is parameterised with the title and a program defined function. This function is applied after the button has been pressed and the notice has been closed.

The `Notice` type constructor will become an instance of the `Dialogs` type constructor class. This would actually be sufficient to allow a program to create both non modal and modal notices, using the class member functions `openDialog` and `openModalDialog`. To simplify creation of notices even further, we want to make the following additional function with a simpler type:

```
openNotice :: (Notice .lst (PSt .l .p)) (PSt .l .p) -> PSt .l .p
```

This specification is bundled in the notice definition module, which now looks as follows:

```
definition module notice

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program defines a new instance of the Dialogs class to create notices.
//  ********************************************************************************

import StdWindow

::  Notice lst pst
 =  Notice [TextLine] (NoticeButton *(lst,pst)) [NoticeButton *(lst,pst)]
::  NoticeButton st
 =  NoticeButton TextLine (IdFun st)

instance Dialogs Notice

openNotice :: (Notice .lst (PSt .l .p)) (PSt .l .p) -> PSt .l .p
```

Let's work out the notice implementation module. Because a notice is a specialised dialog, it is sufficient to map a notice definition into a dialog definition. This mapping is done by the function `noticeToDialog`. For programming convenience, we will convert the text lines into *text controls*, and the notice buttons into *button controls*. Although controls are not properly introduced yet, we hope this does not frustate your attempt to understand what's going on here.

Each notice text line is mapped to a text control. A text control is parameterised with a string, and can have a position attribute. All texts will be left-aligned. This is expressed by setting the (`ControlPos (Left,zero)`) attribute. For any list of strings `texts` this mapping can be conveniently expressed via a list comprehension (the `ListLS` data constructor is required for reasons explained in Chapter 7):

```
ListLS [TextControl text [ControlPos (Left,zero)] \\ text<-texts]
```

We want to place the lines together in a box without additional item space (the default item space and margin values of a dialog are platform dependent). To

make sure that we get our preferred values we override the dialogue item space and margin attributes with our own (using zero margins and an item space of three in both directions). This is done via a `CompoundControl`. The final mapping of the text lines results in:

```
texts‘= CompoundControl
        ( ListLS [  TextControl text [ControlPos (Left,zero)]
                 \\ text<-texts
                 ]
        )
        [ControlHMargin 0 0,ControlVMargin 0 0,ControlItemSpace 3 3]
```

Let's first introduce a function `noticebutton` that maps one notice button to a button control. When a notice button is selected, it should first close the notice (using `closeWindow`), and then apply the argument function. This action is handled by the argument function of the `ControlFunction` attribute. We assume that `wid` refers to the `Id` of the notice dialog. We allow other attributes such as placement and so on to be appended to the button control by passing it as an additional argument of `noticebutton`.

```
noticebutton (NoticeButton text f) atts
 = ButtonControl text [ControlFunction f‘:atts]
where
 f‘ (lst,pst) = f (lst,closeWindow wid pst)
```

Using `noticebutton` we can map the notice buttons to button controls. The mandatory notice button is mapped to the "confirm" control using the window attribute `WindowOk` (Section 6.2.1). We assume that `okid` refers to this `Id`. In addition, the mandatory button is right-aligned in the dialog. So, if `ok` is the mandatory notice button, then it is mapped to the following button control:

```
ok‘ = noticebutton ok [ControlPos (Right,zero),ControlId okid]
```

The optional notice buttons each must appear to the left of the previous notice button. Again, we can use a list comprehension to do this mapping for a list of notice buttons `buttons`:

```
buttons‘= ListLS [noticebutton button [ControlPos (LeftOfPrev,zero)]
                 \\ button<-buttons
                 ]
```

Given a `Notice` value (`Notice texts ok buttons`), `noticeToDialog` creates a titleless `Dialog` and glues the mapped notice controls using the expression (`texts‘ :+: ok‘ :+: buttons‘`) (glueing controls is also discussed in the next chapter):

```
noticeToDialog wid okid (Notice texts ok buttons)
 = Dialog "" (texts‘:+:ok‘:+:buttons‘)
    [   WindowId wid
    ,   WindowOk okid
    ]
```

Given the mapping function `noticeToDialog` it has become a trivial task to define the new instance declaration of the type constructor class `Dialogs`:

```
instance Dialogs Notice where
 openDialog :: .lst (Notice .lst (PSt .l .p)) (PSt .l .p)
                                -> (ErrorReport,PSt .l .p)
 openDialog lst notice pst
  # (wId, pst) = accPIO openId pst
  # (okId,pst) = accPIO openId pst
  = openDialog lst (noticeToDialog wId okId notice) pst

  openModalDialog :: .lst (Notice .lst (PSt .l .p)) (PSt .l .p)
                                     -> (ErrorReport,PSt .l .p)
  openModalDialog lst notice pst
   # (wId, pst) = accPIO openId pst
   # (okId,pst) = accPIO openId pst
   = openModalDialog lst (noticeToDialog wId okId notice) pst

  getDialogType :: (Notice .lst .pst) -> WindowType
  getDialogType _
   = "Notice"
```

Also the definition of the convenience function `openNotice` now becomes trivial:

```
openNotice :: (Notice .lst (PSt .l .p)) (PSt .l .p) -> PSt .l .p
openNotice notice pst
 = snd (openModalDialog undef notice pst)
```

For completeness, the `notice` implementation module is given here.

```
implementation module notice

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program defines a new instance of the Dialogs class to create notices.
//  ********************************************************************************

import StdTuple, StdMisc, StdFunc
import StdId, StdPSt, StdWindow

::  Notice lst pst
 =  Notice [TextLine] (NoticeButton *(lst,pst)) [NoticeButton *(lst,pst)]
::  NoticeButton st
 =  NoticeButton TextLine (IdFun st)

instance Dialogs Notice where
    openDialog :: .lst (Notice .lst (PSt .l .p)) (PSt .l .p)
                                  -> (ErrorReport,PSt .l .p)
    openDialog lst notice pst
        # (wId, pst)    = accPIO openId pst
        # (okId,pst)    = accPIO openId pst
        = openDialog lst (noticeToDialog wId okId notice) pst

    openModalDialog :: .lst (Notice .lst (PSt .l .p)) (PSt .l .p)
                                     -> (ErrorReport,PSt .l .p)
    openModalDialog lst notice pst
        # (wId, pst)    = accPIO openId pst
        # (okId,pst)    = accPIO openId pst
        = openModalDialog lst (noticeToDialog wId okId notice) pst

    getDialogType :: (Notice .ls .ps) -> WindowType
    getDialogType _
```

Figure 6.7: The hello world program, now using a notice.

```
        = "Notice"

openNotice :: (Notice .lst (PSt .l .p)) (PSt .l .p) -> PSt .l .p
openNotice notice pst
    = snd (openModalDialog undef notice pst)

noticeToDialog :: Id Id !(Notice .lst (PSt .l .p))
    ->  Dialog  (:+: (CompoundControl (ListLS TextControl))
                (:+: ButtonControl    (ListLS ButtonControl)
                )) .lst (PSt .l .p)
noticeToDialog wid okid (Notice texts ok buttons)
    = Dialog "" (texts‘:+:ok‘:+:buttons‘)
        [   WindowId    wid
        ,   WindowOk    okid
        ]
where
    texts‘   = CompoundControl
                (  ListLS
                [  TextControl text [ControlPos (Left,zero)]
                \\ text<-texts
                ]
                ) [ControlHMargin 0 0,ControlVMargin 0 0,ControlItemSpace 3 3]

    ok‘      = noticebutton ok [ControlPos (Right,zero),ControlId okid]

    buttons‘ = ListLS
                [  noticebutton button [ControlPos (LeftOfPrev,zero)]
                \\ button<-buttons
                ]

    noticebutton (NoticeButton text f) atts
        = ButtonControl text [ControlFunction f‘:atts]
    where
        f‘ (lst,pst) = f (lst,closeWindow wid pst)
```

Hello world revisited

We can now use our new notice instance to make another implementation of the “Hello world” program. It is very similar to the original version in Section 2.4 but instead of a `Dialog` it now uses a `Notice`. Figure 6.7 shows the notice. The program code is given below.

```
module usenotice

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program shows "Hello world!" using a Notice.
//  ********************************************************************************
```

```
import StdEnv, StdIO
import notice

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    = startIO NoState NoState [openNotice hello] [] world
where
    hello   = Notice ["Hello world!"] (NoticeButton "Quit" (noLS closeProcess)) []
```

Chapter 7

Control handling

The previous chapter introduced windows and dialogues. These top level interface elements can contain *controls*, which are handled in this chapter. Control structures can be hierarchical, i.e. they can be composed of controls themselves. Using controls helps a program to provide a consistent and structured user interface. There are a lot of issues involved when working with controls.

First of all we introduce each standard control object in Section 7.1. Then the glue is introduced to build larger control structures in Section 7.2. An important aspect of controls is to manage their layout, presented in Section 7.3. Related to layout is what should happen in case a window containing controls is resized. This is discussed in Section 7.4. Finally, Section 7.5 contains a number of examples that demonstrate the use of controls.

7.1 The standard controls

Table 2.1 (page 12) shows the standard set of object I/O library controls. They can be divided into three groups:

Platform standard controls: these are the controls that exist on all platforms and that have a well-defined behaviour and look that is platform defined. In the table these are the first seven controls (`RadioControl` . . . `ButtonControl`).

Customised controls: the look and feel of these controls is completely or partially defined by the program. In the table these are the `CustomButtonControl` (look is defined by the program, but it feels like a button) and `CustomControl` (look and feel defined by the program).

Hierarchical controls: there is actually only one such control, the `CompoundControl`. It contains other controls which will be positioned relative to this control. Except that it is hierarchical it can also have a look and feel.

There is an extensive set of control attributes that are shared by most controls. `CompoundControls` have an additional set of attributes that are strongly related to window and dialogue attributes. These will be discussed in Section 7.1.11. Before we discuss each of the controls individually we pay some attention to the shared control attributes.

7.1.1 The shared control attributes

The `ControlId` attribute identifies the control to which it is associated. If you do not provide a `ControlId`, the control can not be modified.

The `ControlPos` attribute determines its layout position (see Section 7.3). If you do not provide a `ControlPos`, the control will be placed right next to the previous control (if it happens to be the first control, it will be positioned at the left-top).

The `ControlSize` attribute provides the initial size of the control. For all platform standard controls except the slider control, the size can be derived from their demanded attributes. So if you do not provide a `ControlSize`, then this is calculated by the system. If you do provide one then this value will override the system calculated size (again, except for slider controls). For customised controls the size is mandatory. For compound controls the size can be calculated given the control elements of the compound control. So, if no `ControlSize` is given, the compound control will exactly fit its element controls. If a `ControlSize` is given, then this value overrides the system calculated size.

The `ControlMinimumSize` attribute defines the minimum size of the control. This value is relevant in case of resizing (see Section 7.4). If no `ControlMinimumSize` is provided, the default value zero is chosen.

The `ControlResize` attribute defines that the control is resizeable. If no `ControlResize` is provided, the control is not resizeable. See Section 7.4 about the resizing behaviour of controls.

The `ControlSelectState` attribute defines whether the control can be used by the user (`Able`) or not (`Unable`). Usually this will affect the look of the control. The default value is `Able`.

The `ControlHide` attribute defines that the control is initially invisible. It does occupy space. If you do not provide this attribute then the control is visible.

The `ControlFunction` and `ControlModsFunction` attributes are the primary callback function attributes of controls. The difference between these two attributes is that the former is simply evaluated whenever the control is selected, and that the latter also provides the callback function with the modifier keys that have been pressed at the moment of selecting the control (for the definition of the modifier keys, see definition module `StdIOCommon`). In an attribute list the first of the two attributes is chosen.

The `ControlMouse` and `ControlKeyboard` attributes add mouse and keyboard handling callback functions to the control. This is only possible for non platform standard controls because platform standard controls have a predefined behaviour.

7.1.2 The RadioControl

A *radio control* is a group of radio control items of which exactly one item is selected. All alternatives are visible. The definition of a radio control is as follows:

```
::  RadioControl lst pst
 =  RadioControl [RadioControlItem *(lst,pst)] RowsOrColumns Index
                 [ControlAttribute *(lst,pst)]
::  RowsOrColumns
 =  Rows    Int
 |  Columns Int
::  RadioControlItem st :== (TextLine,IdFun st)
::  IdFun            st :== st -> st
```

```
::  TextLine            :== String
::  Index               :== Int
```

The items are ordered rowwise (the `Rows` alternative of `RowsOrColumns`) or columnwise (the `Columns` alternative of `RowsOrColumns`). The initially selected item is indicated by the `Index` value. As a convention in the object I/O library, when indicating elements indices range from 1 upto the number of elements. So n elements are indexed by $1 \ldots n$. In case the index is out of range, i.e. less than 1 or larger than n, it is set to 1 and n respectively. Valid radio control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	√
ControlPos	√	ControlFunction	
ControlSize		ControlModsFunction	
ControlMinimumSize		ControlMouse	
ControlResize		ControlKeyboard	
ControlSelectState	√		

When a radio control item is selected the previously selected radio control item will be unchecked, and the new radio control item gets the check mark. The corresponding callback function is then evaluated. The callback function is also evaluated if the currently selected radio control item is selected.

Example This `RadioControl` consists of five items, laid out in two rows. The first item is initially selected.

```
radiocontrol
  = RadioControl
    [("Radio item "+++toString i,id)\\i<-[1..5]] (Rows 2) 1 []
```

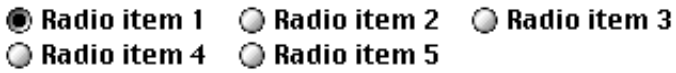

7.1.3 The CheckControl

A *check control* is a group of check control items of which an arbitrary number of items can be selected. All alternatives are visible. The definition of a check control is as follows:

```
::  CheckControl lst pst
 =  CheckControl [CheckControlItem *(lst,pst)] RowsOrColumns
                 [ControlAttribute *(lst,pst)]
::  RowsOrColumns
 =  Rows    Int
 |  Columns Int
::  CheckControlItem st :== (TextLine,MarkState,IdFun st)
::  IdFun            st :== st -> st
::  TextLine            :== String
```

The items are ordered rowwise (the `Rows` alternative of `RowsOrColumns`) or columnwise (the `Columns` alternative of `RowsOrColumns`). The initially selected items are indicated by their `MarkState` value (`Mark` if checked, `Nomark` if not checked). Valid control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	√
ControlPos	√	ControlFunction	
ControlSize		ControlModsFunction	
ControlMinimumSize		ControlMouse	
ControlResize		ControlKeyboard	
ControlSelectState	√		

When a check control item is selected its mark state will be toggled (from `Mark` to `NoMark` and vice versa). No other check control items are affected. The corresponding callback function is then evaluated.

Example This `CheckControl` consists of five items, laid out in two columns. Each odd numbered item has a check mark.

```
checkcontrol
  = CheckControl
    [ ("Check item "+++toString i,if (isOdd i) Mark NoMark,id)
    \\i<-[1..5]
    ] (Columns 2) []
```

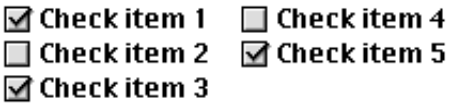

7.1.4 The PopUpControl

A *pop up control* is a group of pop up control items of which exactly one item is selected. The items of a pop up control are presented in a pop up menu. Usually only the currently selected item is displayed in the title of that pop up menu. For this reason pop up controls consume much less space than the functionally equivalent radio controls. The definition of a pop up control is as follows:

```
::  PopUpControl lst pst
 =  PopUpControl [PopUpControlItem *(lst,pst)] Index
                 [ControlAttribute *(lst,pst)]
::  PopUpControlItem st :== (TextLine,IdFun st)
::  IdFun               st :== st -> st
::  TextLine               :== String
```

The initially selected item is indicated by the `Index` value. As a convention in the object I/O library, when indicating elements indices range from 1 upto the number of elements. So n elements are indexed by $1 \ldots n$. In case the index is out of range, i.e. less than 1 or larger than n, it is set to 1 and n respectively. Valid control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	√
ControlPos	√	ControlFunction	
ControlSize		ControlModsFunction	
ControlMinimumSize		ControlMouse	
ControlResize		ControlKeyboard	
ControlSelectState	√		

When a pop up control item is selected the previously selected item is unchecked and the new item becomes the selected item. The corresponding callback function is then evaluated. The callback function is also evaluated if the currently selected radio control item is selected.

Example This `PopUpControl` consists of five items. The first item is the initially selected item.

```
popupcontrol
  = PopUpControl
    [("PopUp item "+++toString i,id)\\i<-[1..5]] 1 []
```

7.1.5 The SliderControl

Slider controls are used to change the view to space consuming visual data in one particular dimension. The definition of a slider control is as follows:

```
::  SliderControl lst pst
 =  SliderControl Direction Length SliderState
                  (SliderAction     *(lst,pst))
                  [ControlAttribute *(lst,pst)]
::  Direction
 =  Horizontal | Vertical
::  Length :== Int
::  SliderState
 =  { sliderMin  :: !Int
    , sliderMax  :: !Int
    , sliderThumb:: !Int
    }
::  SliderAction st :== SliderMove -> st -> st
::  SliderMove
 =  SliderIncSmall | SliderDecSmall
 |  SliderIncLarge | SliderDecLarge
 |  SliderThumb Int
```

A slider control can be laid out in a horizontal direction (the `Horizontal` alternative of `Direction`) or a vertical direction (the `Vertical` alternative of `Direction`). In this direction it can have a certain length. Its width is platform dependent.

The scrolling range of a slider control is defined by the `SliderState` record. The initial slider state determines the integer range: `sliderMin` gives the minimum value, `sliderMax` gives the maximum value. In case these values are given in the wrong order, they will be ordered properly. The initially chosen value is given by the `sliderThumb` value. This value must be inclusively between `sliderMin` and `sliderMax`. If a value smaller than the minimum range is given, then it is set to the minimum. If a value larger than the maximum range is given, then it is set to the maximum.

Valid control attributes for the `SliderControl` are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	√
ControlPos	√	ControlFunction	
ControlSize		ControlModsFunction	
ControlMinimumSize		ControlMouse	
ControlResize	√	ControlKeyboard	
ControlSelectState	√		

A slider control typically has five regions that can be selected by the user. These regions are shown in Figure 7.1

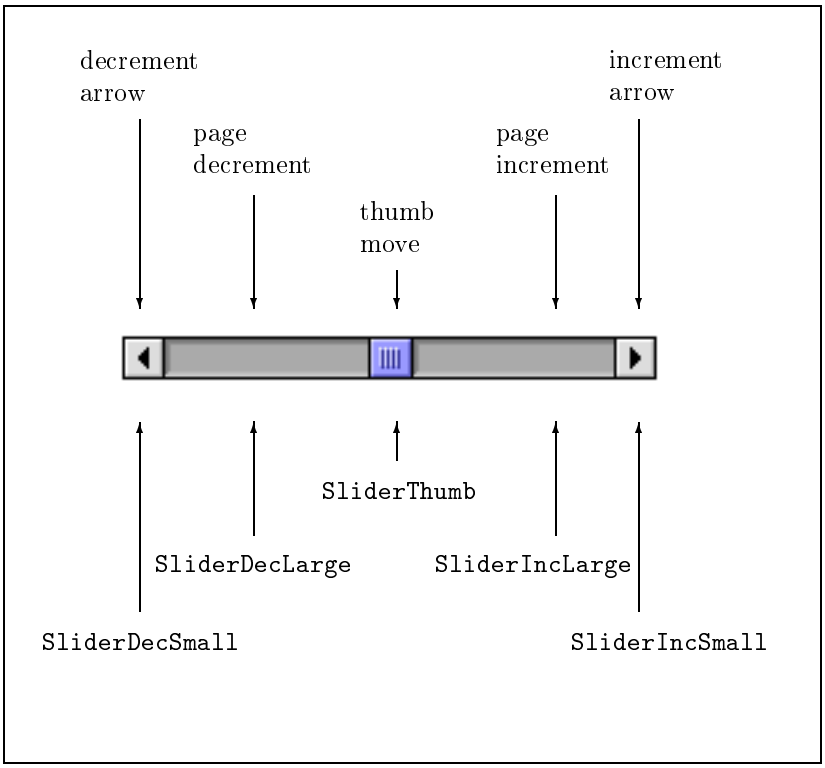

Figure 7.1: The regions of the `SliderControl`.

When the user is working with the slider control, its callback function is evaluated. The algebraic data type `SliderMove` has an alternative constructor for each of the regions of the slider control:

`SliderDecSmall`	*decrement arrow*
`SliderIncSmall`	*increment arrow*
`SliderDecLarge`	*page down region*
`SliderIncLarge`	*page up region*
`SliderThumb`	*thumb move*

The program can decide what to do with this information. It is the responsibility of the programmer that the application responds the way the user expects.

Example This `SliderControl` is orientied horizontally and has a length of 200 pixels.

```
slidercontrol
  = SliderControl Horizontal 200
              {sliderMin=(-100),sliderMax=100,sliderThumb=0}
              (\_ st->st) []
```

7.1.6 The TextControl

A *text control* displays one line of text that can not be changed by the user. The definition of a text control is as follows:

```
::  TextControl lst pst
 =  TextControl TextLine [ControlAttribute *(lst,pst)]
::  TextLine
:== String
```

If the textline contains newlines then these are interpreted as line breaks. This implies that only the initial part of the text is displayed, since text controls consist of one line only. Valid control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	
ControlPos	√	ControlFunction	
ControlSize	√	ControlModsFunction	
ControlMinimumSize		ControlMouse	
ControlResize		ControlKeyboard	
ControlSelectState	√		

The initial size of a text control is determined by its initial text line. If the text control has a `ControlSize` attribute that is larger than its initial size, then that value becomes its initial size. A text control is not resizeable, and it will also not change in size in case its text line is modified.

Example A text control that contains a text with newline characters.

```
textcontrol
  = TextControl "This is a TextControl" []
```

This is a TextControl

7.1.7 The EditControl

The *edit control* is applied to provide the user with an interface to edit (typically small amounts of) textual data. The definition of an edit control is as follows:

```
::  EditControl lst pst
 =  EditControl TextLine Width NrLines
                [ControlAttribute *(lst,pst)]
::  TextLine :== String
::  Width    :== Int
::  NrLines  :== Int
```

An edit control initially displays some text line. If the textline contains newlines then these are interpreted as line breaks. The edit control has an initial interior width (in terms of pixels) and shows an integral number of lines. Valid control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	√
ControlPos	√	ControlFunction	
ControlSize		ControlModsFunction	
ControlMinimumSize		ControlMouse	
ControlResize		ControlKeyboard	√
ControlSelectState	√		

If the `SelectState` of an edit control is `Unable`, the user can not type in text. The current content of an edit control can be retrieved using the `getControlTexts` function of module `StdControl` (Appendix A.4). The program can also keep track of the inserted text by setting the *control keyboard* attribute. For each typed key the keyboard function is evaluated.

Example An `EditControl` that contains text with newline characters. Its interior width is 80 pixels and it shows three lines of text.

```
editcontrol
  = EditControl "This is an \nEditControl" 80 3 []
```

7.1.8 The ButtonControl

The *button control* represents an action that should occur given the current state of the window or dialogue. The definition of a button control is as follows:

```
::  ButtonControl lst pst
 =  ButtonControl TextLine [ControlAttribute *(lst,pst)]
::  TextLine :== String
```

A button control has a title, given by a text line. If the textline contains newlines then these are interpreted as line breaks. Valid control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	√
ControlPos	√	ControlFunction	√
ControlSize	√	ControlModsFunction	√
ControlMinimumSize		ControlMouse	
ControlResize		ControlKeyboard	
ControlSelectState	√		

The initial size of a button control is determined by its initial text line. If the button control has a `ControlSize` attribute that is larger than its initial size, then that value becomes its initial size. The callback function of a button control is the

function that is evaluated in case the button control was selected by the user and its `SelectState` was `Able`. If the name of the button control is modified to a new text line, then its size is not changed.

A button control can be the confirm or cancel button of a window or dialogue by setting their `WindowOk` or `WindowCancel` attribute (see Section 6.2.1).

Example A `ButtonControl` with a title that contains newlines.

```
buttoncontrol
  = ButtonControl "This is a ButtonControl" []
```

7.1.9 The CustomButtonControl

A *custom button control* is a control which *feels* like a button control, but which *look* is customised by the program. The definition of a custom button control is as follows:

```
:: CustomButtonControl lst pst
 = CustomButtonControl Size Look [ControlAttribute (lst,pst)]
:: Size
 = {w::!Int,h::!Int}
:: SelectState
 = Able | Unable
:: UpdateState
 = { oldFrame :: !ViewFrame
   , newFrame :: !ViewFrame
   , updArea  :: !UpdateArea
   }
:: ViewFrame  :== Rectangle
:: UpdateArea :== [ViewFrame]
:: Look       :== SelectState -> UpdateState -> *Picture -> *Picture
```

Both the initial size and look of a custom button control are defined by the program. The look of a custom button control is identical to the look of a window as discussed in Section 6.4.1. The `UpdateState` argument contains zero based rectangles of the same size as the custom button control itself. Every custom button control has a `*Picture` environment to which the look drawing functions are applied. Valid control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	√
ControlPos	√	ControlFunction	√
ControlSize		ControlModsFunction	√
ControlMinimumSize	√	ControlMouse	
ControlResize	√	ControlKeyboard	
ControlSelectState	√		

The callback function of a custom button control is the function that is evaluated in case the custom button control was selected by the user and its `SelectState` was `Able`.

A custom button control can be the confirm or cancel button of a window or dialogue by setting their `WindowOk` or `WindowCancel` attribute (see Section 6.2.1).

Example A `CustomButtonControl` which look depends on its `SelectState`. The picture on the left shows the custom button control in `Able` state, the picture on the right in `Unable` state.

```
custombuttoncontrol
  = CustomButtonControl {w=50,h=50} look []
where
  look :: SelectState UpdateState *Picture -> *Picture
  look Able {newFrame} picture
    # picture = fill newFrame (setPenColour DarkGrey picture)
    # picture = draw newFrame (setPenColour Black    picture)
    = picture
  look Unable {newFrame} picture
    = fill newFrame (setPenColour LightGrey picture)
```

7.1.10 The CustomControl

A *custom control* is a control of which both the look and feel are program defined. The definition of a custom control is as follows:

```
:: CustomControl lst pst
 = CustomControl Size Look [ControlAttribute *(lst,pst)]
:: Size
 = {w::!Int,h::!Int}
:: SelectState
 = Able | Unable
:: UpdateState
 = { oldFrame :: !ViewFrame
   , newFrame :: !ViewFrame
   , updArea  :: !UpdateArea
   }
:: ViewFrame  :== Rectangle
:: UpdateArea :== [ViewFrame]
:: Look       :== SelectState -> UpdateState -> *Picture -> *Picture
```

Both the initial size and look of a custom control are defined by the program. The look of a custom control is identical to the look of a window as discussed in Section 6.4.1. The `UpdateState` argument contains zero based rectangles of the same size as the custom control itself. Every custom control has a `*Picture` environment to which the look drawing functions are applied. Valid control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlHide	√
ControlPos	√	ControlFunction	
ControlSize		ControlModsFunction	
ControlMinimumSize	√	ControlMouse	√
ControlResize	√	ControlKeyboard	√
ControlSelectState	√		

The *feel* of a custom control is defined by its mouse and keyboard callback functions. If the user selects the custom control with the mouse, then the mouse callback function handles all feedback. If the custom control has the keyboard input focus, and the user is typing, then the keyboard callback function handles all feedback.

Example A `CustomControl` which look depends on its `SelectState`. The picture on the left shows the custom control in `Able` state, the picture on the right in `Unable` state.

```
customcontrol
  = CustomControl {w=50,h=50} look []
where
  look :: SelectState UpdateState *Picture -> *Picture
    # picture = fill newFrame (setPenColour DarkGrey picture)
    # picture = draw newFrame (setPenColour Black    picture)
    = picture
  look Unable {newFrame} picture
    = fill newFrame (setPenColour LightGrey picture)
```

7.1.11 The CompoundControl

The *compound control* is a control that contains other controls. A compound control is actually a window within a window. It introduces a new *layout scope*: i.e. controls inside it are positioned relative to the bounds of the compound control. The compound control can have an additional look and feel. The definition of a compound control and its `Controls` class instance declaration is as follows:

```
::  CompoundControl  c lst pst
 =  CompoundControl (c lst pst) [ControlAttribute *(lst,pst)]

instance Controls (CompoundControl c) | Controls c
```

The `c` parameter of the `CompoundControl` type constructor is a type constructor variable that corresponds with the control elements of the compound controls. Any composition of controls that is an instance of the `Controls` type constructor class is a valid argument of `CompoundControl`.

Compared with the previous controls, compound controls have an additional number of control attributes that are irrelevant to the other controls. These attributes are similar to some attributes of windows and dialogues, and of course they intend to have the same meaning. Valid control attributes are:

ControlAttribute:	Valid:	ControlAttribute:	Valid:
ControlId	√	ControlItemSpace	√
ControlPos	√	ControlHMargin	√
ControlSize	√	ControlVMargin	√
ControlMinimumSize	√	ControlLook	√
ControlResize	√	ControlViewDomain	√
ControlSelectState	√	ControlOrigin	√
ControlHide	√	ControlHScroll	√
ControlFunction		ControlVScroll	√
ControlModsFunction			
ControlMouse	√		
ControlKeyboard	√		

The size of a compound control, if not provided as a `ControlSize` attribute, is derived by the system from its control elements (see Section 7.3 for more information on the layout of controls). Related to the size of a compound control and layout of its elements are a number of attributes. The `ControlMinimumSize` attribute determines the minimum size of the compound control (see resizing controls, Section 7.4), the `ControlResize` attribute controls the resize behaviour (see also Section 7.4), the `ControlItemSpace`, `ControlHMargin`, and `ControlVMargin` attributes define the distance between the elements themselves and the horizontal and vertical distance of the elements to the border of the compound control respectively. If a compound control does not specify any of these attributes, it obtains the same attribute values as its parent object (which can be a compound control, a window, or a dialogue).

The compound control anatomy is the same as that of a window (Section 6.1.1). It consists of the same three layers as a window except that we call its top layer the *compound frame* rather than window frame. The compound frame has no title nor features like resize controls and so on.

Analogous to windows, compound controls have a view domain if the `ControlViewDomain` is given. Otherwise they obtain the same default view domain. As explained in Section 6.1.1, a view domain defines a finite area in which can be drawn (Chapter 5), it can also be used as an area to place controls (Section 7.3). Also for compound controls scrolling attributes can be added: `ControlHScroll` and `ControlVScroll`. These attributes control the current view frame orientation of the compound control. The left top point of the view frame that is currently visible is called the *origin*. This value can be set initially with the `ControlOrigin` attribute.

The document layer of a `CompoundControl` can be drawn into *only if* a `ControlLook` attribute is given. In that case the system provides the compound control with a `Picture` drawing environment. If a `ControlLook` is absent, then the compound control is *transparant*.

The feel of a compound control is of course partially determined by its element controls. If the `ControlMouse` (`ControlMouse`) attribute is given, then the compound control can handle all mouse (keyboard) events that are directed to it.

7.2 Control glue

In the previous section the standard set of controls has been discussed. This list does not cover all controls class instances. In the library module `StdControlClass` (Appendix A.5) a number of additional instances are defined, namely the type constructors `:+:`, `ListLS`, `NilLS`, and `AddLS`, `NewLS` (their definition can be found

in module StdIOBasic, Appendix A.12). These additional instances are required to *glue* controls. They are treated below.

7.2.1 :+:

The most common constructor to glue controls is :+:. Its type constructor definition and Controls class instance declaration are as follows:

```
::  :+: t1 t2 local context
 = (:+:) infixr 9 (t1 local context) (t2 local context)

instance Controls ((:+:) c1 c2) | Controls c1 & Controls c2
```

Given two Controls instances c1 and c2, working on the same local state of type local and context state context, the expression c1:+:c2 is also a Controls instance working on the same local state and context state. Because :+: is right associative, the expression c1:+:c2:+:c3 should be read as c1:+:(c2:+:c3).

7.2.2 ListLS and NilLS

In principle the :+: glue is sufficient to create all required control structures. In case of working with a number of control instances of the *same type*, it is much more convenient to use lists and list comprehensions. This glue is provided by the type constructor ListLS. The type constructor NilLS is a shorthand for ListLS []. It can also be conveniently used to state that a CompoundControl or window or dialogue has no controls. Their type constructor definitions and Controls class instance declarations are as follows:

```
::  ListLS t local context = ListLS [t local context]
::  NilLS    local context = NilLS

instance Controls (ListLS c) | Controls c
instance Controls NilLS
```

Given a list of Controls instances cs = [c_1 ...c_n], working on the same local state of type local and context state context, the expression ListLS cs is also a Controls instance working on the same local state and context state.

7.2.3 AddLS and NewLS

The previously discussed glueing type constructors always glue controls that work on the same local state and context state. Two other glueing constructors are defined to extend and change the local state, AddLS and NewLS.

Given a Controls instance c1 that works on a local state of type local and a context state of type context, one can add another Controls instance c2 that works on an *extended* local state of type (add,local) and the same context state of type context. Let x be a value of type add, then this is done by the expression c1 :+: {addLS=x, addDef=c2}.

```
:: AddLS t local context
 = E..add:
```

```
   { addLS :: add
   , addDef:: t *(add,local) context
   }

instance Controls (AddLS c) | Controls c
```

Given a `Controls` instance `c1` that works on a local state of type `local` and a context state of type `context`, one can add another `Controls` instance `c2` that works on a *new* local state of type `new` and the same context state of type `context`. Let `x` be a value of type `new`, then this is done by the expression `c1 :+: {newLS=x, newDef=c2}`.

```
:: NewLS t local context
 = E..new:
   { newLS :: new
   , newDef:: t new context
   }

instance Controls (NewLS c) | Controls c
```

In both cases the extended part of the local state and the new local state are encapsulated completely from the external context using existential quantification.

7.2.4 Example: a counter control

In this section we show an example of glueing controls to form a new `Controls` class instance. A `CompoundControl` is defined that consists of three other controls. It implements a manually incrementable counter. To display the current count value it uses an `Unable EditControl`. Two `ButtonControls` are used to decrement and increment the counter.

```
counter windowid displayid
  = { newLS =initcount
    , newDef=CompoundControl
              (   EditControl (toString initcount) (hmm 50.0) 1
                              [ControlSelectState Unable
                              ,ControlPos         (Center,zero)
                              ,ControlId          displayid
                              ]
              :+: ButtonControl "-" [ControlFunction    (count (-1))
                              ,ControlPos         (Center,zero)
                              ]
              :+: ButtonControl "+" [ControlFunction    (count 1)]
              )   []
    }
where
  initcount = 0

  count :: Int (Int,PSt .l .p) -> (Int,PSt .l .p)
  count dx (count,pst)
    = (count+dx, setText windowid displayid (count+dx) pst)

  setText :: Id Id x (PSt .l .p) -> PSt .l .p | toString x
```

```
setText wid cid x pst
  = appPIO
      (setWindow wid (setControlTexts [(cid,toString x)])) pst
```

Figure 7.2 shows the counter after some user manipulations that resulted in the counter value -15.

Figure 7.2: The counter control.

7.3 Control layout

The object I/O library offers the programmer an expressive layout mechanism to define the layout of controls. Controls can be element of windows, dialogues, and compound controls. The layout rules that are discussed in this section are followed in all these cases. Some layout rules refer to the view domain and frame of the parent object. In case of dialogues these two are identical. They are zero based rectangles with a size equal to the dialogue interior.

As we have seen, every control can have a `ControlPos` attribute. This attribute is defined as follows:

```
::  ControlAttribute ps
 =  ... | ControlPos ItemPos | ...
::  ItemPos
:== (ItemLoc,ItemOffset)
::  ItemLoc
 =  Fix Point
 |  LeftTop    | RightTop    | LeftBottom | RightBottom
 |  Left       | Center      | Right
 |  LeftOf Id  | RightTo Id  | Above Id   | Below Id
 |  LeftOfPrev | RightToPrev | AbovePrev  | BelowPrev
::  ItemOffset
:== Vector
```

The layout position of a control consists of two values: an `ItemLoc` value and an `ItemOffset` value which is a vector. The `ItemLoc` actually determines the location of the control, the `ItemOffset` value adds an offset to this position (which of course influences the position of other controls). The `ItemLoc` values can be divided into four groups:

Fixed position: this is only the `Fix` alternative of `ItemLoc`. `Fix point` places the left top corner of the control of which it is the attribute at the specified `point`. The `point` is given in the coordinate system of the *view domain* of the parent object.

Boundary aligned: these are the alternatives `LeftTop`, `RightTop`, `LeftBottom`, and `RightBottom`. When applied their control is placed at the left-top, right-

top, left-bottom, or right-bottom of the current *view frame* of the parent object.

Line aligned: these are the alternatives `Left`, `Center`, and `Right`. When applied their control is placed below all previous line aligned controls and either left-aligned, centered, or right-aligned with respect to the current *view frame* of the parent object.

Relative position: these are the alternatives `LeftOf`, `RightTo`, `Above`, `Below`, `LeftOfPrev`, `RightToPrev`, `AbovePrev`, and `BelowPrev`. The first four alternatives must be parameterised with the `Id` of a control that is element of the same parent object, otherwise a runtime error will occur. The latter four alternatives can be defined in terms of the first four but have the advantage that you do not have to think of `Id`s for controls that you only want to relatively place other controls to. Placing controls relatively to other controls must construct a *tree* of related controls: cyclic references are not allowed and result also in a runtime error.

Controls that are laid out relatively form a *layout tree* with one *layout root* control. The layout attribute of the layout root control determines the layout positions of the whole layout tree. If it is at a fixed position, the layout tree obtains a fixed position. If it is boundary aligned, the layout tree is aligned at the same boundary. If it is line aligned, the layout tree is line aligned.

Except for the *first* layout root control the default layout attribute for controls is `(RightToPrev,zero)`. For the first layout root control the default attribute is `(Left,zero)`. Consequently, the default layout order is from left to right in one single row.

Controls are allowed to overlap partially or completely. This is particularly useful in case of combinations of hidden and visible controls when at all times only one is visible. It allows the program to change the control structure in an easy way by hiding and showing controls.

In the remaining part of this section a number of examples are given to illustrate control layout. In each of the examples we assume that the controls that are being laid out are placed in a parent object with a view *domain* and view *frame* as given in Figure 7.3. The x axis (the horizontal arrow) and y axis (the vertical arrow) intersect at the coordinate `zero`.

Controls are being displayed as boxes. The control configuration shown in Figure 7.4 occurs frequently in the examples. It consists of five equally sized controls, $c_0 \ldots c_4$. The layout root control is c_0. The controls $c_1 \ldots c_4$ are laid out relatively to c_0 and have the layout attributes `LeftOf`, `RightTo`, `Above`, and `Below` respectively with `zero` offsets.

7.3.1 Layout at fixed position

Controls and layout trees that have a `Fix` layout attribute are being placed relative to the view domain of the parent object. So their visibility depends on the current orientation of the parent view frame (recall that the view frame clips everything that is outside of it). Figure 7.5 shows the control configuration of Figure 7.4 when the layout root control has the attribute `(Fix zero,zero)`.

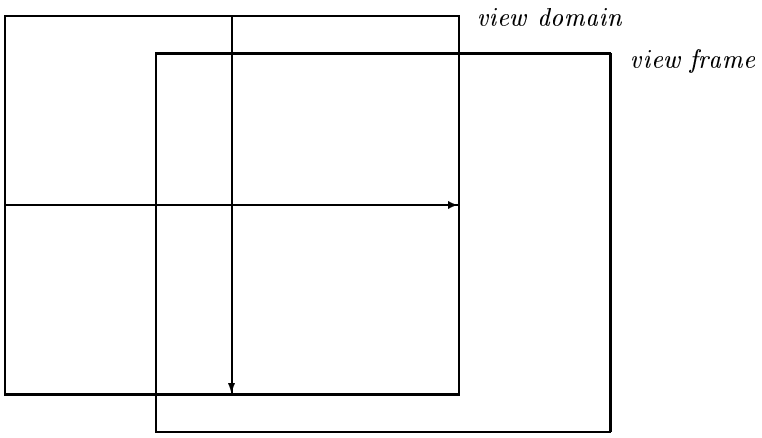

Figure 7.3: View domain and view frame.

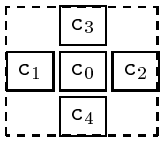

Figure 7.4: A layout tree of five controls.

7.3.2 Layout at view frame boundary

Controls and layout trees that are laid out relative to the parent view frame boundary will always be visible, provided of course that the view frame is suffiently large. If the view frame is not large enough, these controls may become overlapped. Figure 7.6 shows the positions of the control configuration of Figure 7.4 when positioned at every corner of the view frame (using `LeftTop`, `RightTop`, `LeftBottom`, and `RightBottom` with `zero` offsets.

7.3.3 Layout in lines

Laying out controls and layout trees in lines is similar to writing characters in an English piece of text: each new character is placed right next to the previous character until a new line is started. A new line starts below the previous line. This new line can be *left* aligned, *centered*, or *right* aligned. The layout attributes `Left`, `Center`, and `Right` introduce both a new line and its alignment. Figure 7.7 shows the positions of the control configuration of Figure 7.4 when positioned at `Left`, `Center`, and `Right` respectively, using `zero` offsets. It also illustrates that if the view frame is not big enough, controls may become partially invisible.

7.3.4 Layout offsets

So far we have used `zero` offsets in the layout attribute examples. The layout position of a control is changed by an offset vector value `v = {vx,vy}` as follows: first, the layout position of the control is calculated as explained above, using a `zero` offset. Now assume that this results in the *exact* location `pos = {x,y}`. Then

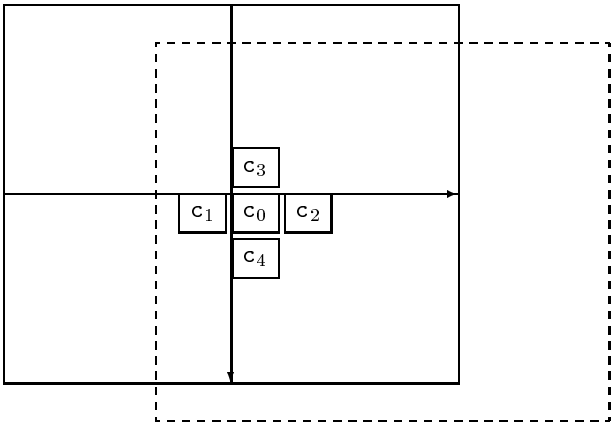

Figure 7.5: The layout tree at `(Fix zero,zero)`.

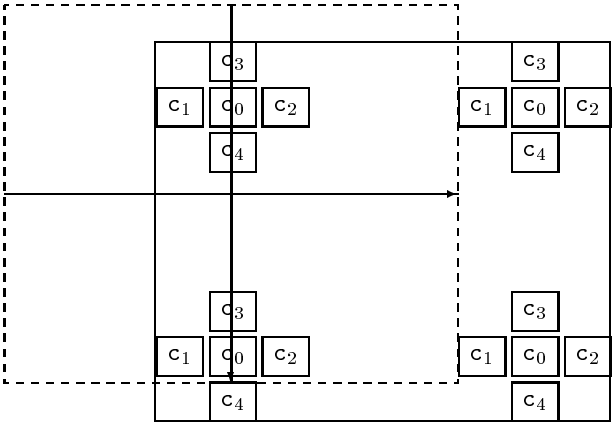

Figure 7.6: The layout tree at `LeftTop`, `RightTop`, `LeftBottom`, and `RightBottom`.

the real position of the control is `{x=x+vx, y=y+vy}`. Figure 7.8 illustrates this. Given two controls c_0 and c_1 it shows the result of placing c_1 at `RightTo` control c_0 with an offset value `v = {vx,vy}`. The dashed box shows the location of c_1 using a `zero` offset.

7.3.5 Layout relative to the previous control

As explained earlier in this section, the default layout attribute of a control is `(RightToPrev,zero)`. The other layout attributes that refer to the previous control are `LeftOfPrev`, `AbovePrev`, and `BelowPrev`. In this section we explain what the previous control is.

Section 7.2 introduced the glue to create control structures. The best way to look at such a control structure is to have a look at its *numbered* graph structure. Consider the following expression: `(a:+:b:+:c)` with `a`, `b`, and `c` standard `Controls` class instances as introduced in Section 7.1. Figure 7.9 shows the graph structure (recall that `:+:` is right associative).

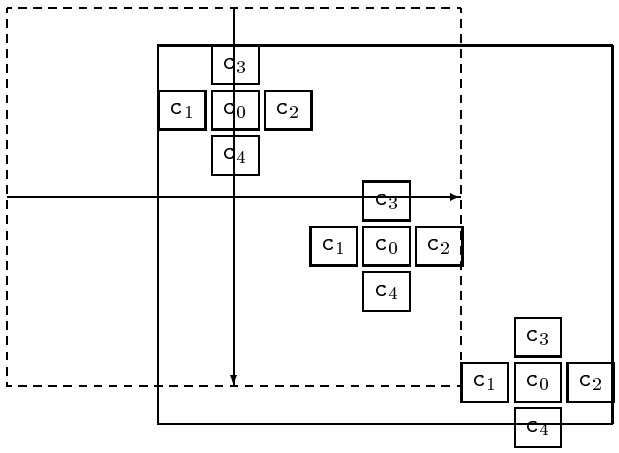

Figure 7.7: The layout tree at `Left`, `Center`, and `Right`.

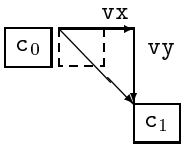

Figure 7.8: Laying out controls using an offset vector.

Each node in the graph has an index. If a node is a glue node, then first number the left sub tree, then the node itself, then the right sub tree. If a node is a standard `Controls` class instance, then number it. The nodes of the sub tree of a `CompoundControl` are not numbered. Proceeding in this way, one obtains the index figures at each of the nodes in Figure 7.9. If a node in the graph with index i represents one of the standard `Controls` class instances then its *previous* control is represented by that node in the graph that has the highest index less than i and represents also one of the standard `Controls` class instances. So the previous control of `c` is not $\texttt{:+:}_4$ but $\texttt{b}_3$ because we assumed that `b` is an instance of the standard `Controls` class. Analogously, the previous element of `b` is neither one of the two `:+:` nodes, but $\texttt{a}_1$. Finally, `a` has no previous control.

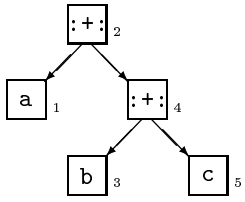

Figure 7.9: The numbered graph of `(a:+:b:+:c)`.

7.4 Resizing controls

The object I/O system has a simple mechanism to let controls respond to resize actions of their parent interface element (window, dialogue, or compound control). If a control wants to respond to resize events, it should have a `ControlResize` attribute. It is defined as follows:

```
::  ControlAttribute st
 =  ... | ControlResize ControlResizeFunction | ...
::  ControlResizeFunction
:== Size -> Size -> Size -> Size
```

The control resize function is applied to its current size, old size of its parent, and the new size of its parent. It returns its own new size. This calculation is performed for all controls that are part of the control that is being resized. If a `CompoundControl` has a resize function, and the new size is different from its previous size, then this computation continues recursively, otherwise the layout of its elements is not recalculated. Given the new sizes of the controls, the layout is recalculated and adjusted accordingly. The effect of this strategy is that the relative layout of controls is never changed in case of resizing a window, dialogue, or compound control.

As an example, consider one wants to have a `CompoundControl` that always displays three `CustomControl`s next to each other at the top of its view frame. The `CompoundControl` takes care that it always has a width dividable by 3, using its own `ControlResize` function `compoundresize`. Its `ControlLook` function draws a rectangle fitting its current view frame.

```
compound     = CompoundControl
                 (ListLS [custom,custom,custom])
                 [ControlResize     compoundresize
                 ,ControlSize       compoundsize
                 ,ControlLook       (\_ {newFrame}->draw newFrame)
                 ,CompoundHMargin   0 0
                 ,CompoundVMargin   0 0
                 ,CompoundItemSpace 0 0
                 ]
compoundsize = {w=60,h=75}
compoundresize _ _ newparentsize=:{w}
             = {newparentsize & w=w/3*3}
```

The `CustomControl`s resize their widths according to the new width of their parent control, using the `ControlResize` function `customresize`. The look of the `CustomControl` simply draws a rectangle fitting its current view frame and the two diagonals.

```
custom
 = CustomControl {w=compoundsize.w/3,h=6} look
                 [ControlResize customresize]
customresize customsize _ {w}
 = {customsize & w=w/3}
look _ {newFrame} picture
 # picture = draw     newFrame                           picture
 # picture = drawLine newFrame.corner1 newFrame.corner2 picture
```

```
 = drawLine {newFrame.corner1 & y=newFrame.corner2.y}
            {newFrame.corner2 & y=newFrame.corner1.y}   picture
```

Figure 7.10 shows what happens with the controls when the parent object is resized. At the left the initial state of the `CompoundControl` and its `CustomControls` is displayed. As explained in Section 7.3, the three custom controls form a layout tree with the layout root control having the layout attribute `(Left,zero)` and the other `CustomControls` `(RightToPrev,zero)`. In the middle, the `CompoundControl` is resized to the right and bottom. This resize action causes first recalculation of the size of the `CompoundControl`, using `compoundresize`. Because this value differs from the old size, recalculation continues for each `CustomControl`. The final result is shown at the right.

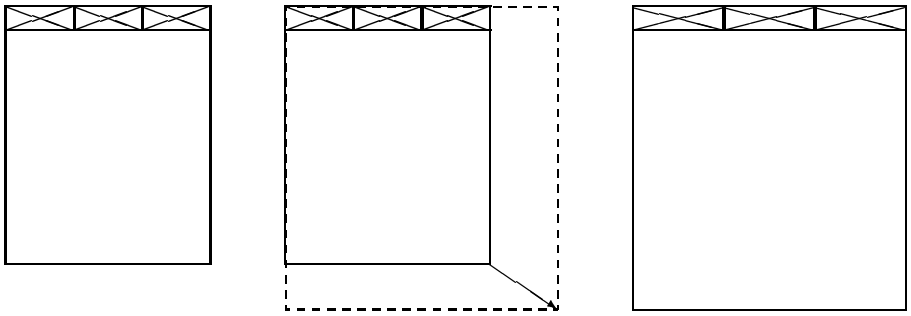

Figure 7.10: Resizing a `CompoundControl` with three `CustomControls`.

7.5 Examples

In this section a number of examples are given to clarify the use of controls.

7.5.1 Keyspotting revisited

In this example we extend the keyspotting example in Section 6.7.1 with the possibility to monitor also the keyboard input of controls. We create a `Window` that contains a `CompoundControl` that contains a `CustomControl`. Before we discuss each of the components below, we have a look at the way they handle keyboard input. A screenshot of the program is given in Figure 7.11.

Each of the components is keyboard sensitive and uses the same `KeyboardFunction` `spotting`. This function is almost identical to the one presented in Section 6.7.1. The only difference is that the keyboard input is now shown in the `CustomControl` instead of the window. For this purpose `spotting` is also parameterised with the `Id` of the `CustomControl`. It is also parameterised with a string that states who currently has the input focus. Now `spotting` changes the `Look` function of the `CustomControl` and forces its update (by using a `True` boolean).

```
spotting (wid,cid) who x pst
=appPIO (setWindow wid [setControlLooks [(cid,True,look text)]]) pst
where
 text = who+++":"+++toString x
```

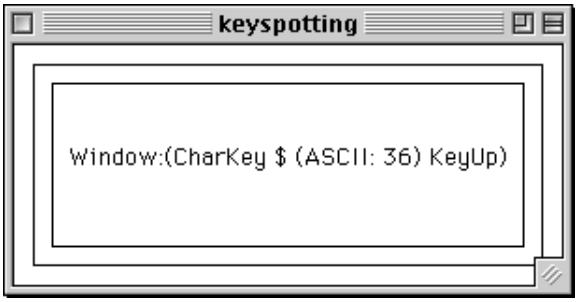

Figure 7.11: The keyspotting program in action.

The `Look` function `look` of the `CustomControl` is almost identical to the original look function: except that it centers its argument string it also draws a framed box around it, so that we can easily see where the components are.

```
look text _ {newFrame} picture
  # picture          = unfill newFrame picture
  # picture          = draw   newFrame picture
  # (width,picture) = getPenFontStringWidth text picture
  = drawAt {x=(w-width)/2,y=h/2} text picture
where
  {w,h}              = rectangleSize newFrame
```

The `CustomControl custom` displays which component is currently receiving what keyboard input. The control is identified by `cid`. It parameterises its `KeyboardFunction spotting` with the required `Id`s and the string `"Control"`. Its initial look draws a box around itself. The control is resizable, specified by adding the `ControlResize` attribute. Whenever the parent object is resized, `custom` changes its size in exactly the same amount.

```
custom
= CustomControl customsize (look "")
  [ControlKeyboard (const True) Able
                   (noLS1 (spotting ids "Control"))
  ,ControlId       cid
  ,ControlResize   resize
  ]

resize oldCSize oldParentSize newParentSize
= { w = oldCSize.w+newParentSize.w-oldParentSize.w
  , h = oldCSize.h+newParentSize.h-oldParentSize.h
  }
```

The `CompoundControl compound` contains only `custom`. It parameterises its `KeyboardFunction spotting` with the required `Id`s and the string `"Compound"`. Its initial size is chosen such that it is large enough to display `custom` completely. It conveniently uses the `look` function to draw a box around itself. The compound control is also resizable, and uses the same resize function as `custom`.

```
compound
= CompoundControl custom
```

```
  [ControlKeyboard (const True) Able
                   (noLS1 (spotting ids "Compound"))
  ,ControlSize     {w=customsize.w+2*margin
                   ,h=customsize.h+2*margin
                   }
  ,ControlResize   resize
  ,ControlLook     (look "")
  ]
```

Finally, the `Window window` contains only `compound`. It parameterises its `spotting` function with the required `Ids` and the string `"Window"`. Its initial size is chosen such that it is large enough to display `compound` completely. For this purpose also the margin layout attributes are set. Termination of the program is taken care of by having the program quit when the user closes the window.

```
window
= Window "keyspotting" compound
  [WindowKeyboard (const True) Able (noLS1 (spotting ids "Window"))
  ,WindowId       windowid
  ,WindowSize     {w=customsize.w+4*margin
                  ,h=customsize.h+4*margin
                  }
  ,WindowResize
  ,WindowHMargin  margin margin
  ,WindowVMargin  margin margin
  ,WindowClose    (noLS closeProcess)
  ]
```

Remaining details that need to be defined are the actual creation of the interactive program and its window. For completeness, we include the program code of `keyspotting` below.

```
module keyspotting

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program monitors keyboard input that is sent to a Window which consists
//  of a CompoundControl which consists of a CustomControl.
//  ********************************************************************************

import StdEnv,StdIO

::  NoState
    = NoState

Start :: *World -> *World
Start world
    # (wid,world)   = openId world
    # (cid,world)   = openId world
    # ids           = (wid,cid)
    # custom        = CustomControl customsize (look "")
                      [   ControlKeyboard (const True) Able
                                          (noLS1 (spotting ids "Control"))
                      ,   ControlId       cid
                      ,   ControlResize   resize
                      ]
    # compound      = CompoundControl custom
                      [   ControlKeyboard (const True) Able
```

```
                                        (noLS1 (spotting ids "Compound"))
                        ,   ControlSize     {w=customsize.w+2*margin
                                            ,h=customsize.h+2*margin
                                            }
                        ,   ControlResize   resize
                        ,   ControlLook     (look "")
                        ]
    # window        = Window "keyspotting" compound
                        [   WindowKeyboard  (const True) Able
                                            (noLS1 (spotting ids "Window"))
                        ,   WindowId        wid
                        ,   WindowSize      {w=customsize.w+4*margin
                                            ,h=customsize.h+4*margin
                                            }
                        ,   WindowResize
                        ,   WindowHMargin   margin margin
                        ,   WindowVMargin   margin margin
                        ,   WindowClose     (noLS closeProcess)
                        ]
    = startIO NoState NoState [snd o openWindow NoState window]
                              [ProcessClose closeProcess] world
where
    customsize  = {w=550,h=100}
    margin      = 10
    resize oldCSize oldParentSize newParentSize
        = { w = oldCSize.w+newParentSize.w-oldParentSize.w
          , h = oldCSize.h+newParentSize.h-oldParentSize.h
          }
    spotting (wid,cid) who x pst
        = appPIO (setWindow wid [setControlLooks [(cid,True,look text)]]) pst
    where
        text    = who+++":"+++toString x
    look text _ {newFrame} picture
        # picture           = unfill newFrame picture
        # picture           = draw   newFrame picture
        # (width,picture)   = getPenFontStringWidth text picture
        = drawAt {x=(w-width)/2,y=h/2} text picture
    where
        {w,h}               = rectangleSize newFrame
```

7.5.2 Mousespotting revisited

In this example we extend the mouse spotting example of Section 6.7.2. Just like then, the revised mouse spotting example is almost identical to the revised keyspotting example discussed above. The only differences are the title of the window and the replacement of the keyboard attributes by mouse attributes. A screenshot of the program is given in Figure 7.12. For completeness, we show the code of the revised mouse spotting example below.

```
module mousespotting

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program monitors mouse input that is sent to a Window which consists
//  of a CompoundControl which consists of CustomControl.
//  ********************************************************************************

import StdEnv,StdIO

::  NoState
    =   NoState

Start :: *World -> *World
```

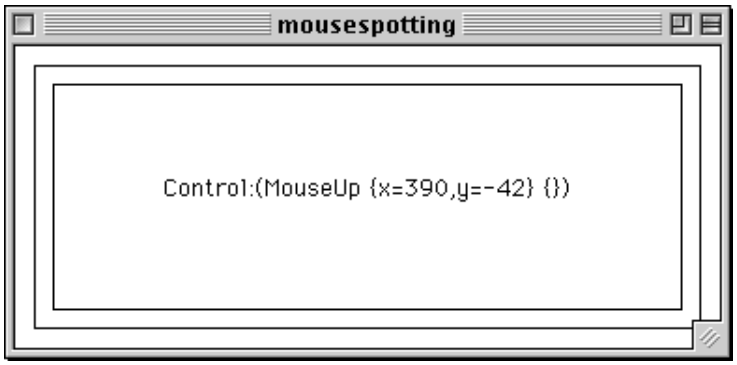

Figure 7.12: The mousespotting program in action.

```
Start world
    # (wid,world)   = openId world
    # (cid,world)   = openId world
      ids           = (wid,cid)
    # custom        = CustomControl customsize (look "")
                        [   ControlMouse    (const True) Able
                                            (noLS1 (spotting ids "Control"))
                        ,   ControlId       cid
                        ,   ControlResize   resize
                        ]
    # compound      = CompoundControl custom
                        [   ControlMouse    (const True) Able
                                            (noLS1 (spotting ids "Compound"))
                        ,   ControlSize     {w=customsize.w+2*margin
                                            ,h=customsize.h+2*margin
                                            }
                        ,   ControlResize   resize
                        ,   ControlLook     (look "")
                        ]
    # window        = Window "mousespotting" compound
                        [   WindowMouse     (const True) Able
                                            (noLS1 (spotting ids "Window"))
                        ,   WindowId        wid
                        ,   WindowSize      {w=customsize.w+4*margin
                                            ,h=customsize.h+4*margin
                                            }
                        ,   WindowResize
                        ,   WindowHMargin   margin margin
                        ,   WindowVMargin   margin margin
                        ,   WindowClose     (noLS closeProcess)
                        ]
    = startIO NoState NoState [snd o openWindow NoState window]
                              [ProcessClose closeProcess] world
where
    customsize  = {w=550,h=100}
    margin      = 10
    resize oldCSize oldParentSize newParentSize
        = { w = oldCSize.w+newParentSize.w-oldParentSize.w
          , h = oldCSize.h+newParentSize.h-oldParentSize.h
          }
    spotting (wid,cid) who x pst
        = appPIO (setWindow wid [setControlLooks [(cid,True,look text)]]) pst
    where
        text    = who+++":"+++toString x
    look text _ {newFrame} picture
        # picture           = unfill newFrame picture
        # picture           = draw   newFrame picture
        # (width,picture)   = getPenFontStringWidth text picture
```

```
    = drawAt {x=(w-width)/2,y=h/2} text picture
where
    {w,h}               = rectangleSize newFrame
```

Chapter 8

Menus

Many interactive applications allow a user to manipulate a number of documents. As we saw in the previous chapter, the user can issue these manipulations by means of the keyboard and mouse. Another common source of manipulations is by issueing *commands* to the application. This is where *menus* come in. Menus help a program to structure the set of available commands. To the user of an application, the use of menus provides a consistent and easily browsable graphical display of the set of available commands. For these reasons it is recommended to use menus in a program.

In Section 8.1 we introduce the standard set of menu definitions that are at a programmers disposal. Then the glue is introduced to create larger menu structures in Section 8.2. One special menu is available for interactive processes that have the multiple document interface (MDI) attribute, the *windows* menu. This menu enumerates the current open and visible windows of that process and give some commands to organise them. This is treated in Section 8.3. Menus provide a consistent graphical interface to users. To enhance the consistency, a number of programming conventions have evolved. These are discussed in Section 8.4.

8.1 Menus and menu elements

Menus and menu elements can be defined by means of the type definitions in module `StdMenuDef` (Appendix A.16). Analogous to `Windows`, `Dialogs`, and `CompoundControls`, the type constructor class `Menus` is parameterised with a type constructor variable.

```
class Menus mdef where
    openMenu :: .lst !(mdef .lst (PSt .l .p)) !(PSt .l .p)
                              -> (!ErrorReport,!PSt .l .p)

instance Menus (Menu m) | MenuElements m
```

The admissable instances have to belong to the `MenuElements` type constructor class. Below we introduce each of the components.

8.1.1 The menu attributes

The menu attributes are used by both menus and menu elements. They are the following:

```
:: MenuAttribute st
 = MenuId           Id
 | MenuSelectState  SelectState
 | MenuIndex        Int
 | MenuShortKey     Char
 | MenuMarkState    MarkState
 | MenuFunction     (IdFun             st)
 | MenuModsFunction (ModifiersFunction st)
```

The `MenuId` attribute identifies the menu or menu element to which it is associated. If you do not provide a `MenuId` the menu (element) can not be modified.

The `MenuSelectState` attribute defines whether the menu (element) can be used by the user (`Able`) or not (`Unable`). Usually this will affect the look of the menu (element). The default value is `Able`.

The `MenuIndex` attribute defines the index position of a menu. Index positions range from one (for the first menu) upto the number of menus. A negative or zero `MenuIndex` attribute value will place the menu in front of all current menus. A `MenuIndex` attribute value that is larger than the current number of menus will place the menu behind all current menus. Other `MenuIndex` attribute values place the menu *behind* the menu with that index value.

The `MenuShortKey` attribute defines a character that can be used by user to select the menu element to which the character is associated by means of the keyboard. The menu element can be selected by pressing that character and some special, platform dependent meta key.

The `MenuMarkState` attribute can add a check mark symbol (`Mark`) or leave it out (`NoMark`) to a menu element. The default value is `NoMark`.

The `MenuFunction` and `MenuModsFunction` attributes add callback functions to menu elements that are evaluated when the menu element to which they are associated is selected by the user. The difference between these two attributes is that the former is simply evaluated whenever the menu element is selected, and that the latter also provides the callback function with the modifier keys that have been pressed at the moment of selecting the menu element (for the definition of the modifier). In a `MenuAttribute` list, the first of these two attributes is chosen.

8.1.2 The Menu

A *menu* is a top level interface element that contains a group of related commands. The definition of a menu is as follows:

```
:: Menu m lst pst
 = Menu Title (m lst pst) [MenuAttribute *(lst,pst)]
```

The usual appearance of a menu is by its title. The user can browse through its commands by mouse or by a platform dependent keyboard interface. Valid menu attributes are:

MenuAttribute:	Valid:
MenuId	√
MenuSelectState	√
MenuIndex	√
MenuShortKey	
MenuMarkState	
MenuFunction	
MenuModsFunction	

Example The left picture shows the menu when not selected by the user, the right picture when selected by the user.

```
menu
  = Menu "Menu"
      (   MenuItem "Open..." [MenuShortKey    'o']
      :+: MenuItem "Close"   [MenuSelectState Unable
                             ,MenuShortKey    'w'
                             ]
      :+: MenuSeparator      []
      :+: MenuItem "Quit"    [MenuShortKey    'q']
      )   []
```

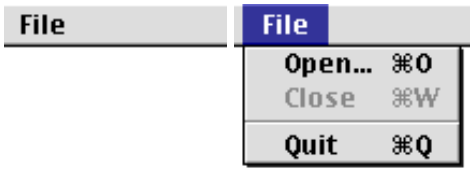

8.1.3 The MenuItem

The *menu item* is the standard element that refers to a command. The definition of a menu item is as follows:

```
:: MenuItem lst pst = MenuItem Title [MenuAttribute *(lst,pst)]
```

The title of a menu element is displayed as a member of its parent menu. When the user selects the menu item its `Menu(Mods)Function` attribute is evaluated if the menu item, and all of its parent menus are `Able`. The appearence of the menu item reflects this state. Valid menu item attributes are:

MenuAttribute:	Valid:
MenuId	√
MenuSelectState	√
MenuIndex	
MenuShortKey	√
MenuMarkState	√
MenuFunction	√
MenuModsFunction	√

An example of a menu item is given in the previous section.

8.1.4 The MenuSeparator

The *menu separator* is a menu element that is only used to separate groups of related menu elements within a parent menu. Graphically, a menu separator usually inserts some vertical space and a horizontal divider within a menu. Menu separators have no further functionality. The definition of a menu separator is as follows:

```
:: MenuSeparator lst pst = MenuSeparator [MenuAttribute *(lst,pst)]
```

Because menu separators have no other purpose than providing some 'white space' between menu elements, the only valid menu separator attribute is the `MenuId` (which can be used to remove a separator dynamically):

MenuAttribute:	Valid:
MenuId	√
MenuSelectState	
MenuIndex	
MenuShortKey	
MenuMarkState	
MenuFunction	
MenuModsFunction	

In the menu example of Section 8.1.2 a menu separator has been used.

8.1.5 The RadioMenu

A *radio menu* element is a group of menu items of which exactly one menu item is selected. All alternatives are visible. The definition of a radio menu is as follows:

```
:: RadioMenu  lst pst = RadioMenu [MenuRadioItem *(lst,pst)] Index
                                  [MenuAttribute *(lst,pst)]
:: MenuRadioItem st :== (Title,Maybe Id,Maybe Char,IdFun st)
```

The initially selected item is indicated by the `Index` value. As a convention in the object I/O library, when indicating elements indices range from 1 upto the number of elements. So n elements are indexed by $1 \ldots n$. In case the index is out of range, i.e. less than 1 or larger than n, it is set to 1 and n respectively. Valid radio menu attributes are:

MenuAttribute:	Valid:
MenuId	√
MenuSelectState	√
MenuIndex	
MenuShortKey	√
MenuMarkState	
MenuFunction	
MenuModsFunction	

When an item of the radio menu is selected the previously selected radio menu item will be unchecked, and the new radio menu item gets the check mark. The corresponding callback function is then evaluated. The callback function is also evaluated if the currently selected radio menu item is selected.

Example A radio menu and its initial look (it is instructive to compare this with the radio *control* example at page 71).

```
radiomenu
  = RadioMenu
    [  ("Radio item "+++toString i,Nothing,Just (iChar i),id)
    \\ i<-[1..5]
    ]  1 []
where
  iChar i = toChar (toInt '1'+i-1)
```

8.1.6 The SubMenu

The *sub menu* is a menu element that contains other menu elements. So it is a menu within a menu, and of course can contain sub menus as well. The definition of a sub menu is as follows:

```
:: SubMenu m lst pst
 = SubMenu Title (m lst pst) [MenuAttribute *(lst,pst)]
```

The usual appearance of a sub menu is by its title. The user can browse through its elements by mouse or by a platform dependent keyboard interface. Valid sub menu attributes are:

MenuAttribute:	Valid:
MenuId	√
MenuSelectState	√
MenuIndex	
MenuShortKey	
MenuMarkState	
MenuFunction	
MenuModsFunction	

Example This sub menu has the same definition as the menu at page 97 except that it has a `SubMenu` data constructor rather than the `Menu` data constructor and a different title. The left picture shows the sub menu when not selected by the user, the right picture when selected by the user.

```
submenu
  = SubMenu "Sub Menu"
      (   MenuItem "Open..." [MenuShortKey    'o']
      :+: MenuItem "Close"   [MenuSelectState Unable
                             ,MenuShortKey    'w'
                             ]
      :+: MenuSeparator      []
      :+: MenuItem "Quit"    [MenuShortKey    'q']
      )   []
```

8.2 Menu glue

In the previous section the standard set of menus and menu elements has been discussed. This list is not complete. In the library module `StdMenuElementClass` (Appendix A.18) a number of additional instances are defined, namely the type constructors `:+:`, `ListLS`, `NilLS`, and `AddLS`, `NewLS` (their definition can be found in Appendix A.13). These additional instances are required to *glue* menus. They are treated below.

8.2.1 `:+:`

The most common constructor to glue menu elements is `:+:`. Its type constructor definition and `MenuElements` class instance declaration are as follows:

```
::  :+: t1 t2 local context
 = (:+:) infixr 9 (t1 local context) (t2 local context)

instance MenuElements ((:+:) m1 m2) | MenuElements m1
                                    & MenuElements m2
```

Given two `MenuElements` instances `m1` and `m2`, working on the same local state of type `local` and context state `context`, the expression `m1:+:m2` is also a `MenuElements` instance working on the same local state and context state. Because `:+:` is right associative, the expression `m1:+:m2:+:m3` should be read as `m1:+:(m2:+:m3)`.

8.2.2 `ListLS` and `NilLS`

In principle the `:+:` glue is sufficient to create all required menu element structures. In case of working with a number of menu instances of the *same type*, it is much more convenient to use lists and list comprehensions. This glue is provided by the type constructor `ListLS`. The type constructor `NilLS` is a shorthand for `ListLS []`. It can also be conveniently used to state that a `SubMenu` or `Menu` has no menu elements. Their type constructor definitions and `MenuElements` class instance declarations are as follows:

```
::  ListLS t local context = ListLS [t local context]
::  NilLS    local context = NilLS

instance MenuElements (ListLS m) | MenuElements m
instance MenuElements NilLS
```

Given a list of `MenuElements` instances `ms = [`m_1 `...`m_n`]`, working on the same local state of type `local` and context state `context`, the expression `ListLS ms` is also a `MenuElements` instance working on the same local state and context state.

8.2.3 AddLS and NewLS

The previously discussed glueing type constructors always glue menu elements that work on the same local state and context state. Two other glueing constructors are defined to extend and change the local state, `AddLS` and `NewLS`.

```
:: AddLS t local context
 = E..add:
   { addLS :: add
   , addDef:: t *(add,local) context
   }

instance MenuElements (AddLS m) | MenuElements m
```

Given a `MenuElements` instance `m1` that works on a local state of type `local` and a context state of type `context`, one can add another `MenuElements` instance `m2` that works on an *extended* local state of type `(add,local)` and the same context state of type `context`. Let `x` be a value of type `add`, then this is done by the expression `m1 :+: {addLS=x, addDef=m2}`.

```
:: NewLS t local context
 = E..new:
   { newLS :: new
   , newDef:: t new context
   }

instance MenuElements (NewLS m) | MenuElements m
```

Given a `MenuElements` instance `m1` that works on a local state of type `local` and a context state of type `context`, one can add another `MenuElements` instance `m2` that works on a *new* local state of type `new` and the same context state of type `context`. Let `x` be a value of type `new`, then this is done by the expression `m1 :+: {newLS=x, newDef=m2}`.

In both cases the extended part of the local state and the new local state are encapsulated completely from the external context using existential quantification.

8.3 The Windows menu

For programs using interactive processes that have the multiple document interface (MDI processes, see Section 11.1) one special menu is added by the object I/O system to the menu system of such an interactive process. This is the *Windows menu*. This menu contains a number of commands to arrange the current set of visible windows, and a list of all window titles. The set of arrange commands is platform dependent, but contains at least the following three commands:

Cascade: this command arranges all windows to have equal size. They will be placed in diagonal order: from the left top to the right bottom of the process window. The windows will always be shown completely inside the process window. If there are to many windows, the arrangements starts over again for the remaining windows, which will therefore overlap the other windows.

Tile Horizontally: this command arranges all windows in rows without overlapping. It is tried to give all windows the same size.

Tile Vertically: this command arranges all windows in columns without overlapping. It is tried to give all windows the same size.

Separated from the window arrange commands by a `MenuSeparator` follows the list of open windows in lexicographical order on the window title. The currently active window is checked. Selecting any of these windows makes that window the active window. (This causes the previously active window to become deactivated, and the new window to become active.)

8.4 Menu conventions

The use of menus provides application users with a consistent and uniform access to the available set of commands. In this section we discuss a number of conventions that are usually followed to increase the level of consistency.

8.4.1 Subsetting the available commands

In general when using an interactive application, the application will move through several states. In each state a particular subset of the complete set of available commands will be applicable to the user while the remaining commands should not be selected. A well designed application should make this clear to the user by *subsetting* the available commands.

The easiest way to subset commands is by *disabling* and *enabling* the menu elements that should be unselectable and selectable respectively. For this purpose the functions `enableMenuElements` and `disableMenuElements` (module `StdMenuElement`, Appendix A.17) are available to enable and disable individual menu elements. Complete menus can be enabled and disabled using the functions `enableMenus` and `disableMenus` (module `StdMenu` (Appendix A.15). This module also contains two functions to enable and disable the whole current set of menus: `enableMenuSystem` and `disableMenuSystem`.

8.4.2 Command conventions

In this section we discuss some conventions that are found frequently in many applications with respect to commands.

Clipboard commands

Applications that support the use of the clipboard (Chapter 12) to *cut*, *copy*, and *paste* private and external data usually are found in an `"Edit"` menu. Conventions are:

cut This command should be enabled only if the application is in a state that an object has been selected that can be transfered to the clipboard. Issueing this command should remove that object from its context and place it in the clipboard. Its name should be `"Cut"` and it should have the shortkey attribute `'x'`.

copy This command should be enabled only if the application is in a state that an object has been selected that can be transfered to the clipboard. Issueing this command should place it in the clipboard, but not remove it from its context. Its name should be `"Copy"` and it should have the shortkey attribute `'c'`.

paste This command should be enabled only if the clipboard contains an object that can be currently incorporated in the application. Issueing this command should read the clipboard and put that object in the application. Its name should be `"Paste"` and it should have the shortkey attribute `'v'`.

Undo command

Applications that allow users to manipulate documents by sequences of commands can support an *undo* command. The undo command can also be undone by the *redo* command. These commands are usually found in an `"Edit"` menu. Conventions are:

undo This command should be enabled only if the user has issued a sequence of commands that can be undone. The number of undoable commands depends on the sophistication of the application. The name of this command is `"Undo"` and it has the shortkey attribute `'z'`.

redo This command should be enabled only if a (sequence of) undo command has been issued. At each selection it restores the changes of the undo command. The name of this command is `"Redo"` and it has the shortkey attribute `'y'`.

Document commands

The document commands are frequently found commands to create new documents, open existing documents, and save and close open documents. These commands are usually found in a `"File"` menu. Conventions are:

new This command should be enabled only if the application can modify a new document. Issueing this command should create or reuse a window containing the new document. Its name should be `"New"` and it should have the shortkey attribute `'n'`.

open This command should be enabled only if the application can modify an additional, existing document. Issueing this command should give the user the opportunity to search for a file that will be opened by the application. For this purpose the `StdFileSelect` function `selectInputFile` can be used (Appendix A.9). The name of the command should be `"Open..."` and it should have the shortkey attribute `'o'`.

close This command should be enabled only if the application has an open document window or dialogue. Issueing this command should close the currently active window or dialogue. It is good programming practice to check if the document has been recently saved. If this is not the case, then the user should be asked if the document should be saved before closing. The name of this command should be `"Close"` and it should have the shortkey attribute `'w'`.

save This command should be enabled only if the currently active document window version differs from a (possibly not present) file version. Issueing this command should save the current state of the document in the active window to file. If there is no file associated yet, then the application should first ask for a file name. For this purpose the `StdFileSelect` function `selectOutputFile` can be used (Appendix A.9). The name of the command should be `"Save"` and it should have the shortkey attribute `'s'`.

save as This command should be enabled only if the application has an open document window. Issueing this command should give the user the possibility to browse the file system and provide a file name. For this purpose the `StdFileSelect` function `selectOutputFile` can be used (Appendix A.9). The name of the command should be `"Save As..."`

Quit command

Users can leave an application using the *quit* command. A user should always be allowed to quit the application. It is good programming practice to check if there are any unsaved documents in the application. If this is the case then the user should be asked if these documents should be saved before closing. The name of the quit command is usually `"Quit"` and has the shortkey attribute `'q'`. This command is usually found in a `"File"` menu.

8.5 Example: a small menu system

In this section we show a program that creates an interactive process with a multiple document interface. It can serve as a framework for writing your own multiple document interface programs. The interactive process contains one menu, titled `"File"`, which consists of three commands: a command to open a new window, titled `"New"`; a command to close the active window, titled `"Close"`; and a command to terminate the interactive process, titled `"Quit"`. The `"New"` and `"Quit"` commands are always available, the `"Close"` command will be subsetted (as described in Section 8.4.1). We start with the commands, and proceed with the menu and process definition.

The `"New"` command is a `MenuItem` with a local state of type `Int`. We use it to give every new window a new title by appending its current value to the string `"Window"`. The initial value is one. Here is the definition of the command:

```
{ newLS  = 1
, newDef = MenuItem "New" [MenuShortKey 'n',MenuFunction new]
}
```

The callback function of the `"New"` command, `new`, uses the local integer state. For each successfully created window the value is incremented. In addition, the `"Close"` command is enabled. If the window can not be created a notice is opened, using the notice library developed in Section 6.8.1.

```
new :: (Int,PSt .l .p) -> (Int,PSt .l .p)
new (i,pst)
# (error,pst) = openWindow 0 window pst
| error<>NoError
  # notice   = Notice ["MDI could not open new window"]
               ( NoticeButton "Ok" id
               ) []
  = (i,openNotice notice pst)
| otherwise
  = (i+1,appPIO (setMenu fileid [enableMenuElements [closeid]]) pst)
```

The window itself is kept very simple. If the user requests to close the window, the `"Close"` command callback function `close` is evaluated. This is done by setting

the (WindowClose (noLS close)) attribute. The look of the window draws the current view frame rectangle and the lines between the diagonally opposite corner points.

```
window = Window ("Window "+++toString i)
           NilLS
           [   WindowClose (noLS close)
           ,   WindowSize  {w=300,h=300}
           ,   WindowLook  look
           ]

look :: SelectState UpdateState *Picture -> *Picture
look _ {newFrame=frame=:{corner1,corner2}} picture
 # picture = draw frame picture
 # picture = drawLine corner1 corner2 picture
 = drawLine {corner1 & x=corner2.x} {corner2 & x=corner1.x} picture
```

The "Close" command is identified by the Id closeid. Because the application initially has no open windows, the initial SelectState is Unable:

```
 MenuItem "Close" [MenuShortKey 'w',MenuFunction (noLS close)
                  ,MenuId closeid,  MenuSelectState Unable
                  ]
```

The callback function of the "Close" command, close, has to close the active window. It retrieves the Id of the active window using the StdWindow function getActiveWindow and closes that window using closeWindow. If the active window was the last open window, then close should disable the "Close" command. The function getWindowsStack returns the list of Ids of all currently open windows. So if this list is empty, then there are no more windows open. In that case the "Close" command is disabled. We get the following definition:

```
close :: (PSt .l .p) -> PSt .l .p
close pst
 # (maybeId,pst)  = accPIO getActiveWindow pst
 | isNothing maybeId
   = abort "Fatal error in MDI: getActiveWindow returned Nothing."
 # pst            = closeWindow (fromJust maybeId) pst
 # (rest,pst)     = accPIO getWindowsStack pst
 | isEmpty rest
   = appPIO (setMenu fileid [disableMenuElements [closeid]]) pst
 | otherwise
   = pst
```

It should be noted that because of subsetting the "Close" command, it is not possible that getActiveWindow returns Nothing. It is good programming practice to check this situation and handle it even though it should not occur. In this case we have chosen to abort the program.

The "Quit" command has a straightforward definition:

```
 MenuItem "Quit" [MenuShortKey 'q',MenuFunction (noLS quit)]
```

The callback function of "Quit", quit, asks the user to confirm this choice. For this purpose we again use a notice. If the user confirms his choice, then the application is terminated, otherwise nothing happens. Figure 8.1 shows the notice.

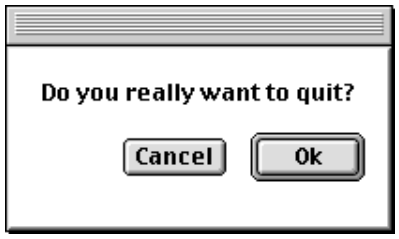

Figure 8.1: The termination notice.

Figure 8.2: The initial menu.

```
quit :: (PSt .l .p) -> PSt .l .p
quit pst
    = openNotice notice pst
where
    notice  = Notice ["Do you really want to quit?"]
                (   NoticeButton "Ok"     (noLS closeProcess))
                [   NoticeButton "Cancel" id
                ]
```

The definition of the `"File"` menu simply glues the definitions of the commands as described above into one menu definition (Figure 8.2 shows the menu when opened).

```
menu
 = Menu "File"
   ( { newLS = 1
     , newDef= MenuItem "New"
                        [MenuShortKey 'n',MenuFunction (noLS new)]
     }
   :+: MenuItem "Close" [MenuShortKey 'w',MenuFunction (noLS close)
                        ,MenuId closeid,  MenuSelectState Unable
                        ]
   :+: MenuSeparator    []
   :+: MenuItem "Quit"  [MenuShortKey 'q',MenuFunction (noLS quit)]
   )
   [   MenuId fileid
   ]
```

The remaining details that need to be arranged are the creation of the two `Id` values `fileid` and `closeid`, the creation of the menu, and the creation of the parent interactive process of the menu. We conveniently reuse the `quit` function to handle user requests to close the interactive process by setting the (`ProcessClose quit`) attribute. Here is the full program code.

```
module MDI
```

```
//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates a Multiple Document Interface process with a Window menu.
//  ********************************************************************************

import StdEnv, StdIO
import notice

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    # (ids,world)   = openIds 2 world
    = startIO NoState NoState [initialise ids] [ProcessClose quit] world

quit :: (PSt .l .p) -> PSt .l .p
quit pst
    = openNotice notice pst
where
    notice  = Notice ["Do you really want to quit?"]
                (   NoticeButton "Ok"     (noLS closeProcess))
                [   NoticeButton "Cancel" id
                ]

initialise :: [Id] (PSt .l .p) -> PSt .l .p
initialise ids=:[fileid,closeid] pst
    # (error,pst)   = openMenu 0 menu pst
    | error<>NoError= abort "MDI could not open File Menu"
    | otherwise     = pst
where
    menu= Menu "File"
            ( { newLS  = 1
              , newDef = MenuItem "New" [MenuShortKey 'n',MenuFunction new]
              }
            :+: MenuItem "Close"        [MenuShortKey 'w',MenuFunction (noLS close)
                                        ,MenuId closeid,  MenuSelectState Unable
                                        ]
            :+: MenuSeparator           []
            :+: MenuItem "Quit"         [MenuShortKey 'q',MenuFunction (noLS quit)]
            )
            [   MenuId fileid
            ]

    new :: (Int,PSt .l .p) -> (Int,PSt .l .p)
    new (i,pst)
        # (error,pst)   = openWindow 0 window pst
        | error<>NoError
            # notice    = Notice ["MDI could not open new window"]
                            (   NoticeButton "Ok" id
                            )   []
            = (i,openNotice notice pst)
        | otherwise
            = (i+1,appPIO (setMenu fileid [enableMenuElements [closeid]]) pst)
    where
        window  = Window ("Window "+++toString i)
                    NilLS
                    [   WindowClose (noLS close)
                    ,   WindowSize  {w=300,h=300}
                    ,   WindowLook  look
                    ]

        look :: SelectState UpdateState *Picture -> *Picture
        look _ {newFrame=frame=:{corner1,corner2}} picture
            # picture   = draw frame picture
```

```
        # picture   = drawLine corner1 corner2 picture
        = drawLine {corner1 & x=corner2.x} {corner2 & x=corner1.x} picture

close :: (PSt .l .p) -> PSt .l .p
close pst
    # (maybeId,pst) = accPIO getActiveWindow pst
    | isNothing maybeId
        = abort "Fatal error in MDI: getActiveWindow returned Nothing."
    # pst           = closeWindow (fromJust maybeId) pst
    # (rest,pst)    = accPIO getWindowsStack pst
    | isEmpty rest
        = appPIO (setMenu fileid [disableMenuElements [closeid]]) pst
    | otherwise
        = pst
```

Chapter 9

Timers

Timers provide interactive programs with a tool to let actions occur at regular time intervals. These actions respond to *timer events*, and so they can be properly defined as callback functions. Typical examples of timer uses are blinking cursors, clocks, and time-out mechanisms.

The definition types of timers can be found in module `StdTimerDef`, Appendix A.36. The main type definitions are as follows:

```
::  Timer t lst pst
 =  Timer TimerInterval (t lst pst) [TimerAttribute *(lst,pst)]
::  TimerInterval
:== Int

::  TimerAttribute   st
 =  TimerId          Id
 |  TimerSelectState SelectState
 |  TimerFunction    (TimerFunction st)

::  TimerFunction    st
:== NrOfIntervals->st->st
```

A `TimerInterval` is an integer value that must be at least zero. The time unit is platform dependent and is defined by the function `ticksPerSecond` (module `StdSystem`, Appendix A.30).

The `TimerAttribute`s are the following:

`TimerId` This attribute identifies the timer. If you do not provide a `TimerId` the timer can not be modified.

`TimerSelectState` This attribute defines whether the timer will respond to timer events (`Able`) or not (`Unable`). The default value is `Able`.

`TimerFunction` This attribute is the callback function that is evaluated when a timer event is handled. Its first argument is the number of whole timer intervals that have elapsed since its previous evaluation, so this value is at least 1. If the timer interval is zero, then this number is always 1. Enabling a timer from a disabled state resets the last evaluation time.

The `Timer` type constructor is parameterised with a type constructor variable. Analogous to menus that contain menu elements, timers can contain *timer elements*. The

instances must be member of the `TimerElements` class. This is expressed by the `Timers` timer *creation* member function, `openTimer` which can be found in module `StdTimer` (Appendix A.35):

```
class Timers tdef where
  openTimer :: .lst !(tdef .lst (PSt .l .p)) !(PSt .l .p)
               -> (!ErrorReport,!PSt .l .p)

instance Timers (Timer t) | TimerElements t
```

Currently, the instances of the `TimerElements` class are *receivers* and the usual *glueing* type constructors that we have already encountered in controls (Section 7.2) and menu elements (Section 8.2), namely `:+:`, `ListLS`, `NilLS`, and `AddLS`, `NewLS`. The receiver instances are declared in module `StdTimerReceiver` (Appendix A.38). Receivers are handled in Chapter 10. The glueing constructors are declared in the same module `StdTimerElementClass` that contains the `TimerElements` class (Appendix A.37).

9.1 Examples

In this section we give some examples to illustrate the use of timers.

9.1.1 Expanding circles

In this example we create a program that uses a timer to draw a number of growing concentric circles in a window periodically. It also opens a menu containing only the quit command. Figure 9.1 shows the program in action. We discuss these object I/O components in reverse order (menu, window, timer).

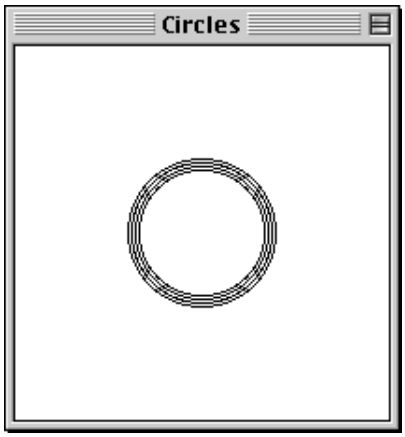

Figure 9.1: The circles program in action.

The menu definition is very simple, as it contains only one command to terminate the example program. Recall that to terminate an interactive process the `StdProcess` function `closeProcess` must be used. The `StdIOBasic` function `noLS` suitably turns `closeProcess` into the desired type. Here is the menu definition.

```
mdef = Menu "Circles"
```

```
        (MenuItem "Quit" [MenuFunction (noLS closeProcess)
                         ,MenuShortKey 'q'
                         ]
        ) []
```

The initial size of the window in which the circles are to be drawn is `windowEdge` by `windowEdge` (200 in the example). To ease drawing, we take care that the view domain of the window has the origin `zero` exactly in the center of the view domain. This can be done conveniently by defining the `WindowViewDomain` attribute to be:

```
viewDomain
 = { corner1 = {x= ~windowEdge/2,y= ~windowEdge/2}
   , corner2 = {x=  windowEdge/2,y=  windowEdge/2}
   }
```

The window does not contain controls (expressed by the `Controls` type instance `NilLS`) and is identified by the value `wid`. The window definition is as follows:

```
wdef = Window "Circles" NilLS
         [WindowId         wid
         ,WindowSize       (rectangleSize viewDomain)
         ,WindowViewDomain viewDomain
         ]
```

The timer draws a number of concentric circles that have an increasing radius. For this purpose it uses a local state of the following type and initial value:

```
:: TimerState
 = { nrCircles    :: Int
   , equiDistance :: Int
   , minRadius    :: Int
   }
initTimerState = { nrCircles=4, equiDistance=2, minRadius=0 }
```

The `nrCircles` field contains the number of circles that are drawn. The `equiDistance` field is the difference of radius between two neighbouring circles. The `minRadius` field keeps track of the radius of the smallest visible circle.

The timer interval is set to a twentieth of a second (`ticksPerSecond/20`). The timer contains no timer elements. Its definition is as follows:

```
tdef = Timer (ticksPerSecond/20) NilLS [TimerFunction timer]
```

The timer function `timer` will be evaluated by the object I/O system every twentieth of a second (if possible). The timer function actually ignores the number of elapsed intervals and simply draws the next sequence of circles. There are two cases to distinguish:

If the smallest circle still fits entirely inside the window (tested by `minRadius < windowEdge/2`), then `timer` *undraws* the smallest circle and *draws* the new circle which should have a radius equal to `minRadius + nrCircles * equiDistance`. Finally, the `minRadius` field is changed to reflect the fact that the smallest visible circle now has radius `minRadius + equiDistance`.

If the smallest circle does not fit entirely inside the window, then `timer` completely *unfills* the window. By setting the new local `TimerState` back to `initTimerState` the circles are drawn again from the center.

```
timer _ (lst=:{nrCircles,equiDistance,minRadius},pst)
 | minRadius<windowEdge/2
   # lst      = {lst & minRadius=minRadius+equiDistance}
     newRadius= minRadius+nrCircles*equiDistance
   # pst      = appPIO
                (appWindowPicture wid
                [ undraw {oval_rx=minRadius,oval_ry=minRadius}
                , draw   {oval_rx=newRadius,oval_ry=newRadius}
                ]
                ) pst
   = (lst,pst)
 | otherwise
   # pst = appPIO (appWindowPicture wid [unfill viewDomain]) pst
   = (initTimerState,pst)
```

The last details that remain to be defined are the actual opening of the menu, window, and timer, the opening of the interactive process, and the creation of `wid`. For completeness we show the complete program code.

```
module circles

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates a window that displays growing concentric circles.
//  For this purpose it uses a timer.
//  ********************************************************************************

import StdEnv, StdIO

::  NoState = NoState

::  TimerState
    =   {   nrCircles       :: Int
        ,   equiDistance    :: Int
        ,   minRadius       :: Int
        }

Start :: *World -> *World
Start world
    = circles (openId world)

circles :: (Id,*World) -> *World
circles (wid,world)
    = startIO NoState
              NoState
              [snd o seqList [openWindow NoState        wdef
                             ,openMenu   NoState        mdef
                             ,openTimer  initTimerState tdef
                             ]
              ] []
              world
where
    windowEdge      = 200
    viewDomain      = { corner1={x= ~windowEdge/2,y= ~windowEdge/2}
                      , corner2={x=  windowEdge/2,y=  windowEdge/2}
                      }
    wdef            = Window "Circles" NilLS
                        [   WindowId         wid
                        ,   WindowSize       (rectangleSize viewDomain)
                        ,   WindowViewDomain viewDomain
                        ]
    mdef            = Menu "Circles"
```

```
                    (   MenuItem "Quit" [MenuFunction (noLS closeProcess)
                                        ,MenuShortKey 'q'
                                        ]
                    )   []
tdef            = Timer (ticksPerSecond/20) NilLS
                    [   TimerFunction   timer
                    ]
initTimerState  = { nrCircles   = 4
                  , equiDistance= 2
                  , minRadius   = 0
                  }
timer _ (lst=:{nrCircles,equiDistance,minRadius},pst)
    | minRadius<windowEdge/2
        # lst         = {lst & minRadius=minRadius+equiDistance}
          newRadius   = minRadius+nrCircles*equiDistance
        # pst         = appPIO
                            (drawInWindow wid
                            [ undraw {oval_rx=minRadius,oval_ry=minRadius}
                            , draw   {oval_rx=newRadius,oval_ry=newRadius}
                            ]
                            ) pst
        = (lst,pst)
    | otherwise
        # pst = appPIO (drawInWindow wid [unfill viewDomain]) pst
        = (initTimerState,pst)
```

9.1.2 Internal clock

In this example we create a program that uses three timers to track the elapsed time since program startup. The timers track the elapsed seconds, minutes, and hours respectively. A dialogue is used to provide visual feedback. Figure 9.2 shows the application in action. We first have a look at the dialogue, and then the three timers.

Figure 9.2: The timing program in action.

The dialogue `ddef` simply uses `TextControl`s to display the hours, minutes, and seconds. These controls are identified by the `Id` values `hoursId`, `minutesId`, and `secondsId` respectively. The dialogue itself is identified by the value `dialogId`. `CompoundControl`s are used to get the layout in the desired way. Observe the use of list comprehensions and the `ListLS Controls` class instance. In this example closing the dialogue also terminates the application. Here is the dialogue definition:

```
ddef
= Dialog "Stopwatch"
  (   CompoundControl
```

```
  (   ListLS [TextControl text [ControlPos (Left,zero)]
             \\ text<-["Hours:","Minutes:","Seconds:"]
             ]
  )   []
  :+: CompoundControl
  (   ListLS [TextControl "00" [ControlPos (Left,zero)
                               ,ControlId  id
                               ]
             \\ id<-[hoursId,minutesId,secondsId]
             ]
  )   []
  )
  [   WindowClose (noLS closeProcess)
  ,   WindowId    dialogId
  ]
```

Because the operation of each of the three timers is very similar, we use one function `tdef` to define them and parameterise it with their respective `TimerIntervals`. The timers do not contain any timer elements. Each timer has a local integer state (initially zero) that keeps the current number of evaluated time units modulo their maximum number. This number is derived from the `TimerInterval` using the straightforward function `maxunit`. The timerfunction `tick` now relies on the `NrOfIntervals` parameter. Given the current number of evaluated time units in its local integer state, each timer function adds the `NrOfIntervals` value to it. Depending on the `TimerInterval` value, the proper modulo value is taken (using `maxunit`). This new value is the new local state and is also drawn in the dialogue, using a local convenience function `setText` and `textid` to determine the `Id` of the corresponding text control.

```
tdef timerInterval
  = Timer timerInterval NilLS [TimerFunction tick]
where
  tick nrElapsed (time,pst)
    # time = (time+nrElapsed) mod (maxunit timerInterval)
    = (time,setText (textid timerInterval) (toString time) pst)

  setText id text pst
    = appPIO (setWindow dialogId [setControlTexts [(id,text)]]) pst

textid interval
  | interval==second = secondsId
  | interval==minute = minutesId
  | interval==hour   = hoursId
maxunit interval
  | interval==second = 60
  | interval==minute = 60
  | interval==hour   = 24
```

The last details that remain to be defined are the actual opening of the three timers, the dialogue, the opening of the interactive process, and the creation of the proper `Ids`. For completeness we show the complete program code.

```
module stopwatch
```

```
//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates a window that tracks the elapsed time since startup.
//  For this purpose it uses three timers to track the seconds, minutes, and hours
//  separately.
//  ********************************************************************************

import StdEnv,StdIO

::  NoState
    =   NoState
::  DialogInfo
    =   {   secondsId   :: Id
        ,   minutesId   :: Id
        ,   hoursId     :: Id
        ,   dialogId    :: Id
        }

second  :== ticksPerSecond
minute  :== 60*second
hour    :== 60*minute

openDialogInfo :: *env -> (DialogInfo,*env) | Ids env
openDialogInfo env
    # ([secondsid,minutesid,hoursid,dialogid:_],env) = openIds 4 env
    = ({ secondsId=secondsid
       , minutesId=minutesid
       , hoursId  =hoursid
       , dialogId =dialogid
       }
      ,env
      )

Start :: *World -> *World
Start world
    # (dialogInfo,world)    = openDialogInfo world
    = startIO NoState NoState [initialise dialogInfo] [] world

initialise :: DialogInfo (PSt .l .p) -> (PSt .l .p)
initialise {secondsId,minutesId,hoursId,dialogId} pst
    # (errors,pst)  = seqList [  openTimer 0 (tdef timerinfo)
                              \\ timerinfo<-[second,minute,hour]
                              ]  pst
    | any ((<>) NoError) errors
        = closeProcess pst
    # (error,pst)   = openDialog NoState ddef pst
    | error<>NoError= closeProcess pst
    | otherwise     = pst
where
    tdef timerInterval
        = Timer timerInterval NilLS [TimerFunction tick]
    where
        tick nrElapsed (time,pst)
            # time  = (time+nrElapsed) mod (maxunit timerInterval)
            = (time,setText (textid timerInterval) (toString time) pst)

        setText id text pst
            = appPIO (setWindow dialogId [setControlTexts [(id,text)]]) pst

        textid interval
            | interval==second   = secondsId
            | interval==minute   = minutesId
            | interval==hour     = hoursId
        maxunit interval
            | interval==second   = 60
            | interval==minute   = 60
```

```
        | interval==hour     = 24

ddef= Dialog "Stopwatch"
        (   CompoundControl
        (   ListLS  [  TextControl text [ControlPos (Left,zero)]
                    \\ text<-["Hours:","Minutes:","Seconds:"]
                    ]
        )   []
        :+: CompoundControl
        (   ListLS  [  TextControl "00" [ControlPos (Left,zero)
                                        ,ControlId  id]
                    \\ id<-[hoursId,minutesId,secondsId]
                    ]
        )   []
        )
        [   WindowClose (noLS closeProcess)
        ,   WindowId    dialogId
        ]
```

Chapter 10

Receivers

All the interactive object I/O components discussed so far have in common that the events to which they respond are *abstract*. In this context abstract means that it is not specified in detail what *concrete* events cause a specific callback function to be evaluated. For instance, a callback function associated with a `MenuItem` (Section 8.1.3) is evaluated when it has been selected by the user. How this selection takes place is not specified.

In this section we discuss an interactive object I/O component that responds to program defined events, or rather *messages*. This component is the *receiver*. It plays an important role in the construction of interactive components. There are *no restrictions* on the kind of messages that can be *sent* or *received*, provided that they are type correct. The latter is obtained by using special identification values for receivers, the *receiver ids* (Chapter 4).

There are also receivers that are able to receive data from other programs via a network. These receivers are not topic of this chapter but of Chapter 14.

We will first have a look at the definition of receivers in Section 10.1. Receivers can be opened as top level object I/O components, but also as elements of windows, dialogues, menus, and timers. This is discussed in Section 10.2. Knowing how to define and open receivers, we show which functions are available to send messages in Section 10.3. Section 10.4 contains a number of examples to demonstrate the use of receivers.

10.1 Receiver definitions

There are two kinds of receivers. *Uni-directional* receivers respond only to messages. *Bi-directional* receivers respond to messages and also reply with a message. The types needed to define receivers can be found in the module `StdReceiverDef`, Appendix A.28.

A uni-directional receiver that responds to messages of type `msg` and which has a local state of type `lst` and process state of type `pst`) is defined by an expression of type (`Receiver msg lst pst`):

```
:: Receiver  msg   lst pst
 = Receiver  (RId  msg)      (ReceiverFunction  msg      *(lst,pst))
                             [ReceiverAttribute          *(lst,pst)]
:: ReceiverFunction msg st
```

```
:== msg -> st -> st
```

A bi-directional receiver that responds to messages of type `msg` and returns a response message of type `resp` and which has a local state of type `lst` and process state `pst`) is defined by an expression of type (`Receiver2 msg resp lst pst`):

```
:: Receiver2 msg resp lst pst
 = Receiver2 (R2Id msg resp) (Receiver2Function msg resp *(lst,pst))
                             [ReceiverAttribute          *(lst,pst)]
:: Receiver2Function msg resp st
:== msg -> st -> (resp,st)
```

The set of receiver attributes is currently limited to the `ReceiverSelectState` which default value is `Able`.

```
::  ReceiverAttribute ps
 =  ReceiverSelectState SelectState
```

10.2 Receiver creation

Receivers can be opened as top level interface elements. This is done in the usual, overloaded way (module `StdReceiver`, Appendix A.27) :

```
class Receivers rdef where
  openReceiver   :: .lst !(rdef .lst (PSt .l .p)) !(PSt .l .p)
                                  -> (!ErrorReport,!PSt .l .p)
  ...

instance Receivers (Receiver  msg)
instance Receivers (Receiver2 msg resp)
```

Receivers can also be opened as elements of windows, dialogues, menus, and timers. So, in a window or dialogue one can not only add the set of controls as discussed in Chapter 7 but also receivers. This is accomplished by declaring receivers to be instances of the `Controls` type constructor class in module `StdControlReceiver` (Appendix A.7). Analogously, receivers can be *menu elements* (module `StdMenuReceiver`, Appendix A.19), and *timer elements* (module `StdTimerReceiver`, Appendix A.38).

In all cases, when opening a receiver, its `R(2)Id` value must not be used by another receiver (Chapter 4).

10.3 Message passing

In contrast with the previously discussed object I/O components, receivers *must* have an identification value. The *message passing* functions require this identification value to ensure that a message of the correct type is sent. The message passing functions can be found in module `StdReceiver`, Appendix A.27.

All message passing functions return a report about the message passing action. This report is an algebraic type `SendReport`, which has the following alternatives:

```
::  SendReport
 =  SendOk
 |  SendUnknownReceiver
 |  SendUnableReceiver
 |  SendDeadlock
```

For all functions, the alternative value `SendOk` is returned in case message passing was successful. The alternative `SendUnknownReceiver` is returned in case the indicated receiver is not open at the moment of sending the message. The other `SendReport` alternatives are discussed below.

We start with message passing to uni-directional receivers in Section 10.3.1. Bi-directional message passing is discussed in Section 10.3.2.

10.3.1 Uni-directional message passing

There are two functions a programmer can use to send a message to a uni-directional receiver: `asyncSend` and `syncSend` which have the same function types:

```
asyncSend :: !(RId msg) msg !(PSt .l .p) -> (!SendReport,!PSt .l .p)
syncSend  :: !(RId msg) msg !(PSt .l .p) -> (!SendReport,!PSt .l .p)
```

Using `asyncSend`, a message is placed at the end of the asynchronous message queue of the indicated receiver and will, at some point, be handled by that receiver. It is unspecified *when* that event occurs. Asynchronous messages that are sent in sequence will be evaluated in that sequence. Asynchronous message passing only fails in case the indicated receiver is not open, as discussed above. One should observe that successfully queueing an asynchronous message does not guarantee that the message will be handled. The receiver might be disabled or closed before all of its messages have been handled.

If one wants to enforce a receiver to handle a message, one should use `syncSend`. This function does not place the message in the message queue of the indicated receiver, but evaluates its receiver function given the message. This implies that `syncSend` has to *switch context* because the indicated receiver may be part of another object I/O component. Another consequence is that synchronous messages may overtake asynchronous messages. However, synchronous messages that are sent in sequence will be evaluated in that sequence.

Sending a synchronous message fails in case the indicated receiver does not exist. If the indicated receiver does exist, but its `ReceiverSelectState` attribute is `Unable`, then the message is also not handled. In this case, the `SendReport` alternative `SendUnableReceiver` is returned. Finally, even if the indicated receiver exists and is `Able` communication can still fail. This is possible when the indicated receiver itself is involved in a synchronous message passing transaction and waiting for termination. If this transaction involves the original receiver, then a deadlock situation is detected. In that case, `syncSend` also fails and returns with `SendDeadlock`.

Examples of uni-directional message passing are given in Section 10.4. The first example, in Section 10.4.1, demonstrates the use of asynchronous message passing, while the second example, in Section 10.4.2, demonstrates synchronous message passing.

10.3.2 Bi-directional message passing

Bi-directional message passing is *synchronous*. A message is sent using the function `syncSend2`:

```
syncSend2 :: !(R2Id msg resp) msg !(PSt .l .p)
   -> (!(!SendReport,!Maybe resp), !PSt .l .p)
```

Analogous to `syncSend`, `syncSend2` locates the indicated bi-directional receiver and applies the argument message immediately to the corresponding receiver function. If this receiver could be found and happens to be `Able`, evaluation of the receiver function will yield a response message `resp`. In this case the `SendReport` result of `syncSend2` is `SendOk`, and the response value is returned as `(Just resp)`. In all exceptional cases, there is no response value and `Nothing` is returned. To evaluate the receiver function, `syncSend2` has to switch context as well.

Bi-directional receivers can be used to retrieve local encapsulated data. Example 10.4.3 demonstrates this.

10.4 Examples

In this section we give some examples to illustrate the use of receivers.

10.4.1 Talk windows

In this example we create a program that uses receivers to send keyboard input from one window to another window, and vice versa. This results in a talk like application (although it is not very useful as a talk application because it runs on one computer). Figure 10.1 shows the application in action.

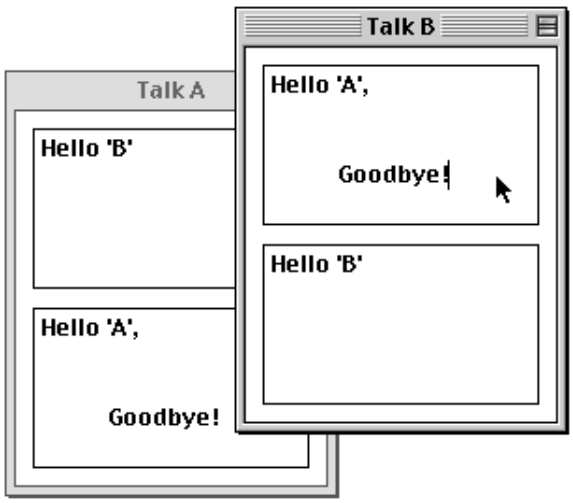

Figure 10.1: The talk program in action.

The initialisation action of the talk program, defined by the function `initialise`, first creates a menu that has only one quit menu item. Then two `RId` values are generated that will be used to identify the two receivers. The function `openTalkWindow`

is then applied twice to create the two windows. Its parameters are the name of the window and the two receiver ids. The first receiver id parameter identifies the private receiver, while the second identifies the other receiver.

```
initialise :: (PSt .l .p) -> PSt .l .p
initialise pst
  # menu        = Menu "Talk"
                  (MenuItem "Quit" [MenuShortKey 'q'
                                   ,MenuFunction (noLS closeProcess)
                                   ]
                  ) []
  # (error,pst) = openMenu undef menu pst
  | error<>NoError
    = abort "talk could not open menu."
  | otherwise
    # (a,pst)   = accPIO openRId pst
    # (b,pst)   = accPIO openRId pst
    # pst       = openTalkWindow "A" a b pst
    # pst       = openTalkWindow "B" b a pst
    = pst
```

`openTalkWindow` creates a window in which the user can type text and see the messages of the other talk window. The window consists of three components: in the first `EditControl`, identified by `inId`, the user can type text. The `ControlKeyboard` attribute takes care that the program can respond to keyboard input. The second `EditControl`, identified by `outId`, is used to present the messages coming from the other talk window. To prevent the user from typing text in this control its initial `ControlSelectState` attribute is `Unable`. The `Receiver` component is the one to which the messages are being sent.

```
openTalkWindow :: String (RId String) (RId String) (PSt .l .p)
                                                    -> PSt .l .p
openTalkWindow name me you pst
  # (wId,  pst)= accPIO openId pst
  # (inId, pst)= accPIO openId pst
  # (outId,pst)= accPIO openId pst
    input      = EditControl "" (hmm 50.0) 5
                     [ ControlId       inId
                     , ControlKeyboard inputfilter Able
                                       (noLS1 (input wId inId you))
                     ]
    output     = EditControl "" (hmm 50.0) 5
                     [ ControlId       outId
                     , ControlPos      (BelowPrev,zero)
                     , ControlSelectState Unable
                     ]
  # (size,pst) = accPIO (controlSize (input:+:output) True
                            Nothing Nothing Nothing) pst
  # receiver   = Receiver me (noLS1 (receive wId outId)) []
  # wdef       = Window ("Talk "+++name) (input:+:output:+:receiver)
                     [ WindowId    wId
                     , WindowSize  size
                     ]
  # (error,pst)= openWindow undef wdef pst
```

```
  | error<>NoError
    = abort "talk could not open window."
  | otherwise
    = pst
```

The input `EditControl` has a keyboard filter, `inputfilter`, that accepts only `KeyDown` keyboard input.

```
inputfilter :: KeyboardState -> Bool
inputfilter keystate = getKeyboardStateKeyState keystate<>KeyUp
```

For accepted keyboard input value of type `KeyboardState`, the callback function `input` is evaluated. It first gets the current state of the window, using the `StdControl` library function `getWindow` (Appendix A.4). From this value the current content of the input `EditControl` can be retrieved, using the function `getControlTexts`. This new content, which is a `String`, is being sent asynchronously to the other receiver. Note that `input` assumes that both `getWindow` and `asyncSend` never fail.

```
input :: Id Id (RId String) KeyboardState (PSt .l .p) -> PSt .l .p
input wId inId you _ pst
  # (Just wst,pst)
        = accPIO (getWindow wId) pst
    text = fromJust (snd (hd (getControlTexts [inId] wst)))
  = snd (asyncSend you text pst)
```

For every string message received from the other talk window, the receiver function `receive` is evaluated. It simply replaces the current content of the output `EditControl` with the new text. This is done using the function `setControlTexts`. The function `setEditControlCursor` makes sure that the end of the text is visible.

```
receive :: Id Id String (PSt .l .p) -> PSt .l .p
receive wId outId text pst
  = appPIO (setWindow wId [setControlTexts [(outId,text)]
                          ,setEditControlCursor outId (size text)
                          ]) pst
```

For completeness the whole program is shown here.

```
module talk

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates two windows that communicate with each other using message
//  passing. Text that has been typed in one window is being sent to the other, and
//  vice versa.
//  ********************************************************************************

import  StdEnv, StdIO

::  NoState
    =   NoState

Start :: *World -> *World
```

```
Start world
    = startIO NoState NoState [initialise] [] world
where
    initialise :: (PSt .l .p) -> PSt .l .p
    initialise pst
        # menu          = Menu "Talk"
                            (   MenuItem "Quit" [   MenuShortKey 'q'
                                                ,   MenuFunction (noLS closeProcess)
                                                ]
                            )   []
        # (error,pst)   = openMenu undef menu pst
        | error<>NoError
            = abort "talk could not open menu."
        | otherwise
            # (a,pst)   = accPIO openRId pst
            # (b,pst)   = accPIO openRId pst
            # pst       = openTalkWindow "A" a b pst
            # pst       = openTalkWindow "B" b a pst
            = pst

openTalkWindow :: String (RId String) (RId String) (PSt .l .p) -> PSt .l .p
openTalkWindow name me you pst
    # (wId,  pst)   = accPIO openId pst
    # (inId, pst)   = accPIO openId pst
    # (outId,pst)   = accPIO openId pst
      input         = EditControl "" (hmm 50.0) 5
                        [   ControlId       inId
                        ,   ControlKeyboard inputfilter Able
                                            (noLS1 (input wId inId you))
                        ]
      output        = EditControl "" (hmm 50.0) 5
                        [   ControlId       outId
                        ,   ControlPos      (BelowPrev,zero)
                        ,   ControlSelectState Unable
                        ]
    # (size,pst)    = accPIO (controlSize (input:+:output)
                                Nothing Nothing Nothing) pst
    # receiver      = Receiver me (noLS1 (receive wId outId)) []
    # wdef          = Window ("Talk "+++name) (input:+:output:+:receiver)
                        [   WindowId    wId
                        ,   WindowSize  size
                        ]
    # (error,pst)   = openWindow undef wdef pst
    | error<>NoError= abort "talk could not open window."
    | otherwise     = pst
where
    inputfilter :: KeyboardState -> Bool
    inputfilter keystate
        = getKeyboardStateKeyState keystate<>KeyUp

    input :: Id Id (RId String) KeyboardState (PSt .l .p) -> PSt .l .p
    input wId inId you _ pst
        # (Just wst,pst)    = accPIO (getWindow wId) pst
          text              = fromJust (snd (hd (getControlTexts [inId] wst)))
        = snd (asyncSend you text pst)

    receive :: Id Id String (PSt .l .p) -> PSt .l .p
    receive wId outId text pst
        = appPIO (setWindow wId [ setControlTexts [(outId,text)]
                                , setEditControlCursor outId (size text)
                                ])  pst
```

10.4.2 Resetting the counter

In this example we extend the example counter in Section 7.2.4 (page 82) with a means to reset the counter to zero. We proceed in a bottom-up style: the counter control is a compound control extended with a receiver that, when it receives a message, will reset the counter to zero. This control encapsulates its local counter state. Then a button control is defined that, when selected, sends a message to the receiver component of the counter control. The whole is being placed in a dialogue. Figure 10.2 gives a snapshot of the program.

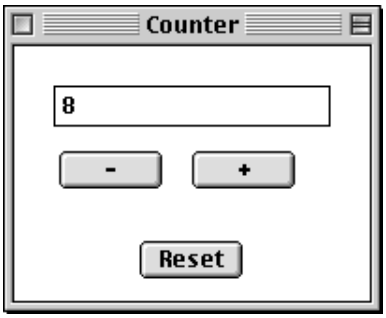

Figure 10.2: The counter control with reset button.

The main component is of course the counter control, defined by `counter`. It encapsulates an integer local state with initial value `initcount`. Its definition is almost identical to the one shown on page 82 except that a `Receiver` has been added.

```
counter
  = {newLS = initcount
    ,newDef= CompoundControl
             (   EditControl (toString initcount) (hmm 50.0) 1
                             [ControlSelectState Unable
                             ,ControlPos         (Center,zero)
                             ,ControlId          displayid
                             ]
             :+: ButtonControl "-" [ControlFunction    (count (-1))
                             ,ControlPos         (Center,zero)
                             ]
             :+: ButtonControl "+" [ControlFunction    (count   1 )]
             :+: Receiver resetid reset []
             )   []
    }
```

The only purpose of the receiver is to reset the current local counter value to `initcount` and show this by changing the content of the `EditControl`. Note that the receiver function `reset` is not interested at all in the message.

```
reset :: m (Int,PSt .l .p) -> (Int,PSt .l .p)
reset _ (_,ps)
  = (initcount,setText windowid displayid initcount ps)
```

The reset button is a straightforward `ButtonControl`. It is centered below the counter control. Its `ControlFunction` is just to send a message synchronously to

the receiver component of the counter control. Because this component does not care about the message, the button function can be as bold to send the `StdMisc` library function `undef`. This function, when evaluated, aborts the application. This demonstrates that message passing is truely lazy in the message argument.

```
resetbutton
  = ButtonControl "Reset"
       [ControlFunction (noLS (snd o syncSend resetid undef))
       ,ControlPos      (Center,zero)
       ]
```

The final details of the program are to generate the proper identification values and to create the initial process and dialogue. For completeness, the program code is given here.

```
module counterreset

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program defines a Controls component that implements a manually settable
//  counter. A receiver is used to add a reset option.
//  ********************************************************************************

import StdEnv, StdIO

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    =   startNDI NoState NoState initialise [] world
where
    initialise ps
        #   (windowid, ps)  = accPIO openId  ps
        #   (displayid,ps)  = accPIO openId  ps
        #   (resetid,  ps)  = accPIO openRId ps
        #   (error,ps)      = openDialog NoState
                                  (dialog windowid displayid resetid) ps
        |   error<>NoError
            =   abort "counter could not open Dialog."
        |   otherwise
            =   ps

dialog windowid displayid resetid
    =   Dialog "Counter"
            (   counter
            :+: resetbutton
            )
            [   WindowId    windowid
            ,   WindowClose (noLS closeProcess)
            ]
where
    counter
        =   {   newLS   = initcount
            ,   newDef  = CompoundControl
                          (   EditControl (toString initcount) (hmm 50.0) 1
                                          [ControlSelectState   Unable
                                          ,ControlPos           (Center,zero)
                                          ,ControlId            displayid
                                          ]
                          :+: ButtonControl "-" [ControlFunction      (count (-1))
                                          ,ControlPos           (Center,zero)
                                          ]
```

```
                        :+: ButtonControl "+" [ControlFunction      (count   1 )]
                        :+: Receiver resetid reset []
                        )   []
        }
where
    initcount   = 0

    count :: Int (Int,PSt .l .p) -> (Int,PSt .l .p)
    count dx (count,ps)
        =   (count+dx,setText windowid displayid (count+dx) ps)

    reset :: m (Int,PSt .l .p) -> (Int,PSt .l .p)
    reset _ (_,ps)
        =   (initcount,setText windowid displayid initcount ps)

    setText :: Id Id x (PSt .l .p) -> PSt .l .p | toString x
    setText wid cid x ps
        =   appPIO (setWindow wid (setControlTexts [(cid,toString x)])) ps

resetbutton
    =   ButtonControl "Reset"
            [ControlFunction (noLS (snd o syncSend resetid undef))
            ,ControlPos      (Center,zero)
            ]
```

10.4.3 Reading the counter

In this example we extend the counter example once more and add a dialogue that reads the local counter value. To be able to do this a bi-directional receiver is added to the counter. Figure 10.3 gives a snapshot of the program.

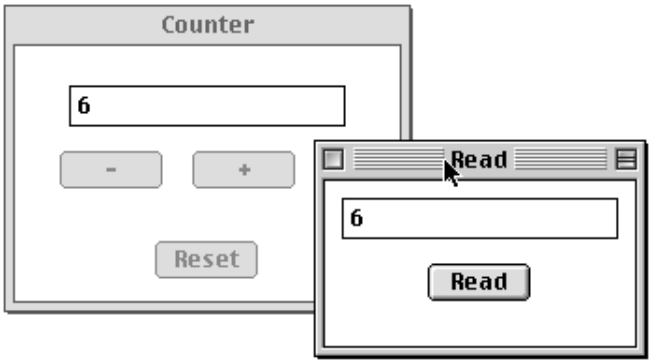

Figure 10.3: Reading the counter control with reset button.

The first change is to extend the counter component with a bi-directional receiver component defined by the expression `(Receiver2 readid read [])`, where `readid` is a `R2Id` identification value. The receiver function `read` is very straightforward: it returns the current local counter value without changing anything (also in this case `read` is not interested in the input message type):

```
read :: m (Int,PSt .l .p) -> (Int,(Int,PSt .l .p))
read _ (count,pst)
    =  (count,(count,pst))
```

The initialisation action opens another dialogue defined by `display` that obtains and shows the counter value on request. This request can be done by a user by

pressing the ButtonControl labeled "Read". Again, an Unable EditControl is used to display the read value.

```
display windowid displayid readid
  = Dialog "Read"
    (   EditControl   "" (hmm 50.0) 1
                            [ControlSelectState Unable
                            ,ControlId          displayid
                            ]
    :+: ButtonControl "Read" [ControlFunction    (noLS read)
                            ,ControlPos         (Center,zero)
                            ]
    )
    [   WindowId    windowid
    ,   WindowClose (noLS closeProcess)
    ]
```

When the "Read" button has been selected, its callback function read is evaluated. It sends a message to the bi-directional receiver component of the counter control using syncSend2. If this action fails it aborts the program. Although this situation will never occur, this check has been added for program hygiene. In a successfull communication, value will be (Just count), and it is this value that is then displayed in the EditControl of the dialogue.

```
read pst
  # ((error,value),pst) = syncSend2 readid undef pst
  | error<>SendOk
    = abort "could not read counter value"
  | otherwise
    = setText windowid displayid (fromJust value) pst
```

The initialisation action creates the necessary identification values. Because of the Id assignment rules (Chapter 4) the Id displayid can be used in both dialogues. Here is the complete program code.

```
module counterread

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program defines a Controls component that implements a manually settable
//  counter.
//  A bi-directional receiver is added to give external access to the counter value.
//  ********************************************************************************

import StdEnv, StdIO

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    = startIO NoState NoState [initialise] [] world

initialise pst
    # (windowid, pst)= accPIO openId   pst
    # (displayid,pst)= accPIO openId   pst
    # (resetid,  pst)= accPIO openRId  pst
```

```
    # (readid,   pst)= accPIO openR2Id pst
    # (error,    pst)= openDialog NoState
                          (dialog windowid displayid resetid readid) pst
    | error<>NoError
        = abort "counter could not open counter dialog."
    # (windowid, pst)= accPIO openId pst
    # (error,    pst)= openDialog NoState (display windowid displayid readid) pst
    | error<>NoError = abort "counter could not open display dialog."
    | otherwise      = pst
where
    dialog windowid displayid resetid readid
        = Dialog "Counter"
            (   counter
            :+: resetbutton
            )
            [   WindowId    windowid
            ,   WindowClose (noLS closeProcess)
            ]
    where
        counter
            = { newLS   = initcount
              , newDef  = CompoundControl
                          (   EditControl (toString initcount) (hmm 50.0) 1
                                          [ControlSelectState Unable
                                          ,ControlPos      (Center,zero)
                                          ,ControlId       displayid
                                          ]
                          :+: ButtonControl "-" [ControlFunction (count (-1))
                                          ,ControlPos      (Center,zero)
                                          ]
                          :+: ButtonControl "+" [ControlFunction (count   1 )]
                          :+: Receiver  resetid reset []
                          :+: Receiver2 readid  read  []
                          )   []
              }
        where
            initcount = 0

            count :: Int (Int,PSt .l .p) -> (Int,PSt .l .p)
            count dx (count,pst)
                # count = count+dx
                = (count,setText windowid displayid count pst)

            reset :: m (Int,PSt .l .p) -> (Int,PSt .l .p)
            reset _ (_,pst)
                = (initcount,setText windowid displayid initcount pst)

            read :: m (Int,PSt .l .p) -> (Int,(Int,PSt .l .p))
            read _ (count,pst)
                = (count,(count,pst))

        resetbutton
            = ButtonControl "Reset" [ControlFunction (noLS reset)
                                    ,ControlPos      (Center,zero)
                                    ]
        where
            reset pst
                = snd (syncSend resetid undef pst)

    display windowid displayid readid
        = Dialog "Read"
            (   EditControl   "" (hmm 50.0) 1
                                    [ControlSelectState Unable
                                    ,ControlId          displayid
                                    ]
            :+: ButtonControl "Read" [ControlFunction   (noLS read)
                                    ,ControlPos         (Center,zero)
```

```
                                ]
        )
        [   WindowId    windowid
        ,   WindowClose (noLS closeProcess)
        ]
where
    read pst
        # ((error,value),pst)   = syncSend2 readid undef pst
        | error<>SendOk
            = abort "could not read counter value"
        | otherwise
            = setText windowid displayid (fromJust value) pst

setText :: Id Id x (PSt .l .p) -> PSt .l .p | toString x
setText wid cid x pst
    = appPIO (setWindow wid [setControlTexts [(cid,toString x)]]) pst
```

Chapter 11

Interactive processes

The examples discussed so far used the `StdProcess` library function `startIO`. This function creates one interactive process that engages in some graphical user interface actions with a user and then terminates, using the `StdProcess` function `closeProcess`. In this chapter we show how an interactive process can spawn new interactive processes that will run *interleaved*. It is also possible to create a number of processes at once.

We start the discussion by looking at the ways to define interactive processes in Section 11.1. This is followed by the functions to open interactive processes. Finally we give some examples to illustrate their application.

11.1 Defining interactive processes

The types to define interactive processes can be found in the module `StdProcessDef` (Appendix A.25). In the object I/O library, interactive processes are distinguished by their *document interface*. There are three kinds of document interfaces, and there are also three corresponding type constructors to define an individual interactive process:

No Document Interface (NDI): An interactive process with this document interface does not present a document to a user. It has no menus. It is typically used for 'background' interactive processes. The interactive process is allowed to open dialogues, timers, and receivers. The type constructor that defines a NDI process is `NDIProcess`:

```
:: NDIProcess p
 = E. .l: NDIProcess l (ProcessInit      (PSt l p))
                       [ProcessAttribute (PSt l p)]
```

Single Document Interface (SDI): An interactive process with this document interface presents exactly one document to the user. All menu commands are associated with this document. The interactive process is allowed to open dialogues, timers, and receivers. The type constructor that defines a SDI process is `SDIProcess`:

```
:: SDIProcess wdef p
 = E. .l lst: SDIProcess l lst (wdef         lst (PSt l p))
```

```
                                (ProcessInit      (PSt l p))
                                [ProcessAttribute (PSt l p)]
```

Multiple Document Interface (MDI): An interactive process with this document interface can present an arbitrary number of documents to the user (even zero). In its menu system all available commands are presented. The interactive process is allowed to open dialogues, timers, and receivers. The type constructor that defines a MDI process is `MDIProcess`:

```
:: MDIProcess p
 = E. .l: MDIProcess l (ProcessInit      (PSt l p))
                       [ProcessAttribute (PSt l p)]
```

The function `startIO` that we have used so far in our examples actually creates a MDI process.

The type constructor definitions of each of the three kinds of interactive processes are identical in case of NDI and MDI processes. A SDI process has two additional parameters: `lst` and `(wdef lst (PSt l p))`. The type constructor variable `wdef` must be a `Windows` type constructor class instance. The variable `lst` is the local state of that window. A SDI process can not close this window.

The interactive process type constructors are polymorphic in the public process state `p`. Each constructor introduces an initial local process state `l` and encapsulates it using existential quantification. The initialisation action is given by a process state transition function:

```
::  ProcessInit pst :== [IdFun pst]
::  IdFun       st  :== st -> st
```

Process definitions can be attributed with the following alternatives (defined in module `StdIOCommon`):

ProcessWindowPos, ProcessWindowSize, ProcessWindowResize: A *process window* is the root window in which all top-level user interface elements are created. On a Macintosh the process window is simply the screen. On Windows(95/NT) it is also the screen in case of NDI processes. For SDI and MDI processes it is a window that can be resized by the user[1]. If this happens, and a `ProcessWindowResize` attribute has been specified, then a function of type `ProcessWindowResizeFunction` is evaluated:

```
::  ProcessWindowResizeFunction pst
:== Size -> Size -> pst -> pst
```

The two `Size` parameters are the size of the process window before and after resizing.

ProcessHelp, ProcessAbout: These two attributes contain callback functions that are evaluated whenever the user issues a request for help or wants information about your application.

[1] This has not yet been implemented in this version of the object I/O library.

ProcessActivate, ProcessDeactivate: These two attributes correspond closely to the `WindowActivate` and `WindowDeactivate` attributes. Recall that keyboard and mouse input is always directed to the active window. The parent interactive process that contains this window is the *active process*. If the input focus is moved to a window that is not owned by the interactive process then the `ProcessDeactivate` attribute function is evaluated to inform the program that the interactive process has become inactive. If an inactive process obtains the input focus, its `ProcessActivate` attribute function is evaluated to inform the program that it has become active again.

ProcessClose: This attribute corresponds closely to the `WindowClose` attribute. If for some reason the interactive process is requested to be closed, its `ProcessClose` attribute function is evaluated. It can take the opportunity to save data to disk and to ask the user if the process can be closed safely. It is however the responsibility of the program to terminate the process. It is good programming practice to always include this attribute because the user of your program expects to be able to close your application in this way.

ProcessShareGUI: When this attribute is set, the interactive process *shares* the graphical user interface elements with its parent. This means that the windows and dialogues share the window stack of the parent, and menus share the menu area. By default, this attribute is not set. This means that normally an interactive process behaves like a separate program.

ProcessNoWindowMenu: This attribute is valid only for MDI processes. In the default case every MDI process has a special "Windows" menu (discussed in Section 8.3). If this attribute is set, then this menu is not added to the menu system of the interactive process (which is the case for MDI processes created by `startIO`).

11.2 Interactive process creation

As mentioned briefly in the beginning of this chapter, interactive processes can be created one by one, or by groups. In this section we will first discuss individual creation of interactive processes in Section 11.2.1, and then have a look at the creation of multiple interactive processes in Section 11.2.2. Finally, we discuss the relations between interactive processes in Section 11.2.3.

11.2.1 Creating single processes

The basic function to create individual processes from a `World` environment is `startIO`, which we have encountered many times. For completeness we repeat its type definition once more:

```
startIO ::!.l !.p !(ProcessInit (PSt .l .p))
            ![ProcessAttribute (PSt .l .p)] !*World -> *World
```

Interactive processes can also create individual processes. They can not use `startIO` because the `World` environment is not retrievable from their context, which is `PSt`. For this purpose the overloaded function `shareProcesses` is available:

```
class shareProcesses pdef :: !(pdef .p) !(PSt .l .p) -> PSt .l .p
```

```
instance shareProcesses  NDIProcess
instance shareProcesses (SDIProcess wdef) | Windows wdef
instance shareProcesses  MDIProcess
```

When applied to a NDI, SDI, or MDI process definition, `shareProcesses` adds an initial version of that interactive process to the process state administration. At some point in time that interactive process will be initialised and joins the game. So process creation is asynchronous.

From the type definition of `shareProcesses` one can see that the public process state components of the process states have the same type. One should also observe that no new public process state value is introduced. Interactive processes that are created by `shareProcesses` *share* the public state component of the process that spawned them. Interactive processes that share the same public process state value constitute a *process group*. Because interactive processes run interleaved, the public process state component can be changed in turn by each of the members of that group. This explains why it is called 'public'.

If one wants to spawn an interactive process that has a public process state of its own using only the `shareProcesses` class instances is not sufficient because they assume a public process state is already given. For this purpose the `ProcessGroup` type constructor (module `StdProcessDef`) can be used:

```
:: ProcessGroup pdef
 = E. .p: ProcessGroup p (pdef p)
```

`ProcessGroup` introduces a public process state value and encapsulates it using existential quantification. The type constructor variable `pdef` can be any of the instances of the `shareProcesses` class defined above. In this way we obtain two overloaded functions that can spawn an interactive process with an independent public process state:

```
class Processes pdef where
    startProcesses :: !pdef !*World      -> *World
    openProcesses  :: !pdef !(PSt .l .p) -> PSt .l .p

instance Processes (ProcessGroup pdef) | shareProcesses pdef
```

The two constructor class functions spawn an interactive process that does not share its public process state with the process that created it (in case of `openProcesses`). We also see the other process creation function that operates on a `World` environment, `startProcesses`. It is this function that is used by `startIO` to do its job. Here is its implementation:

```
startIO local public init atts world
= startProcesses (ProcessGroup public
                  (MDIProcess local init [ProcessNoWindowMenu:atts])
                 ) world
```

11.2.2 Creating multiple processes

In the previous section we have shown all functions that are needed to create individual interactive processes. Because these functions are overloaded they are also suited to create a number of processes within one function application. This is

achieved by declaring suitable *glueing* instances for the type constructor classes `Processes` (`:^:` and `[]`) and `shareProcesses` (`:~:`, and `ListCS`). Except for lists, these type definitions can be found in module `StdIOBasic` (Appendix A.12).

`:^:` The type constructor `:^:` glues two arbitrary type constructors. Its type constructor definition and `Processes` class instance declaration are as follows:

```
:: :^: t1 t2 = (:^:) infixr 9 t1 t2

instance Processes (:^: pdef1 pdef2) | Processes pdef1
                                     & Processes pdef2
```

Given two `Processes` instances `p1` and `p2`, then the expression `p1:^:p2` is also a `Processes` instance. Because `:^:` is right associative, an expression such as `p1 :^: p2 :^: p3` should be read as `p1 :^: (p2 :^: p3)`.

`[]` In principle the `:^:` glue is sufficient to create all required process structures. In case of working with a number of process instances of the *same type*, it is much more convenient to use lists and list comprehensions. Because the process type constructors are not parameterised we can use the normal lists for this purpose. So we get the following `Processes` class instance declaration:

```
instance Processes [pdef] | Processes pdef
```

Given a list of `Processes` instances `ps = [p1 ...pn]`, then the expression `ps` itself is also a `Processes` instance.

`:~:` The type constructor `:~:` glues two type constructors that work on the same context. Its type constructor definition and `shareProcesses` class instance declaration are as follows:

```
:: :~: t1 t2 context = (:~:) infixr 9 (t1 context) (t2 context)

instance shareProcesses (:~: pdef1 pdef2) | shareProcesses pdef1
                                          & shareProcesses pdef2
```

Given two `shareProcesses` instances `p1` and `p2` working on the same context state of type `context`, then the expression `p1 :~: p2` is also a `shareProcesses` instance working on the same context state. Because `:~:` is right associative, an expression such as `p1 :~: p2 :~: p3` should be read as `p1 :~: (p2 :~: p3)`.

`ListCS` In principle the `:~:` glue is sufficient to create all required process group structures. In case of working with a number of process group instances of the *same type*, it is much more convenient to use lists and list comprehensions. This glue is provided by the type constructor `ListCS`. Its type constructor definition and `shareProcesses` class instance declaration is as follows:

```
:: ListCS t context = ListCS [t context]

instance shareProcesses (ListCS pdef) | shareProcesses pdef
```

Given a list of `shareProcesses` instances `ps = [p1 ...pn]`, working on the same context state of type `context`, then the expression `ListCS ps` is also a `shareProcesses` instance working on the same context state.

11.2.3 Process relations

An interactive program in general consists of a number of process groups, each of which consists of a number of interactive processes. As explained above, interactive processes and process groups can be created dynamically.

Except for interactive processes created with the `ProcessShareGUI` attribute, there is no special parent-child relationship between an interactive process and the interactive processes that it creates. For instance, termination of one process (using `closeProcess`) has no consequence for the other processes. Termination of a process also terminates all child processes that have been created with the `ProcessShareGUI` attribute.

Process groups exist by virtue of their element processes. As soon as all interactive processes of one process group have been closed, the process group vanishes and the shared public state becomes garbage.

The process creation functions that work on the `World` environment, `startProcesses` and its derived function `startIO`, terminate as soon as all of the process groups that have been created during their life cycle have been closed.

11.3 Examples

In this section we give some examples of the use of interactive processes.

11.3.1 Talk revisited

In this example we have a new look at the talk example of Section 10.4.1. In that version, the program created one interactive process, using `startIO`, which opened the two talk windows. In the new version for each talk window we create a SDI process. Receivers are still used to send the user typed messages to each of the talk windows. The menu is now created for both processes. Below we discuss the differences.

The initialisation of the new talk program is of course different. The function `talk` defines the SDI process that contains the talk window and menu. The two talk processes are going to be created in one process group. This implies that the `RIds` are also created earlier. We obtain the following `Start` rule:

```
Start :: *World -> *World
Start world
# (a,    world) = openRId world
# (b,    world) = openRId world
# (talkA,world) = talk "A" a b world
# (talkB,world) = talk "B" b a world
= startProcesses (ProcessGroup NoState (ListCS [talkA,talkB])) world
```

The menu definition that is moved to `talk` is almost identical to the old version. The only difference is termination of the application. In the old version termination was no issue because there was only one interactive process. In the new version, closing one talk process does not close the other. Instead, before closing its parent process, the `quit` function sends a new message to the other process to request termination. To do this the message type needs to be changed. The message type is now the following algebraic data type:

```
:: Message
 = NewLine String
 | Quit
```

Instead of sending a `String s` in the old version we send the value `(NewLine s)` in the new version. The request to terminate is done by sending the `Quit` message. The `quit` function is now defined as:

```
quit :: (RId Message) (PSt .l .p) -> PSt .l .p
quit you pst
 = closeProcess (snd (syncSend you Quit pst))
```

Of course the definition of the receiver function `receive` must be changed accordingly:

```
receive :: Id Id Message (PSt .l .p) -> PSt .l .p
receive wId outId (NewLine text) pst=:{io}
 = {pst & io=setWindow wId [ setControlTexts [(outId,text)]
                           , setEditControlCursor outId (size text)
                           ] io}
receive _ _ Quit pst
 = closeProcess pst
```

On receipt of the `Quit` message, the receiver only needs to terminate its parent process, knowing that the requesting process will terminate itself.

These are the major differences. Below is the complete program.

```
module talk

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates two interactive processes that communicate via message
//  passing.
//  In a future distributed version this program can be used as a graphical talk
//  application.
//
//  ********************************************************************************

import  StdEnv, StdIO

::  Message
    =   NewLine String                      // Transmit a line of text
    |   Quit                                // Request termination
::  NoState
    =   NoState                             // The singleton data type

Start :: *World -> *World
Start world
    # (a,    world) = openRId world
    # (b,    world) = openRId world
    # (talkA,world) = talk "A" a b world
    # (talkB,world) = talk "B" b a world
    = startProcesses (ProcessGroup NoState (ListCS [talkA,talkB])) world

talk :: String (RId Message) (RId Message) *env
    -> (SDIProcess (Window (:+: EditControl
                           (:+: EditControl
                               (Receiver Message)
                           ))) .p,*env)
```

```
    | Ids env
talk name me you env
    # (wId,  env)   = openId env
    # (outId,env)   = openId env
    # (inId, env)   = openId env
      input         = EditControl   "" (hmm 50.0) 5
                          [   ControlId       inId
                          ,   ControlKeyboard inputfilter Able
                                              (noLS1 (input wId inId you))
                          ]
      output        = EditControl   "" (hmm 50.0) 5
                          [   ControlId       outId
                          ,   ControlPos      (BelowPrev,zero)
                          ,   ControlSelectState Unable
                          ]
      receiver      = Receiver me (noLS1 (receive wId outId)) []
      talkwindow    = Window ("Talk "+++name) (input:+:output:+:receiver)
                          [   WindowId    wId
                          ,   WindowSize  {w=hmm 50.0,h=120}
                          ]
      menu          = Menu ("&Talk "+++name)
                          (   MenuItem "&Quit"
                              [   MenuShortKey 'q',MenuFunction (noLS (quit you))]
                          )   []
    = ( SDIProcess  NoState
                    undef
                    talkwindow
                    [snd o openMenu undef menu]
                    [ProcessClose (quit you)]
      , env
      )

inputfilter :: KeyboardState -> Bool
inputfilter keystate
    = getKeyboardStateKeyState keystate<>KeyUp

input :: Id Id (RId Message) KeyboardState (PSt .l .p) -> PSt .l .p
input wId inId you _ pst
    # (Just window,pst) = accPIO (getWindow wId) pst
      text              = fromJust (snd (hd (getControlTexts [inId] window)))
    = snd (asyncSend you (NewLine text) pst)

receive :: Id Id Message (PSt .l .p) -> PSt .l .p
receive wId outId (NewLine text) pst=:{io}
    = {pst & io=setWindow wId [ setControlTexts [(outId,text)]
                              , setEditControlCursor outId (size text)
                              ] io}
receive _ _ Quit pst
    = closeProcess pst

quit :: (RId Message) (PSt .l .p) -> PSt .l .p
quit you pst
    = closeProcess (snd (syncSend you Quit pst))
```

11.3.2 Clock revisited

In this example we are going to turn the clock example of Section 9.1.2 into a stopwatch component that can be added in an arbitrary interactive process. The stopwatch commands will be to *reset* timing, *pause* timing, *continue* timing, and *close* the stopwatch component. All stopwatch definitions are placed in the module `stopwatch.icl`. The function `stopwatch` defines the stopwatch component. A main module, `usestopwatch.icl` that opens and controls the stopwatch is also defined. We first look at the stopwatch and then at the main program.

The stopwatch component

The original clock program created an interactive process with three timers and a dialogue. Each of the three timers changes a *local* state that keeps track of the elapsed seconds, minutes, and hours. This situation is schematised in Figure 11.1.

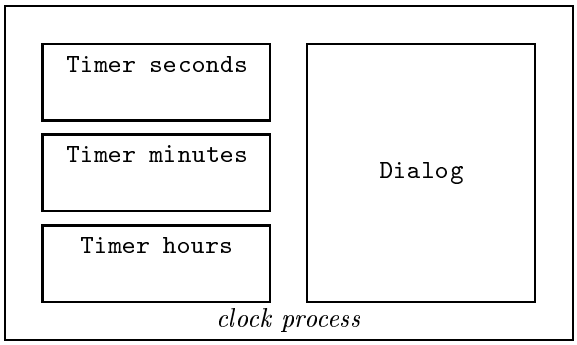

Figure 11.1: The structure of the clock process.

The stopwatch process is controlled by sending messages to a 'gateway' receiver. The timers are extended with a receiver component that handle the commands *reset*, *pause*, and *continue*. These commands will be sent to them by the gateway receiver. This gateway receiver handles the *close* command. Finally, the stopwatch process creates the same dialogue as the original clock program. The stopwatch process is schematised in Figure 11.2.

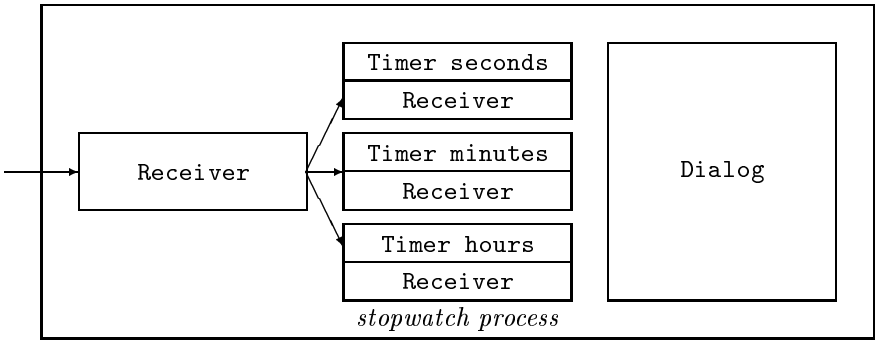

Figure 11.2: The structure of the stopwatch process.

So the new components are the gateway receiver and the receiver components of the timers. We first look at the gateway receiver and then timer receivers. Both receivers accept messages of the following type:

```
:: StopwatchCommands
 = Reset
 | Pause
 | Continue
 | Close
```

The alternatives of the algebraic type `StopwatchCommands` correspond of course with the stopwatch commands *reset*, *pause*, *continue*, and *close*. The gateway receiver function `receive`, on receiving the `Close` message simply terminates the

interactive process by applying `closeProcess` to its process state. Every other message is routed to the timer receiver components which are identified by the list `timerinfos`. Here is the function definition of `receive`:

```
receive :: StopwatchCommands (PSt .l .p) -> PSt .l .p
receive Close pst
 = closeProcess pst
receive msg pst
 = snd (seqList [syncSend timerRId msg\\{timerRId}<-timerinfos] pst)
```

The timer receiver components receive only the `StopwatchMessage` alternatives `Reset`, `Pause`, and `Continue`. In the clock example, the timer was parameterised with its timer interval. In this example, we also need to identify both the timer and its receiver component. So the definition of a stopwatch timer now is:

```
tdef :: TimerInfo
     -> Timer (Receiver StopwatchCommands) Int (PSt .l .p)
tdef {timerId,timerRId,timerInterval}
 = Timer timerInterval (Receiver timerRId receive [])
     [ TimerId       timerId
     , TimerFunction tick
     ]
```

The *reset* command should set the timer back to zero. One might suppose that it is sufficient to change only the value of the local state to zero, but that is not completely true. Resetting the stopwatch can occur at any moment. At that moment the timer should be synchronised with its local state. This can be done by first *disabling* and then *enabling* the timer. For this purpose the `StdTimer` functions `disableTimer` and `enableTimer` should be used. The reason that it works is because `enableTimer`, when applied to a disabled timer, synchronises the timer with the moment of evaluation. It does nothing in case the indicated timer was already enabled. Finally, on receiving the `Reset` message, the timer receiver component must set the corresponding text field of the dialogue to zero. This gives the following definition of the `Reset` alternative.

```
receive Reset (time,pst)
 # pst = appListPIO [disableTimer timerId,enableTimer timerId] pst
 # pst = setText (textid timerInterval) "00" pst
 = (0,pst)
```

The *pause* command should halt the timer until further notice (either *reset* or *continue*). This is easily done by disabling the timer:

```
receive Pause (time,pst)
 = (time,appPIO (disableTimer timerId) pst)
```

The *continue* command should let the timer continue from where it was paused. This is easily done by enabling the timer:

```
receive Continue (time,pst)
 = (time,appPIO (enableTimer timerId) pst)
```

The final details of the stopwatch component are to create the proper identification values for the timers and their receiver components, and to export its definition as an interactive process. This is done by the function `stopwatch`. The stopwatch process is defined as a process group with no interesting local or public process state (using the ubiquitous `NoState` singleton type constructor). Its initialisation action first creates the required `Id`s for the dialogue, and then the parameters required for the timers. Then initialisation proceeds as described above.

```
stopwatch :: (RId StopwatchCommands) -> ProcessGroup NDIProcess
stopwatch rid
 = ProcessGroup NoState (NDIProcess NoState [initialise‘] [])
where
 initialise‘ pst
  # (dialogIds, pst) = accPIO openDialogIds  pst
  # (timerInfos,pst) = accPIO openTimerInfos pst
  = initialise rid dialogIds timerInfos pst
```

For completeness, the definition module and implementation module of the stopwatch component are given below.

```
implementation module stopwatch

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program defines a stopwatch process component.
//  It uses three timers to track the seconds, minutes, and hours separately.
//  Message passing is used to reset, pause, and continue timing.
//  The current time is displayed using a dialogue.
//  ********************************************************************************

import StdEnv,StdIO

::  NoState
    =   NoState
::  DialogIds
    =   {   secondsId       :: Id
        ,   minutesId       :: Id
        ,   hoursId         :: Id
        ,   dialogId        :: Id
        }
::  TimerInfo
    =   {   timerId         :: Id
        ,   timerRId        :: RId StopwatchCommands
        ,   timerInterval   :: TimerInterval
        }
::  StopwatchCommands
    =   Reset
    |   Pause
    |   Continue
    |   Close

second  :== ticksPerSecond
minute  :== 60*second
hour    :== 60*minute

openDialogIds :: *env -> (DialogIds,*env) | Ids env
openDialogIds env
    # ([secondsid,minutesid,hoursid,dialogid:_],env) = openIds 4 env
    = ( { secondsId=secondsid
        , minutesId=minutesid
        , hoursId  =hoursid
```

```
        , dialogId =dialogid
        }
      , env
      )

openTimerInfos :: *env -> ([TimerInfo],*env) | Ids env
openTimerInfos env
    # (tids,env)  = openIds  3 env
    # (rids,env)  = openRIds 3 env
    # intervals   = [second,minute,hour]
    = ( [ {timerId=tid,timerRId=rid,timerInterval=i}
        \\ tid<-tids & rid<-rids & i<-intervals
        ]
      , env
      )

stopwatch :: (RId StopwatchCommands) -> ProcessGroup NDIProcess
stopwatch rid
    = ProcessGroup NoState (NDIProcess NoState [initialise`] [])
where
    initialise` pst
        # (dialogIds, pst) = accPIO openDialogIds  pst
        # (timerInfos,pst) = accPIO openTimerInfos pst
        = initialise rid dialogIds timerInfos pst

initialise :: (RId StopwatchCommands) DialogIds [TimerInfo]
              (PSt .l .p) -> (PSt .l .p)
initialise rid {secondsId,minutesId,hoursId,dialogId} timerinfos pst
    # (errors,pst)  = seqList [  openTimer 0 (tdef timerinfo)
                              \\ timerinfo<-timerinfos
                              ]  pst
    | any ((<>) NoError) errors
        = closeProcess pst
    # (error,pst)   = openDialog NoState ddef pst
    | error<>NoError
        = closeProcess pst
    # (error,pst)   = openReceiver NoState rdef pst
    | error<>NoError
        = closeProcess pst
    | otherwise
        = pst
where
    tdef {timerId,timerRId,timerInterval}
        = Timer timerInterval (Receiver timerRId receive [])
            [   TimerId         timerId
            ,   TimerFunction   tick
            ]
    where
        tick nrElapsed (time,pst)
            # time  = (time+nrElapsed) mod (maxunit timerInterval)
            = (time,setText (textid timerInterval) (toString time) pst)

        setText id text pst
            = appPIO (setWindow dialogId [setControlTexts [(id,text)]]) pst

        receive Reset (time,pst)
            # pst   = appListPIO [disableTimer timerId,enableTimer timerId] pst
            # pst   = setText (textid timerInterval) "00" pst
            = (0,pst)
        receive Pause (time,pst)
            = (time,appPIO (disableTimer timerId) pst)
        receive Continue (time,pst)
            = (time,appPIO (enableTimer timerId) pst)

        textid interval
            | timerInterval==second = secondsId
            | timerInterval==minute = minutesId
```

```
            | timerInterval==hour   = hoursId
        maxunit interval
            | timerInterval==second = 60
            | timerInterval==minute = 60
            | timerInterval==hour   = 24

    ddef    = Dialog "Stopwatch"
                (   CompoundControl
                (   ListLS  [  TextControl text [ControlPos (Left,zero)]
                            \\ text<-["Hours:","Minutes:","Seconds:"]
                            ]
                )   []
                :+: CompoundControl
                (   ListLS  [  TextControl "00" [ControlPos (Left,zero)
                                                ,ControlId  id
                                                ]
                            \\ id<-[hoursId,minutesId,secondsId]
                            ]
                )   []
                )
                [   WindowClose (noLS closeProcess)
                ,   WindowId    dialogId
                ]
    rdef    = Receiver rid (noLS1 receive) []
    where
        receive Close pst
            = closeProcess pst
        receive msg pst
            = snd (seqList [syncSend timerRId msg \\ {timerRId}<-timerinfos] pst)

implementation module stopwatch

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program defines a stopwatch process component.
//  It uses three timers to track the seconds, minutes, and hours separately.
//  Message passing is used to reset, pause, and continue timing.
//  The current time is displayed using a dialogue.
//  ********************************************************************************

import StdEnv,StdIO

::  NoState
    =   NoState
::  DialogIds
    =   {   secondsId       :: Id
        ,   minutesId       :: Id
        ,   hoursId         :: Id
        ,   dialogId        :: Id
        }
::  TimerInfo
    =   {   timerId         :: Id
        ,   timerRId        :: RId StopwatchCommands
        ,   timerInterval   :: TimerInterval
        }
::  StopwatchCommands
    =   Reset
    |   Pause
    |   Continue
    |   Close

second  :== ticksPerSecond
minute  :== 60*second
hour    :== 60*minute

openDialogIds :: *env -> (DialogIds,*env) | Ids env
openDialogIds env
```

```
    # ([secondsid,minutesid,hoursid,dialogid:_],env) = openIds 4 env
    = ( { secondsId=secondsid
        , minutesId=minutesid
        , hoursId  =hoursid
        , dialogId =dialogid
        }
      , env
      )

openTimerInfos :: *env -> ([TimerInfo],*env) | Ids env
openTimerInfos env
    # (tids,env)  = openIds  3 env
    # (rids,env)  = openRIds 3 env
    # intervals   = [second,minute,hour]
    = ( [ {timerId=tid,timerRId=rid,timerInterval=i}
        \\ tid<-tids & rid<-rids & i<-intervals
        ]
      , env
      )

stopwatch :: (RId StopwatchCommands) -> ProcessGroup NDIProcess
stopwatch rid
    = ProcessGroup NoState (NDIProcess NoState [initialise`] [])
where
    initialise` pst
        # (dialogIds, pst) = accPIO openDialogIds  pst
        # (timerInfos,pst) = accPIO openTimerInfos pst
        = initialise rid dialogIds timerInfos pst

initialise :: (RId StopwatchCommands) DialogIds [TimerInfo]
              (PSt .l .p) -> (PSt .l .p)
initialise rid {secondsId,minutesId,hoursId,dialogId} timerinfos pst
    # (errors,pst)  = seqList [  openTimer 0 (tdef timerinfo)
                              \\ timerinfo<-timerinfos
                              ]  pst
    | any ((<>) NoError) errors
        = closeProcess pst
    # (error,pst)   = openDialog NoState ddef pst
    | error<>NoError
        = closeProcess pst
    # (error,pst)   = openReceiver NoState rdef pst
    | error<>NoError
        = closeProcess pst
    | otherwise
        = pst
where
    tdef {timerId,timerRId,timerInterval}
        = Timer timerInterval (Receiver timerRId receive [])
            [   TimerId         timerId
            ,   TimerFunction   tick
            ]
    where
        tick nrElapsed (time,pst)
            # time  = (time+nrElapsed) mod (maxunit timerInterval)
            = (time,setText (textid timerInterval) (toString time) pst)

        setText id text pst
            = appPIO (setWindow dialogId [setControlTexts [(id,text)]]) pst

        receive Reset (time,pst)
            # pst   = appListPIO [disableTimer timerId,enableTimer timerId] pst
            # pst   = setText (textid timerInterval) "00" pst
            = (0,pst)
        receive Pause (time,pst)
            = (time,appPIO (disableTimer timerId) pst)
        receive Continue (time,pst)
            = (time,appPIO (enableTimer timerId) pst)
```

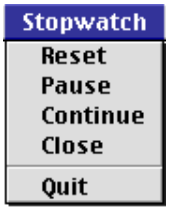

Figure 11.3: The stopwatch menu.

```
        textid interval
            | timerInterval==second = secondsId
            | timerInterval==minute = minutesId
            | timerInterval==hour   = hoursId
        maxunit interval
            | timerInterval==second = 60
            | timerInterval==minute = 60
            | timerInterval==hour   = 24

    ddef    = Dialog "Stopwatch"
                (   CompoundControl
                (   ListLS  [  TextControl text [ControlPos (Left,zero)]
                            \\ text<-["Hours:","Minutes:","Seconds:"]
                            ]
                )   []
                :+: CompoundControl
                (   ListLS  [  TextControl "00" [ControlPos (Left,zero)
                                                ,ControlId  id
                                                ]
                            \\ id<-[hoursId,minutesId,secondsId]
                            ]
                )   []
                )
                [   WindowClose (noLS closeProcess)
                ,   WindowId    dialogId
                ]
    rdef    = Receiver rid (noLS1 receive) []
    where
        receive Close pst
            = closeProcess pst
        receive msg pst
            = snd (seqList [syncSend timerRId msg \\ {timerRId}<-timerinfos] pst)
```

Using the stopwatch

Given the stopwatch component module `stopwatch`, we can create a program that opens, uses, and closes the stopwatch. The program will be as simple as possible. As its initialisation action it will create a `RId` needed to open the stopwatch component. It also opens a menu, `mdef`, that triggers the stopwatch commands *reset*, *pause*, *continue*, and *close*. Finally, it contains the *quit* command to terminate the whole program. Its definition is as follows (Figure 11.3 shows the menu in its initial state):

```
mdef
 = Menu "&Stopwatch"
   (   MenuItem "&Reset"    [MenuFunction (noLS (send Reset))]
   :+: MenuItem "&Pause"    [MenuFunction (noLS (send Pause))]
   :+: MenuItem "C&ontinue" [MenuFunction (noLS (send Continue))]
   :+: MenuItem "&Close"    [MenuFunction (noLS (send Close))]
```

```
    :+: MenuSeparator        []
    :+: MenuItem "&Quit"     [MenuFunction (noLS (closeProcess o
                                                  (send Close)))]
    )   []
```

Each of the stopwatch command menu functions is defined by the `send` function which is parameterised with the corresponding `StopwatchCommands` message alternative. Its purpose is to send its argument message to the gateway receiver of the stopwatch process, identified by `stopwatchid`. If this fails it also emits a system beep. Here is its definition:

```
send msg pst
 # (error,pst)   = syncSend stopwatchid msg pst
 | error<>SendOk = appPIO beep pst
 | otherwise     = pst
```

Here is the complete code of the main program.

```
module usestopwatch

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates a simple program that uses the stopwatch process.
//  The program only has a menu to open the stopwatch and control it.
//  ********************************************************************************

import StdEnv, StdIO
import stopwatch

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    # (stopwatchid,world) = openRId world
    = startIO NoState NoState [initialise stopwatchid] [] world

initialise :: (RId StopwatchCommands) (PSt .l .p) -> PSt .l .p
initialise stopwatchid pst
    # pst           = openProcesses (stopwatch stopwatchid) pst
    # (error,pst)   = openMenu NoState mdef pst
    | error<>NoError= closeProcess pst
    | otherwise     = pst
where
    mdef    = Menu "&Stopwatch"
                (   MenuItem "&Reset"    [MenuFunction (noLS (send Reset))]
                :+: MenuItem "&Pause"    [MenuFunction (noLS (send Pause))]
                :+: MenuItem "C&ontinue" [MenuFunction (noLS (send Continue))]
                :+: MenuItem "&Close"    [MenuFunction (noLS (send Close))]
                :+: MenuSeparator        []
                :+: MenuItem "&Quit"     [MenuFunction (noLS (closeProcess o
                                                              (send Close)))]
                ) []
    send msg pst
        # (error,pst)   = syncSend stopwatchid msg pst
        | error<>SendOk = appPIO beep pst
        | otherwise     = pst
```

Chapter 12

Clipboard handling

The *clipboard* is a universal simple communication metaphor between applications and within applications. An application can write some data to the clipboard (typically text or pictures) which can be read at a later point of time by the same or another application. This mechanism is supported by the functions in the definition module `StdClipboard`. At the moment only text can be handled. Because we intend to incorporate pictures as well in the near future the current version is set up in such a way that it can be extended upward compatibly. For this purpose an abstract data type, `ClipboardItem`, is defined. Two overloaded functions from the type constructor class `Clipboard` take care that data types can be converted to and from `ClipboardItems`:

```
:: ClipboardItem

class Clipboard item where
   toClipboard   :: !item          -> ClipboardItem
   fromClipboard :: !ClipboardItem -> Maybe item

instance Clipboard {#Char}
```

A convention of using the clipboard is that applications should provide several 'popular' data formats for the same content in descending order of accuracy. For instance, a text processor can first store its private format for the laid out text including font and style information, followed by an ASCII version of the same text, followed by a picture of the laid out text. For this reason the function `setClipboard` that writes the clipboard is not applied to one single clipboard data item but a list of them. The previous content will be destroyed completely. Because programs are supposed to provide only one data item of each format from this list duplicate types of clipboard items are removed. Note that providing `setClipboard` with an empty list will clear the clipboard.

```
setClipboard :: ![ClipboardItem] !(PSt .l .p) -> PSt .l .p
```

The function that reads the clipboard, `getClipboard`, simply gets a list of the current content of the clipboard in descending order of accuracy. With the conversion function `fromClipboard` an application can determine easily if some data item is present that it can handle.

```
getClipboard :: !(PSt .l .p) -> (![ClipboardItem],!PSt .l .p)
```

Finally, because reading in a complete clipboard can be time-consuming or space-consuming the function `clipboardHasChanged` is provided that checks whether the clipboard has been updated since the last time the program checked it.

```
clipboardHasChanged :: !(PSt .l .p) -> (!Bool,!PSt .l .p)
```

12.1 Example: a clipboard editor

To illustrate the use of the clipboard, we construct a small program that shows the current content of the clipboard and that can write some text to the clipboard (see the picture below). It will create only one dialogue with two text fields. In the first text field, the *show* field, the content of the clipboard can be *loaded* by pressing a button. In the second text field, the *set* field, the content of the clipboard can be *stored* by pressing a button.

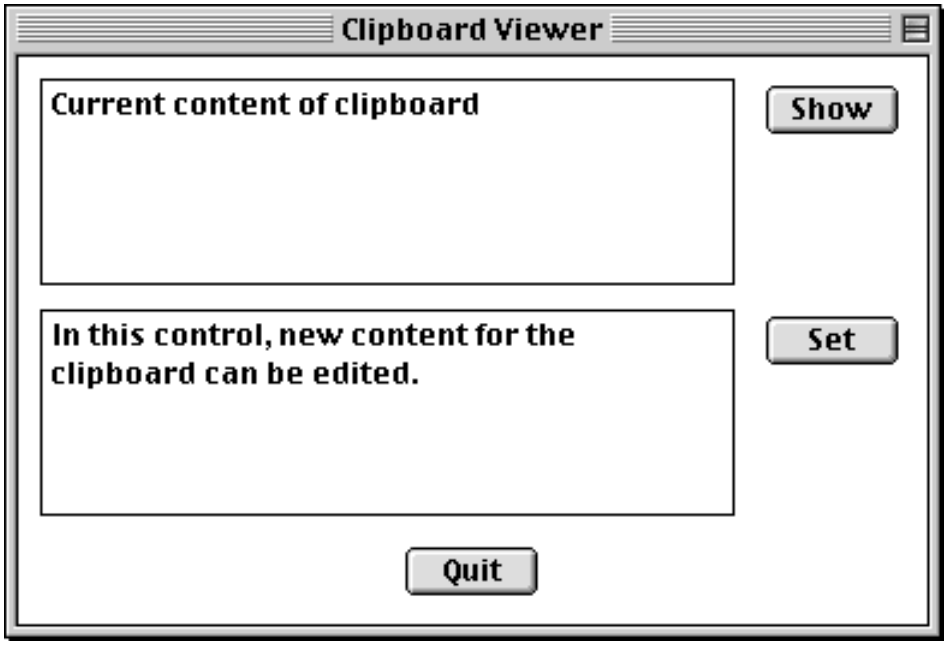

The show text field and its activating button can be defined as follows:

```
showclip
  = EditControl   "" width nrlines [ControlSelectState Unable
                                   ,ControlId  showid
                                   ,ControlPos (Left,zero)
                                   ]
    :+:
    ButtonControl "Show" [ControlFunction (noLS show)]
```

The show text field is an edit text control that will not respond to keyboard input (because its `SelectState` attribute is `Unable`). It is identified by some `Id` of value `showid`. It will have some `width` and a height defined by the number of lines `nrlines`. The "Show" button, when selected, must read the content of the clipboard and figure out if there was a text clipboard item. It then sets the text of the show text field to the loaded clipboard content (an empty string if nothing was found). This action can be defined as follows:

```
show pst
  # (content,pst) = getClipboard pst
    text          = getString content
  = appPIO (setWindow viewid [setControlTexts [(showid,text)]]) pst

getString [clip:clips]
  | isNothing item= getString clips
  | otherwise     = fromJust item
where
  item            = fromClipboard clip
getString []
  = ""
```

The set text field and its activating button are defined as follows:

```
setclip
  = EditControl   "" width nrlines [ControlId  setid
                                   ,ControlPos (Left,zero)
                                   ]
    :+:
    ButtonControl "Set"  [ControlFunction (noLS set)]
```

The set text field is an edit text control which accepts keyboard input. It is identified by some `Id` of value `setid`. It has the same dimensions as the show text control. The "Set" button, when selected, must get the content of the set text control and write this to the clipboard. This action is defined as follows:

```
set pst
  # (wst,pst)= accPIO (getWindow viewid) pst
    text     = fromJust (snd (hd (getControlTexts [setid]
                          (fromJust wst))))
  = setClipboard [toClipboard text] pst
```

The definition of the clipboard viewing dialogue simply summaries these elements and adds a "Quit" button to terminate the program:

```
clipview = Dialog "Clipboard Viewer"
                  (showclip :+: setclip :+: quit)
                  [WindowId viewid]
quit     = ButtonControl "Quit"
                  [ControlFunction (noLS closeProcess)
                  ,ControlPos      (Center,zero)
                  ]
```

The last details that remain to be defined are the opening of the interactive program which is very analogous to the "Hello world!" of Section 2.4. For completeness we show the complete program code.

```
module clipboardview

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program creates a dialog to display and change the current content of the
```

```
//  clipboard.
//  ********************************************************************************

import StdEnv, StdIO

::  NoState
    =   NoState

Start :: *World -> *World
Start world
    # (ids,world)   = openIds 3 world
    = startProcesses (ProcessGroup NoState
                        (NDIProcess NoState [initialise ids] [])
                     ) world

initialise ids pst
    # (error,pst)   = openDialog NoState clipview pst
    | error<>NoError= closeProcess pst
    | otherwise     = pst
where
    (viewid,showid,setid)   = (ids!!0,ids!!1,ids!!2)

    clipview= Dialog "Clipboard Viewer"
                ( showclip :+: setclip :+: quit )
                [ WindowId viewid ]
    showclip= EditControl "" width nrlines [ ControlSelectState Unable
                                           , ControlId          showid
                                           , ControlPos         (Left,zero)
                                           ]
              :+:
              ButtonControl "Show"         [ ControlFunction    (noLS show)
                                           ]
    setclip = EditControl "" width nrlines [ ControlId          setid
                                           , ControlPos         (Left,zero)
                                           ]
              :+:
              ButtonControl "Set"          [ ControlFunction    (noLS set)
                                           ]
    quit    = ButtonControl "Quit"         [ ControlFunction    (noLS closeProcess)
                                           , ControlPos         (Center,zero)
                                           ]
    width   = hmm 100.0
    nrlines = 5

    show pst
        # (content,pst) = getClipboard pst
          text          = getString content
        = appPIO (setWindow viewid [setControlTexts [(showid,text)]]) pst

    set pst
        # (wst,pst)     = accPIO (getWindow viewid) pst
          text          = fromJust (snd (hd (getControlTexts [setid]
                                (fromJust wst))))
        = setClipboard [toClipboard text] pst

    getString [clip:clips]
        | isNothing item= getString clips
        | otherwise     = fromJust item
    where
        item            = fromClipboard clip
    getString []
        = ""
```

Chapter 13

Printing

In this chapter we introduce two modules that enable Clean programs to print: the `StdPrint` module (Appendix A.22) and the `StdPrintText` module (Appendix A.23). At first we will discuss the dialogues with whom the user is confronted when he wants to print something (Section 13.1). The `StdPrint` module can be used to print any arbitrary output. It is illustrated in the Sections 13.2 to 13.5. The `StdPrintText` module can be used when only text should be printed. It's use is explained in Section 13.6.

13.1 The User Interface for Printing

When a user wants to get some printed output, he has to specify several parameters for printing. These parameters can be divided into two groups: the *print setup* parameters determine the printer and the paper format that are used and the page orientation (portrait or landscape). The *print job* parameters include the number of copies and the range of pages to print. According to these two sets there are two dialogues to specify these parameters and two Clean types for objects that store this information. The dialogues are the *print setup dialogue* (Figure 13.1) and the *print job dialogue* (Figure 13.2). The types are the abstract data type `PrintSetup` and the `JobInfo` record (see Section 13.2). The print job dialogue launches the print job if the user does not cancel. In most applications the user has to answer this dialogue just before printing, but using the print setup dialogue is optional.

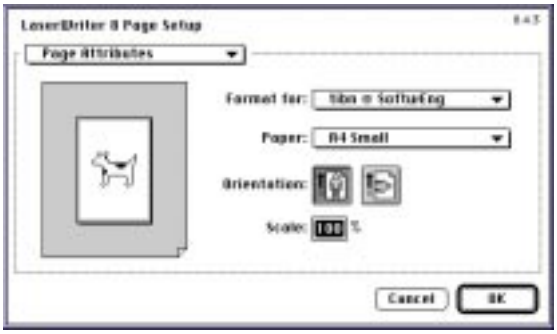

Figure 13.1: a print setup dialogue

There is always a default `PrintSetup` which is returned by the `defaultPrintSetup`

Figure 13.2: a print job dialogue

function. Such a `PrintSetup` can be edited with the print setup dialogue when it is applied to the `printSetupDialog` function. Both these functions are members of the `PrintSetupEnvironments` class:

```
::  PrintSetup

class PrintSetupEnvironments env where
    defaultPrintSetup ::              !*env -> (!PrintSetup,!*env)
    printSetupDialog  :: !PrintSetup !*env -> (!PrintSetup,!*env)
```

The `env` parameter can either be the `World` or the `IOSt`.

On Macintosh computers the print setup dialogue and the print job dialogue are typically reachable from two different menu items. On Windows(95/NT) it is typical, that the print setup dialogue can only be reached by clicking on a button in the print job dialogue. The `StdSystem` module's function `printSetupTypical` can be used to determine whether a seperate menu item for the print setup dialogue is typical on the used platform.

13.2 The print function

The `print` function in the `StdPrint` module has the following signature:

```
print   ::  !Bool !Bool
            .(PrintInfo !*Picture -> ([IdFun *Picture],!*Picture))
            !PrintSetup !*printEnv
        ->  (!PrintSetup,!*printEnv)
        |   PrintEnvironments printEnv

::  PrintInfo
    =   {   printSetup  ::  PrintSetup
        ,   jobInfo     ::  JobInfo
        }
::  JobInfo
    =   {   range   ::  !(!Int,!Int)
        ,   copies  ::  !Int
        }
```

The parameters of

```
print doDialog emulateScreen pages printSetup printEnv
```

have the following meaning:

1. `doDialog`: if True, the print job dialogue will pop up, otherwise printing will happen in the default way.
2. `emulateScreen`: iff True, the screen resolution will be emulated, see Section 13.3.
3. `pages`: this function determines the printed output. Therefore it generates a list of drawing functions. Each drawing function corresponds to one page. The function can access a `PrintInfo` record and a `Picture`. The `Picture` parameter can be used to retrieve information like font metrics or string widths from a printer `Picture`.
4. `printSetup`: the print setup. As mentioned before, the user can change this print setup from the print job dialogue under Windows(95/NT).
5. `printEnv`: a print environment is either the `PSt` or the `Files` sub world. If the print environment is the `PSt`, then the print function can update the contents of windows during printing (to be more exact: between the rendering of two pages). This is important on the Windows(95/NT) platforms, since printing can take some time. On the Macintosh platform printing is a blocking operation, so no window updating is necessary. If the print environment is the `Files` sub world, then no window updating will happen, and, on the Windows(95/NT) platform, no cancel dialogue for printing will pop up.

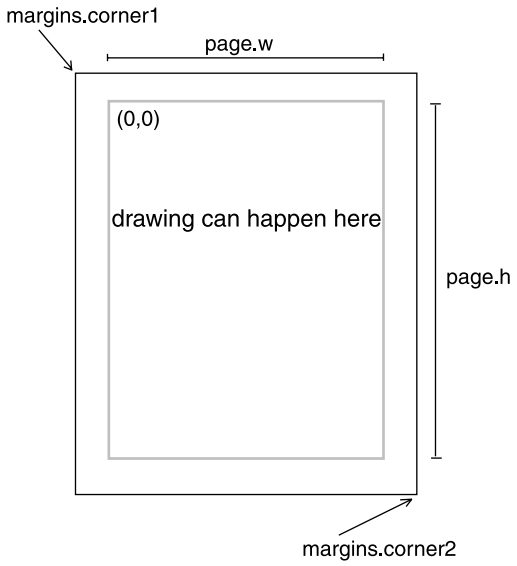

Figure 13.3: Page measurements

The `PrintInfo` record contains information that will be needed for generating the list of drawing functions. The fields of the `JobInfo` part of this record have the following meaning:

1. `range`: this field contains the numbers of the first and last page which the user has chosen via the print dialogue. If the user has chosen "all pages", then the first page will be one, and the last page will be 9999.

2. `copies`: the number of copies to generate. This will not necessarily be equal to the number of copies as specified in the print job dialogue. Some printer drivers take themselves care of producing the appropriate number of copies. In that case the contents of this field will be one.

The `PrintInfo` record also contains a `PrintSetup` field. This print setup will be used for printing The `printSetup` field can be passed to the `getPageDimensions` function:

```
getPageDimensions   ::  !PrintSetup !Bool-> PageDimensions

::  PageDimensions
    =   {   page        :: !Size
        ,   margins     :: !Rectangle
        ,   resolution  :: !(!Int,!Int)
        }
```

The record fields have the following meaning:

1. `page`: this field contains the size of the drawable area of a sheet in pixels, see Figure 13.3. The `print` function sets the origin of such a `Picture` to `(0,0)`. So any drawing in a rectangle from `(0,0)` to `(page.w-1,page.h-1)` will be seen on the printed paper.

2. `margins`: this field contains information about the size of a whole sheet of paper. This is not only the area in which the printer can draw, but also the margins of the paper. Both coordinates of `corner1` are negative and the coordinates of `corner2` are greater than the corresponding values of the `page` field, see Figure 13.3.

3. `resolution`: the horizontal and vertical resolution of the printer in *dpi* (dots per inch).

These measurements of a printer picture depend on whether the screen resolution is emulated or not. Iff the Boolean parameter of the `getPageDimensions` function is True, then the values for a printer picture that emulates the screen resolution are returned.

The `print` function returns the `PrintSetup` that is stored in the `PrintInfo` record.

Here is an easy example program, `printExample1`, that uses the `print` function:

```
module printExample1

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program demonstrates the use of the "print" function.
//  It prints two pages, one with the text "Hello Printer", and one, that informs
//  about the printer's resolution.
//  ********************************************************************************

import StdEnv, StdPrint
```

```
Start world
    # (defaultPS, world) = defaultPrintSetup world
    = snd (accFiles (print True False pages defaultPS) world)
  where
    pages :: PrintInfo *Picture -> ([IdFun *Picture],*Picture);
    pages { printSetup, jobInfo={ range=(first,last), copies } } picture
      # {resolution=(xRes,_)} = getPageDimensions printSetup False
        bothPages
            = [ drawAt { x=100, y=100 } "Hello Printer",
                drawAt { x=100, y=100 }
                       ("horizontal Resolution: "+++toString xRes
                                                 +++" dpi.") ]
        oneCopy = bothPages % (first-1,last-1)
      = ( flatten (repeatn copies oneCopy), picture)
```

On the Macintosh platform, a Clean program needs additional "extra memory" for printing. In the "Application Options" dialogue of the CleanIDE, this value should be set to about 200K, otherwise the program may crash.

If `printExample1` is executed, the print job dialogue will pop up. Let's assume that the user chooses to print only one copy of the first two pages. These values will be passed to the `pages` function, via its first argument of type `PrintInfo`. So the values of `first` and `copies` are one, and the value of `last` is two. The list of drawing functions, which is returned by the `pages` function is then equal to `bothPages`. It contains two drawing functions, so two pages will be printed. The first printed page will contain the text "Hello Printer", and the second page will contain the printer resolution.

If the user choses to print only the first page, then the value of `last` is also one. The `%` operator selects only the first element of `bothPages` and so only one page is printed. If the user choses to print all pages, then the value of `last` is 9999. The `%` operator then selects both pages, which would subsequently be printed. The application of `repeatn` ensures that the right number of copies is generated.

In general, printing is done similar like drawing on the screen, namely by applying drawing functions on a `Picture`. In the rest of this chapter printer `Pictures` and screen `Pictures` are distinguished. There are two differences between these two kinds of `Pictures`:

1. Functions which hilite parts of a `Picture`, or perform a XOR operation, will have no effect on a printer `Picture`.
2. Bitmapped fonts cannot be drawn on both kinds of `Pictures` with the same sizes. This is because the pixel resolutions of both output devices are, in general, different. TrueType fonts work well for both kinds.

13.3 Reusing drawing functions

In many cases an application should be able to draw a certain content on printer `Pictures` as well as on screen `Pictures`. Ideally, such a program should use the same code for both output devices. But a problem arises, because in Clean *two* different units of measurement are used for drawing, namely *Points* and *Pixels* (see also Chapter 5).

The physical extent and position of a graphical object on the paper or on the screen always depends on `Int` values, which are passed to a drawing function. These values are interpreted either as point values or as pixel values. Point values occur only when the size of a font is specified. In all other cases pixel values are used, e.g. in `drawAt {x=100, y=100}...`.

A point is approximately $\frac{1}{72}$ of an inch. A pixel has of course also a physical extent, which can be measured in inches. But unlike a point, the size of a pixel varies from output device to output device: the size of a printer pixel can be about 3 to 4 times smaller than the size of a screen pixel. This implies that the result of applying a drawing function on a printer `Picture` will result in smaller images compared with screen `Pictures`. Only when text is drawn, this shrinking does not happen, because the size of the used font is specified in points, a resolution independent specification of physical extent.

Consider for example the following function:

```
myDrawfunction picture
    # string  = "This text is boxed"
      box     = {box_w=120,box_h=15}
    # picture = drawAt {x=20,y=80} string picture
    # picture = drawAt {x=18,y=68} box    picture
    = picture
```

could cause a screen to contain the following:

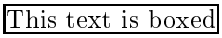

but the printed output could look like this:

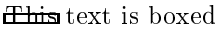

The reason is that the upper code uses constants to express pixel values. There a two ways to deal with this problem: *implicit* and *explicit* scaling.

Implicit scaling Setting the `print` function's second Boolean argument to `True` enables implicit scaling. This causes printer `Pictures` to emulate the screen resolution, so that the shrinking does not occur. This approach has the advantage that it's very easy to write a function that can be used for drawing on both a screen and printer `Picture`. But not everything can be printed very well by using this option. For instance lines on the paper can not be thinner than on the screen. Further problems arise when the printer resolution is not a whole multiple of the screen resolution. This causes rounding errors in the emulation process. If the ratio between printer and screen resolution is for instance 3.5, then the emulation of line drawing with a pen size of one screen pixel will result in a line on the paper that is only 3 printer pixels thick. The following drawing function draws five hundred horizontal lines:

```
seq [drawAt {x=0,y=y} {vx=500,vy=0} \\ y<-[0..500]]
```

On paper space appears between the lines, but not on the screen.

Explicit scaling If the disadvantages of implicit scaling are not acceptable, the application has to perform the scaling itself. One way to obtain this is to avoid constant pixel values. Consider for instance a program that draws a string on a certain place in a window or sheet. The position of the string could be specified by using values that are related to string widths or font metrics. Since string width and font metrics always depend on a `Picture`, they will reflect the resolution of the output device. The upper `myDrawfunction` on page 156 could be rewritten to:

```
myDrawfunction picture
  # (metrics,picture)= getPenFontMetrics picture
    (mwidth,descent) = (metrics.fMaxWidth,metrics.fDescent)
    height           = fontLineHeight metrics
  # (swidth,picture) = getPenFontStringWidth string picture
    string  = "This text is boxed"
    box     = {box_w=swidth+12*mwidth/10,box_h=0-height}
  # picture = drawAt {x=20,            y=80}      string picture
  # picture = drawAt {x=20-2*mwidth/10,y=80+descent} box picture
  = picture
```

The constants in this example are not pixel values, because they are *multiplied* with pixel values like `height`. This example also shows that writing resolution independent drawing functions results in more complex code.

Another way to perform scaling is to use the printer and screen resolution explicitly. Then the code fragment on page 156 could be rewritten to:

```
myDrawfunction (sclNom,sclDenom) picture
  # string  = "This text is boxed"
    box     = {box_w=scl 120,box_h=scl 15}
  # picture = drawAt {x=scl 20,y=scl 80} string picture
  # picture = drawAt {x=scl 18,y=scl 68} box    picture
  = picture
where
  scl x     = sclNom*x/sclDenom
```

For printing, `sclNom` should be the printer resolution, and `sclDenom` should be the resolution of the screen. For drawing on the screen, both values should be one. The `StdPicture` module's function `getResolution` will return the horizontal and vertical resolution of a screen or printer picture:

```
getResolution :: !*Picture -> (!(!Int,!Int),!*Picture)
```

The `accScreenPicture` function (see Appendix A.26) can be handy in this context. The printer resolution can of course also be found in the `PageDimensions` record.

13.4 The printUpdateFunction function

The `printUpdateFunction` function in `StdPrint` is a good example of using implicit scaling. This function offers a very easy way to print the contents of a window. It has the following signature:

```
printUpdateFunction ::
     Bool
     (UpdateState -> *Picture -> *Picture)
     [Rectangle]
     !PrintSetup
     *printEnv
  -> (!PrintSetup, !*printEnv)
  | PrintEnvironments printEnv
```

The first and the last two parameters are identical to the first and last two parameters of the `print` function. Also the result has the same semantics as shown with the `print` function. Let's assume that an application has opened a window. The contents of the window is defined by its `WindowLook` attribute. This attribute contains a `Look` function which updates the contents of a window everytime this is required (see also Section 6.4.1). The `Look` is defined as:

```
::  Look :== SelectState -> UpdateState -> *Picture -> *Picture
```

If such a `Look` function is curried with a `SelectState`, it can be passed to the `printUpdateFunction` function as the second argument.

The list of rectangles parameter specifies the parts of the view domain that should be printed. If such a rectangle is too big for one sheet of paper, then it will be distributed on several pages.

13.5 The printPagePerPage function

The third function in the `StdPrint` module is the `printPagePerPage` function. The `printUpdateFunction` and the `print` functions are both specialisations of this function. This function performs a state transition function on a polymorphic state. The state transition function itself has to be passed to the `printPagePerPage` function. With each application of this state transition function the contents of the actually printed page is drawn. Furtheron this state transition function returns a Boolean value which indicates whether there are further pages to print. So it will be applied repeatedly, until this value becomes `True`.

`printPagePerPage` has the following signature (slightly altered for simplicity):

```
printPagePerPage :: !Bool !Bool
    .unq
    (.unq PrintInfo *Picture  -> ((Bool,Point),(.state,*Picture)))
    ((.state,       *Picture) -> ((Bool,Point),(.state,*Picture)))
    !PrintSetup
    !*printEnv
 -> (Alternative .unq .state,!*printEnv)
 | PrintEnvironments printEnv

:: Alternative unq state = Cancelled unq | StartedPrinting state

instance PrintEnvironments (PSt .l .p)
instance PrintEnvironments Files
```

The arguments of

```
printPagePerPage doDialog emulateScreen unq prepare stateTrans
                 printEnv
```

have the following meaning:

`doDialog`, `emulateScreen`, `printSetup` **and** `printEnv`**:** analogous to those in the `print` function.

unq: if a unique object *unq*, e.g. a text file, has to be accessed and passed back by `printPagePerPage` it should be passed via this parameter. If the user cancels printing via the print job dialogue, `(Cancelled` *unq*`)` will be returned. Otherwise the `prepare` function is applied to this object.

prepare: calculates the initial state. In order to do so it can access the `PrintInfo` record and a printer `Picture`. Drawing in that `Picture` will have no effect on the printed output. The `prepare` also returns a Boolean and a `Point` value. The Boolean value is `False` iff there is another page to print. The `Point` will be the origin of the next page, if there is one.

stateTrans: is the state transition function that generates the pages to be printed. Like the `prepare` function, each application of `stateTrans` returns a Boolean and a `Point` value with the same meaning. If the user did not cancel printing via the cancel dialogue, the `printPagePerPage` function will return the final `state` in the `StartedPrinting` alternative constructor of `Alternative`. Since this `state` is polymorphic, any arbitrary value can be returned, e.g. the `PrintInfo` record.

13.6 Printing text

The module `StdPrintText` (Appendix A.23) offers functions to print text. These functions are overloaded in the argument that specifies the text to print. In this way the source of the text can be an arbitrary data structure, e.g. a textfile or a list of characters.

The `printText1` function has the following signature:

```
printText1 :: !Bool        !WrapMode
              !FontDef     !Int
              !*charStream
              !PrintSetup !*printEnv
           -> ((!*charStream,!PrintSetup),!*printEnv)
           |  CharStreams charStream & PrintEnvironments printEnv

class CharStreams cs where
    getChar    :: !*cs -> (!Bool,!Char,!*cs)
    savePos    :: !*cs -> *cs
    restorePos :: !*cs -> *cs
    eos        :: !*cs -> (!Bool,!*cs)

::  WrapMode   :== Int

NoWrap         :== 0
LeftJustify    :== 1
RightJustify   :== 2
```

This function simply prints the text contained in the `charStream` object.

The first and the two last parameters are identical to the first and last parameters of the `print` function. Also the returned `PrintSetup` and print environment have their usual semantics. The other parameters have the following meaning:

WrapMode: this parameter determines how lines are handled that are too long to fit on the paper. `NoWrap` suppresses wrapping, while `LeftJustify` and

`RightJustify` wrap long lines and adjust the rest of the line to the left and right margin respectively.

`FontDef:` this parameter determines the font of the text.

`Int:` this parameter controls the tab width, measured in the number of space characters.

A `CharStreams` object contains the text to print. Its functionality is analogous to the behaviour of a text file. So a `CharStream` contains a finite sequence of characters and a position pointer which points to the actual character. Applying `getChar` retrieves this actual character and increases the position pointer. The Boolean return value is True iff this operation was successful. Characters can be read sequentially, until the `eos` (end of stream) function returns `True`. The actual position of a `CharStream` can be saved with the `savePos` function, and the `restorePos` function resets the position pointer to the previously saved position.

The following example illustrates how to print a text that is stored in a list of characters. Therefore the `CharStreams` class is instanciated with a proper type `ListCharStream`.

```
module printCharList

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program demonstrates the use of the function "printText1".
//  It instantiates the CharStreams class with the ListCharStream type. Objects
//  of type ListCharstream contain a list of characters, which should be printed.
//  ********************************************************************************

import StdEnv, StdPrintText

:: *ListCharStream = { list :: [Char] , savedPos :: [Char] }

instance CharStreams ListCharStream
  where
    getChar sc=:{list}
        # empty = isEmpty list
        = ( not empty, if empty ' ' (hd list),
            { sc & list=if empty list (tl list) })
    savePos sc
        = { sc & savedPos=sc.list }
    restorePos sc
        = { sc & list=sc.savedPos }
    eos sc=:{list}
        = (isEmpty list,sc)

fontDef = { fName="Courier New", fStyles=[], fSize=9 }

Start world
 # (defaultPS, world) = defaultPrintSetup world
 = accFiles (printText1 True NoWrap fontDef 4
                        { list=['Hello printer again'], savedPos=[] }
                        defaultPS)
            world
```

The `StdPrintText` module instanciates the `CharStreams` class with the `*FileCharStream` type. The following example illustrates how to print a text file, and how to use the functions `fileToCharStream` and `charStreamToFile`.

```
module printFile
```

```
//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program demonstrates to print a text file by using the function
//   "printText1".
//  ********************************************************************************

import StdEnv, StdPrintText

fileName = "printFile08.icl"
fontDef = { fName="Courier New", fStyles=[], fSize=9 }

Start world
    # (ok,file,world) = fopen fileName FReadData world
    | not ok
        = abort ("file "+++fileName+++" not found")
    # (defaultPS,world) = defaultPrintSetup world
      ((charStream,_),world)
        = accFiles (printText1 True NoWrap fontDef 4
                               (fileToCharStream file) defaultPS)
                   world
      (ok,world) = fclose (charStreamToFile charStream) world
    | not ok
        = abort "can't close file"
    = world
```

The `printText2` function allows you to print text with a header on each page. This function expects the same parameters as `printText1`, but takes in addition two strings. The first string will appear on the left corner of each header, and the second string, concatenated with the current page number, will appear on the right side of each header. The previous example can be altered with:

```
module printFileWithHeader

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program demonstrates to print a text file by using the function
//   "printText2".
//  ********************************************************************************

import StdEnv, StdPrintText

fileName = "printFileWithHeader.icl"
fontDef = { fName="Courier New", fStyles=[], fSize=9 }

Start world
    # (ok,file,world) = fopen fileName FReadData world
    | not ok
        = abort ("file "+++fileName+++" not found")
    # (defaultPS,world) = defaultPrintSetup world
      ((charStream,_),world)
            = accFiles (printText2 "Deze file is geprint door Clean"
                                   "pagina "
                                   True NoWrap fontDef 4
                                   (fileToCharStream file)
                                   defaultPS)
                       world
      (ok,world) = fclose (charStreamToFile charStream) world
    | not ok
        = abort "can't close file"
    = world


// Deze file is geprint door Clean = This file is printed with Clean
```

```
// pagina = page
```

Finally, the function `printText3` permits full control over the look of a header and/or trailer (or footer) on each page. It's signature is:

```
printText3 ::!Bool !WrapMode !FontDef !Int
    (PrintInfo *Picture -> (userInfo, (Int,Int), *Picture))
    (userInfo Int PrintInfo *Picture -> *Picture)
    !*charStream
    !PrintSetup !*printEnv
->  (!(!*charStream,!PrintSetup),!*printEnv)
| CharStreams charStream & PrintEnvironments printEnv
```

The fifth and sixth parameter of

```
printText3 doDialog wrapMode fontParams spacesPerTab
           textRangeFunc
           eachPageDrawFunc
           charStream
           printSetup
           printEnv
```

have the following meaning:

`textRangeFunc`**:** this function takes a `PrintInfo` record and a printer `Picture` and returns a triple. The first result can be any data of arbitrary type. This data will be passed by `printText3` to `eachPageDrawFunc`. The second triple result is a pair `(top,bottom)`, where `top < bottom`. (`printText3` aborts if this is not the case.) The printed text will appear within these y-coordinates only, leaving you room to print a header and a trailer for each page.

`eachPageDrawFunc`**:** this function is responsible for drawing a header and trailer. As parameters it takes the data produced by `textRangeFunc`, the actual page number, the `PrintInfo` record and a `Picture`. This function will be called from `printText3` for each page. It should return the `Picture` in which the header and/or trailer for the current page are drawn. Drawing is not restricted to any particular area on the printer `Picture`.

Chapter 14

TCP

In this chapter we cover Clean's TCP interface. Blocking is an inherent problem with network I/O: A function *blocks*, when it can not be reduced further, until some external condition changes. A program can not do anything, while a function blocks[1]. An example of blocking is a program, that tries to receive some data. Such a program could block until the data arrives.

It makes a difference whether a Clean program engages in event driven I/O or not. The event driven part will be started up by functions like `startIO` or `startProcesses`. If a program calls this function then we call it an "I/O" program, otherwise it is a "World" program. Principally it does not matter when a World program blocks. But care has to be taken with I/O programs: they should not block, or at least not block a "long" time. The "$\frac{1}{10}$ second rule" says, that an I/O program should not take longer than $\frac{1}{10}$ second to process an event. Hence an I/O program should not block longer than this time (if the programmer respects that rule). The reason for this rule is, that the user of an I/O program should have the impression, that the program immediately responds to his actions. Furtheron, window update events should also be handled by the program immediately. Otherwise the windows might get looking ugly.

After a short introduction to TCP in Section 14.1 some basic ideas are discussed in Section 14.2. Section 14.3 is about the blocking approach, which is applicable in World programs. The non blocking approach for I/O programs is discussed in Section 14.4.

14.1 Introduction to TCP

TCP (Transmission Control Protocol) offers a possibility to transfer data between programs which are running on different computers via a network. TCP is a connection oriented protocol. Before data can be transferred a connection between two running programs has to be established. A TCP connection is always a duplex connection. That means, that both sides can send and receive data when they are connected. After all data has been sent and received, the connection has to be teard down.

Two programs that establish a connection with each other are categorized due to their role in establishing the connection: one program is the server and the other one

[1]The Clean compiler produces *sequential* code. This is the reason, why a program can not do anything else, while it is blocking. If the Clean compiler would be able to produce non sequential code, blocking would be much less of a problem.

the client. The client has to know the address of the server. The address consists of two parts: an IP address and a port number. The IP address uniquely identifies every computer on the internet. An IP address is a 32 bit quantity, which is often written down in "dotted decimal form" like e.g. "131.174.33.11". It is possible to establish different connections to one computer at the same time. The port number discriminates these connections. The port number is a 16 bit quantity. For some services some port numbers are reserved. E.g. "HTTP" is reachable via the port number 80 and "telnet" via port number 23. When a connection is established, the following happens: first the server has to *listen* on a certain port. Listening on a port means being prepared for receiving connection requests from clients. The client knows both this port number and the IP address of the computer, where the server is running. So he issues a connection request, which will reach the server via the internet. If the server accepts the connection request, the connection is established.

Now data can be transferred. The sent data will reach the other side without being duplicated or altered in its chronological order.

A service called DNS (Dynamic Name System) helps us human beings to avoid having to remind IP addresses. The DNS service translates alphanumerical aliases like "www.cs.kun.nl" into IP adresses.

14.2 Basic Ideas

In Clean's TCP interface *channels* play an important role. Intuitively a channel is like a pipe: What you fill in on one side comes out on the other end after a while. We do not transport material goods like water or gas with our pipes, but rather information. So we call the things that we stuff in our channels *messages*. It is well possible, that many messages reside within a channel at the same time while they are on their way to the other end. But it is not possible, that they come out in another order than they were put in (we do not support out of band data). The most important things we can do with a channel are of course sending and receiving.

In our approach each channel type has a message type. This means, that we can only send and receive messages with that message type on a certain channel. But of course it is possible to define several channel types. For instance the following channel type is defined:

```
::  *TCP_RCharStream    :== TCP_RCharStream_ Char
```

The `TCP_RCharStream` type is a channel type with the message type `Char`. It is not necessary here to know, how `TCP_RCharStream_` is defined (note that underscore). Generally a channel type is a type constructor with one argument, the message type. If we wanted to receive on such a channel we could apply the receive function:

```
receive     ::          !*(*ch .a)  !*env
            -> (!.a,    !*(*ch .a), !*env)
            |  ChannelEnv env & Receive ch
```

To help understanding that seemingly weird definition we give an example, how this type could be specialized. It is necessary to know that `TCP_CharStream` is an instance of the `Receive` class, and that a `ChannelEnv` can be either the `World`, the `IOSt` or the `PSt`. We defined the type of `receive` as general as possible. The following function definition has a type that is more restrictive than it could be:

```
receive2    ::           TCP_RCharStream  *World
            -> (Char,    TCP_RCharStream, *World)
receive2 ch env = receive ch env
```

We gain the type of `receive2` by uniform substitution of `!*(*ch .a)` with `TCP_RCharStream` and `!.a` with `Char`. From the type of `receive2` we see that we can receive a `Char`, if we apply the `receive2` function on a `TCP_RCharStream` in a `World`.

Channel types are uniquely attributed.

14.3 Blocking TCP in World programs

Since this section is about blocking use of TCP, the described methods are applicable for World programs. Section 14.4 discusses non blocking use of TCP in I/O programs.

The library modules that are relevant for this section are `StdChannels`, `StdTCPDef` and `StdTCPChannels` (see Appendixes A.2, A.33 and A.32). The module `StdTCP` imports all these modules similar to the `StdEnv` and `StdIO` modules.

Generally a TCP session is divided into three phases: the connection establishment phase, the data transfer phase and the disconnect phase. These three phases will be covered in the Sections 14.3.1 to 14.3.3. In these sections small example functions are described. These functions are combined to give an example for a client and a server program in Section 14.3.4. The technique of multiplexing is discussed in Section 14.3.5 which includes a fancy chat server example program. Section 14.3.6 introduces some further channels that can be used.

14.3.1 Establishing a connection

We show the process of establishing a connection first for client programs, and afterwards for server programs.

The client program has to know the IP address and port number of the server program. There is a function that uses the DNS to get an IP Address:

```
lookupIPAddress ::  !String !*env
                ->  (!Maybe IPAddress, !*env)
                |   ChannelEnv env
```

The `String` that is passed can be an alphanumerical internet address like "martinpc.cs.kun.nl", but also a dotted decimal notation like "131.174.33.11". To establish a new connection, the client simply calls the `connectTCP_MT` function:

```
connectTCP_MT   :: !(Maybe !Timeout) !(!IPAddress,!Port) !*env
                -> (!TimeoutReport, !Maybe TCP_DuplexChannel, !*env)
                |  ChannelEnv env

::  Port        :== Int
::  Timeout     :== Int
::  TimeoutReport
                =   TR_Expired
                |   TR_Success
                |   TR_NoSuccess
```

The optional `Timeout` value is measured in (platform dependent) ticks. The `TimeoutReport` informs about success or failure of the attempt to connect. Iff this value is `TR_Success` then `Just` a `TCP_DuplexChannel` is returned. A `TCP_DuplexChannel` is a record that contains two channels: one TCP channel to send and one to receive:

```
::  *TCP_DuplexChannel
       :== DuplexChannel *TCP_SChannel_ *TCP_RChannel_ ByteSeq

::  DuplexChannel sChannel rChannel a
    =   {   sChannel    ::  sChannel a
        ,   rChannel    ::  rChannel a
        }
::  *TCP_SChannel       :== TCP_SChannel_ ByteSeq
::  *TCP_RChannel       :== TCP_RChannel_ ByteSeq
```

The `TCP_SChannel` and the `TCP_RChannel` types are channel types. The type of the messages, which can be received is the abstract data type `ByteSeq`. This type resembles sequences of bytes. So with one message a sequence of bytes can be sent or received (see Appendix A.33).

Example We define a function, that looks up an IP address and tries to connect to that machine on port 2000. For simplicity we let the program abort, if any of the operations fails. Otherwise the result is a `TCP_DuplexChannel`.

```
clientConnect :: !*World -> (!TCP_DuplexChannel, !*World)
clientConnect world
  # (mbIPAddr, world)    = lookupIPAddress "martinpc.cs.kun.nl"
                                           world
  | isNothing mbIPAddr
    = abort "DNS lookup failed"
  # ipAddr               = fromJust mbIPAddr
    (tReport, mbDuplex, world)
                         = connectTCP_MT Nothing (ipAddr, 2000)
                                         world
  | tReport<>TR_Success
    = abort "can't connect to port 2000"
  # duplexChannel        = fromJust mbDuplex
  = (duplexChannel, world)
```

Now we discuss the process of establishing a connection for a server program.

The server has to call the `openTCP_Listener` function to listen on a given port:

```
openTCP_Listener    ::  !Port !*env
                    ->  (!OkBool, !Maybe TCP_Listener, !*env)
                    |   ChannelEnv env

::  *TCP_Listener   :== TCP_Listener_ (IPAddress, TCP_DuplexChannel)
::  OkBool          :== Bool
```

The `Port` parameter should be the desired port number. Iff listening on the given port succeeds, then the `OkBool` result will be `True` and the `Maybe` result will be `Just` a `TCP_Listener`. If another program on the same machine already listens on the given port, this attempt will fail. A `TCP_Listener` is a channel on which connection

requests from remote clients can be received. So the server does this by applying the `receive` function on the `TCP_Listener`:

```
receive     ::          !*(*ch .a)  !*env
            -> (!.a,    !*(*ch .a), !*env)
            |  ChannelEnv env & Receive ch
```

From the type definition of `TCP_Listener` we see, that the message type `a` is a pair, which consists of an `IPAddress` and a `TCP_DuplexChannel`. Receiving a message of such a type establishes a new connection. The `IPAddress` is the IP address of the remote client, which wants to connect. The `TCP_DuplexChannel` is the new connection itself.

So a `TCP_Listener` is a *channel* on which *channels* can be received. With one `receive` on a `TCP_Listener` a `TCP_SChannel` (to send) *and* a `TCP_RChannel` (to receive) are gained. (People familiar with the sockets API will recognize that a `receive` on a `TCP_Listener` resembles the `accept` call.)

Example The following function listens on port 2000 and accepts one connection:

```
serverConnect :: !*World -> (!TCP_DuplexChannel, !*World)
serverConnect world
    #   (ok, mbListener, world)
                            = openTCP_Listener 2000 world
    |   not ok
        = abort "can't open Listener on port 2000"
    #   listener            = fromJust mbListener
        ((_, duplexChannel), listener, world)
                            = receive listener world
        world               = closeRChannel listener world
    = (duplexChannel, world)
```

The `closeRChannel` function is discussed in Section 14.3.3.

14.3.2 Basic Operations on Channels

In the previous section we came in contact with three channel types: the `TCP_-Listener`, the `TCP_SChannel` and the `TCP_RChannel`. These types are instances of type constructor classes, which are defined in the module `StdChannels` (Appendix A.2). The `TCP_Listener` and the `TCP_RChannel` are instances of the `Receive` class:

```
class Receive ch
where
    receive_MT    ::  !(Maybe !Timeout)           !*(*ch .a)  !*env
                  ->  (!TimeoutReport, !Maybe !.a,!*(*ch .a), !*env)
                                                  | ChannelEnv  env
    available     ::                              !*(*ch .a)  !*env
                  ->  (!Bool,                     !*(*ch .a), !*env)
                                                  | ChannelEnv  env
    eom           ::                              !*(*ch .a)  !*env
                  ->  (!Bool,                     !*(*ch .a), !*env)
                                                  | ChannelEnv  env
    ...
```

Each of these functions gets and returns a channel and an environment. "MT" is a shorthand for "Maybe Timeout". The `receive_MT` function tries to receive a message of type `a` on the channel. This function might block until the message arrives. A timeout value can be given to limit this time. The timeout value is given in (platform dependent) ticks. The returned `TimeoutReport` determines whether the receive was succesful (`TR_Success`), whether the timeout expired (`TR_Expired`) or whether the other side teard down the connection (`TR_NoSuccess`). The `available` function polls on the channel, whether a message is currently available. The `eom` ("End of Messages") returns `True`, iff it is sure that available will never become `True` anymore. To understand these functions, the following *model* of receive channels is handy:

A receive channel consists of a message queue and is always in one of three states: the "idle" state, the "available" state, or the "EOM" state, as shown in Figure 14.1.

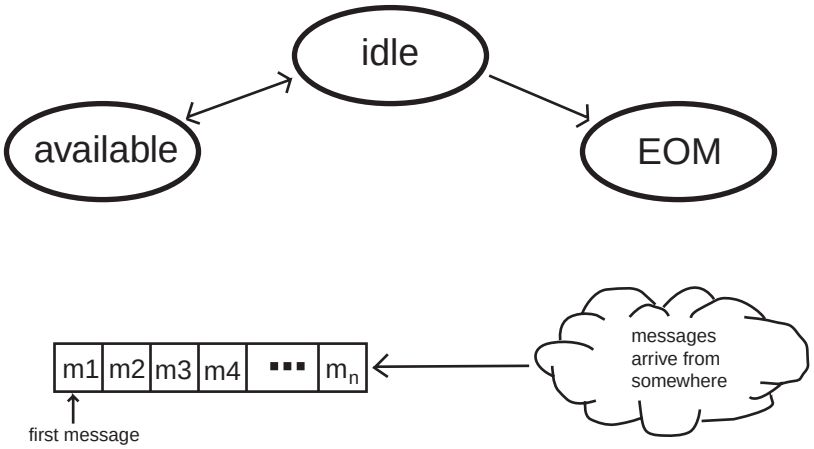

Figure 14.1: model of a receive channel

Iff the receive channel is in the available state, then there is a message in the message queue. Iff the channel is in the idle or EOM state, then the message queue is empty. If a message arrives, then the state can change from idle to available, but if the channel is in the EOM state, no message will arrive anymore. In this case the state can not change anymore, and this happens, when the other side tears down the connection.

To ease programming there are some specialisations of the functions above:

```
receive         ::          !*(*ch .a)  !*env
                -> (!.a,    !*(*ch .a), !*env)
                |   ChannelEnv env & Receive ch
nreceive        ::  !Int    !*(*ch .a)  !*env
                ->  (![.a], !*(*ch .a), !*env)
                |   ChannelEnv env & Receive ch
```

`receive` receives one message from the channel and `(nreceive n)` receives `n` messages. These functions are only partially defined. If the receive channel state is EOM or becomes EOM while the function blocks, then the program will abort! These functions should only be used if it is sure that the specified amount of data will be received.

Example The following function receives a byte sequence and compares it with a given string. If the strings are different, the program aborts.

```
serverReceive :: String TCP_RChannel  *World
              ->       (TCP_RChannel, *World)
serverReceive expectedMessage rChannel world
  # (message, rChannel, world)       = receive rChannel world
  | toString message<>expectedMessage
    = abort "received wrong message"
  = (rChannel, world)
```

A special role play channels on which channels can be received, currently only the `TCP_Listener`. If a message is `available` on such a channel it might be possible that a following `receive_MT` will return `Nothing`. This happens when a client tries to connect to a server, but disconnects again before the connection is accepted (received) from the server. Furtheron `eom` will never get `True` for this kind of channels.

On all other channels it holds, that if `available` is True, receiving on that channel will not block and will return `Just` a message.

Before we discuss the basic operations on send channels, we introduce a model for these channels, see Figure 14.2.

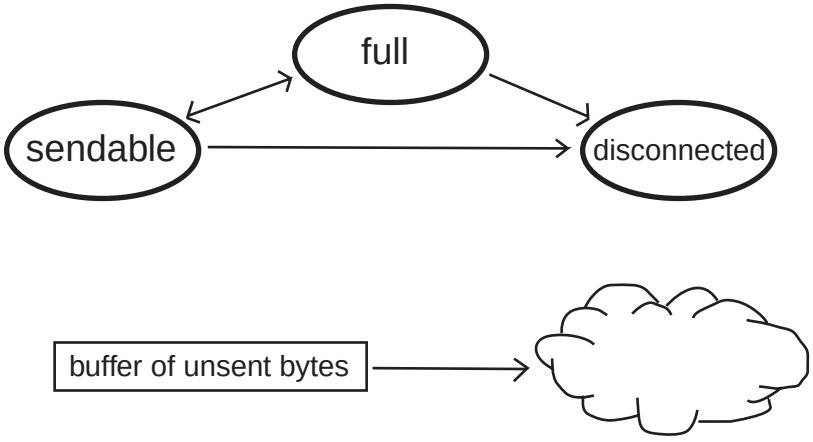

Figure 14.2: model of a send channel

A send channel has also three states: the "sendable" state, the "full" state and the "disconnected" state. Furtheron there is a buffer of unsent bytes. If the channel is in the disconnected state, then no sending of data is possible. This happens, when the remote side tears down the connection. The disconnected state is the final state (similar to the EOM state for receive channels).

Sending on a channel might block as well as receiving does. Since this is not as obvious as in the receive case, here is a scenario where sending becomes blocking: Let's assume a TCP connection, where data is sent only in one direction, from a "producer" program to a "consumer" program. Let's assume further, that the producer sends much faster than the consumer, e.g. the producer sends one MByte per second but the consumer receives only one KByte per second, because he doesn't want to receive faster. It is obvious, that this does not work forever! At a certain point in time all buffers in the two participating computers and in the computers

between them get full. The producer has to be stopped in his overproduction then, which means, that the producer's sending operation will block.

This is the reason, why there are the sendable and the full states and the message queue. Iff a channel is in the full state, then trying to send will not pump any data into the internet, the data will be queued locally in the channels internal buffer instead. Iff the channel state changes from full to sendable then it's reasonable to flush this buffer. The buffer is a part of the Clean heap and does not have a limited size (apart from the usual memory limitations).

Now let's have a look at the basic operations on send channels:

```
class Send ch
where
    send_MT       ::  !(Maybe !Timeout) !.a       !*(*ch .a)  !*env
                  ->  (!TimeoutReport, !Int,      !*(*ch .a), !*env)
                                                  | ChannelEnv  env
    nsend_MT      ::  !(Maybe !Timeout) ![.a]     !*(*ch .a)  !*env
                  ->  (!TimeoutReport, !Int,      !*(*ch .a), !*env)
                                                  | ChannelEnv  env
    flushBuffer_MT::  !(Maybe !Timeout)           !*(*ch .a)  !*env
                  ->  (!TimeoutReport, !Int,      !*(*ch .a), !*env)
                                                  | ChannelEnv  env
    disconnected  ::                              !*(*ch .a)  !*env
                  ->  (!Bool,                     !*(*ch .a), !*env)
                                                  | ChannelEnv  env
    bufferSize    ::                              !*(*ch .a)
                  ->  (!Int,                      !*(*ch .a))
    ...
```

The `send_MT` function sends a message and the `nsend_MT` function sends a list of messages. Both functions take a `(Maybe Timeout)` argument. If not all of the data could be sent within the timeout period, the rest will be queued in the send channels internal buffer. This internal buffer will not be sent automatically, it can be sent (flushed) with the `flushBuffer_MT` function. This function tries to send as much as possible from the internal buffer. The `Int` value that is returned by these three functions is the number of sent bytes. If the timeout expires before everything could be sent, then the returned `TimeoutReport` will be `TR_Expired`. If the send channel's state changed to disconnected during the send operation the `TimeoutReport` will be `TR_NoSucces`. Only if all data was sent, this value will be `TR_Success`. The size of the internal buffer in bytes is returned by the `bufferSize` function. The `disconnected` function is used to poll whether the channel is in the disconnected state.

The following functions are to avoid having to specify timeouts for sending. They are specialisations of the functions in the `Send` class.

```
send            :: !.a          !*(*ch .a)  !*env
                ->              (!*(*ch .a), !*env)
                |   ChannelEnv env & Send ch
nsend           :: ![.a]        !*(*ch .a)  !*env
                ->              (!*(*ch .a), !*env)
                |   ChannelEnv env & nsend_MT ch
```

Example The following function sends the message "hello server".

```
clientSend :: TCP_SChannel *World -> (TCP_SChannel, *World)
clientSend sChannel world
    = send (toByteSeq "hello server") sChannel world
```

The `MaxSize` class allows to limit the size of the received byte sequences for receive channels. For further details see the `StdChannels` module.

Any object whose type is an instance of the `toString` class can be converted into a `ByteSeq`:

```
toByteSeq :: x -> ByteSeq | toString x
```

And vice versa: `ByteSeq`s can be converted into strings, because they are an instance of `toString`.

14.3.3 Tearing down a connection

Using TCP there are principally two ways to tear down a connection: *graceful* and *abortive* disconnects. The abortive disconnect will happen when the `abortConnection` function is applied on a send channel. If the channel should be closed gracefully, then the `closeChannel_MT` function should be applied. These two functions are both members of the `Send` class:

```
class Send ch
where
    closeChannel_MT ::  !(Maybe !Timeout)         !*(*ch .a)  !*env
                    ->  (!TimeoutReport, !Int,                !*env)
                                                  | ChannelEnv  env
    abortConnection ::                            !*(*ch .a)  !*env
                    ->                                        !*env
                                                  | ChannelEnv  env
    ...
```

The `closeChannel_MT` function takes a `(Maybe Timeout)` argument, because it tries to send the channel's internal buffer before closing the channel. This function has also a specialized version without a timeout:

```
closeChannel :: !*(*ch .a) !*env -> !*env
              |  ChannelEnv env & Send ch
```

The internal buffer will not be sent, when the `abortConnection` function is used. Furtheron it is possible that data that is already sent will not reach the other side, when this function is used.

To close receive channels the `closeRChannel` function should be used:

```
class closeRChannel ch :: !*(*ch .a) !*env -> !*env
                       |  ChannelEnv  env
```

This function will not block.

14.3.4 Putting it together

The example functions that were introduced in Section 14.3.1 and 14.3.2 are now assembled together into two programs: a `client` and a `server` program. The client will attempt to connect to the server and send the message "hello server". The server will check, whether the received message matches that text. Both sides will close the connection gracefully afterwards. It is assumed that the server runs on a machine called "martinpc.cs.kun.nl".

Here is the client program:

```
module client

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program implements a minimal TCP session (client).
//
//  ********************************************************************************

import StdEnv, StdTCP

clientConnect :: !*World -> (!TCP_DuplexChannel, !*World)
clientConnect world
    #   (mbIPAddr, world)               = lookupIPAddress "martinpc.cs.kun.nl" world
    |   isNothing mbIPAddr
        = abort "DNS lookup failed"
    #   ipAddr                          = fromJust mbIPAddr
        (tReport, mbDuplex, world)      = connectTCP_MT Nothing (ipAddr, 2000) world
    |   tReport<>TR_Success
        = abort "can't connect to port 2000"
    #   duplexChannel                   = fromJust mbDuplex
    = (duplexChannel, world)

clientSend :: TCP_SChannel *World -> (TCP_SChannel, *World)
clientSend sChannel world
    = send (toByteSeq "hello server") sChannel world

Start world
    #   ({sChannel, rChannel}, world)   = clientConnect world
        (sChannel, world)               = clientSend sChannel  world
        world                           = closeChannel  sChannel world
        world                           = closeRChannel rChannel world
    = world
```

And here is the server:

```
module server

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program implements a minimal TCP session (server).
//
//  ********************************************************************************

import StdEnv, StdTCP

serverConnect :: !*World -> (!TCP_DuplexChannel, !*World)
serverConnect world
    #   (ok, mbListener, world)                     = openTCP_Listener 2000 world
    |   not ok
        = abort "can't open Listener on port 2000"
    #   listener                                    = fromJust mbListener
        ((_, duplexChannel), listener, world)       = receive listener world
```

```
        world                                       = closeRChannel listener world
    = (duplexChannel, world)

serverReceive :: String TCP_RChannel *World -> (TCP_RChannel, *World)
serverReceive expectedMessage rChannel world
    #   (message, rChannel, world)                  = receive rChannel world
    |   toString message<>expectedMessage
        = abort "received wrong message"
    = (rChannel, world)

Start world
    #   ({sChannel, rChannel}, world)  = serverConnect world
        (rChannel, world)              = serverReceive "hello server" rChannel world
        world                          = closeChannel  sChannel world
        world                          = closeRChannel rChannel world
    = world
```

14.3.5 Multiplexing

The `selectChannel_MT` function allows to determine, on which of a set of channels operations can be used in a non blocking way. This is very useful to determine the channel in the set, for which data has arrived first. It's signature is:

```
selectChannel_MT :: !(Maybe !Timeout)
                    !*r_channels  !*s_channels  !*World
                 -> (![(!Int, !SelectResult)],
                    !*r_channels, !*s_channels, !*World)
                |   SelectReceive r_channels & SelectSend s_channels

::  SelectResult
    =   SR_Available
    |   SR_EOM
    |   SR_Sendable
    |   SR_Disconnected
```

The `r_channels` parameter can be a set of lists of receive channels. These lists can be combined with the `:^:` constructor (Appendix A.13). For instance there are defined:

```
:: *TCP_RChannels       =   TCP_RChannels [TCP_RChannel]
:: *TCP_Listeners       =   TCP_Listeners [TCP_Listener]
```

If `tcp_RChannels` is a list of `TCP_RChannels` and `tcp_Listeners` is a list of type `TCP_Listener` then the following expression is a valid value for the `r_channels` parameter:

```
TCP_RChannels tcp_RChannels :^: TCP_Listeners tcp_Listeners
```

The `:^:` constructor allows us to combine lists of channels with different types. Similar sets of send channels can be created for the `s_channels` argument. To specify an empty set of channels, it's also possible to pass

```
Void
```

as a value for `r_channels` or `s_channels`.

If the timeout expires, the function will return an empty list. Otherwise the `Int` part of each result pair will identify one of the channels out of `r_channels` or `s_channels` for which the `SelectResult` holds. For an example let's suppose that both lists `tcp_RChannels` and `tcp_Listeners` contain two elements. In the fragment

```
# r_channels                 =     TCP_RChannels tcp_RChannels
                                :^: TCP_Listeners tcp_Listeners
  ([(who, what):_], r_channels, _, world)
                             = selectChannel_MT Nothing r_channels
                                                Void    world
  (    TCP_RChannels tcp_RChannels
   :^: TCP_Listeners tcp_Listeners)
                             = r_channels
```

a set of four receive channels is passed via the `r_channels` parameter to the select function. So the `who` value will be inbetween zero and three. Let's assume here for simplicity that the list that is returned by the `selectChannel_MT` function only contains one element. If at first some data would have arrived for the first channel in the `tcp_RChannels` list, then `who` would equal zero, for the other element of that list `who` would equal one. Since the `TCP_Listeners` data constructor is the right argument of `:^:`, the `who` value would be two or respectively three, if at first a connection request would have arrived for one of the listeners. In all these cases the value for `what` would have been `SR_Available`. But if at first one of the `TCP_RChannels` would get into the EOM state, then the result would be `SR_EOM`. Since only receive channels are passed to the `selectChannel_MT` function, the `SelectResult` result can only be `SR_Available` or `SR_EOM`. The `SR_Sendable` and the `SR_Disconnected` values can only be returned, if some send channels are part of the `s_channels` parameter. As an example we pass a list `tcp_SChannels` of `TCP_SendChannels` to the `selectChannel_MT` function.

```
# ( [(who, what):_], _,
    TCP_SChannels tcp_SChannels,
    world
  )              = selectChannel_MT Nothing Void
                                    (TCP_SChannels tcp_SChannels)
                                    world
```

The `who` value will identify one of the send channels in the list (`tcp_SChannels !! who`). Since no receive channels were passed to the `selectChannel_MT` function, the `what` value can only be `SR_Sendable`, or `SR_Disconnected`. These values indicate, that the channel is in the sendable respectively disconnected state.

Of course it is also possible to pass receive channels and send channels to `selectChannel_MT`. The `SelectResult` result determines, whether a receive channel or a send channel is identified by the `Int` result. The numbering of channels begins with zero for receive channels and for send channels. If more than one channel could be selected by `selectChannel_MT`, then priority is given from left to right.

As an example program we'll discuss a server for a *chat* application. A chat client program (see Figure 14.3) will ask it's user for a nickname and the address of the host, where the server is running. After connecting to that server the client application will pop up a window with two edit controls. In the upper field the user can input some text, which will be broadcasted via the server to all other people, who are currently connected to the server. The client program is not a topic here,

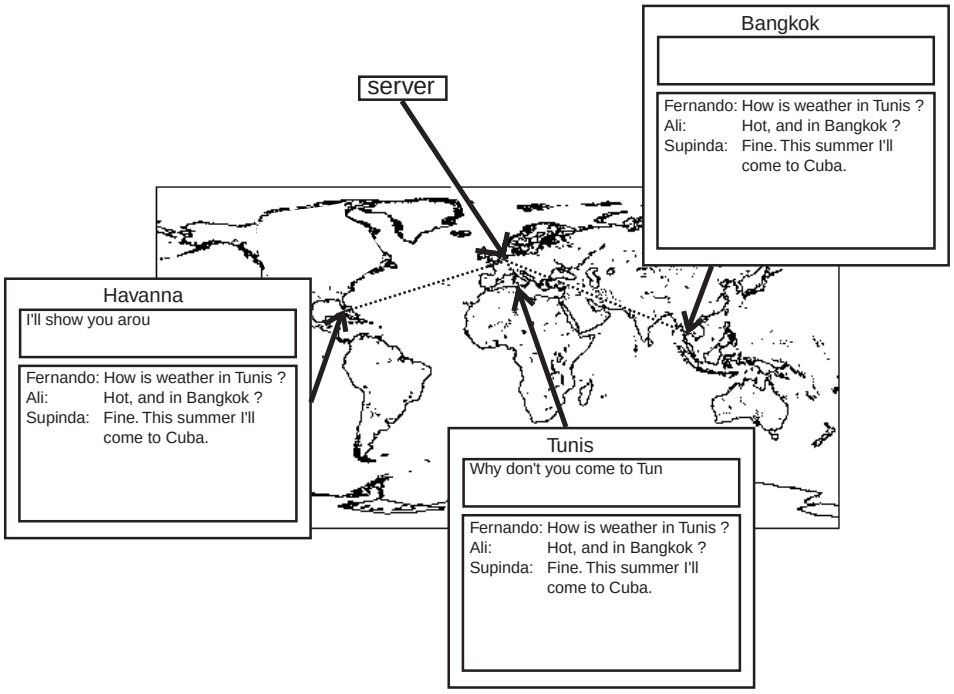

Figure 14.3: chatting via the internet

because it is an I/O program, but it is incorporated in the examples part of the Object I/O library.

The server program at first opens a listener on the chosen port 2000. It uses a list of `ChanInfo` records. Each element of this list corresponds to one connection to a client. Apart from the two channels for the communication, the nickname for the connection is stored in a `ChanInfo` record. The set of receive channels which is passed to the `selectChannel_MT` function consists of the listener and the receive channels for each open connection. There are three cases to handle: a new connection is made, data has been sent, a connection is closed.

```
module chatServer

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program demonstrates the usage of the selectChannel_MT function
//
//  ********************************************************************************

import  StdEnv, StdTCPChannels, StdMaybe

chatPort    :== 2000

::  *ChanInfo
    =   {   sndChan     ::  TCP_SChannel
        ,   rcvChan     ::  TCP_RChannel
        ,   nickname    ::  String
        }

Start world
    # (ok, mbListener, world)   = openTCP_Listener chatPort world
    | not ok
```

```
        = abort ("chatServer: can't listen on port "+++toString chatPort)
    # (console, world)      = stdio world
    = loop (fromJust mbListener) [] console world


loop :: !TCP_Listener ![ChanInfo] !*File !*World -> *World
loop listener channels console world
    # (sChans, rChans, nicknames)
                = unzip3 channels
      glue      = (TCP_Listeners [listener]) :^: (TCP_RChannels rChans)
      ([(who,what):_],glue, _, world)
                = selectChannel_MT Nothing glue Void world
      (TCP_Listeners [listener:_]) :^: (TCP_RChannels rChans)
                = glue
      channels  = zip3 sChans rChans nicknames

    // case 1: someone wants to join the chatroom

    | who==0
        # (tReport, mbNewMember, listener, world)
                        = receive_MT (Just 0) listener world
        | tReport<>TR_Success        // the potential new member changed his mind
            = loop listener channels console world
        # (_,{sChannel,rChannel})   = fromJust mbNewMember
          (byteSeq, rChannel, world)
                        = receive rChannel world
          nickname      = toString byteSeq
          message       = "*** "+++nickname+++" joined the group."
          console       = fwrites (message+++"\n") console
          channel       = {sndChan=sChannel,rcvChan=rChannel,nickname=nickname}
          channels      = [channel:channels]
          (channels, world) = broadcastString message channels [] world
        | nickname % (0,3)=="quit"
            = quit listener channels world
        = loop listener channels console world

    // case 2: somebody has something to say

    | what==SR_Available
        # (channel=:{rcvChan, nickname}, channels)
                        = selectList (who-1) channels
          (byteSeq, rcvChan, world)
                        = receive rcvChan world
          message       = toString byteSeq
          channels      = channels++[{channel & rcvChan=rcvChan}]
          (channels,world)  = broadcastString (nickname+++": "+++message)
                                              channels [] world
        = loop listener channels console world

    // case 3: somebody leaves the group

    | what==SR_EOM
    # ({sndChan, rcvChan, nickname}, channels)
                        = selectList (who-1) channels
      message           = "*** "+++nickname+++" left the group"
      console           = fwrites (message+++"\n") console
      (channels,world)  = broadcastString message channels [] world
      world             = seq [closeChannel sndChan,closeRChannel rcvChan] world
    = loop listener channels console  world

broadcastString :: !String ![ChanInfo] ![ChanInfo] !*World -> ([ChanInfo],!*World)
broadcastString string [] akku world
    = (u_reverse akku, world)
broadcastString string [channel=:{sndChan}:channels] akku world
    # (sndChan, world)      = send (toByteSeq string) sndChan world
    = broadcastString string channels [{channel & sndChan=sndChan}:akku] world
```

```
selectList :: !Int [.a] -> (!.a,![.a])
selectList n l
    # (left, [element:right])   = splitAt n l
    = (element, left++right)

quit listener channels world
    # world = closeRChannel listener world
    = closeChannels channels world

closeChannels [] world
    = world
closeChannels [{sndChan, rcvChan}: channels] world
    # world = seq [closeChannel sndChan, closeRChannel rcvChan] world
    = closeChannels channels world

unzip3 :: ![!ChanInfo] -> (![TCP_SChannel], ![TCP_RChannel], ![String])
unzip3 [] = ([],[],[])
unzip3 [{sndChan, rcvChan, nickname}:t]
    # (a,b,c) = unzip3 t
    = ([sndChan:a], [rcvChan:b], [nickname:c])

zip3 :: ![TCP_SChannel] ![TCP_RChannel] ![String] -> [!ChanInfo]
zip3 [] [] [] = []
zip3 [sndChan:a] [rcvChan:b] [nickname:c]
    = [{sndChan=sndChan, rcvChan=rcvChan, nickname=nickname} : zip3 a b c]

u_reverse list = reverse_ list []
where
    reverse_ [hd:tl] list   = reverse_ tl [hd:list]
    reverse_ [] list        = list
```

Server processes typically run in the backgroud. Macintosh users should be aware that the MacOS was not designed as a multitasking operating system. Hence it is not a default property of programs to be able to run in the background. But the desired behaviour can be achieved if the "Can Backgroud" bit of the "SIZE" resource of the executable file is set. The program that allows to edit these resources of files is called "ResEdit".

14.3.6 More Channels

Sometimes it is very handy, if every single byte can be sent or received atomically. For this purpose we implemented the following two channels:

```
::  *TCP_SCharStream    :== TCP_SCharStream_ Char
::  *TCP_RCharStream    :== TCP_RCharStream_ Char
```

On a `TCP_S(R)CharStream` characters can be sent (received). The functions `toSCharStream` and `toRCharStream` convert `TCP_S(R)Channels` into `TCP_S(R)CharStreams`. The `nsend` and `nreceive` functions are handy to use with this kind of channels. The following function takes a `TCP_SChannel`, converts it into a character stream, and sends the characters "hello partner".

```
s :: !TCP_SChannel !*World -> (!TCP_SCharStream, !*World)
s tcp_SChannel world
    # sCharStream = toSCharStream tcp_SChannel
    = nsend ['hello partner'] sCharStream world
```

If `rCharStream` is the corresponding `TCP_RCharStream` and `world` the `World`, then the other side could receive these thirteen characters with the following function application:

```
nreceive 13 rCharStream world
```

Another kind of channels are the so called string channels. One handicap in using raw TCP is, that the `ByteSeq` packets are not atomic. That means, that sending two byte sequences with sizes of e.g. 10 and 14 bytes could be received on the other side as one byte sequence with a size of 24 bytes. String channels use their own protocol, which is built on top of the TCP protocol. If a string is sent on a string channel, then at first a representation of the size of that string is sent, and afterwards the contents. In this way it is possible to send empty strings as well as strings with sizes of several megabytes. On the corresponding receive channel these strings will be received as a whole. String channels should communicate only with other string channels, since these channels use their very own protocol. It is possible to use the `setMaxSize` function of the `MaxSize` class to limit the size of receivable strings. In this way an application can be protected against bogus programs, which claim to send strings that are to huge to fit in the memory of the used computer.

The definitions and instanciations for string channels can be found in the module `StdStringChannels` (Appendix A.29). We only show here the following:

```
::  *StringSChannel     :== StringSChannel_ String
::  *StringRChannel     :== StringRChannel_ String
```

The character streams and the string channels can be passed to the `selectChannel_MT` function. The naming convention is, that the data constructors for the objects that are passed to the `selectChannel_MT` function end with an additional "`s`". E.g. it is possible to pass `TCP_RCharStreams` and `StringRChannels` to the `selectChannel_MT` function:

```
selectChannel_MT
        Nothing
        (TCP_RCharStreams [ch0,ch1] :^: StringRChannels [ch2,ch3])
        Void
        world
```

14.4 Non Blocking TCP in I/O Programs

As said before, I/O programs should not block for a "long" time. To explain, how this condition can be solved we take a look at a very similar method of handling input: handling keybord input. In World programs we simply can write something like

```
# (line, console) = freadline console
```

The program would block until the user presses the "Enter" key. In an I/O program we should not use the `freadline` function to get keyboard input. We have seen in Section 6.7.1 that we should specify a callback function as a part of the `WindowKeyboard` attribute. When the user presses a key, the runtime system will look up this callback function (which is stored in the `PSt`) and apply it on the key code and the `PSt`. The callback function performs the necessary actions on the program state as a reaction on the user's keyboard input.

With TCP connections similar things happen. One obvious difference is, that a callback function for TCP is not specified as a part of a `WindowAttribute`. Instead

it is specified as a `ReceiverFunction` which is a part of a receiver definition (see Chapter 10). To receive data we have to open a receiver. If via a TCP connection some data arrives, the `ReceiverFunction` of a receiver will be applied to this data. But not only the receiving of data will cause a `ReceiverFunction` to be applied. Generally, if a channel's state changes, the runtime system will generate a certain *event*, on which a `ReceiverFunction` will be applied. There are two events for receive channels:

```
::  ReceiveMsg m = Received m | EOM
```

and two events for send channels:

```
::  SendEvent = Sendable |  Disconnected
```

A `ReceiveMsg m` event informs the application, that the message `m` has arrived. The `EOM` event informs about closure of the channel. The two `SendEvent`s inform the application, that the state of a send channel changed to sendable or disconnected.

Another underlying idea is the following: When receivers are opened, receive channels are *eaten* but send channels are not *eaten*!

Eating has to do with our uniqueness typing system. We say a function eats an object if the type of that object is uniquely attributed, and not returned.

Example Let's have a look on the following two functions:

```
sum_eating :: *{Int} -> Int
sum_eating {[0]=a0, [1]=a1}
    = a0+a1

sum_not_eating :: *{Int} -> (Int, *{Int})
sum_not_eating a=:{[0]=a0, [1]=a1}
    = (a0+a1, a)
```

`sum_eating` eats it's unique array argument, but `sum_not_eating` does not eat it. As a consequence, the passed array can not be used anymore if it is passed to `sum_eating`. The following application would be rejected by the type system:

```
Start = (sum_eating a) + a.[0]
where
    a = {47,11}
```

`a` is not unique on the right hand of the equal sign, because it is used twice. It was tried to use `a` in the expression `a.[0]`, although `sum_eating` has eaten `a`.

Fortunately we can calculate the desired sum with the following rule:

```
Start
    #   (s1, a‘)    = sum_not_eating a
    = s1 + a‘.[0]
where
    a = {47,11}
```

As a consequence, it is impossible to apply any function to a receive channel, after a receiver was opened for such a channel.

Let's examine, how we can open receivers for receive channels.

For each receive channel there is an algebraic receiver definition type, which is used for opening a receiver. An object of such a type is passed to the `openReceiver` function, as shown in Chapter 10. For `TCP_RChannels` there is a `TCP_Receiver`:

```
::  *TCP_Receiver ls ps
    =   TCP_Receiver
            Id TCP_RChannel
            (ReceiverFunction (ReceiveMsg ByteSeq) *(ls,ps))
            [ReceiverAttribute                     *(ls,ps)]

::  ReceiverFunction    m   ps  :== m -> ps -> ps
```

To open such a receiver, an `Id`, a `TCP_RChannel`, a `ReceiverFunction` and `ReceiverAttributes` have to be specified. The `Id` can be used to disable or close the receiver. If data has arrived on the channel, the `ReceiverFunction` will be called with the `Received` alternative. If the connection is teard down by the remote peer, then the `ReceiverFunction` will be called with the `EOM` alternative. Similar definitions are the `TCP_ListenerReceiver` (for `TCP_Listeners`), the `TCP_CharReceiver` (for receiving character by character) and the `StringChannelReceiver` (for `StringRChannels`) (see Appendixes A.33 and A.29).

Example We show a function `f` that opens a receiver for string channels. The `ReceiverFunction rcvFun` will store the received strings in a file, which is also passed to `f` and which is stored in the local state of the receiver. When the `EOM` event happens, the receiver will be closed by the runtime system after evaluation of the `EOM` alternative. It is not necessary to close such a receiver explicitly.

```
f :: StringRChannel *File (PSt .l .p) -> (PSt .l .p)
f stringChannel file pSt
  # (rId, pSt)       = openRId pSt
    (errReport, pSt) = openReceiver
                           file
                           (StringChannelReceiver rId
                                     stringChannel rcvFun [])
                           pSt
  | errReport<>NoError
     = abort "an error occurred"
  = pSt
where
  rcvFun (Received string) (file, pSt)
    = (fwrites string file, pSt)
  rcvFun EOM (file, pSt)
    = (undef, snd (accFiles (fclose file) pSt))
```

Now we turn our focus on receivers for send channels.

Receivers for send channels are called `SendNotifiers`. To open a `SendNotifier` we use the function `openSendNotifier`, which does not eat it's channel argument:

```
openSendNotifier  ::  .ls !(SendNotifier *(*ch .a) .ls (PSt .l .p))
                      !(PSt .l .p)
                  ->  (!ErrorReport,!*(*ch .a),!PSt .l .p)
                  |   accSChannel ch & Send ch

::  SendNotifier sChannel ls ps
    =   SendNotifier
            sChannel
            (ReceiverFunction SendEvent     *(ls,ps))
            [ReceiverAttribute              *(ls,ps)]
```

We see, that no `Id` is used. Indeed we do not need an `Id`. The only thing we want to do with a `SendNotifier` is closing it. This is done automatically, when we close the corresponding send channel.

Since the send channel will not be eaten, it has to be stored somewhere in the program state.

To send some data we can apply the following functions, which are defined in the `StdChannels` module. These functions simply call their `MT` counterpart with a timeout of zero:

```
send_NB         :: !.a          !*(*ch .a)  !*env
                ->             (!*(*ch .a), !*env)
                |   ChannelEnv env & Send ch
flushBuffer_NB  ::              !*(*ch .a)  !*env
                ->             (!*(*ch .a), !*env)
                |   ChannelEnv env & Send ch
```

"`NB`" is a shorthand for "non blocking".

Typically an I/O program uses the `send_NB` function to send in a non blocking way. If not all of the data can be sent immediately because the send channel's state changed to full, the `send_NB` function will store the unsent data in the send channel's internal buffer. When the flow conditions again allow sending, the runtime system will apply the `SendNotifier`'s `ReceiverFunction` on the `Sendable` event. The `ReceiverFunction` should try to flush the internal buffer then by using the `flushBuffer_NB` function.

As an example we will discuss a server program, that accepts one connection and echoes the incoming data. Therefore it opens two receivers in it's initialization function: for the receive channel a receiver will be opened to receive the incoming data and for the send channel a send notifier will be opened. The send notifier allows the application to perform flow control. It should be possible, that the remote side wants to receive the echoed data much slower than it wants to send. The echo server monitors the size of the internal buffer of the send channel. If this buffer size is greater than zero, then the receiver which receives the incoming data will be disabled. This receiver will be enabled only, when due to a `Sendable` event the buffer can be flushed again.

```
module echoServer

//  ********************************************************************************
//  Clean tutorial example program.
//
//  This program demonstrates the usage of functions for event driven TCP.
//  It listens on port 7, accepts a connection and echoes the input
//
```

```
//  ********************************************************************************

import  StdEnv, StdTCP, StdIO

echoPort    :== 7

::  *PState :== PSt TCP_DuplexChannel Bool
//  The Boolean value stores, whether EOM happened on the receive channel.

Start world
    # (_, mbListener, world)    = openTCP_Listener echoPort world
      ((_,duplexChan), listener, world)
                                = receive (fromJust mbListener) world
      world                     = closeRChannel listener world
    = startIO duplexChan False [initialize] [] world

//////////////////////////////////////////////////////////
////    initialize - the function to initialze the PSt ////
//////////////////////////////////////////////////////////

initialize  :: PState -> PState
initialize pSt=:{ ls={rChannel,sChannel}, io }
    # (tcpRcvId, io)    = openId io
      pSt               = { pSt & ls = { rChannel=undef, sChannel=undef }, io=io }

  // open a receiver for the receive channel

      (errReport1, pSt) = openReceiver tcpRcvId
                                    (TCP_Receiver tcpRcvId rChannel rcvFun []) pSt

  // open a receiver for the send channel

      (errReport2, sChannel, pSt)
                        = openSendNotifier tcpRcvId
                                    (SendNotifier sChannel sndFun []) pSt
    | errReport1<>NoError || errReport2<>NoError
        = abort "error: can't open receiver"
    = { pSt & ls={ rChannel=undef, sChannel=sChannel } }

//////////////////////////////////////////////////////////////////////////
////    rcvFun - the callback function for the receive channels receiver ////
//////////////////////////////////////////////////////////////////////////

rcvFun :: (ReceiveMsg ByteSeq) (Id,PState) -> (Id,PState)
rcvFun (Received byteSeq) (tcpRcvId, pSt=:{ ls=ls=:{sChannel}, io})
    # (sChannel, io)        = send_NB byteSeq sChannel io
      (buffSize,sChannel)   = bufferSize sChannel

  //  disable this receiver, if the send channel is full

      io = case buffSize of
            0 -> io
            _ -> disableReceivers [tcpRcvId] io
    = (tcpRcvId, { pSt & ls={ ls & sChannel=sChannel}, io=io })
rcvFun EOM (tcpRcvId, pSt=:{ ls=ls=:{sChannel}, io})
    # (buffSize,sChannel)   = bufferSize sChannel
      pSt   = { pSt & ls  = { ls & sChannel=sChannel}, ps=True, io=io }

  //  close program only, if all data in the send channel's internal buffer has been
  //  sent

      pSt = case buffSize of
                0   -> closeProcess (close pSt)
                _   -> pSt
    = (tcpRcvId, pSt)

//////////////////////////////////////////////////////////////////////
```

```
////    sndFun - the callback function for the send channels receiver ////
//////////////////////////////////////////////////////////////////////////

sndFun :: SendEvent (Id,PState) -> (Id,PState)
sndFun Sendable (tcpRcvId, pSt=:{ ls=ls=:{sChannel}, ps=eomHappened ,io})
    # (sChannel, io)        = flushBuffer_NB sChannel io
      (buffSize,sChannel)   = bufferSize sChannel
      pSt   = { pSt & ls  = { ls & sChannel=sChannel}, io=io }

  //  enable the receive channel's receiver again, if the send channel is still
  //  sendable

      pSt = case (buffSize,eomHappened) of
                (0, False)  -> { pSt & io = enableReceivers [tcpRcvId] pSt.io }
                (0, True )  -> close pSt
                _           -> pSt
    = (tcpRcvId, pSt)
sndFun Disconnected (ls, pSt)
    = (ls, closeProcess pSt)

close   ::  PState -> PState
close pSt=:{ls=ls=:{sChannel}, io}
    #   io              = closeChannel sChannel io
    = { pSt & ls={ ls & sChannel=undef}, io=io }
```

Is it forbidden to use blocking functions in I/O programs? This question arises because using blocking functions mostly results in nicer programs. The answer is: of course it is not forbidden. The only problem that arises when blocking functions are used is, that the program will simply do not anything else while it blocks. In particular it will not update the contents of windows. The programmer has to estimate himself whether this is acceptable or not. Furtheron it depends on many factors, whether the functions will block or not. For example there is no problem to send data in a blocking way to a program which is known to receive fast enough.

Appendix A

I/O library

A.1 StdBitmap

```
definition module StdBitmap


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdBitmap contains functions for reading bitmap files and drawing bitmaps.
//  ********************************************************************************


import  StdMaybe
from    StdFile     import FileSystem
from    osbitmap    import Bitmap
import  StdPicture
export  FileSystem World


openBitmap      :: !{#Char} !*env -> (!Maybe Bitmap,!*env)  | FileSystem env
/*  openBitmap reads in a bitmap from file.
    The String argument must be the file name of the bitmap.
    If the bitmap could be read, then (Just bitmap) is returned, otherwise Nothing
    is returned.
*/

getBitmapSize   :: !Bitmap -> Size
/*  getBitmapSize returns the size of the given bitmap.
    In case the bitmap is the result of an erroneous openBitmap, then the size is
    zero.
*/

resizeBitmap :: !Bitmap !Size -> Bitmap
/*  zooms or stretches a bitmap. The second argument is the size
    of the resulting bitmap
*/

instance Drawables Bitmap
/*  draw       bitmap
        draws the given bitmap with its left top at the current pen position.
    drawAt pos bitmap
        draws the given bitmap with its left top at the given pen position.
    undraw(At)
        equals unfill(At) the box {box_w=w,box_h=h} with {w,h} the size of bitmap.
*/
```

A.2 StdChannels

```
definition module StdChannels

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdChannels defines operations on channels
//  ********************************************************************************

import  StdMaybe
from    StdOverloaded   import ==, toString
from    channelenv      import  ChannelEnv

instance ChannelEnv World
// other instances are IOSt & PSt (see StdPSt)

////////////////////////////////// receive channels ////////////////////////////////

class Receive ch
where
    receive_MT      ::  !(Maybe !Timeout)           !*(*ch .a)  !*env
                    ->  (!TimeoutReport, !Maybe !.a,!*(*ch .a), !*env)
                                                    | ChannelEnv  env
    receiveUpTo     ::  !Int                        !*(*ch .a)  !*env
                    ->  (![.a],                     !*(*ch .a), !*env)
                                                    | ChannelEnv  env
    available       ::                              !*(*ch .a)  !*env
                    ->  (!Bool,                     !*(*ch .a), !*env)
                                                    | ChannelEnv  env
    eom             ::                              !*(*ch .a)  !*env
                    ->  (!Bool,                     !*(*ch .a), !*env)
                                                    | ChannelEnv  env
/*  receive_MT
        tries to receive on a channel. This function will block, until data can be
        received, eom becomes true or the timeout expires.
    receiveUpTo max ch env
        receives messages on a channel until available becomes False or max
        messages have been received.
    available
        polls on a channel, whether some data is ready to be received. If the
        returned Boolean is True, then a following receive_MT will not block and
        return TR_Success.
    eom ("end of messages")
        polls on a channel, whether data can't be received anymore.
*/

class closeRChannel ch  :: !*(*ch .a) !*env -> !*env    | ChannelEnv  env
// closes the channel

//////////////////////////////// send channels /////////////////////////////////////

class Send ch
where
    send_MT         ::  !(Maybe !Timeout) !.a       !*(*ch .a)  !*env
                    ->  (!TimeoutReport, !Int,      !*(*ch .a), !*env)
                                                    | ChannelEnv  env
    nsend_MT        ::  !(Maybe !Timeout) ![.a]     !*(*ch .a)  !*env
                    ->  (!TimeoutReport, !Int,      !*(*ch .a), !*env)
                                                    | ChannelEnv  env
    flushBuffer_MT  ::  !(Maybe !Timeout)           !*(*ch .a)  !*env
                    ->  (!TimeoutReport, !Int,      !*(*ch .a), !*env)
                                                    | ChannelEnv  env
    closeChannel_MT ::  !(Maybe !Timeout)           !*(*ch .a)  !*env
                    ->  (!TimeoutReport, !Int,                  !*env)
                                                    | ChannelEnv  env
    abortConnection ::                              !*(*ch .a)  !*env
```

```
                    ->                                          !*env
                                                          | ChannelEnv  env
    disconnected    ::                                    !*(*ch .a)  !*env
                    ->  (!Bool,                           !*(*ch .a), !*env)
                                                          | ChannelEnv  env
    bufferSize      ::                                    !*(*ch .a)
                    ->  (!Int,                            !*(*ch .a))
/*  send_MT mbTimeout a ch env
        adds the data a to the channels internal buffer and tries to send this buffer
    nsend_MT mbTimeout l ch env
        adds the data l to the channels internal buffer and tries to send this buffer
    flushBuffer_MT
        tries to send the channels internal buffer
    closeSChannel_MT
        first tries to send the channels internal buffer and then closes the channel
    abortConnection
        will cause an abortive disconnect (sent data can be lost)
    disconnected
        polls on a channel, whether data can't be sent anymore. If the returned
        Boolean is True, then a following send_MT will not block and return
        TR_NoSuccess
    bufferSize returns the size of the channels internal buffer in bytes

    The integer value that is returned by send_MT, nsend_MT, flushBuffer_MT &
    closeSChannel_MT is the number of sent bytes.
*/

////////////////////////////////// miscellaneous ////////////////////////////////////

class MaxSize ch
  where
    setMaxSize      ::  !Int !*(*ch .a) -> *(*ch .a)
    getMaxSize      ::       !*(*ch .a) -> (!Int, !*(*ch .a))
    clearMaxSize    ::       !*(*ch .a) -> *(*ch .a)
//  to set,get or clear the maximum size of the data that can be received

:: DuplexChannel sChannel rChannel a
    =   {   sChannel    ::  sChannel a
        ,   rChannel    ::  rChannel a
        }

::  TimeoutReport
    =   TR_Expired
    |   TR_Success
    |   TR_NoSuccess

::  Timeout     :== Int     // timeout in ticks

::  ReceiveMsg m        =   Received m
                        |   EOM
                        //  receiving "EOM" will automatically close the receiver
::  SendEvent           =   Sendable
                        |   Disconnected
                        //  receiving "Disconnected" will automatically close the
                        //  receiver

instance == TimeoutReport
instance toString TimeoutReport

////////////////////////////////// derived functions ////////////////////////////////

nreceive_MT         ::  !(Maybe !Timeout) !Int      !*(*ch .a)  !*env
                    ->  (!TimeoutReport, ![.a],     !*(*ch .a), !*env)
                                                    | Receive ch & ChannelEnv  env
/*  nreceive_MT mbTimeout n ch env
        tries to call receive_MT n times. If the result is (tReport, l, ch2, env2),
        then the following holds:
```

```
            tReport==TR_Succes      <=> length l==n
            tReport==TR_NoSuccess    => length l<n
*/

//  the following two receive functions call their "_MT" counterpart with no
//  timeout. If the data can't be received because eom became True the function will
//  abort.

receive         ::          !*(*ch .a)  !*env
                -> (!.a,    !*(*ch .a), !*env)
                |   ChannelEnv env & Receive ch
nreceive        ::  !Int    !*(*ch .a)  !*env
                ->  (![.a], !*(*ch .a), !*env)
                |   ChannelEnv env & Receive ch

//  the following three send functions call their "_MT" counterpart with no timeout.

send            :: !.a          !*(*ch .a)  !*env
                ->             (!*(*ch .a), !*env)
                |   ChannelEnv env & Send ch
nsend           :: ![.a]        !*(*ch .a)  !*env
                ->             (!*(*ch .a), !*env)
                |   ChannelEnv env & nsend_MT ch
closeChannel    ::              !*(*ch .a)  !*env
                ->                          !*env
                |   ChannelEnv env & Send ch

//  the following two send functions call their "_MT" counterpart with timeout == 0.
//  "NB" is a shorthand for "non blocking"

send_NB         :: !.a          !*(*ch .a)  !*env
                ->             (!*(*ch .a), !*env)
                |   ChannelEnv env & Send ch
flushBuffer_NB  ::              !*(*ch .a)  !*env
                ->             (!*(*ch .a), !*env)
                |   ChannelEnv env & Send ch
```

A.3 StdClipboard

```
definition module StdClipboard


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdClipboard specifies all functions on the clipboard.
//  ********************************************************************************


import  StdMaybe
from    iostate import  PSt, IOSt


//  Clipboard data items:

::  ClipboardItem

class Clipboard item where
    toClipboard     :: !item            -> ClipboardItem
    fromClipboard   :: !ClipboardItem   -> Maybe item
/*  toClipboard
        makes an item transferable to the clipboard.
    fromClipboard
        attempts to retrieve an item of the instance type from the clipboard item.
        If this fails, the result is Nothing, otherwise it is (Just item).
*/

instance Clipboard {#Char}


//  Access to the current content of the clipboard:

setClipboard :: ![ClipboardItem] !(PSt .l .p) -> PSt .l .p
getClipboard :: !(PSt .l .p) -> (![ClipboardItem],!PSt .l .p)
/*  setClipboard
        replaces the current content of the clipboard with the argument list.
        Of the list only the first occurence of a ClipboardItem of the same type
        will be stored in the clipboard.
        Note that setClipboard [] erases the clipboard.
    getClipboard
        gets the current content of the clipboard without changing the content.
*/


clipboardHasChanged :: !(PSt .l .p) -> (!Bool,!PSt .l .p)
/*  clipboardHasChanged holds if the current content of the clipboard is different
    from the last access to the clipboard.
*/
```

A.4 StdControl

```
definition module StdControl


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdControl specifies all control operations.
//  Changing controls in a window/dialogue requires a *WState.
//  Reading the status of controls requires a WState.
//  ********************************************************************************


import  StdControlDef, StdMaybe
from    iostate     import IOSt


::  WState


getWindow               :: !Id !(IOSt .l .p) -> (!Maybe WState, !IOSt .l .p)
/*  getWindow returns a read-only WState for the indicated window.
    In case the indicated window does not exist Nothing is returned.
*/

setWindow               :: !Id ![IdFun *WState] !(IOSt .l .p) -> IOSt .l .p
/*  Apply the control changing functions to the current state of the indicated
    window.
    In case the indicated window does not exist nothing happens.
*/


/*  Functions that change the state of controls.
    When applied to unknown Ids these functions have no effect.
*/
showControls            :: ![Id]                            !*WState -> *WState
hideControls            :: ![Id]                            !*WState -> *WState
/*  (show/hide)Controls makes the indicated controls visible/invisible.
    Hiding a control overrides the visibility of its elements, which become
        invisible.
    Showing a hidden control re-establishes the visibility state of its elements.
*/

enableControls          :: ![Id]                            !*WState -> *WState
disableControls         :: ![Id]                            !*WState -> *WState
/*  (en/dis)ableControls (en/dis)ables the indicated controls.
    Disabling a control overrides the SelectStates of its elements, which become
        unselectable.
    Enabling a disabled control re-establishes the SelectStates of its elements.
*/

markCheckControlItems   :: !Id ![Index]                     !*WState -> *WState
unmarkCheckControlItems :: !Id ![Index]                     !*WState -> *WState
/*  (unm/m)arkCheckControlItems unmarks/marks the indicated check items of the given
    CheckControl. Indices range from 1 to the number of check items. Illegal indices
    are ignored.
*/

selectRadioControlItem  :: !Id  !Index                      !*WState -> *WState
/*  selectRadioControlItem marks the indicated radio item of a RadioControl, causing
    the mark of the previously marked radio item to disappear. The item is given by
    the Id of the RadioControl and its index position (counted from 1).
*/

selectPopUpControlItem  :: !Id  !Index                      !*WState -> *WState
/*  selectPopUpControlItem marks the indicated popup item of a PopUpControl, causing
```

```
    the mark of the previously marked popup item to disappear. The item is given by
    the Id of the PopUpControl and its index position (counted from 1).
*/

moveControlViewFrame    :: !Id Vector                       !*WState -> *WState
/*  moveControlViewFrame moves the orientation of the CompoundControl over the given
    vector, and updates the control if necessary. The control frame is not moved
    outside the ViewDomain of the control. MoveControlViewFrame has no effect if the
    indicated control has no ControlDomain attribute.
*/

setControlTexts         :: ![(Id,String)]                   !*WState -> *WState
/*  setControlTexts sets the text of the indicated (Text/Edit/Button)Controls.
    If the indicated control is a (Text/Button)Control, then AltKey are interpreted
    by the system.
    If the indicated control is an EditControl, then the text is taken as it is.
*/

setEditControlCursor    :: !Id !Int                         !*WState -> *WState
/*  setEditControlCursor sets the cursor at position @2 of the current content of
    the EditControl.
    In case @2<0, then the cursor is set at the start of the current content.
    In case @2>size content, then the cursor is set at the end of the current
    content.
*/

setControlLooks         :: ![(Id,Bool,Look)]                !*WState -> *WState
/*  setControlLooks applied to a CompoundControl turns it into a non-transparant
    CompoundControl.
    Setting the Look only redraws the indicated controls if the corresponding
    Boolean is True.
*/

setSliderStates         :: ![(Id,SliderState->SliderState)] !*WState -> *WState
setSliderThumbs         :: ![(Id,Int)]                      !*WState -> *WState
/*  setSliderStates
        applies the function to the current SliderState of the indicated
        SliderControl and redraws the settings if necessary.
    setSliderThumbs
        sets the new thumb value of the indicated SliderControl and redraws the
        settings if necessary.
*/


drawInControl           :: !Id ![DrawFunction]              !*WState -> *WState
/*  Draw in a (Custom(Button)/Compound)Control. If the CompoundControl is
    transparant then this operation has no effect.
*/


getControlTypes         ::          !WState -> [(ControlType,Maybe Id)]
getCompoundTypes        :: !Id      !WState -> [(ControlType,Maybe Id)]
/*  getControlTypes
        yields the list of ControlTypes of the component controls of this window.
    getCompoundTypes
        yields the list of ControlTypes of the component controls of this
        CompoundControl.
    For both functions (Just id) is yielded if the component control has a
    (ControlId id) attribute, and Nothing otherwise. Component controls are not
    collected recursively through CompoundControls.
    If the indicated CompoundControl is not a CompoundControl, then [] is yielded.
*/

getControlLayouts       :: ![Id] !WState -> [(Bool,(Maybe ItemPos,ItemOffset))]
                                                                // (Nothing,zero)
getControlViewSizes     :: ![Id] !WState -> [(Bool,Size)]               // zero
getControlSelectStates  :: ![Id] !WState -> [(Bool,SelectState)]        // Able
```

```
getControlShowStates    :: ![Id] !WState -> [(Bool,Bool)]               // False
getControlTexts         :: ![Id] !WState -> [(Bool,Maybe String)]       // Nothing
getControlNrLines       :: ![Id] !WState -> [(Bool,Maybe NrLines)]      // Nothing
getControlLooks         :: ![Id] !WState -> [(Bool,Maybe Look)]         // Nothing
getControlMinimumSizes  :: ![Id] !WState -> [(Bool,Maybe Size)]         // Nothing
getControlResizes       :: ![Id] !WState -> [(Bool,Maybe ControlResizeFunction)]
                                                                        // Nothing
getRadioControlItems    :: ![Id] !WState -> [(Bool,Maybe [TextLine])]   // Nothing
getRadioControlSelection:: ![Id] !WState -> [(Bool,Maybe Index)]        // Nothing
getCheckControlItems    :: ![Id] !WState -> [(Bool,Maybe [TextLine])]   // Nothing
getCheckControlSelection:: ![Id] !WState -> [(Bool,Maybe [Index])]      // Nothing
getPopUpControlItems    :: ![Id] !WState -> [(Bool,Maybe [TextLine])]   // Nothing
getPopUpControlSelection:: ![Id] !WState -> [(Bool,Maybe Index)]        // Nothing
getSliderDirections     :: ![Id] !WState -> [(Bool,Maybe Direction)]    // Nothing
getSliderStates         :: ![Id] !WState -> [(Bool,Maybe SliderState)]  // Nothing
getControlViewFrames    :: ![Id] !WState -> [(Bool,Maybe ViewFrame)]    // Nothing
getControlViewDomains   :: ![Id] !WState -> [(Bool,Maybe ViewDomain)]   // Nothing
getControlItemSpaces    :: ![Id] !WState -> [(Bool,Maybe (Int,Int))]    // Nothing
getControlMargins       :: ![Id] !WState -> [(Bool,Maybe ((Int,Int),(Int,Int)))]
                                                                        // Nothing

/*  Functions that return the current state of controls.
    The result list is of equal length as the argument Id list. Each result list
    element corresponds in order with the argument Id list. Of each element the
    first Boolean result is False in case of invalid Ids (if so dummy values are
    returned - see comment).
    Important: controls with no ControlId attribute, or illegal ids, can not be
    found in the WState!

    getControlLayouts
        Yields (Just ControlPos) if the indicated control had a ControlPos attribute
        and Nothing otherwise. The ItemOffset offset is the exact current location
        of the indicated control (LeftTop,offset).
    getControlViewSizes
        Yields the current view frame size of the indicated control. Note that for
        any control other than the CompoundControl this is the exact size of the
        control.
    getControlSelectStates
        Yields the current SelectState of the indicated control.
    getControlShowStates
        Yields True if the indicated control is visible, and False otherwise.
    getControlTexts
        Yields (Just text) of the indicated (Text/Edit/Button)Control.
        If the control is not such a control, then Nothing is yielded.
    getControlNrLines
        Yields (Just nrlines) of the indicated EditControl.
        If the control is not such a control, then Nothing is yielded.
    getControlLooks
        Yields the Look of the indicated (Custom/CustomButton/Compound)Control.
        If the control is not such a control, or is a transparant CompoundControl,
        then Nothing is yielded.
    getControlMinimumSizes
        Yields (Just minimumsize) if the indicated control had a ControlMinimumSize
        attribute and Nothing otherwise.
    getControlResizes
        Yields (Just resizefunction) if the indicated control had a ControlResize
        attribute and Nothing otherwise.
    getRadioControlItems
        Yields the TextLines of the items of the indicated RadioControl.
        If the control is not such a control, then Nothing is yielded.
    getRadioControlSelection
        Yields the index of the selected radio item of the indicated RadioControl.
        If the control is not such a control, then Nothing is yielded.
    getCheckControlItems
        Yields the TextLines of the items of the indicated CheckControl.
        If the control is not such a control, then Nothing is yielded.
    getCheckControlSelection
```

```
        Yields the indices of the selected checkitems of the indicated CheckControl.
        If the control is not such a control, then Nothing is yielded.
    getPopUpControlItems
        Yields the TextLines of the items of the indicated PopUpControl.
        If the control is not such a control, then Nothing is yielded.
    getPopUpControlSelection
        Yields the Index of the indicated PopUpControl.
        If the control is not such a control, then Nothing is yielded.
    getSliderDirections
        Yields (Just Direction) of the indicated SliderControl.
        If the control is not such a control, then Nothing is yielded.
    getSliderStates
        Yields (Just SliderState) of the indicated SliderControl.
        If the control is not such a control, then Nothing is yielded.
    getControlViewFrames
        Yields (Just ViewFrame) of the indicated CompoundControl.
        If the control is not such a control, then Nothing is yielded.
    getControlViewDomains
        Yields (Just ViewDomain) of the indicated CompoundControl.
        If the control is not such a control, then Nothing is yielded.
    getControlItemSpaces
        Yields (Just (horizontal space,vertical space)) of the indicated
        CompoundControl.
        If the control is not such a control, then Nothing is yielded.
    getControlMargins
        Yields (Just (ControlHMargin,ControlVMargin)) of the indicated
        CompoundControl.
        If the control is not such a control, then Nothing is yielded.
*/
```

A.5 StdControlClass

```
definition module StdControlClass


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdControlClass define the standard set of controls instances.
//  ********************************************************************************


import  StdIOCommon, StdControlDef
from    windowhandle    import ControlState
from    StdPSt          import PSt, IOSt


class Controls cdef where
    controlToHandles:: !(cdef .ls (PSt .l .p)) -> [ControlState .ls (PSt .l .p)]
    getControlType  ::  (cdef .ls .ps) -> ControlType

instance Controls (AddLS  c)            | Controls c
instance Controls (NewLS  c)            | Controls c
instance Controls (ListLS c)            | Controls c
instance Controls NilLS
instance Controls ((:+:) c1 c2)         | Controls c1 & Controls c2
instance Controls RadioControl
instance Controls CheckControl
instance Controls PopUpControl
instance Controls SliderControl
instance Controls TextControl
instance Controls EditControl
instance Controls ButtonControl
instance Controls CustomButtonControl
instance Controls CustomControl
instance Controls (CompoundControl c)   | Controls c
```

A.6 StdControlDef

```
definition module StdControlDef


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdControl contains the types to define the standard set of controls.
//  ********************************************************************************


import  StdIOCommon
from    StdPicture  import  DrawFunction, Picture


::  RadioControl        ls ps
 =  RadioControl        [RadioControlItem *(ls,ps)] RowsOrColumns Index
                                            [ControlAttribute *(ls,ps)]
::  CheckControl        ls ps
 =  CheckControl        [CheckControlItem *(ls,ps)] RowsOrColumns
                                            [ControlAttribute *(ls,ps)]
::  PopUpControl        ls ps
 =  PopUpControl        [PopUpControlItem *(ls,ps)] Index
                                            [ControlAttribute *(ls,ps)]
::  SliderControl       ls ps
 =  SliderControl       Direction Length SliderState (SliderAction *(ls,ps))
                                            [ControlAttribute *(ls,ps)]
::  TextControl         ls ps
 =  TextControl         TextLine                [ControlAttribute *(ls,ps)]
::  EditControl         ls ps
 =  EditControl         TextLine Width NrLines  [ControlAttribute *(ls,ps)]
::  ButtonControl       ls ps
 =  ButtonControl       TextLine                [ControlAttribute *(ls,ps)]
::  CustomButtonControl ls ps
 =  CustomButtonControl Size Look               [ControlAttribute *(ls,ps)]
::  CustomControl       ls ps
 =  CustomControl       Size Look               [ControlAttribute *(ls,ps)]
::  CompoundControl      c ls ps
 =  CompoundControl     (c ls ps)               [ControlAttribute *(ls,ps)]

::  TextLine                :== String
::  NrLines                 :== Int
::  Width                   :== Int
::  Length                  :== Int
::  RowsOrColumns
    =   Rows    Int
    |   Columns Int
::  RadioControlItem    ps  :== (TextLine,            IOFunction ps)
::  CheckControlItem    ps  :== (TextLine, MarkState, IOFunction ps)
::  PopUpControlItem    ps  :== (TextLine,            IOFunction ps)
::  Look                    :== SelectState -> UpdateState -> *Picture -> *Picture
::  SliderAction        ps  :== SliderMove  -> ps -> ps
::  SliderMove
    =   SliderIncSmall
    |   SliderDecSmall
    |   SliderIncLarge
    |   SliderDecLarge
    |   SliderThumb Int

::  ControlAttribute ps                         // Default:
    =   ControlId           Id                  // no id
    |   ControlPos          ItemPos             // (RightTo previous,zero)
    |   ControlSize         Size                // system derived/overruled
    |   ControlMinimumSize  Size                // zero
    |   ControlResize       ControlResizeFunction   // no resize
    |   ControlSelectState  SelectState          // control Able
```

```
    |   ControlHide                                  // initially visible
    |   ControlFunction     (IOFunction     ps)     // id
    |   ControlModsFunction (ModsIOFunction ps)     // ControlFunction
    |   ControlMouse        MouseStateFilter SelectState (MouseFunction ps)
                                                    // no mouse input/overruled
    |   ControlKeyboard     KeyboardStateFilter SelectState (KeyboardFunction ps)
                                                    // no keyboard input/overruled
//  For CompoundControls only:
    |   ControlItemSpace    Int Int                 // system dependent
    |   ControlHMargin      Int Int                 // system dependent
    |   ControlVMargin      Int Int                 // system dependent
    |   ControlLook         Look                    // control is transparant
    |   ControlViewDomain   ViewDomain              // {zero,max range}
    |   ControlOrigin       Point                   // Left top of ViewDomain
    |   ControlHScroll      ScrollFunction          // no horizontal scrolling
    |   ControlVScroll      ScrollFunction          // no vertical   scrolling

::  ControlResizeFunction
    :== Size ->                                     // current  control size
        Size ->                                     // old      window  size
        Size ->                                     // new      window  size
        Size                                        // new      control size
::  ScrollFunction
    :== ViewFrame       ->                          // current  view
        SliderState     ->                          // current  state of scrollbar
        SliderMove      ->                          // action of the user
        Int                                         // new thumb value of scrollbar
::  ControlType
    :== String
```

A.7 StdControlReceiver

```
definition module StdControlReceiver


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdControlReceiver defines Receiver(2) controls instances.
//  ********************************************************************************


import  StdReceiverDef, StdControlClass


instance Controls (Receiver  m  )
instance Controls (Receiver2 m r)
```

A.8 StdEventTCP

```
definition module StdEventTCP

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdEventTCP provides functions for using event driven TCP
//  ********************************************************************************

import  StdChannels, StdTCPDef
import  StdReceiver
from    tcp_bytestreams import TCP_SCharStream_

instance Receivers      TCP_ListenerReceiver
instance Receivers      TCP_Receiver
instance Receivers      TCP_CharReceiver

openSendNotifier        ::  .ls !(SendNotifier *(*ch .a) .ls (PSt .l .p))
                            !(PSt .l .p)
                        ->  (!ErrorReport,!*(*ch .a),!PSt .l .p)
                        |   accSChannel ch & Send ch
/*  opens a send notifier, which informs the application, that sending on the
    channel is again possible due to flow conditions. Possible error reports are
    NoError and ErrorNotifierOpen
*/

closeSendNotifier       ::  !*(*ch .a) !(IOSt .l .p)
                        ->  (!*(*ch .a), !IOSt .l .p)
                        |   accSChannel ch
/*  to close a send notifier. This function will be called implicitly if a send
    channel is closed, so there is no need to do it explicitly then.
*/

lookupIPAddress_async   ::  !String !(InetLookupFunction (PSt .l .p)) !(PSt .l .p)
                        ->  (PSt .l .p)
/*  lookupIPAddress_async asynchronously looks up an IP address. The String can be
    in dotted decimal form or alphanumerical. The InetLookupFunction will be called
    with the IP address, if this address was found, otherwise with Nothing.
*/

connectTCP_async        ::  !(!IPAddress,!Port) !(InetConnectFunction (PSt .l .p))
                            !(PSt .l .p)
                        ->  (PSt .l .p)
/*  connectTCP_async asynchronously tries to establish a new connection. The
    InetConnectFunction will be called with the new duplex channel if this attempt
    was succesful, otherwise with Nothing
*/

class accSChannel ch    ::  (TCP_SChannel -> (x, TCP_SChannel)) *(*ch .a)
                        ->  (x, *(*ch .a))
/*  This overloaded function supports the openSendNotifier function. It applies an
    access function on the underlying TCP_SChannel
*/

instance accSChannel TCP_SChannel_
instance accSChannel TCP_SCharStream_
```

A.9 StdFileSelect

```
definition module StdFileSelect


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdFileSelect defines the standard file selector dialogue.
//  ********************************************************************************


import StdMaybe, StdString

class FileSelectEnv env where
    selectInputFile ::                  !*env -> (!Maybe String,!*env)
    selectOutputFile:: !String !String !*env -> (!Maybe String,!*env)
/*  selectInputFile
        opens a dialogue in which the user can browse the file system to select an
        existing file.
        If a file has been selected, the String result contains the complete
        pathname of the selected file.
        If the user has not selected a file, Nothing is returned.
    selectOutputFile
        opens a dialogue in which the user can browse the file system to save a
        file.
        The first argument is the prompt of the dialogue (default: "Save As:")
        The second argument is the suggested filename.
        If the indicated directory already contains a file with the indicated name,
        selectOutputFile opens a new dialogue to confirm overwriting of the existing
        file.
        If either this dialogue is not confirmed or browsing is cancelled then
        Nothing is returned, otherwise the String result is the complete pathname of
        the selected file.
*/

instance FileSelectEnv World
```

A.10 StdId

```
definition module StdId


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdId specifies the generation functions for identification values.
//  ********************************************************************************


from    id      import  Id, RId, R2Id, RIdtoId, R2IdtoId, toString, ==
from    iostate import  IOSt

class Ids env where
    openId      ::      !*env -> (!Id,          !*env)
    openIds     :: !Int !*env -> (![Id],        !*env)

    openRId     ::      !*env -> (!RId  m,      !*env)
    openRIds    :: !Int !*env -> (![RId m],     !*env)

    openR2Id    ::      !*env -> (!R2Id  m r,   !*env)
    openR2Ids   :: !Int !*env -> (![R2Id m r],  !*env)
/*  There are three types of identification values:
    -   RId  m:     for uni-directional message passing (see StdReceiver)
    -   R2Id m r:   for bi-directional  message passing (see StdReceiver)
    -   Id:         for all other Object I/O library components
    Of each generation function there are two variants:
    -   to create exactly one identification value.
    -   to create a number of identification values.
            If the integer argument <=0, then an empty list of identification values
            is generated.
*/

instance Ids World
instance Ids (IOSt .l .p)
```

A.11 StdIO

```
definition module StdIO

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdIO contains all definition modules of the Object I/O library.
//  ********************************************************************************

import
    StdId,                  // The operations that generate identification values
    StdIOCommon,            // Function and type definitions used in the library
    StdMaybe,               // The Maybe data type
    StdPSt,                 // Operations on PSt that are not device related
    StdSystem,              // System dependent operations

    StdFileSelect,          // File selector dialogues

    StdPictureDef,          // Type definitions for picture handling
    StdPicture,             // Picture handling operations
    StdBitmap,              // Defines an instance for drawing bitmaps

    StdProcessDef,          // Type definitions for process handling
    StdProcess,             // Process handling operations

    StdClipboard,           // Clipboard handling operations

    StdControlDef,          // Type definitions for controls
    StdControlClass,        // Standard controls class instances
    StdControlReceiver,     // Receiver controls class instances
    StdControl,             // Control handling operations

    StdMenuDef,             // Type definitions for menus
    StdMenuElementClass,    // Standard menus class instances
    StdMenuReceiver,        // Receiver menus class instances
    StdMenuElement,         // Menu element handling operations
    StdMenu,                // Menu handling operations

    StdReceiverDef,         // Type definitions for receivers
    StdReceiver,            // Receiver handling operations

    StdTimerDef,            // Type definitions for timers
    StdTimerElementClass,   // Standard timer class instances
    StdTimerReceiver,       // Receiver timer class instances
    StdTimer,               // Timer handling operations
    StdTime,                // Time related operations

    StdWindowDef,           // Type definitions for windows
    StdWindow,              // Window handling operations

    StdPrint,               // for printing
    StdPrintText            // for printing text
```

A.12 StdIOBasic

```
definition module StdIOBasic


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdIOBasic defines basic types and access functions for the I/O library.
//  ********************************************************************************


import  StdOverloaded, StdString


/*  General type constructors for composing context-independent data structures.
*/
::  :^:     t1 t2           = (:^:) infixr 9 t1 t2


/*  General type constructors for composing context-dependent data structures.
*/
::  :~:     t1 t2       cs  = (:~:) infixr 9 (t1 cs) (t2 cs)
::  ListCS      t       cs  = ListCS [t cs]
::  NilCS               cs  = NilCS


/*  General type constructors for composing local and context-dependent
    data structures.
*/
::  :+:     t1 t2   ls  cs  = (:+:) infixr 9 (t1 ls cs) (t2 ls cs)
::  ListLS      t   ls  cs  = ListLS [t ls cs]
::  NilLS           ls  cs  = NilLS
::  NewLS       t   ls  cs  = E..new: {newLS::new, newDef:: t   new     cs}
::  AddLS       t   ls  cs  = E..add: {addLS::add, addDef:: t *(add,ls) cs}

noLS ::     (.a->.b)    (.c,.a) -> (.c,.b)  // Lift function    a  ->  b
                                            // to             (c,a)->(c,b)
noLS1:: (.x->.a->.b) .x (.c,.a) -> (.c,.b)  // Lift function x->  a  ->  b
                                            // to            x->(c,a)->(c,b)

::  Void            =   Void
::  Index           :== Int
::  Title           :== String


::  Vector          =   {vx::!Int,vy::!Int}

instance            ==          Vector      // @1-@2==zero
instance            +           Vector      // {vx=@1.vx+@2.vx,vy=@1.vy+@2.vy}
instance            -           Vector      // {vx=@1.vx-@2.vx,vy=@1.vy-@2.vy}
instance            zero        Vector      // {vx=0,vy=0}
instance            ~           Vector      // zero-@1
instance            toString    Vector

class toVector      x :: !x ->  Vector


::  Size            =   {w ::!Int,h ::!Int}

instance            ==          Size        // @1.w==@2.w && @1.h==@2.h
instance            zero        Size        // {w=0,h=0}
instance            toVector    Size        // {w,h}->{vx=w,vy=h}
instance            toString    Size


::  Point
```

```
    =   {   x       :: !Int
        ,   y       :: !Int
        }
::  Rectangle
    =   {   corner1 :: !Point
        ,   corner2 :: !Point
        }

instance            ==          Point       // @1-@2==zero
instance            +           Point       // {x=@1.x+@2.x,y=@1.y+@2.y}
instance            -           Point       // {x=@1.x-@2.x,y=@1.y-@2.y}
instance            zero        Point       // {x=0,y=0}
instance            toVector    Point       // {x,y}->{vx=x,vy=y}
instance            toString    Point

instance            ==          Rectangle   //    @1.corner1==@2.corner1
                                            // && @1.corner2==@2.corner2
instance            zero        Rectangle   // {corner1=zero,corner2=zero}
instance            toString    Rectangle

rectangleSize       :: !Rectangle -> Size   // {w=abs (@1.corner1-@1.corner2).x,
                                            //  h=abs (@1.corner1-@1.corner2).y}
movePoint   :: !Vector !Point -> Point      // {vx,vy} {x,y} -> {vx+x,vy+y}


::  IdFun ps        :== ps -> ps
```

A.13 StdIOCommon

```
definition module StdIOCommon


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdIOCommon defines common types and access functions for the
//  Object I/O library.
//  ********************************************************************************


import  StdOverloaded
import  StdString
import  StdIOBasic
from    id      import  Id, RId, R2Id, RIdtoId, R2IdtoId, toString, ==
from    key     import  SpecialKey,
                        BeginKey,
                        ClearKey,
                        DeleteKey, DownKey,
                        EndKey, EnterKey, EscapeKey,
                        F1Key,  F2Key,  F3Key,  F4Key,  F5Key,
                        F6Key, F7Key, F8Key, F9Key, F10Key,
                        F11Key, F12Key, F13Key, F14Key, F15Key,
                        HelpKey,
                        LeftKey,
                        PgDownKey, PgUpKey,
                        RightKey,
                        UpKey

/*  General type constructors for composing local and context-dependent
    data structures.
*/
::  SelectState     =   Able | Unable
::  MarkState       =   Mark | NoMark

enabled     :: !SelectState -> Bool         // @1 == Able
marked      :: !MarkState   -> Bool         // @1 == Mark

instance    == SelectState                  // Constructor equality
instance    == MarkState                    // Constructor equality
instance    ~  SelectState                  // Able <-> Unable
instance    ~  MarkState                    // Mark <-> NoMark

::  KeyboardState
    =   CharKey     Char        KeyState    // ASCII character input
    |   SpecialKey  SpecialKey  KeyState Modifiers
                                            // Special key input
::  KeyState
    =   KeyDown     IsRepeatKey             // Key is down
    |   KeyUp                               // Key goes up
::  IsRepeatKey                             // Flag on key down:
    :== Bool                                // True iff key is repeating
::  Key
    =   IsCharKey    Char
    |   IsSpecialKey SpecialKey
::  KeyboardStateFilter                     // Predicate on KeyboardState:
    :== KeyboardState -> Bool               // evaluate KeyFunction only if True

getKeyboardStateKeyState:: !KeyboardState -> KeyState
getKeyboardStateKey     :: !KeyboardState -> Key

instance    == KeyState                     // Equality on KeyState


::  MouseState
```

```
    =   MouseMove   Point Modifiers         // Mouse is up     (position,modifiers)
    |   MouseDown   Point Modifiers Int     // Mouse goes down (and nr down)
    |   MouseDrag   Point Modifiers         // Mouse is down   (position,modifiers)
    |   MouseUp     Point Modifiers         // Mouse goes up   (position,modifiers)
::  ButtonState
    =   ButtonStillUp                       // MouseMove
    |   ButtonDown                          // MouseDown _ _ 1
    |   ButtonDoubleDown                    //           _ _ 2
    |   ButtonTripleDown                    //           _ _ >2
    |   ButtonStillDown                     // MouseDrag
    |   ButtonUp                            // MouseUp
::  MouseStateFilter                        // Predicate on MouseState:
    :== MouseState -> Bool                  // evaluate MouseFunction only if True

getMouseStatePos        :: !MouseState  -> Point
getMouseStateModifiers  :: !MouseState  -> Modifiers
getMouseStateButtonState:: !MouseState  -> ButtonState

instance    == ButtonState                  // Constructor equality


::  SliderState
    =   {   sliderMin   :: !Int
        ,   sliderMax   :: !Int
        ,   sliderThumb :: !Int
        }

instance == SliderState                     // @1.sliderMin   == @2.sliderMin
                                            // @1.sliderMax   == @2.sliderMax
                                            // @1.sliderThumb == @2.sliderThumb


::  UpdateState
    =   {   oldFrame    :: !ViewFrame
        ,   newFrame    :: !ViewFrame
        ,   updArea     :: !UpdateArea
        }
::  ViewDomain          :== Rectangle
::  ViewFrame           :== Rectangle
::  UpdateArea          :== [ViewFrame]



::  Modifiers
    =   {   shiftDown   :: !Bool            // True iff shift   down
        ,   optionDown  :: !Bool            // True iff option  down
        ,   commandDown :: !Bool            // True iff command down
        ,   controlDown :: !Bool            // True iff control down
        ,   altDown     :: !Bool            // True iff alt     down
        }

//  Constants to check which of the Modifiers are down.

NoModifiers :== {shiftDown  = False
                ,optionDown = False
                ,commandDown= False
                ,controlDown= False
                ,altDown    = False
                }
ShiftOnly   :== {NoModifiers & shiftDown    = True}
OptionOnly  :== {NoModifiers & optionDown   = True}
CommandOnly :== {NoModifiers & commandDown  = True}
ControlOnly :== {NoModifiers & controlDown  = True}
AltOnly     :== {NoModifiers & altDown      = True}


/*  The layout language used for windows and controls.
```

```
*/
::  ItemPos
    :== (   ItemLoc
        ,   ItemOffset
        )
::  ItemLoc
    //  Absolute:
    =   Fix     Point
    //  Relative to corner:
    |   LeftTop
    |   RightTop
    |   LeftBottom
    |   RightBottom
    //  Relative in next line:
    |   Left
    |   Center
    |   Right
    //  Relative to other item:
    |   LeftOf  Id
    |   RightTo Id
    |   Above   Id
    |   Below   Id
    //  Relative to previous item:
    |   LeftOfPrev
    |   RightToPrev
    |   AbovePrev
    |   BelowPrev
::  ItemOffset
    :== Vector

instance    == ItemLoc                      // Constructor and value equality


/*  The Direction type.
*/
::  Direction
    =   Horizontal
    |   Vertical

instance    == Direction                    // Constructor equality


/*  Document interface type of interactive processes.
*/
::  DocumentInterface
    =   NDI                                 // No       Document Interface
    |   SDI                                 // Single   Document Interface
    |   MDI                                 // Multiple Document Interface

instance    == DocumentInterface            // Constructor equality


/*  Process attributes.
*/
::  ProcessAttribute ps                     // Default:
    =   ProcessWindowPos    ItemPos         // Platform dependent
    |   ProcessWindowSize   Size            // Platform dependent
    |   ProcessWindowResize (ProcessWindowResizeFunction ps)
                                            // Platform dependent
    |   ProcessHelp         (IOFunction ps) // No Help  facility
    |   ProcessAbout        (IOFunction ps) // No About facility
    |   ProcessActivate     (IOFunction ps) // No action on activate
    |   ProcessDeactivate   (IOFunction ps) // No action on deactivate
    |   ProcessClose        (IOFunction ps) // Process is closed
    |   ProcessShareGUI                     // Process does not share parent GUI
//  Attributes for MDI processes only:
    |   ProcessNoWindowMenu                 // Process has WindowMenu
```

```
::  ProcessWindowResizeFunction ps
    :== Size                                // Old ProcessWindow size
     -> Size                                // New ProcessWindow size
     -> ps -> ps


/*  Frequently used function types.
*/
::  IOFunction          ps  :==                     ps -> ps
::  ModsIOFunction      ps  :== Modifiers       ->  ps -> ps
::  MouseFunction       ps  :== MouseState      ->  ps -> ps
::  KeyboardFunction    ps  :== KeyboardState   ->  ps -> ps


/*  Common error report types.
*/
::  ErrorReport                             // Usual cause:
    =   NoError                             // Everything went allright
    |   ErrorViolateDI                      // Violation against DocumentInterface
    |   ErrorIdsInUse                       // Object contains Ids that are bound
    |   ErrorUnknownObject                  // Object can not be found
    |   ErrorNotifierOpen                   // Second send notifier opened

instance    ==          ErrorReport         // Constructor equality
instance    toString    ErrorReport         // Constructor as String

::  OkBool                                  // True iff successful operation
    :== Bool
```

A.14 StdMaybe

```
definition module StdMaybe

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdMaybe defines the Maybe type.
//  ********************************************************************************

::  Maybe x
    =   Just x
    |   Nothing

isJust      :: !(Maybe .x) -> Bool      // case @1 of (Just _) -> True; _ -> False
isNothing   :: !(Maybe .x) -> Bool      // not o isJust
fromJust    :: !(Maybe .x) -> .x        // \(Just x) -> x

// for possibly unique elements

u_isJust    :: !(Maybe .x) -> (!Bool, !Maybe .x)
u_isNothing :: !(Maybe .x) -> (!Bool, !Maybe .x)
```

A.15 StdMenu

```
definition module StdMenu


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdMenu defines functions on menus.
//  ********************************************************************************


import  StdMenuDef, StdMenuElementClass
from    iostate     import PSt, IOSt


//  Operations on unknown Ids are ignored.

class Menus mdef where
    openMenu    :: .ls !(mdef .ls (PSt .l .p)) !(PSt .l .p)
                               -> (!ErrorReport,!PSt .l .p)
    getMenuType ::      (mdef .ls .ps)  -> MenuType
/*  Open the given menu definition for this interactive process.
    openMenu may not be permitted to open a menu depending on its DocumentInterface
        (see the comments at the shareProcesses instances in module StdProcess).
    In case a menu with the same Id is already open then nothing happens. In case
        the menu has the WindowMenuId Id then nothing happens. In case the menu does
        not have an Id, it will obtain an Id which is fresh with respect to the
        current set of menus. The Id can be reused after closing this menu. In case
        menu elements are opened with duplicate Ids, the menu will not be opened.
    In case the menu definition does not have a MenuIndex attribute (see StdMenuDef)
        it will be opened behind the last menu. In case the menu definition has a
        MenuIndex attribute it will be placed behind the menu indicated by the
        integer index.
        The index of a menu starts from one for the first present menu. If the index
        is negative or zero, then the new menu is added before the first menu. If
        the index exceeds the number of menus, then the new menu is added behind the
        last menu.
*/

instance Menus  (Menu m)    | MenuElements m


closeMenu :: !Id !(IOSt .l .p) -> IOSt .l .p
/*  closeMenu closes the indicated Menu and all of its elements.
    The WindowMenu can not be closed by closeMenu (in case the Id argument equals
    WindowMenuId).
*/


openMenuElements    :: !Id     !Index .ls (m .ls (PSt .l .p))      !(IOSt .l .p)
                                                -> (!ErrorReport,!IOSt .l .p)
                                                |  MenuElements m
openSubMenuElements :: !Id !Id !Index .ls (m .ls (PSt .l .p))      !(IOSt .l .p)
                                                -> (!ErrorReport,!IOSt .l .p)
                                                |  MenuElements m
openRadioMenuItems  :: !Id !Id !Index ![MenuRadioItem (PSt .l .p)] !(IOSt .l .p)
                                                -> (!ErrorReport,!IOSt .l .p)
/*  Add menu elements to the indicated Menu, SubMenu, or RadioMenu.
    openMenuElements:
        adds menu elements to the Menu identified by the Id argument.
    openSubMenuElements:
        adds menu elements to the SubMenu identified by the second Id argument,
        which must be contained in the Menu identified by the first Id argument.
    openRadioMenuItems:
        adds menu radio items to the RadioMenu identified by the second Id argument,
        which must be contained in the Menu identified by the first Id argument.
```

```
        If the RadioMenu was empty, then the first item in the list will be checked.
    Menu elements are added after the item with the specified index. The index of a
        menu element starts from one for the first menu element in the indicated
        menu.
        If the index is negative or zero, then the new menu elements are added
        before the first menu element of the indicated menu.
        If the index exceeds the number of menu elements in the indicated menu, then
        the new menu elements are added behind the last menu element of the
        indicated menu.
    No menu elements are added if the indicated menu does not exist.
    open(Sub)MenuElements have no effect in case menu elements with duplicate Ids
    are opened.
*/


closeMenuElements :: !Id ![Id] !(IOSt .l .p) -> IOSt .l .p
/*  closeMenuElements
        closes menu elements of the Menu identified by the first Id argument by
        their Ids. The elements of (Sub/Radio)Menus will be removed first.
*/


closeMenuIndexElements      :: !Id     ![Index] !(IOSt .l .p) -> IOSt .l .p
closeSubMenuIndexElements   :: !Id !Id ![Index] !(IOSt .l .p) -> IOSt .l .p
closeRadioMenuIndexElements :: !Id !Id ![Index] !(IOSt .l .p) -> IOSt .l .p
/*  Close menu elements of the indicated Menu, SubMenu, or RadioMenu by their Index
    position.
    closeMenuIndexElements:
        closes menu elements of the Menu identified by the Id argument.
    closeSubMenuIndexElements:
        closes menu elements of the SubMenu identified by the second Id argument,
        which must be contained in the Menu identified by the first Id argument.
    closeRadioMenuIndexElements:
        closes menu items of the RadioMenu identified by the second Id argument,
        which must be contained in the Menu identified by the first Id argument.
    Analogous to openMenuElements and openRadioMenuItems indices range from one to
        the number of menu elements in a menu. Invalid indices (less than one or
        larger than the number of menu elements of the menu) are ignored.
    If the currently checked element of a RadioMenu is closed, the first remaining
        element of that RadioMenu will be checked.
    Closing a (Sub/Radio)Menu closes the indicated (Sub/Radio)Menu and all of its
    elements.
*/


enableMenuSystem :: !(IOSt .l .p) -> IOSt .l .p
disableMenuSystem:: !(IOSt .l .p) -> IOSt .l .p
/*  Enable/disable the menu system of this interactive process. When the menu system
    is re-enabled the previously selectable menus and elements will become
    selectable again.
    Enable/disable operations on the menu(element)s of a disabled menu system take
    effect when the menu system is re-enabled.
    enableMenuSystem has no effect in case the interactive process has a (number of)
    modal dialogue(s).
*/


enableMenus :: ![Id] !(IOSt .l .p) -> IOSt .l .p
disableMenus:: ![Id] !(IOSt .l .p) -> IOSt .l .p
/*  Enable/disable individual menus.
    The WindowMenu can not be enabled/disabled.
    Disabling a menu overrules the SelectStates of its elements, which become
    unselectable.
    Enabling a disabled menu re-establishes the SelectStates of its elements.
    Enable/disable operations on the elements of a disabled menu take effect when
    the menu is re-enabled.
*/
```

```
getMenuSelectState :: !Id !(IOSt .l .p) -> (!Maybe SelectState,!IOSt .l .p)
/*  getMenuSelectState yields the current SelectState of the indicated menu. In case
    the menu does not exist, Nothing is returned.
*/


getMenus :: !(IOSt .l .p) -> (![(Id,MenuType)],!IOSt .l .p)
/*  getMenus yields the Ids and MenuTypes of the current set of menus of this
    interactive process.
*/


getMenuPos :: !Id !(IOSt .l .p) -> (!Maybe Index,!IOSt .l .p)
/*  getMenuPos yields the index position of the indicated menu in the current list
    of menus.
    In case the menu does not exist, Nothing is returned.
*/


setMenuTitle :: !Id !Title !(IOSt .l .p) -> IOSt .l .p
getMenuTitle :: !Id        !(IOSt .l .p) -> (!Maybe Title,!IOSt .l .p)
/*  setMenuTitle sets the title of the indicated menu.
        In case the menu does not exist or refers to the WindowMenu, nothing
        happens.
    getMenuTitle retrieves the current title of the indicated menu.
        In case the menu does not exist, Nothing is returned.
*/
```

A.16 StdMenuDef

```
definition module StdMenuDef


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdMenu contains the types to define the standard set of menus and their
//  elements.
//  ********************************************************************************


import  StdIOCommon, StdMaybe


/*  Menus:              */
::  Menu        m ls ps = Menu          Title         (m ls ps)
                                        [MenuAttribute *(ls,ps)]

/*  Menu elements:      */
::  SubMenu     m ls ps = SubMenu       Title         (m ls ps)
                                        [MenuAttribute *(ls,ps)]
::  RadioMenu     ls ps = RadioMenu     [MenuRadioItem *(ls,ps)] Index
                                        [MenuAttribute *(ls,ps)]
::  MenuItem      ls ps = MenuItem      Title
                                        [MenuAttribute *(ls,ps)]
::  MenuSeparator ls ps = MenuSeparator [MenuAttribute *(ls,ps)]

::  MenuRadioItem ps  :== (Title,Maybe Id,Maybe Char,IOFunction ps)

::  MenuAttribute       ps                      // Default:
//  Attributes for Menus and MenuElements:
    =   MenuId              Id                  // no Id
    |   MenuSelectState     SelectState         // menu(item) Able
//  Attributes only for Menus:
    |   MenuIndex           Int                 // end of current menu list
//  Attributes ignored by (Sub)Menus:
    |   MenuShortKey        Char                // no ShortKey
    |   MenuMarkState       MarkState           // NoMark
    |   MenuFunction        (IOFunction     ps) // \x->x
    |   MenuModsFunction    (ModsIOFunction ps) // MenuFunction

::  MenuType        :== String
::  MenuElementType :== String
```

A.17 StdMenuElement

```
definition module StdMenuElement


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdMenuElement specifies all functions on menu elements.
//  Changing the status of menu elements requires a *MState.
//  Reading  the status of menu elements requires a  MState.
//  ********************************************************************************


import  StdMenuDef
from    iostate import IOSt


::  MState


getMenu :: !Id !(IOSt .l .p) -> (!Maybe MState, !IOSt .l .p)
/*  getMenu returns a read-only MState for the indicated menu.
    In case the indicated menu does not exist Nothing is returned.
*/


setMenu :: !Id ![IdFun *MState] !(IOSt .l .p) -> IOSt .l .p
/*  Apply the menu element changing functions to the current state of the indicated
    menu.
    In case the indicated menu does not exist nothing happens.
*/


/*  Functions that change the state of menu elements.
    When applied to unknown Ids none of these functions have effect.
*/

enableMenuElements  :: ![Id] !*MState -> *MState
disableMenuElements :: ![Id] !*MState -> *MState
/*  (en/dis)ableMenuElements set the SelectState of the indicated menu elements.
    Disabling a (Sub/Radio)Menu overrules the SelectStates of its elements, which
        become unselectable.
    Enabling a disabled (Sub/Radio)Menu re-establishes the SelectStates of its
        elements.
    (En/Dis)able operations on the elements of a disabled (Sub/Radio)Menu take
        effect when the (Sub/Radio)Menu is re-enabled.
*/

setMenuElementTitles :: ![(Id,Title)] !*MState -> *MState
/*  setMenuElementTitles sets the titles of the indicated menu elements.
*/

markMenuItems   :: ![Id] !*MState -> *MState
unmarkMenuItems :: ![Id] !*MState -> *MState
/*  (un)markMenuItems sets the MarkState of the indicated MenuItems.
*/

selectRadioMenuItem      :: !Id !Id    !*MState -> *MState
selectRadioMenuIndexItem :: !Id !Index !*MState -> *MState
/*  selectRadioMenu(Index)Item
        selects the indicated MenuRadioItem of a RadioMenu, causing the mark of the
        previously marked MenuRadioItem to disappear.
    selectRadioMenuItem
        indicates the MenuRadioItem by the Id of its parent RadioMenu and its Id.
    selectRadioMenuIndexItem
        indicates the MenuRadioItem by the Id of its parent RadioMenu and its index
```

```
        position (counted from 1).
*/


/*  Functions that read the state of menu elements.
*/

getMenuElementTypes         ::     !MState -> [(MenuElementType,Maybe Id)]
getCompoundMenuElementTypes :: !Id !MState -> [(MenuElementType,Maybe Id)]
/*  getMenuElementTypes
        yields the list of MenuElementTypes of all menu elements of this menu.
    getCompoundMenuElementTypes
        yields the list of MenuElementTypes of all menu elements of this
        (Sub/Radio)Menu.
    Both functions return (Just id) if the element has a MenuId attribute, and
    Nothing otherwise.
    Ids are not collected recursively through (Sub/Radio)Menus.
*/

getSelectedRadioMenuItem    :: !Id   !MState -> (!Index,!Maybe Id)
/*  getSelectedRadioMenuItem
        returns the Index and Id, if any, of the currently selected MenuRadioItem of
        the indicated RadioMenu.
    If the RadioMenu does not exist or is empyt, the Index is zero and the Id is
    Nothing.
*/

getMenuElementSelectStates  :: ![Id] !MState -> [(Bool,SelectState)]    // Able
getMenuElementMarkStates    :: ![Id] !MState -> [(Bool,MarkState)]      // NoMark
getMenuElementTitles        :: ![Id] !MState -> [(Bool,Maybe String)]   // Nothing
getMenuElementShortKey      :: ![Id] !MState -> [(Bool,Maybe Char)]     // Nothing
/*  The  result list is of equal length as the argument Id list.
    Each result list element corresponds in order with the argument Id list.
    Of each element the first Boolean result is False in case of invalid id
        (if so dummy values are returned - see comment).
    -   getMenuElementSelectStates
            yield the SelectState of the indicated elements.
    -   getMenuElementMarkStates
            yield the MarkState of the indicated elements.
    -   getMenuElementTitles
            yields (Just title) of the indicated (SubMenu/MenuItem),
            Nothing otherwise.
    -   getMenuElementShortKey
            yields (Just key) of the indicated MenuItem, Nothing otherwise.
*/
```

A.18 StdMenuElementClass

```
definition module StdMenuElementClass


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdMenuElementClass defines the standard set of menu element instances.
//  ********************************************************************************


import  StdMenuDef
from    menuhandle  import MenuElementState


class MenuElements m
where
    menuElementToHandles    :: !(m .ls .ps) -> [MenuElementState .ls .ps]
    getMenuElementType      ::  (m .ls .ps) -> MenuElementType

instance MenuElements (AddLS    m)  | MenuElements m    // getMenuElementType==""
instance MenuElements (NewLS    m)  | MenuElements m    // getMenuElementType==""
instance MenuElements (ListLS   m)  | MenuElements m    // getMenuElementType==""
instance MenuElements NilLS                             // getMenuElementType==""
instance MenuElements ((:+:) m1 m2) | MenuElements m1
                                    & MenuElements m2   // getMenuElementType==""
instance MenuElements (SubMenu  m)  | MenuElements m
instance MenuElements RadioMenu
instance MenuElements MenuItem
instance MenuElements MenuSeparator
```

A.19 StdMenuReceiver

```
definition module StdMenuReceiver


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdMenuReceiver defines Receiver(2) menu element instances.
//  ********************************************************************************


import  StdReceiverDef, StdMenuElementClass


//  Receiver components:
instance MenuElements (Receiver  m  )
instance MenuElements (Receiver2 m r)
```

A.20 StdPicture

```
definition module StdPicture


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdPicture contains the drawing operations and access to Pictures.
//  ********************************************************************************


from    StdFunc     import  St
from    osfont      import  Font
from    ospicture   import  Picture
import  StdPictureDef


::  DrawFunction
    :== *Picture -> *Picture

//  The picture attributes:
::  PictureAttribute                        // Default:
    =   PicturePenSize      Int             // 1
    |   PicturePenPos       Point           // zero
    |   PicturePenColour    Colour          // Black
    |   PicturePenFont      Font            // DefaultFont


//  Pen position attributes:
setPenPos               :: !Point               !*Picture -> *Picture
getPenPos               ::                      !*Picture -> (!Point,!*Picture)

class movePenPos figure :: !figure              !*Picture -> *Picture
//  Move the pen position as much as when drawing the figure.
instance movePenPos Vector
instance movePenPos Curve

//  PenSize attributes:
setPenSize              :: !Int                 !*Picture -> *Picture
getPenSize              ::                      !*Picture -> (!Int,!*Picture)

setDefaultPenSize       ::                      !*Picture -> *Picture
//  setDefaultPenSize = setPenSize 1

//  Colour attributes:
setPenColour            :: !Colour              !*Picture -> *Picture
getPenColour            ::                      !*Picture -> (!Colour,!*Picture)

setDefaultPenColour     ::                      !*Picture -> *Picture
//  setDefaultPenColour = setPenColour BlackColour

//  Font attributes:
setPenFont              :: !Font                !*Picture -> *Picture
getPenFont              ::                      !*Picture -> (!Font,!*Picture)

setDefaultPenFont       ::                      !*Picture -> *Picture


/*  Font operations:
*/
openFont        :: !FontDef !*Picture -> (!(!Bool,!Font),!*Picture)
openDefaultFont ::          !*Picture -> (!Font,         !*Picture)
openDialogFont  ::          !*Picture -> (!Font,         !*Picture)
/*  openFont
        creates the font as specified by the name, stylistic variations, and size.
        The Boolean result is True only if the font is available and need not be
```

```
        scaled.
        In all other cases, an existing font is returned (depending on the system).
    openDefaultFont
        returns the font used by default by applications.
    openDialogFont
        returns the font used by default by the system.
*/

getFontNames    ::                      !*Picture -> (![FontName],   !*Picture)
getFontStyles   ::           !FontName  !*Picture -> (![FontStyle],  !*Picture)
getFontSizes    :: !Int !Int !FontName  !*Picture -> (![FontSize],   !*Picture)
/*  getFontNames
        returns the FontNames  of all available fonts.
    getFontStyles
        returns the FontStyles of all available styles of a particular FontName.
    getFontSizes
        returns all FontSizes in increasing order of a particular FontName that are
        available without scaling. The sizes inspected are inclusive between the two
        Integer arguments. (Negative values are set to zero.)
        In case the requested font is unavailable, the styles or sizes of the
        default font are returned.
*/

getFontDef :: !Font -> FontDef
/*  getFontDef returns the name, stylistic variations and size of the argument Font.
*/

getPenFontCharWidth     ::       ! Char     !*Picture -> (! Int,         !*Picture)
getPenFontCharWidths    ::       ![Char]    !*Picture -> (![Int],        !*Picture)
getPenFontStringWidth   ::       ! String   !*Picture -> (! Int,         !*Picture)
getPenFontStringWidths  ::       ![String]  !*Picture -> (![Int],        !*Picture)
getPenFontMetrics       ::                  !*Picture -> (!FontMetrics, !*Picture)

getFontCharWidth        :: !Font ! Char     !*Picture -> (!Int,          !*Picture)
getFontCharWidths       :: !Font ![Char]    !*Picture -> (![Int],        !*Picture)
getFontStringWidth      :: !Font ! String   !*Picture -> (!Int,          !*Picture)
getFontStringWidths     :: !Font ![String]  !*Picture -> (![Int],        !*Picture)
getFontMetrics          :: !Font            !*Picture -> (!FontMetrics, !*Picture)
/*  get(Pen)Font(Char/String)Width(s)
        return the width of the argument (Char/String)(s) given the Font argument
        or current PenFont attribute.
    get(Pen)FontMetrics
        returns the FontMetrics of the Font argument or current PenFont attribute.
*/


/*  Region functions.
    A Region is defined by a collection of shapes.
*/
::  Region

//  Basic access functions on Regions:

isEmptyRegion   :: !Region -> Bool
getRegionBound  :: !Region -> Rectangle
/*  isEmptyRegion
        holds if the argument region covers no pixels (it is empty).
    getRegionBound
        returns the smallest enclosing rectangle of the argument region.
        If the region is empty, zero is returned.
*/

//  Constructing a region:

class toRegion area :: !area -> Region

::  PolygonAt
```

```
    =   {   polygon_pos :: !Point
        ,   polygon     :: !Polygon
        }

instance toRegion Rectangle
instance toRegion PolygonAt
instance toRegion [area]            | toRegion area
instance toRegion (:^: area1 area2) | toRegion area1 & toRegion area2


//  Drawing and restoring picture attributes:

appPicture      :: !(IdFun *Picture) !*Picture -> *Picture
accPicture      :: !(St *Picture .x) !*Picture -> (.x,!*Picture)
/*  (app/acc)Picture f pict
        apply f to pict. After drawing,  the picture attributes of the result
        picture are restored to those of pict.
*/


//  Drawing within in a clipping region:

appClipPicture :: !Region !(IdFun *Picture) !*Picture -> *Picture
accClipPicture :: !Region !(St *Picture .x) !*Picture -> (.x,!*Picture)


//  Drawing in 'exclusive or' mode:

appXorPicture   :: !(IdFun *Picture) !*Picture -> *Picture
accXorPicture   :: !(St *Picture .x) !*Picture -> (.x,!*Picture)
/*  (app/acc)XorPicture f pict
        apply f to pict in the appropriate platform xor mode.
*/


//  Drawing in 'hilite' mode:

class Hilites figure where
    hilite  ::        !figure !*Picture -> *Picture
    hiliteAt:: !Point !figure !*Picture -> *Picture
/*  hilite
        draws figures in the appropriate 'hilite' mode at the current  pen position.
    hiliteAt
        draws figures in the appropriate 'hilite' mode at the argument pen position.
    Both functions reset the 'hilite' mode after drawing.
*/

instance Hilites Box        // Hilite a box
instance Hilites Rectangle  // Hilite a rectangle (note: hiliteAt pos r = hilite r)


//  Drawing points:

drawPoint       ::               !*Picture -> *Picture
drawPointAt     :: !Point        !*Picture -> *Picture
/*  drawPoint
        plots a point at the current  pen position p and moves to p+{vx=1,vy=0}
    drawPointAt
        plots a point at the argument pen position, but retains the pen position.
*/

//  Drawing lines:

drawLineTo      :: !Point        !*Picture -> *Picture
drawLine        :: !Point !Point !*Picture -> *Picture
/*  drawLineTo
        draws a line from the current pen position to the argument point which
```

```
        becomes the new pen position.
    drawLine
        draws a line between the two argument points, but retains the pen position.
*/


/*  Drawing and filling operations.
    These functions are divided into the following classes:
    Drawables:
        draw     'line-oriented' figures at the current  pen position.
        drawAt   'line-oriented' figures at the argument pen position.
        undraw     f = appPicture (draw     f o setPenColour background)
        undrawAt x f = appPicture (drawAt x f o setPenColour background)
    Fillables:
        fill     'area-oriented' figures at the current  pen position.
        fillAt   'area-oriented' figures at the argument pen position.
        unfill     f = appPicture (fill     f o setPenColour background)
        unfillAt x f = appPicture (fillAt x f o setPenColour background)
*/
class Drawables figure where
    draw    ::        !figure !*Picture -> *Picture
    drawAt  :: !Point !figure !*Picture -> *Picture
    undraw  ::        !figure !*Picture -> *Picture
    undrawAt:: !Point !figure !*Picture -> *Picture

class Fillables figure where
    fill    ::        !figure !*Picture -> *Picture
    fillAt  :: !Point !figure !*Picture -> *Picture
    unfill  ::        !figure !*Picture -> *Picture
    unfillAt:: !Point !figure !*Picture -> *Picture


//  Text drawing operations:
//  Text is always drawn with the baseline at the y coordinate of the pen.

instance Drawables  Char
instance Drawables  {#Char}
/*  draw     text:
        draws the text starting at the current pen position.
        The new pen position is directly after the drawn text including spacing.
    drawAt p text:
        draws the text starting at p.
*/


//  Vector drawing operations:
instance Drawables Vector
/*  draw     v:
        draws a line from the current pen position pen to pen+v.
    drawAt p v:
        draws a line from p to p+v.
*/


/*  Oval drawing operations:
    An Oval o is a transformed unit circle
    with    horizontal radius rx    o.oval_rx
            vertical   radius ry    o.oval_ry
    Let (x,y) be a point on the unit circle:
        then (x',y') = (x*rx,y*ry) is a point on o.
    Let (x,y) be a point on o:
        then (x',y') = (x/rx,y/ry) is a point on the unit circle.
*/
instance Drawables Oval
instance Fillables Oval
/*  draw     o:
        draws an oval with the current pen position being the center of the oval.
```

```
    drawAt p o:
        draws an oval with p being the center of the oval.
    fill     o:
        fills an oval with the current pen position being the center of the oval.
    fillAt p o:
        fills an oval with p being the center of the oval.
    None of these functions change the pen position.
*/


/*  Curve drawing operations:
    A Curve c is a slice of an oval o
    with    start angle a   c.curve_from
            end   angle b   c.curve_to
            direction   d   c.curve_clockwise
    The angles are taken in radians (counter-clockwise).
    If d holds then the drawing direction is clockwise, otherwise drawing occurs
    counter-clockwise.
*/
instance Drawables Curve
instance Fillables Curve
/*  draw     c:
        draws a curve with the starting angle a at the current pen position.
        The pen position ends at ending angle b.
    drawAt p c:
        draws a curve with the starting angle a at p.
    fill     c:
        fills the figure obtained by connecting the endpoints of the drawn curve
        (draw c) with the center of the curve oval.
        The pen position ends at ending angle b.
    fillAt p c:
        fills the figure obtained by connecting the endpoints of the drawn curve
        (drawAt p c) with the center of the curve oval.
*/


/*  Box drawing operations:
    A Box b is a horizontally oriented rectangle
    with    width  w        b.box_w
            height h        b.box_h
    In case w==0 (h==0),   the Box collapses to a vertical (horizontal) vector.
    In case w==0 and h==0, the Box collapses to a point.
*/
instance Drawables  Box
instance Fillables  Box
/*  draw     b:
        draws a box with left-top corner at the current pen position p and
        right-bottom corner at p+(w,h).
    drawAt p b:
        draws a box with left-top corner at p and right-bottom corner at p+(w,h).
    fill     b:
        fills a box with left-top corner at the current pen position p and
        right-bottom corner at p+(w,h).
    fillAt p b:
        fills a box with left-top corner at p and right-bottom corner at p+(w,h).
    None of these functions change the pen position.
*/


/*  Rectangle drawing operations:
    A Rectangle r is always horizontally oriented
    with    width  w        abs (r.corner1.x-r.corner2.x)
            height h        abs (r.corner1.y-r.corner2.y)
    In case w==0 (h==0),   the Rectangle collapses to a vertical (horizontal) vector.
    In case w==0 and h==0, the Rectangle collapses to a point.
*/
instance Drawables Rectangle
```

```
instance Fillables Rectangle
/*  draw     r:
        draws a rectangle with diagonal corners r.corner1 and r.corner2.
    drawAt p r:
        draw r
    fill     r:
        fills a rectangle with diagonal corners r.corner1 and r.corner2.
    fillAt p r:
        fill r
    None of these functions change the pen position.
*/


/*  Polygon drawing operations:
    A Polygon p is a figure
    with    shape   p.polygon_shape
    A polygon p at a point base is drawn as follows:
        drawPicture [setPenPos base:map draw shape]++[drawToPoint base]
*/
instance Drawables Polygon
instance Fillables Polygon
/*  None of these functions change the pen position.
*/


getResolution :: !*Picture -> (!(!Int,!Int),!*Picture)
/*  getResolution returns the horizontal and vertical resolution of a Picture in dpi
    (dots per inch).
    In case of a printer Picture:
        the return values are the printer resolution (if it is not emulating the
        screen resolution).
    In case of a screen Picture:
        the resolution is the number of pixels that fit in one "screen inch".
        A "screen inch" is the physical size of a 72 point font on the screen.
        Although some screens allow the user to alter the screen resolution,
        getResolution always returns the same value for the same screen.
        The reason is that if a user for instance increases the screen resolution,
        not only the size of a pixel decreases, but also the size of a "screen
        inch". So a 12 point font will not appear with 12 point size, but smaller
        (A point is a physical unit, defined as 1 point is approximately 1/72 inch.)
*/
```

A.21 StdPictureDef

```
definition module StdPictureDef


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdPictureDef contains the predefined figures that can be drawn.
//  ********************************************************************************


import  StdIOBasic


::  Box                                     // A box is a rectangle
    =   {   box_w           :: !Int         // The width  of the box
        ,   box_h           :: !Int         // The height of the box
        }
::  Oval                                    // An oval is a stretched unit circle
    =   {   oval_rx         :: !Int         // The horizontal radius (stretch)
        ,   oval_ry         :: !Int         // The vertical   radius (stretch)
        }
::  Curve                                   // A curve is a slice of an oval
    =   {   curve_oval      :: !Oval        // The source oval
        ,   curve_from      :: !Real        // Starting angle (in radians)
        ,   curve_to        :: !Real        // Ending   angle (in radians)
        ,   curve_clockwise :: !Bool        // Direction: True iff clockwise
        }
::  Polygon                                 // A polygon is an outline shape
    =   {   polygon_shape   :: ![Vector]    // The shape of the polygon
        }
::  FontDef
    =   {   fName           :: !FontName    // Name of the font
        ,   fStyles         :: ![FontStyle] // Stylistic variations
        ,   fSize           :: !FontSize    // Size in points
        }
::  FontMetrics
    =   {   fAscent         :: !Int         // Distance between top    and base line
        ,   fDescent        :: !Int         // Distance between bottom and base line
        ,   fLeading        :: !Int         // Distance between two text lines
        ,   fMaxWidth       :: !Int         // Max character width including spacing
        }
::  FontName    :== String
::  FontStyle   :== String
::  FontSize    :== Int
::  Colour
    =   RGB RGBColour
    |   Black       | White
    |   DarkGrey    | Grey      | LightGrey // 75%, 50%, and 25% Black
    |   Red         | Green     | Blue
    |   Cyan        | Magenta   | Yellow
::  RGBColour
    =   {   r   :: !Int                     // The contribution of red
        ,   g   :: !Int                     // The contribution of green
        ,   b   :: !Int                     // The contribution of blue
        }

//  Colour constants:
BlackRGB                :== {r=MinRGB,g=MinRGB,b=MinRGB}
WhiteRGB                :== {r=MaxRGB,g=MaxRGB,b=MaxRGB}
MinRGB                  :== 0
MaxRGB                  :== 255

//  Font constants:
SerifFontDef            :== {fName="Times",      fStyles=[],fSize=10}
SansSerifFontDef        :== {fName="Arial",      fStyles=[],fSize=10}
```

```
SmallFontDef            :== {fName="Small Fonts",fStyles=[],fSize=7 }
NonProportionalFontDef  :== {fName="Courier",    fStyles=[],fSize=10}
SymbolFontDef           :== {fName="Symbol",     fStyles=[],fSize=10}

//  Font style constants:
ItalicsStyle            :== "Italic"
BoldStyle               :== "Bold"
UnderlinedStyle         :== "Underline"

//  Standard lineheight of a font is the sum of its leading, ascent and descent:
fontLineHeight fMetrics :== fMetrics.fLeading + fMetrics.fAscent + fMetrics.fDescent

//  Useful when working with Ovals and Curves:
PI                      :== 3.1415926535898
```

A.22 StdPrint

```
definition module StdPrint

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdPrint specifies general printing functions.
//  Related functions and modules:
//  -   getResolution (StdPicture)
//  -   resizeBitmap  (StdBitmap)
//  -   StdPrintText to print text.
//  ********************************************************************************


import  StdPicture, StdIOCommon
from    StdOverloaded   import  ==
from    osprint         import  PrintSetup, JobInfo, PrintInfo, Alternative,
                                Cancelled, StartedPrinting, PrintEnvironments
from    iostate         import  IOSt, PSt

::  PageDimensions
    =   {   page        :: !Size        // size of the drawable area of the page
        ,   margins     :: !Rectangle   // This field contains information about the
                                        // size of the margins on a sheet in pixels.
                                        // Drawing can't occur within these margins.
                                        // The margin Rectangle is bigger than the
                                        // page size. Its values are:
                                        // corner1.x<=0 && corner1.y<=0 &&
                                        // corner2.x>=page.w && corner2.y>=page.h
        ,   resolution  :: !(!Int,!Int) // horizontal and vertical printer
                                        // resolution in dpi
        }

class PrintSetupEnvironments env
  where
    defaultPrintSetup   ::  !*env -> (!PrintSetup, !*env)
        // returns a default print setup
    printSetupDialog    ::  !PrintSetup !*env -> (!PrintSetup, !*env)
        // lets the user choose a print setup via the print setup dialog

instance PrintSetupEnvironments World
instance PrintSetupEnvironments (IOSt .l .p)

getPageDimensions   ::  !PrintSetup !Bool-> PageDimensions
instance == PageDimensions

fwritePrintSetup    ::  !PrintSetup !*File -> *File
//  writes PrintSetup to file (text or data)

freadPrintSetup     ::  !*File !*env -> (!Bool, !PrintSetup, !*File, !*env)
                    |   PrintSetupEnvironments env
//  reads PrintSetup from File (text or data). If resulting Boolean is True:success,
//  otherwise the default PrintSetup is returned


print   ::  !Bool !Bool
            .(PrintInfo !*Picture -> ([IdFun *Picture],!*Picture))
            !PrintSetup !*printEnv
        ->  (!PrintSetup,!*printEnv)
        |   PrintEnvironments printEnv

/*  print doDialog emulateScreen pages printSetup env
    sends output to the printer and returns the used print setup, which can differ
    from the input print setup
    doDialog:
        if True a dialog will pop up that lets the user choose all printing options,
```

```
        otherwise printing will happen in the default way.
    emulateScreen:
        if True, the printing routine will emulate the resolution of the screen.
        That means that a pixel on paper has the same dimension as on screen.
        Otherwise, the used resolution will be the printer resolution, with the
        effect that coordinates get much "tighter".
    pages:
        this function should calculate a list of functions, each function
        representing one page to be printed. Each of these drawing functions is
        applied to an initial printer Picture.
    env:
        a PrintEnvironment is either the PSt or the Files system.
*/

printUpdateFunction
        ::  !Bool (UpdateState -> *Picture -> *Picture) [Rectangle]
            !PrintSetup !*printEnv
        ->  (!PrintSetup, !*printEnv)
        | PrintEnvironments printEnv

/*  printUpdateFunction doDialog update area printSetup env
    sends the content of the update function of a given area to the printer:
    doDialog:
        identical to print.
    update:
        this function will be applied to an UpdateState of value
            {oldFrame=area,newFrame=area,updArea=[area]}.
    area:
        the area to be sent to the printer. If a rectangle of this area does not
        fit on one sheet, it will be distributed on several sheets.
    printSetup,env,result value:
        identical to print.
*/

printPagePerPage
    ::  !Bool !Bool
        .x
        .(.x -> .(PrintInfo -> .(*Picture -> ((.Bool,Point),(.state,*Picture)))))
        ((.state,*Picture) -> ((.Bool,Point),(.state,*Picture)))
        !PrintSetup !*printEnv
    ->  (Alternative .x .state,!*printEnv)
    |   PrintEnvironments printEnv

/*  printPagePerPage doDialog emulateScreen x prepare pages printSetup env
    sends output to the printer.
    This function can be used more efficiently than print. The major difference is
    that the pages function is a state transition function instead of a page list
    producing function. Each page transition function generates one page for the
    printer. An additional feature of printPagePerPage is that it is possible to
    set the origin of the printer Pictures.

    doDialog:
        identical to print.
    emulateScreen:
        identical to print.
    x:
        this value is passed to the prepare function.
    prepare:
        this function calculates the initial page print state.
        Iff there are no pages to print, the return Boolean must be True.
        The returned Point is the Origin of the first printer Picture.
    pages:
        this state transition function produces the printed pages.
        The state argument consists of the state information and an initial printer
        Picture which Origin has been set by the previous return Point value.
        If there are no more pages to print, the return Boolean must be True. In
        that case the result of printPagePerPage is (StartedPrinting state),
```

```
        with state the current state value. If printing should continue, the
        return Boolean is False.
        The returned Point is the Origin of the next printer Picture.
    printSetup, env:
        identical to print.

    If printing is cancelled via the print dialog, then (Cancelled x) will be
    returned, otherwise (StartedPrinting ...)
*/

instance PrintEnvironments World
// other instances are the Files subworld and PSt (see osprint.dcl)
```

A.23 StdPrintText

```
definition module StdPrintText

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdPrintText specifies functions to print text.
//  ********************************************************************************

import StdPicture,StdPrint
from    StdString import String


:: WrapMode :== Int

NoWrap          :== 0
LeftJustify     :== 1
RightJustify    :== 2

class CharStreams cs
where
    getChar     :: !*cs -> (!Bool,!Char,!*cs)
        //  returns the charstreams next character, and whether this operation was
        //  successful
    savePos     :: !*cs -> *cs
        //  saves actual position of charstream to be enable the restorePos function
        //  to restore it then.
    restorePos  :: !*cs -> *cs
    eos         :: !*cs -> (!Bool,!*cs)
        //  end of stream

instance CharStreams FileCharStream

:: *FileCharStream

fileToCharStream :: !*File -> *FileCharStream
charStreamToFile :: !*FileCharStream -> *File




printText1 :: !Bool !WrapMode !FontDef !Int !*charStream !PrintSetup !*printEnv
         ->   (!(!*charStream,!PrintSetup),!*printEnv)
         | CharStreams charStream & PrintEnvironments printEnv

/*  printText1 doDialog wrapMode font spacesPerTab charStream printSetup env
    prints a CharStream:
    doDialog:
        identical to print (StdPrint)
    wrapMode:
        controls word wrapping in case lines do not fit. NoWrap suppresses wrapping.
        LeftJustify and RightJustify wrap text to the left and right respectively.
    font:
        the text will be printed in this font.
    spacesPerTab:
        the number of spaces a tab symbol represents.
    charStream:
        the charStream to be printed.
    printSetup, env:
        identical to print (StdPrint)
*/



printText2 :: !String !String !Bool !WrapMode !FontDef !Int !*charStream !PrintSetup
```

```
            !*printEnv
          -> (!(!*charStream,!PrintSetup),!*printEnv)
          |  CharStreams charStream & PrintEnvironments printEnv

/*  printText2 titleStr pageStr doDialog wrapMode fontParams spacesPerTab charStream
             printSetup env
    prints a charStream with a header on each page.
    titleStr:
        this String will be printed on each page at the left corner of the header
    pageStr:
        this String and the actual page number are printed on the right corner of
        the header
    The other parameters are identical to printText1.
*/



printText3 ::!Bool !WrapMode !FontDef !Int
             .(PrintInfo *Picture -> (state, (Int,Int), *Picture))
             (state Int *Picture -> *Picture)
             !*charStream
             !PrintSetup !*printEnv
         ->  (!(!*charStream,!PrintSetup),!*printEnv)
         | CharStreams charStream & PrintEnvironments printEnv

/*  printText3 doDialog wrapMode font spacesPerTab textRange eachPageDraw charStream
             printSetup env
    prints a charStream with a header and trailer on each page.
    textRange:
        this function takes a PrintInfo record and the printer Picture on which the
        text will be printed. It returns a triple (state,range,picture):
        state:
            a value of arbitrary type that can be used to pass data to the page
            printing function pages.
        range:
            a pair (top,bottom), where top<bottom. The printed text will appear
            within these y-coordinates only, so a header and a trailer can be
            printed for each page.
    eachPageDraw:
        this function draws the header and/or trailer for the current page. Its
        arguments are the data produced by textRange, the actual page number, and an
        initial printer Picture. This function is applied by printText3 before each
        new page receives its text.
    The other parameters are identical to printText1.
*/


/*  If a file is openend with FReadData, then all possible newline conventions
    (unix,mac,dos) will be recognized. All these printing functions will replace
    nonprintable characters of the font with ASCII spaces. Exceptions are: newline,
    formfeed and tab. So the ASCII space has to be a printable character in the used
    font. A form feed character will cause a form feed, and it will also end a line.
*/
```

A.24 StdProcess

```
definition module StdProcess


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdProcess contains the process creation and manipulation functions.
//  ********************************************************************************


import  StdProcessDef, StdWindow


/*  General process topology creation functions:
*/

class Processes pdef where
    startProcesses          :: !pdef      !*World      -> *World
    openProcesses           :: !pdef      !(PSt .l .p) -> PSt .l .p

class shareProcesses pdef   :: !(pdef .p) !(PSt .l .p) -> PSt .l .p

/*  (start/open/share)Processes creates an interactive process topology specified by
        the pdef argument.
        All interactive processes can communicate with each other by means of  the
        file system or by message passing.
        By default, processes obtain a private ProcessWindow. However, if a process
        has the ProcessShareGUI attribute, then the process will share the
        ProcessWindow of the current interactive process. Every interactive process
        can create processes in this way. This results in a tree of processes (see
        also the notes of termination at closeProcess).
    startProcesses aborts the application if the argument world does not contain a
        file system.
        startProcesses terminates as soon as all interactive processes that are
        created by startProcesses and their child processes have terminated. It
        returns the final world, consisting of the final file system and event
        stream.
    shareProcesses adds the interactive processes specified by the pdef argument to
        the process group of the current interactive process. The new interactive
        processes can communicate with all interactive processes of the current
        process group by means of the public process state component.
*/

instance Processes      (ProcessGroup pdef) | shareProcesses pdef
instance Processes      [pdef]              | Processes      pdef
instance Processes      (:^: pdef1 pdef2)   | Processes      pdef1
                                            & Processes      pdef2

instance shareProcesses  NDIProcess
instance shareProcesses (SDIProcess wdef)   | Windows        wdef
instance shareProcesses  MDIProcess
instance shareProcesses (ListCS    pdef )   | shareProcesses pdef
instance shareProcesses (:~: pdef1 pdef2)   | shareProcesses pdef1
                                            & shareProcesses pdef2


//  Convenience process creation functions:

startIO :: !.l !.p !(ProcessInit (PSt .l .p)) ![ProcessAttribute (PSt .l .p)]
           !*World -> *World
/*  startIO creates one process group of one interactive MDI process which is
    initialised with the ProcessInit argument.
*/
```

```
//  Process access operations:

closeProcess    :: !(PSt .l .p) -> PSt .l .p
/*  closeProcess removes all abstract devices that are held in the interactive
    process.
    If the interactive process has processes that share its GUI then these will also
    be closed recursively. As a result evaluation of this interactive process
    including GUI sharing processes will terminate.
*/


hideProcess     :: !(IOSt .l .p) -> IOSt .l .p
showProcess     :: !(IOSt .l .p) -> IOSt .l .p
/*  If the interactive process is active, hideProcess hides the interactive process,
    and showProcess makes it visible. Note that hiding an interactive process does
    NOT disable the process but simply makes it invisible.
*/

getProcessWindowPos :: !(IOSt .l .p) -> (!Point,!IOSt .l .p)
/*  getProcessWindowPos returns the current position of the ProcessWindow.
*/


getProcessWindowSize:: !(IOSt .l .p) -> (!Size,!IOSt .l .p)
/*  getProcessWindowSize returns the current size of the ProcessWindow.
*/
```

A.25 StdProcessDef

```
definition module StdProcessDef


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.0.1
//
//  StdProcessDef contains the types to define interactive processes.
//  ********************************************************************************


import  StdIOCommon
from    iostate import PSt, IOSt


::  ProcessGroup pdef
    =   E..p:ProcessGroup p (pdef p)
::  NDIProcess      p                                       // DocumentInterface: NDI
    =   E..l:   NDIProcess  l
                            (ProcessInit      (PSt l p))
                            [ProcessAttribute (PSt l p)]
::  SDIProcess wdef p                                       // DocumentInterface: SDI
    =   E..l ls:SDIProcess  l ls
                            (wdef          ls (PSt l p))
                            (ProcessInit      (PSt l p))
                            [ProcessAttribute (PSt l p)]
::  MDIProcess      p                                       // DocumentInterface: MDI
    =   E..l:   MDIProcess  l
                            (ProcessInit      (PSt l p))
                            [ProcessAttribute (PSt l p)]

/*  NDIProcesses can't open windows and menus.
    SDIProcesses can't open windows, except for their argument window.
    MDIProcesses can open an arbitrary number of device instances.
*/

::  ProcessInit ps
    :== [IdFun ps]
```

A.26 StdPSt

```
definition module StdPSt


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdPSt defines operations on PSt and IOSt that are not abstract device related.
//  ********************************************************************************


import  StdFunc, StdFile
import  StdFileSelect, StdIOCommon, StdPicture, StdTime
from    iostate import  PSt, IOSt
from    channelenv  import  ChannelEnv


/*  PSt is an environment instance of the following classes:
    - FileEnv       (see StdFile)
    - FileSelectEnv (see StdFileSelect)
    - TimeEnv       (see StdTime)
    - ChannelEnv    (see StdChannels)
*/
instance FileEnv        (PSt .l .p)
instance FileSelectEnv  (PSt .l .p)
instance TimeEnv        (PSt .l .p)
instance ChannelEnv     (PSt .l .p)
instance Ids            (PSt .l .p)


/*  IOSt is an environment instance of the following classes:
    - FileEnv
    - TimeEnv
    - ChannelEnv
*/

instance FileEnv    (IOSt .l .p)
instance ChannelEnv (IOSt .l .p)
instance TimeEnv    (IOSt .l .p)

/*  accScreenPicture provides access to an initial Picture as it would be created in
    a window or control.
*/
class accScreenPicture env :: !(St *Picture .x) !*env -> (!.x,!*env)

instance accScreenPicture World
instance accScreenPicture (IOSt .l .p)


beep :: !(IOSt .l .p) -> IOSt .l .p
/*  beep emits the alert sound.
*/


//  Operations on the global cursor:

/* RWS ---
setCursor       :: !CursorShape !(IOSt .l .p) -> IOSt .l .p
resetCursor     ::              !(IOSt .l .p) -> IOSt .l .p
obscureCursor   ::              !(IOSt .l .p) -> IOSt .l .p
/*  setCursor       overrules the shape of the cursor of all windows.
    resetCursor     removes the overruled cursor shape of all windows.
    obscureCursor   hides the cursor until the mouse is moved.
*/
```

```
//  Operations on the DoubleDownDistance:

setDoubleDownDistance :: !Int !(IOSt .l .p) -> IOSt .l .p
/*  setDoubleDownDistance sets the maximum distance the mouse is allowed to move to
    generate a ButtonDouble(Triple)Down button state. Negative values are set to
    zero.
*/
--- RWS */

//  Operations on the DocumentInterface of an interactive process:

getDocumentInterface :: !(IOSt .l .p) -> (!DocumentInterface, !IOSt .l .p)
/*  getDocumentInterface returns the DocumentInterface of the interactive process.
*/


//  Operations on the attributes of an interactive process:

setProcessActivate  :: !(IdFun (PSt .l .p)) !(IOSt .l .p) -> IOSt .l .p
setProcessDeactivate:: !(IdFun (PSt .l .p)) !(IOSt .l .p) -> IOSt .l .p
setProcessHelp      :: !(IdFun (PSt .l .p)) !(IOSt .l .p) -> IOSt .l .p
setProcessAbout     :: !(IdFun (PSt .l .p)) !(IOSt .l .p) -> IOSt .l .p
/*  These functions set the ProcessActivate, ProcessDeactivate, ProcessHelp, and
    ProcessAbout attribute of the interactive process respectively.
*/


//  Coercing PSt component operations to PSt operations.

appListPIO  :: ![.IdFun (IOSt .l .p)]   !(PSt .l .p) ->          PSt .l .p
appListPLoc :: ![.IdFun .l]             !(PSt .l .p) ->          PSt .l .p
appListPPub :: ![.IdFun .p]             !(PSt .l .p) ->          PSt .l .p

appPIO      :: !.(IdFun (IOSt .l .p))   !(PSt .l .p) ->          PSt .l .p
appPLoc     :: !.(IdFun .l)             !(PSt .l .p) ->          PSt .l .p
appPPub     :: !.(IdFun .p)             !(PSt .l .p) ->          PSt .l .p

//  Accessing PSt component operations.

accListPIO  :: ![.St (IOSt .l .p) .x]   !(PSt .l .p) -> (![.x], !PSt .l .p)
accListPLoc :: ![.St .l           .x]   !(PSt .l .p) -> (![.x], !PSt .l .p)
accListPPub :: ![.St .p           .x]   !(PSt .l .p) -> (![.x], !PSt .l .p)

accPIO      :: !.(St (IOSt .l .p) .x)   !(PSt .l .p) -> (!.x,   !PSt .l .p)
accPLoc     :: !.(St .l           .x)   !(PSt .l .p) -> (!.x,   !PSt .l .p)
accPPub     :: !.(St .p           .x)   !(PSt .l .p) -> (!.x,   !PSt .l .p)
```

A.27 StdReceiver

```
definition module StdReceiver


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdReceiver specifies all receiver operations.
//  ********************************************************************************


import  StdReceiverDef, StdMaybe
from    iostate import  PSt, IOSt
from    id      import  RId, R2Id, RIdtoId, R2IdtoId, ==


//  Operations on the ReceiverDevice.

//  Open uni- and bi-directional receivers:

class Receivers rdef where
    openReceiver    :: .ls !*(*rdef .ls (PSt .l .p)) !(PSt .l .p)
                    -> (!ErrorReport,!PSt .l .p)
    getReceiverType :: *(*rdef .ls .ps)
                    -> ReceiverType
/*  openReceiver
        opens the given receiver if no receiver currently exists with the given
        R(2)Id. The R(2)Id has to be used to send messages to this receiver.
    getReceiverType
        returns the type of the receiver (see also getReceivers).
*/

instance Receivers (Receiver  msg)
instance Receivers (Receiver2 msg resp)


closeReceiver       :: !Id !(IOSt .l .p) -> IOSt .l .p
/*  closeReceiver closes the indicated uni- or bi-directional receiver.
    Invalid Ids have no effect.
*/


getReceivers        :: !(IOSt .l .p) -> (![(Id,ReceiverType)], !IOSt .l .p)
/*  getReceivers returns the Ids and ReceiverTypes of all currently open uni- or
    bi-directional receivers of this interactive process.
*/


enableReceivers         :: ![Id] !(IOSt .l .p) ->                      IOSt .l .p
disableReceivers        :: ![Id] !(IOSt .l .p) ->                      IOSt .l .p
getReceiverSelectState  :: ! Id  !(IOSt .l .p) -> (!Maybe SelectState,!IOSt .l .p)
/*  (en/dis)ableReceivers
        (en/dis)able the indicated uni- or bi-directional receivers.
        Note that this implies that in case of synchronous message passing messages
        can fail (see the comments of syncSend and syncSend2 below). Invalid Ids
        have no effect.
    getReceiverSelectState
        yields the current SelectState of the indicated receiver. In case the
        receiver does not exist, Nothing is returned.
*/


//  Inter-process communication:

//  Message passing status report:
::  SendReport
```

```
    =   SendOk
    |   SendUnknownReceiver
    |   SendUnableReceiver
    |   SendDeadlock

instance ==       SendReport
instance toString SendReport

asyncSend :: !(RId msg) msg !(PSt .l .p) -> (!SendReport, !PSt .l .p)
/*  asyncSend posts a message to the receiver indicated by the argument RId. In case
    the indicated receiver belongs to this process, the message is simply buffered.
    asyncSend is asynchronous: the message will at some point be received by the
    indicated receiver.
    The SendReport can be one of the following alternatives:
    -   SendOk: No exceptional situation has occurred. The message has been sent.
                Note that even though the message has been sent, it cannot be
                guaranteed that the message will actually be handled by the
                indicated receiver because it might become closed, forever disabled,
                or flooded with synchronous messages.
    -   SendUnknownReceiver:
                The indicated receiver does not exist.
    -   SendUnableReceiver:
                Does not occur: the message is always buffered, regardless whether
                the indicated receiver is Able or Unable. Note that in case the
                receiver never becomes Able, the message will not be handled.
    -   SendDeadlock:
                Does not occur.
*/

syncSend :: !(RId msg) msg !(PSt .l .p) -> (!SendReport, !PSt .l .p)
/*  syncSend posts a message to the receiver indicated by the argument RId. In case
    the indicated receiver belongs to the current process, the corresponding
    ReceiverFunction is applied directly to the message argument and current process
    state.
    syncSend is synchronous: this interactive process blocks evaluation until the
    indicated receiver has received the message.
    The SendReport can be one of the following alternatives:
    -   SendOk: No exceptional situation has occurred. The message has been sent and
                handled by the indicated receiver.
    -   SendUnknownReceiver:
                The indicated receiver does not exist.
    -   SendUnableReceiver:
                The receiver exists, but its ReceiverSelectState attribute is Unable.
                Message passing is halted. The message is not sent.
    -   SendDeadlock:
                The receiver is involved in a synchronous, cyclic communication
                with the current process. Blocking the current process would result
                in a deadlock situation. Message passing is halted to circumvent the
                deadlock. The message is not sent.
*/

syncSend2 :: !(R2Id msg resp) msg  !(PSt .l .p)
    -> (!(!SendReport,!Maybe resp), !PSt .l .p)
/*  syncSend2 posts a message to the receiver indicated by the argument R2Id. In
    case the indicated receiver belongs to the current process, the corresponding
    Receiver2Function is applied directly to the message argument and current
    process state.
    syncSend2 is synchronous: this interactive process blocks until the indicated
    receiver has received the message.
    The SendReport can be one of the following alternatives:
    -   SendOk: No exceptional situation has occurred. The message has been sent and
                handled by the indicated receiver. The response of the receiver is
                returned as well as (Just response).
    -   SendUnknownReceiver:
                The indicated receiver does not exist.
    -   SendUnableReceiver:
                The receiver exists, but its ReceiverSelect attribute is Unable.
```

```
                Message passing is halted. The message is not sent.
    -   SendDeadlock:
                The receiver is involved in a synchronous, cyclic communication
                with the current process. Blocking the current process would result
                in a deadlock situation. Message passing is halted to circumvent the
                deadlock. The message is not sent.
    In all other cases than SendOk, the optional response is Nothing.
*/
```

A.28 StdReceiverDef

```
definition module StdReceiverDef


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdReceiverDef contains the types to define the standard set of receivers.
//  ********************************************************************************


import  StdIOCommon


::  Receiver  m   ls ps = Receiver  (RId  m)   (ReceiverFunction  m   *(ls,ps))
                                                [ReceiverAttribute     *(ls,ps)]
::  Receiver2 m r ls ps = Receiver2 (R2Id m r) (Receiver2Function m r *(ls,ps))
                                                [ReceiverAttribute     *(ls,ps)]

::  ReceiverFunction  m   ps :== m -> ps -> ps
::  Receiver2Function m r ps :== m -> ps -> (r,ps)

::  ReceiverAttribute   ps                          // Default:
    =   ReceiverSelectState SelectState             // receiver Able
    |   ReceiverConnectedReceivers  [Id]            // []
::  ReceiverType
    :== String
```

A.29 StdStringChannels

```
definition module StdStringChannels

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdStringChannels provides channel instanciates to send and receive Strings
//  These channels use their own protocol above TCP
//  ********************************************************************************

import  StdString
import  StdChannels, StdTCPDef, StdTCPChannels, StdReceiver, StdEventTCP

/*  If a string via a StringChannel is sent, then first the length of the string is
    sent, and then the string itself, e.g. sending the string "abc" will result in
    "3 abc\xD"
*/

////////////////////// StringChannels to receive ////////////////////////////////////

::  *StringRChannel_ a
::  *StringRChannel     :== StringRChannel_ String
::  *StringRChannels    =   StringRChannels [StringRChannel]
::  *StringChannelReceiver ls ps
    =   StringChannelReceiver
            (RId (ReceiveMsg String)) StringRChannel
            (ReceiverFunction (ReceiveMsg String)       *(ls,ps))
            [ReceiverAttribute                          *(ls,ps)]

toStringRChannel    ::  TCP_RChannel -> StringRChannel

instance Receivers          StringChannelReceiver
instance Receive            StringRChannel_
instance closeRChannel      StringRChannel_
instance MaxSize            StringRChannel_

////////////////////// StringChannels to send ////////////////////////////////////

::  *StringSChannel_ a
::  *StringSChannel     :== StringSChannel_ String
::  *StringSChannels    =   StringSChannels [StringSChannel]

toStringSChannel    ::  TCP_SChannel -> StringSChannel

instance Send               StringSChannel_

////////////////////////////////////////////////////////////////////////////////////

// for openSendNotifier, closeSendNotifier
instance accSChannel        StringSChannel_

// for selectChannel
instance SelectSend         StringSChannels
instance SelectReceive      StringRChannels
instance getNrOfChannels    StringRChannels
```

A.30 StdSystem

```
definition module StdSystem


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdSystem defines platform dependent constants and functions.
//  ********************************************************************************


import  StdIOCommon


//  System dependencies concerning the file system.

dirseparator    :: Char     // Separator between folder- and filenames in a pathname
homepath        :: !String -> String
applicationpath :: !String -> String
newlineChars    :: !String
/*  dirseparator
        is the separator symbol used between folder- and filenames in a file path.
    homepath
        prefixes the 'home' directory file path to the given file name.
    applicationpath
        prefixes the 'application' directory file path to the given file name.
    Use these directories to store preference/options/help files of an application.

    newlineChars
        the newline characters in a textfile
*/


//  System dependencies concerning the time resolution

ticksPerSecond  :: Int
/*  ticksPerSecond returns the maximum timer resolution per second.
*/


//  System dependencies concerning the screen resolution.

mmperinch       :== 25.4

hmm             :: !Real -> Int
vmm             :: !Real -> Int
hinch           :: !Real -> Int
vinch           :: !Real -> Int
/*  h(mm/inch) convert millimeters/inches into pixels, horizontally.
    v(mm/inch) convert millimeters/inches into pixels, vertically.
*/

maxScrollWindowSize :: Size
maxFixedWindowSize  :: Size
/*  maxScrollWindowSize
        yields the range at which scrollbars are inactive.
    maxFixedWindowSize
        yields the range at which a window still fits on the screen.
*/

printSetupTypical   :: Bool
```

A.31 StdTCP

```
definition module StdTCP

import  StdChannels,
        StdTCPChannels,
        StdEventTCP,
        StdStringChannels
```

A.32 StdTCPChannels

```
definition module StdTCPChannels

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdTCPChannels provides instanciates to use TCP
//  ********************************************************************************

import  StdTCPDef, StdChannels
from    StdIOCommon     import OkBool
from    StdIOBasic      import Void
from    tcp_bytestreams import TCP_SCharStream_, TCP_RCharStream_

////////////////////////////////// Listeners //////////////////////////////////////

instance Receive            TCP_Listener_
instance closeRChannel      TCP_Listener_
//  receiving on a listener will accept a TCP_DuplexChannel. eom never becomes True
//  for listeners.

//////////////////////////////// TCP send channels ////////////////////////////////

instance Send               TCP_SChannel_

//////////////////////////////// TCP receive channels ////////////////////////////////

instance Receive            TCP_RChannel_
instance closeRChannel      TCP_RChannel_
instance MaxSize            TCP_RChannel_

////////////////////////////// TCP char streams to receive ////////////////////////////

::  *TCP_RCharStream    :== TCP_RCharStream_ Char
::  *TCP_RCharStreams   =   TCP_RCharStreams [TCP_RCharStream]

toRCharStream   ::  !TCP_RChannel -> TCP_RCharStream

instance Receive            TCP_RCharStream_
instance closeRChannel      TCP_RCharStream_

////////////////////////////// TCP char streams to send ////////////////////////////

::  *TCP_SCharStream    :== TCP_SCharStream_ Char
::  *TCP_SCharStreams   =   TCP_SCharStreams [TCP_SCharStream]

toSCharStream   ::  !TCP_SChannel -> TCP_SCharStream

instance Send               TCP_SCharStream_

//////////////////////////// establishing connections ////////////////////////////////

lookupIPAddress     ::  !String !*env
                    ->  (!Maybe IPAddress, !*env)
                    |   ChannelEnv env
connectTCP_MT       ::  !(Maybe !Timeout) !(!IPAddress,!Port) !*env
                    ->  (!TimeoutReport, !Maybe TCP_DuplexChannel, !*env)
                    |   ChannelEnv env
openTCP_Listener    ::  !Port !*env
                    ->  (!OkBool, !Maybe TCP_Listener, !*env)
                    |   ChannelEnv env
tcpPossible         ::  !*env
                    ->  (!Bool, !*env)
                    |   ChannelEnv env

/*  lookupIPAddress
```

```
        input String can be in dotted decimal form or alphanumerical. In the latter
        case the DNS is called
    connectTCP
        tries to establish a TCP connection
    openTCP_Listener
        to listen on a certain port
    tcpPossible
        whether tcp can be started on this computer
*/

////////////////////////////// multiplexing  ////////////////////////////////////

selectChannel_MT    :: !(Maybe !Timeout) !*r_channels  !*s_channels  !*World
                    -> (![(!Int, !SelectResult)],
                                        !*r_channels, !*s_channels, !*World)
                    |   SelectReceive r_channels & SelectSend s_channels
/*  selectChannel_MT mbTimeout r_channels s_channels world
        determines on which channels first "something happens". If the result is
        an empty list, then the timeout expired. Otherwise each (who, what) element
        of the result identifies one channel in r_channels or s_channels. The
        what value determines whether available/eom/disconnected on the identified
        channel would have returned True. what==SR_Sendable indicates, that it is
        possible to send non blocking on the identified channel. If r_channels
        contains r channels and if s_channels contains s channels, then the
        following holds:
            isMember what [SR_Available ,SR_EOM]            =>  0<=who<r
            isMember what [SR_Sendable ,SR_Disconnected]    =>  0<=who<s
*/

instance == SelectResult
instance toString SelectResult

//  the following classes support the selectChannel_MT function

class SelectReceive channels
  where
    accRChannels    :: (PrimitiveRChannel -> (x, PrimitiveRChannel)) !*channels
                    -> (![x], !*channels)
    getRState       :: !Int !*channels !*World
                    -> (!Maybe !SelectResult, !*channels, !*World)
/*  accRChannels f channels
        applies a function on each channel in channels and returns a list, which
        contains the result for each application
    getRState
        applies available and eom on the channel which is identified by the Int
        parameter and returns SR_Available or SR_EOM or Nothing
*/

class SelectSend channels
  where
    accSChannels    :: (TCP_SChannel -> (x, TCP_SChannel)) !*channels
                    -> (![x], !*channels)
    appDisconnected :: !Int !*channels !*World
                    -> (!Bool, !*channels, !*World)
/*  accSChannels
        applies a function on each channel in channels and returns a list, which
        contains the result for each application
    appDisconnected
        returns whether disconnected is True for the channel which is identified by
        the Int parameter
*/

class getNrOfChannels channels  :: !*channels
                                -> (!Int, !*channels)
//  getNrOfChannels channels
//      returns the number of channels in channels
```

```
instance SelectReceive TCP_RChannels,   TCP_Listeners,  TCP_RCharStreams,   Void
instance SelectReceive (:^: *x *y)      | SelectReceive, getNrOfChannels x
                                        & SelectReceive y

instance SelectSend TCP_SChannels, TCP_SCharStreams, Void
instance SelectSend (:^: *x *y)         | SelectSend, getNrOfChannels x
                                        & SelectSend y

instance getNrOfChannels TCP_RChannels, TCP_Listeners, TCP_RCharStreams,
                         TCP_SChannels, TCP_SCharStreams, Void
instance getNrOfChannels (:^: *x *y)    | getNrOfChannels x & getNrOfChannels y
```

A.33 StdTCPDef

```
definition module StdTCPDef

//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdTCPDef provides basic definitions for using TCP
//  ********************************************************************************

import  StdString
import  StdReceiverDef, StdMaybe
from    StdChannels     import  DuplexChannel, ReceiveMsg, SendEvent
from    tcp             import  TCP_SChannel_,TCP_RChannel_,TCP_Listener_,IPAddress

::  *TCP_SChannel       :== TCP_SChannel_ ByteSeq
::  *TCP_RChannel       :== TCP_RChannel_ ByteSeq
::  *TCP_Listener       :== TCP_Listener_ (IPAddress, TCP_DuplexChannel)

::  Port                :== Int

::  *TCP_DuplexChannel  :== DuplexChannel *TCP_SChannel_ *TCP_RChannel_ ByteSeq

::  ByteSeq
//  a sequence of bytes

instance toString   ByteSeq
instance ==         ByteSeq
toByteSeq       ::  !x          -> ByteSeq  | toString x
byteSeqSize     ::  !ByteSeq    -> Int
// byteSeqSize returns the size in bytes

instance toString IPAddress
//  returns ip address in dotted decimal form

//////////////////////////// for event driven processing ////////////////////////////

//  to receive byte sequences

::  *TCP_Receiver ls ps
    =   TCP_Receiver
            Id TCP_RChannel
            (ReceiverFunction (ReceiveMsg ByteSeq)  *(ls,ps))
            [ReceiverAttribute                      *(ls,ps)]

::  SendNotifier sChannel ls ps
    =   SendNotifier
            sChannel
            (ReceiverFunction SendEvent     *(ls,ps))
            [ReceiverAttribute              *(ls,ps)]

//  to accept new connections

::  *TCP_ListenerReceiver ls ps
    =   TCP_ListenerReceiver
            Id TCP_Listener
            ((ReceiveMsg (IPAddress,TCP_DuplexChannel)) -> *(*(ls,ps) -> *(ls,ps)))
            [ReceiverAttribute                              *(ls,ps)]

//  to receive characters

::  *TCP_CharReceiver ls ps
    =   TCP_CharReceiver
            Id TCP_RChannel (Maybe NrOfIterations)
            (ReceiverFunction (ReceiveMsg Char)     *(ls,ps))
            [ReceiverAttribute                      *(ls,ps)]
```

```
/*  For efficency the receiver function of a TCP_CharReceiver will be called from
    a loop. Within this loop no other events can be handled. The NrOfIterations
    parameter limits the maximum number of iterations.
*/

::  NrOfIterations          :== Int
::  InetLookupFunction ps   :== (Maybe IPAddress) -> *(ps -> ps)
::  InetConnectFunction ps  :== (Maybe TCP_DuplexChannel) -> *(ps -> ps)

////////////////////////////// for multiplexing ////////////////////////////////

:: *TCP_RChannels       =   TCP_RChannels [TCP_RChannel]
:: *TCP_SChannels       =   TCP_SChannels [TCP_SChannel]
:: *TCP_Listeners       =   TCP_Listeners [TCP_Listener]

::  *PrimitiveRChannel
    =   TCP_RCHANNEL TCP_RChannel
    |   TCP_LISTENER TCP_Listener

::  SelectResult
    =   SR_Available
    |   SR_EOM
    |   SR_Sendable
    |   SR_Disconnected
```

A.34 StdTime

```
definition module StdTime


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdTime contains time related operations.
//  ********************************************************************************

import  StdOverloaded
from    ostick import Tick

::  Time
    =   {   hours   :: !Int     // hours        (0-23)
        ,   minutes :: !Int     // minutes      (0-59)
        ,   seconds :: !Int     // seconds      (0-59)
        }
::  Date
    =   {   year    :: !Int     // year
        ,   month   :: !Int     // month        (1-12)
        ,   day     :: !Int     // day          (1-31)
        ,   dayNr   :: !Int     // day of week  (1-7, Sunday=1, Saturday=7)
        }

wait                :: !Int .x -> .x

/*  wait n x suspends the evaluation of x modally for n ticks.
    If n<=0, then x is evaluated immediately.
*/

instance < Tick
intPlusTick     ::  !Int  !Tick -> Tick
tickDifference  ::  !Tick !Tick -> Int

class TimeEnv env where
    getBlinkInterval:: !*env -> (!Int,  !*env)
    getCurrentTime  :: !*env -> (!Time, !*env)
    getCurrentDate  :: !*env -> (!Date, !*env)
    getCurrentTick  :: !*env -> (!Tick, !*env)
/*  getBlinkInterval
        returns the time interval in ticks that should elapse between blinks of
        e.g. a cursor. This interval may be changed by the user while the
        interactive process is running!
    getCurrentTime
        returns the current Time.
    getCurrentDate
        returns the current Date.
    getCurrentTick
        returns the current Tick.
*/

instance TimeEnv World
```

A.35 StdTimer

```
definition module StdTimer


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdTimer specifies all timer operations.
//  ********************************************************************************


import  StdTimerDef, StdTimerElementClass, StdMaybe
from    StdSystem   import ticksPerSecond
from    iostate     import PSt, IOSt


class Timers tdef where
    openTimer   :: .ls !(tdef .ls (PSt .l .p)) !(PSt .l .p)
                               -> (!ErrorReport,!PSt .l .p)
    getTimerType::      (tdef .ls .ps) -> TimerType
/*  Open a new timer.
    This function has no effect in case the interactive process already contains a
    timer with the same Id. In case TimerElements are opened with duplicate Ids, the
    timer will not be opened. Negative TimerIntervals are set to zero.
    In case the timer does not have an Id, it will obtain an Id which is fresh with
    respect to the current set of timers. The Id can be reused after closing this
    timer.
*/

instance Timers (Timer t)   | TimerElements t


closeTimer :: !Id !(IOSt .l .p) -> IOSt .l .p
/*  closeTimer closes the timer with the indicated Id.
*/


getTimers  ::     !(IOSt .l .p) -> ([(Id,TimerType)],!IOSt .l .p)
/*  getTimers returns the Ids and TimerTypes of all currently open timers.
*/


enableTimer         :: !Id !(IOSt .l .p) -> IOSt .l .p
disableTimer        :: !Id !(IOSt .l .p) -> IOSt .l .p
getTimerSelectState :: !Id !(IOSt .l .p) -> (!Maybe SelectState,!IOSt .l .p)
/*  (en/dis)ableTimer (en/dis)ables the indicated timer.
    getTimerSelectState yields the SelectState of the indicated timer. If the timer
    does not exist, then Nothing is yielded.
*/


setTimerInterval    :: !Id !TimerInterval   !(IOSt .l .p) -> IOSt .l .p
getTimerInterval    :: !Id                  !(IOSt .l .p)
                    -> (!Maybe TimerInterval,!IOSt .l .p)
/*  setTimerInterval
        sets the TimerInterval of the indicated timer.
        Negative TimerIntervals are set to zero.
    getTimerInterval
        yields the TimerInterval of the indicated timer.
        If the timer does not exist, then Nothing is yielded.
*/
```

A.36 StdTimerDef

```
definition module StdTimerDef


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdTimerDef contains the types to define the standard set of timers.
//  ********************************************************************************


import  StdIOCommon


::  Timer t ls ps   = Timer TimerInterval (t ls ps) [TimerAttribute *(ls,ps)]

::  TimerInterval
    :== Int

::  TimerAttribute  ps                              // Default:
    =   TimerId             Id                      // no Id
    |   TimerSelectState    SelectState             // timer Able
    |   TimerFunction       (TimerFunction ps)      // \_ x->x

::  TimerFunction   ps  :== NrOfIntervals->ps->ps
::  NrOfIntervals       :== Int

::  TimerType           :== String
::  TimerElementType    :== String
```

A.37 StdTimerElementclass

```
definition module StdTimerElementClass


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdTimerElementClass define the standard set of timer element instances.
//  ********************************************************************************


import  StdIOCommon, StdTimerDef
from    timerhandle import TimerElementState


class TimerElements t where
    timerElementToHandles   :: !(t .ls .ps) -> [TimerElementState .ls .ps]
    getTimerElementType     ::  (t .ls .ps) -> TimerElementType

instance TimerElements (NewLS   t)    | TimerElements t  // getTimerElementType==""
instance TimerElements (AddLS   t)    | TimerElements t  // getTimerElementType==""
instance TimerElements (ListLS  t)    | TimerElements t  // getTimerElementType==""
instance TimerElements NilLS                             // getTimerElementType==""
instance TimerElements ((:+:) t1 t2) | TimerElements t1
                                     & TimerElements t2 // getTimerElementType==""
```

A.38 StdTimerReceiver

```
definition module StdTimerReceiver


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdTimerReceiver defines Receiver(2) timer element instances.
//  ********************************************************************************


import  StdReceiverDef, StdTimerElementClass


//  Receiver components for timers:
instance TimerElements  (Receiver  m  )
instance TimerElements  (Receiver2 m r)
```

A.39 StdWindow

```
definition module StdWindow


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdWindow defines functions on windows and dialogues.
//  ********************************************************************************


import  StdMaybe, StdWindowDef
from    StdPSt          import PSt, IOSt
from    StdControlClass import Controls, ControlState


//  Functions applied to non-existent windows or unknown ids have no effect.
class Windows wdef
where
    openWindow      :: .ls !(wdef .ls (PSt .l .p)) !(PSt .l .p)
                                -> (!ErrorReport,!PSt .l .p)
    getWindowType   ::      (wdef .ls .ps)  -> WindowType

class Dialogs wdef
where
    openDialog      :: .ls !(wdef .ls (PSt .l .p)) !(PSt .l .p)
                                -> (!ErrorReport,!PSt .l .p)
    openModalDialog :: .ls !(wdef .ls (PSt .l .p)) !(PSt .l .p)
                                -> (!ErrorReport,!PSt .l .p)
    getDialogType   ::      (wdef .ls .ps)  -> WindowType

/*  open(Window/Dialog) opens the given window(dialog).
    If the Window(Dialog) has no WindowIndex attribute (see StdWindowDef), then the
    new window is opened frontmost.
    If the Window(Dialog) has a WindowIndex attribute, then the new window is
    opened behind the window indicated by the integer index:
        Index value 1 indicates the top-most window.
        Index value M indicates the bottom-most modal window, if there are M modal
            windows.
        Index value N indicates the bottom-most window, if there are N windows.
        If index<M, then the new window is added behind the bottom-most modal window
            (at index M).
        If index>N, then the new window is added behind the bottom-most window
            (at index N).
    openModalDialog always opens a window at the front-most position.
    openWindow may not be permitted to open a window depending on its
        DocumentInterface (see the comments at the ShareProcesses instances in
        module StdProcess).
    In case the window does not have an Id, it will obtain an Id which is fresh with
    respect to the current set of windows. The Id can be reused after closing this
    window.
    In case a window with the same Id is already open the window will not be opened.
    In case controls are opened with duplicate Ids, the window will not be opened.
    openModalDialog terminates when:
        the window has been closed (by means of closeWindow), or the process has
        been terminated (by means of closeProcess).
*/

instance Windows (Window c) | Controls c
instance Dialogs (Dialog c) | Controls c


closeWindow :: !Id !(PSt .l .p) -> PSt .l .p
/*  If the indicated window is not an inactive modal window, then closeWindow closes
    the window.
    In case the Id of the window was generated by open(Window/Dialog), it will
```

```
    become reusable for new windows/dialogues.
    In case of unknown Id, closeWindow does nothing.
*/

closeControls :: !Id [Id] !Bool !(IOSt .l .p) -> IOSt .l .p
/*  closeControls removes the indicated controls (second argument) from the
    indicated window (first argument) and recalculates the layout iff the Boolean
    argument is True.
*/


openControls        :: !Id     .ls (cdef .ls (PSt .l .p)) !(IOSt .l .p)
                                        -> (!ErrorReport,!IOSt .l .p)
                                        |  Controls cdef
openCompoundControls:: !Id !Id .ls (cdef .ls (PSt .l .p)) !(IOSt .l .p)
                                        -> (!ErrorReport,!IOSt .l .p)
                                        |  Controls cdef
/*  openControls
        adds the given controls argument to the indicated window.
    openCompoundControls
        adds the given controls argument to the indicated compound control.
    Both functions have no effect in case the indicated window/dialog/compound
    control could not be found (ErrorUnknownObject) or if controls are opened with
    duplicate Ids (ErrorIdsInUse).
*/


setControlPos :: !Id !Id !ItemPos   !(IOSt .l .p) -> (!Bool,!IOSt .l .p)
/*  setControlPos changes the current layout position of the indicated control to
    the new position.
    If there are relatively laynout controls, then their layout also changes. The
    window is not resized.
    The Boolean result is False iff the window or control id are unknown, or if the
    new ItemPos refers to an unknown control.
*/


controlSize :: !(cdef .ls (PSt .l .p))
                    !(Maybe (Int,Int)) !(Maybe (Int,Int)) !(Maybe (Int,Int))
                    !(IOSt .l .p)
            -> (!Size,!IOSt .l .p)
            |  Controls cdef
/*  controlSize calculates the size of the given control definition as it would be
    opened as an element of a window/dialog.
    The Maybe arguments are the prefered horizontal margins, vertical margins, and
    item spaces (see also the (Window/Control)(H/V)Margin and
    (Window/Control)ItemSpace attributes). If Nothing is specified, their default
    values are used.
*/


hideWindows     :: ![Id]    !(IOSt .l .p) -> IOSt .l .p
showWindows     :: ![Id]    !(IOSt .l .p) -> IOSt .l .p
getHiddenWindows::          !(IOSt .l .p) -> (![Id],!IOSt .l .p)
getShownWindows ::          !(IOSt .l .p) -> (![Id],!IOSt .l .p)
/*  (hide/show)Windows hides/shows the indicated modeless windows (modal dialogues
    can not be hidden).
    get(Hidden/Shown)Windows yields the list of currently visible/invisible windows.
*/


activateWindow  :: !Id      !(IOSt .l .p) -> IOSt .l .p
/*  activateWindow makes the indicated window the active window.
    If the window was hidden, then it will become shown.
    If there are modal dialogues, then the window will be placed behind the last
    modal dialog.
    activateWindow has no effect in case the window is unknown or is a modal dialog.
```

```
*/

getActiveWindow ::          !(IOSt .l .p) -> (!Maybe Id,!IOSt .l .p)
/*  getActiveWindow returns the Id of the window that currently has the input focus
    of the interactive process.
    Nothing is returned if there is no such window.
    Note that hidden windows never are active windows, and that modal windows never
    are hidden.
*/


stackWindow     :: !Id !Id  !(IOSt .l .p) -> IOSt .l .p
/*  stackWindow id1 id2 places the window with id1 behind the window with id2.
    If id1 or id2 is unknown, or id1 indicates a modal window, stackWindow does
    nothing.
    If id2 indicates a modal window, then the window with id1 is placed behind the
    last modal window.
*/

getWindowStack  :: !(IOSt .l .p) -> (![(Id,WindowType)],!IOSt .l .p)
getWindowsStack :: !(IOSt .l .p) -> (![Id],             !IOSt .l .p)
getDialogsStack :: !(IOSt .l .p) -> (![Id],             !IOSt .l .p)
/*  getWindowStack returns the Ids and WindowTypes of all currently open windows
    (including the hidden windows), in the current stacking order starting with the
    active window.
    get(Windows/Dialogs)Stack is equal to getWindowStack, restricted to Windows
    instances and Dialogs instances respectively.
*/


getDefaultHMargin   ::      !(IOSt .l .p) -> ((Int,Int),        !IOSt .l .p)
getDefaultVMargin   ::      !(IOSt .l .p) -> ((Int,Int),        !IOSt .l .p)
getDefaultItemSpace ::      !(IOSt .l .p) -> ((Int,Int),        !IOSt .l .p)
getWindowHMargin    :: !Id  !(IOSt .l .p) -> (!Maybe (Int,Int), !IOSt .l .p)
getWindowVMargin    :: !Id  !(IOSt .l .p) -> (!Maybe (Int,Int), !IOSt .l .p)
getWindowItemSpace  :: !Id  !(IOSt .l .p) -> (!Maybe (Int,Int), !IOSt .l .p)
/*  getDefault((H/V)Margin)/ItemSpace) return the default values for the horizontal
    and vertical window/dialogue margins and item spaces.
    getWindow((H/V)Margin/ItemSpace) return the current horizontal and vertical
    margins and item spaces of the indicated window. These will have the default
    values in case they are not specified.
    In case the window does not exist, Nothing is yielded.
*/


enableWindow            :: !Id  !(IOSt .l .p) -> IOSt .l .p
disableWindow           :: !Id  !(IOSt .l .p) -> IOSt .l .p
enableWindowMouse       :: !Id  !(IOSt .l .p) -> IOSt .l .p
disableWindowMouse      :: !Id  !(IOSt .l .p) -> IOSt .l .p
enableWindowKeyboard    :: !Id  !(IOSt .l .p) -> IOSt .l .p
disableWindowKeyboard   :: !Id  !(IOSt .l .p) -> IOSt .l .p
/*  (en/dis)ableWindow
        (en/dis)ables the indicated window.
    (en/dis)ableWindowMouse
        (en/dis)ables mouse handling of the indicated window.
    (en/dis)ableWindowKeyboard
        (en/dis)ables keyboard handling of the indicated window.
    Disabling a window overrules the SelectStates of its elements, which all become
    Unable.
    Reenabling the window reestablishes the SelectStates of its elements.
    The functions have no effect in case of invalid Ids or Dialogs instances.
    The latter four functions also have no effect in case the Window does not have
    the indicated attribute.
*/


getWindowSelectState        :: !Id  !(IOSt .l .p) ->(!Maybe SelectState,!IOSt .l .p)
```

```
getWindowMouseSelectState   :: !Id  !(IOSt .l .p) ->(!Maybe SelectState,!IOSt .l .p)
getWindowKeyboardSelectState:: !Id  !(IOSt .l .p) ->(!Maybe SelectState,!IOSt .l .p)
/*  getWindowSelectState
        yields the current SelectState of the indicated window.
    getWindow(Mouse/Keyboard)SelectState
        yields the current SelectState of the mouse/keyboard of the indicated
        window.
    The functions return Nothing in case of invalid Ids or Dialogs instances or if
    the Window does not have the indicated attribute.
*/


getWindowMouseStateFilter   :: !Id                          !(IOSt .l .p)
                            -> (!Maybe MouseStateFilter,    ! IOSt .l .p)
getWindowKeyboardStateFilter:: !Id                          !(IOSt .l .p)
                            -> (!Maybe KeyboardStateFilter, ! IOSt .l .p)
setWindowMouseStateFilter   :: !Id !MouseStateFilter        !(IOSt .l .p)
                                                           -> IOSt .l .p
setWindowKeyboardStateFilter:: !Id !KeyboardStateFilter     !(IOSt .l .p)
                                                           -> IOSt .l .p
/*  getWindow(Mouse/Keyboard)StateFilter yields the current
        (Mouse/Keyboard)StateFilter of the indicated window. Nothing is yielded in
        case the window does not exist or has no Window(Mouse/Keyboard) attribute.
    setWindow(Mouse/Keyboard)StateFilter replaces the current
        (Mouse/Keyboard)StateFilter of the indicated window. If the indicated window
        does not exist the function has no effect.
*/


drawInWindow    :: !Id ![DrawFunction]      !(IOSt .l .p) -> IOSt .l .p
/*  drawInWindow applies the list of drawing functions in left-to-right order to the
    picture of the indicated window (behind all controls).
    drawInWindow has no effect in case the window is unknown or is a Dialog.
*/

updateWindow    :: !Id !(Maybe ViewFrame)   !(IOSt .l .p) -> IOSt .l .p
/*  updateWindow applies the WindowLook attribute function of the indicated window.
    The SelectState argument of the Look attribute is the current SelectState of the
    window.
    The UpdateState argument of the Look attribute is
        {oldFrame=frame,newFrame=frame,updArea=[frame]}
    where frame depends on the optional ViewFrame argument:
        in case of (Just rectangle):
            the intersection of the current ViewFrame of the window and rectangle.
        in case of Nothing:
            the current ViewFrame of the window.
    updateWindow has no effect in case of unknown windows, or if the indicated
    window is a Dialog, or the window has no WindowLook attribute, or the optional
    viewframe argument is empty.
*/

setWindowLook   :: !Id !Bool !Look  !(IOSt .l .p) -> IOSt .l .p
getWindowLook   :: !Id              !(IOSt .l .p) -> (!Maybe Look,!IOSt .l .p)
/*  setWindowLook sets the Look of the indicated window.
        The window is redrawn only if the Boolean argument is True.
        setWindowLook has no effect in case the window does not exist, or is a
        Dialog.
    getWindowLook returns the (Just Look) of the indicated window.
        In case the window does not exist, or is a Dialog, or has no WindowLook
        attribute, the result is Nothing.
*/


setWindowPos    :: !Id !ItemPos !(IOSt .l .p) -> IOSt .l .p
getWindowPos    :: !Id          !(IOSt .l .p) -> (!Maybe ItemOffset,!IOSt .l .p)
/*  setWindowPos places the window at the indicated position.
        If the ItemPos argument refers to the Id of an unknown window (in case of
```

```
        LeftOf/RightTo/Above/Below), setWindowPos has no effect.
        If the ItemPos argument is one of (LeftOf/RightTo/Above/Below)Prev, then the
        previous window is the window that is before the window in the current
        stacking order.
        If the window is frontmost, setWindowPos has no effect. setWindowPos also
        has no effect if the window would be moved outside the screen, or if the Id
        is unknown or refers to a modal Dialog.
    getWindowPos returns the current item offset position of the indicated window.
        The corresponding ItemPos is (LeftTop,offset). Nothing is returned in case
        the window does not exist.
*/


moveWindowViewFrame :: !Id Vector !(IOSt .l .p) -> IOSt .l .p
/*  moveWindowViewFrame moves the orientation of the view frame of the indicated
    window over the given vector, and updates the window if necessary. The view
    frame is not moved outside the ViewDomain of the window.
    In case of unknown Id, or of Dialogs, moveWindowViewFrame has no effect.
*/

getWindowViewFrame  :: !Id !(IOSt .l .p) -> (!ViewFrame,!IOSt .l .p)
/*  getWindowViewFrame returns the current view frame of the window in terms of the
    ViewDomain. Note that in case of a Dialog, getWindowViewFrame returns
    {zero,size}.
    In case of unknown windows, the ViewFrame result is zero.
*/

setWindowViewSize   :: !Id Size !(IOSt .l .p) -> IOSt .l .p
getWindowViewSize   :: !Id      !(IOSt .l .p) -> (!Size,!IOSt .l .p)
/*  setWindowViewSize
        sets the size of the view frame of the indicated window as given, and
        updates the window if necessary. The size is fit between the minimum size
        and the screen dimensions.
        In case of unknown Ids, or of Dialogs, setWindowViewSize has no effect.
    getWindowViewSize yields the current size of the view frame of the indicated
        window. If the window does not exist, zero is returned.
*/

setWindowViewDomain :: !Id ViewDomain   !(IOSt .l .p) -> IOSt .l .p
getWindowViewDomain :: !Id              !(IOSt .l .p)
                    -> (!Maybe ViewDomain,!IOSt .l .p)
/*  setWindowViewDomain
        sets the view domain of the indicated window as given. The window view frame
        is moved such that a maximum portion of the view domain is visible. The
        window is not resized.
        In case of unknown Ids, or of Dialogs, setWindowViewDomain has no effect.
    getWindowViewDomain
        returns the current ViewDomain of the indicated window.
        Nothing is returned in case the window does not exist or is a Dialog.
*/


setWindowTitle  :: !Id Title       !(IOSt .l .p) -> IOSt .l .p
setWindowOk     :: !Id Id          !(IOSt .l .p) -> IOSt .l .p
setWindowCancel :: !Id Id          !(IOSt .l .p) -> IOSt .l .p
setWindowCursor :: !Id CursorShape !(IOSt .l .p) -> IOSt .l .p
getWindowTitle  :: !Id             !(IOSt .l .p) -> (!Maybe Title,      !IOSt .l .p)
getWindowOk     :: !Id             !(IOSt .l .p) -> (!Maybe Id,         !IOSt .l .p)
getWindowCancel :: !Id             !(IOSt .l .p) -> (!Maybe Id,         !IOSt .l .p)
getWindowCursor :: !Id             !(IOSt .l .p) -> (!Maybe CursorShape,!IOSt .l .p)
/*  setWindow(Title/Ok/Cancel/Cursor) set the indicated window attributes.
        In case of unknown Ids, these functions have no effect.
    getWindow(Title/Ok/Cancel/Cursor) get the indicated window attributes.
        In case of unknown Ids, the result is Nothing.
*/
```

A.40 StdWindowDef

```
definition module StdWindowDef


//  ********************************************************************************
//  Clean Standard Object I/O library, version 1.1
//
//  StdWindowDef contains the types to define the standard set of windows and
//  dialogues.
//  ********************************************************************************


import  StdControlDef


::  Window c ls ps = Window Title (c ls ps) [WindowAttribute *(ls,ps)]
::  Dialog c ls ps = Dialog Title (c ls ps) [WindowAttribute *(ls,ps)]

::  WindowAttribute ps                          // Default:
//  Attributes for Windows and Dialogs:
    =   WindowId            Id                  // system defined id
    |   WindowPos           ItemPos             // system dependent
    |   WindowIndex         Int                 // open front-most
    |   WindowSize          Size                // screen size
    |   WindowHMargin       Int Int             // system dependent
    |   WindowVMargin       Int Int             // system dependent
    |   WindowItemSpace     Int Int             // system dependent
    |   WindowOk            Id                  // no default (Custom)ButtonControl
    |   WindowCancel        Id                  // no cancel  (Custom)ButtonControl
    |   WindowHide                              // initially visible
    |   WindowClose         (IOFunction ps)     // user can't close window
    |   WindowInit          [IdFun ps]          // no actions after opening window
//  Attributes for Windows only:
    |   WindowSelectState   SelectState         // Able
    |   WindowLook          Look                // show system dependent background
    |   WindowViewDomain    ViewDomain          // {zero,max range}
    |   WindowOrigin        Point               // left top of picture domain
    |   WindowHScroll       ScrollFunction      // no horizontal scrolling
    |   WindowVScroll       ScrollFunction      // no vertical   scrolling
    |   WindowMinimumSize   Size                // system dependent
    |   WindowResize                            // fixed size
    |   WindowActivate      (IOFunction ps)     // id
    |   WindowDeactivate    (IOFunction ps)     // id
    |   WindowMouse         MouseStateFilter    SelectState (MouseFunction    ps)
                                                // no mouse input
    |   WindowKeyboard      KeyboardStateFilter SelectState (KeyboardFunction ps)
                                                // no keyboard input
    |   WindowCursor        CursorShape         // no change of cursor
::  CursorShape
    =   StandardCursor
    |   BusyCursor
    |   IBeamCursor
    |   CrossCursor
    |   FatCrossCursor
    |   ArrowCursor
    |   HiddenCursor

::  WindowType
    :== String
```

Index