

FORMAL SPECIFICATION OF WINDOW SYSTEMS

by
Jonathan Bowen

Technical Monograph PRG-74
ISBN 0-902928-56-2

June, 1989

Oxford University Computing Laboratory
Programming Research Group
8-11 Keble Road
Oxford OX1 3QD
England

Copyright © 1989 Jonathan Bowen

Oxford University Computing Laboratory

Programming Research Group

8-11 Keble Road

Oxford OX1 3QD

England

Electronic mail: `bowen@uk.ac.oxford.prg` (JANET)

**FORMAL
SPECIFICATION
OF
WINDOW
SYSTEMS**



Jonathan Bowen

X

Blit

example

WM

To Alice, Emma and Jane

“Oh! I do so wish I could see *that* bit!”

Lewis Carroll, *Through the Looking-Glass*

Formal Specification of Window Systems

Jonathan Bowen

Summary

Window management systems are now used extensively for user interfaces to computer systems. Part I of this monograph introduces some of the fundamental ideas in window systems using a formal notation. Part II outlines three real systems and attempts to capture the essence of each system using the same formal notation and ideas introduced in Part I. Low-level detail is avoided to keep the length to a manageable size.

In Part I, chapter 1 introduces general concepts useful for specifying pixel maps and window systems. Chapter 2 defines the raster-op function which is fundamental to many graphics operations and chapter 3 introduces a simple example window system. In Part II, chapters 4-6 detail three particular window systems. Chapter 7 remarks on experience gained by formally specifying the three window systems.

The formal notation used, **Z**, is based on set theory, and has been developed at the Programming Research Group in Oxford.

Contents

I	An Introduction to Window Systems	1
1	Basic Concepts	3
1.1	Pixel positions	3
1.2	Pixel maps	5
1.3	Windows	7
2	Raster-Op Functions	10
2.1	Operations on Pixel Values	10
2.2	Operations on Pixel Maps	11
2.3	Display Operations	14
2.4	An Example - Swapping Pixel Maps	15
2.5	Conclusion	16
3	An Example Window System	18
3.1	State of the Window System	18
3.2	Operations on Windows	19
3.3	Error Conditions	21
3.4	Status Operations	22
3.5	Upgraded Operations	24
3.6	Conclusion	25

II	A Comparison of Three Real Window Systems	27
	Introduction	29
4	The ITC 'WM' Window Manager (CMU)	30
4.1	State	30
4.2	Window Creation and Deletion	34
4.3	Window Size	35
4.4	Windows Visibility	36
4.5	Other Window Operations	37
4.6	Errors	38
4.7	The ITC Network	40
4.8	Simplifications and Assumptions	42
4.9	Comments	42
5	Blit Windows (AT&T)	44
5.1	State	44
5.2	New Layers	47
5.3	New Processes	48
5.4	Mouse Operations	50
5.5	The 'mux' Multiplexer	51
5.6	Errors	52
5.7	Simplifications, Assumptions and Comments	53
6	X Window System (MIT)	54
6.1	State	54
6.2	Creating and Destroying Windows	57
6.3	Manipulating Windows	59
6.4	Other Window Operations	64
6.5	Errors	65
6.6	Simplifications and Assumptions	66
6.7	Comments and Inconsistencies	67
7	Conclusions	69
7.1	Comparison of Window Systems	69
7.2	Formal Specification of Existing Systems	70
7.3	General Conclusions	71

<i>CONTENTS</i>	iii
Acknowledgements	72
References	73
Appendices	74
A Glossary of Z notation	77
A.1 Abbreviated names	77
A.2 Horizontal paragraphs	78
A.3 Vertical paragraphs	78
A.4 Declarations and operators	78
A.5 Expressions	78
A.6 Predicates	79
A.7 Schema expressions	79
A.8 Sets	80
A.9 Relations	80
A.10 Functions	81
A.11 Numbers and finiteness	81
A.12 Sequences	82
A.13 Bags	82
A.14 Conventions	83
B Index of formal definitions	84

Part I

An Introduction to Window Systems

Chapter 1

Basic Concepts

Before we start to attempt to specify a window system formally, it is helpful to introduce a few fundamental and generally applicable ideas. This section gives a formal framework to aid the description of *pixels*, their organisation into *pixel maps* and a number of *windows*.

The specification language used throughout this monograph is the **Z** notation [Haye87, King88, Loom88, Spiv88a, Spiv89]. This is a typed language based on set theory and first order predicate calculus. It has been developed at Oxford over the past few years. The notation here conforms almost entirely to the notation described in [Spiv89]. The only extension is the use of the schema piping operator (' \gg ') which is included in [King88].

1.1 Pixel positions

A raster graphics display is made up of a set of pixels with positions or coordinates. These are normally defined in X-Y coordinate space. The display is a fixed size bounded rectangle in the X-Y plane.

| $Xsize, Ysize : \mathbf{N}_1$

The offset in a particular direction is specified from zero up by convention. The position of a pixel may be specified by a pair of X-Y coordinates.

$Xrange == 0 .. Xsize - 1$

$Yrange == 0 .. Ysize - 1$

$Pixel == Xrange \times Yrange$

The pixel at (0,0) is normally at the lower left-hand corner of the display and the pixel at ($Xsize - 1, Ysize - 1$) is at the top right by convention.

Many operations are applied to pairs of pixels.

$\begin{array}{l} \text{PixelPair} \\ \text{pix}_1, \text{pix}_2 : \text{Pixel} \\ \text{x}_1, \text{x}_2 : \text{Xrange} \\ \text{y}_1, \text{y}_2 : \text{Yrange} \\ \text{pix}_1 = (\text{x}_1, \text{y}_1) \\ \text{pix}_2 = (\text{x}_2, \text{y}_2) \end{array}$
--

The '+', '-' and '≤' operators may be overloaded to apply to pixel positions. '+' and '-' may be used for moving pixel areas around the display. '≤' can be used to define pixel ordering from the bottom left to top right.

$\begin{array}{l} - + - \\ - - - : (\text{Pixel} \times \text{Pixel}) \rightarrow \text{Pixel} \\ - \leq - : \text{Pixel} \leftrightarrow \text{Pixel} \end{array}$
$\begin{array}{l} \forall \text{PixelPair} \bullet \\ (\text{x}_1 + \text{x}_2 < \text{Xsize} \wedge \text{y}_1 + \text{y}_2 < \text{Ysize}) \Rightarrow \\ \quad \text{pix}_1 + \text{pix}_2 = (\text{x}_1 + \text{x}_2, \text{y}_1 + \text{y}_2) \wedge \\ (\text{x}_2 \leq \text{x}_1 \wedge \text{y}_2 \leq \text{y}_1) \Rightarrow \\ \quad \text{pix}_1 - \text{pix}_2 = (\text{x}_1 - \text{x}_2, \text{y}_1 - \text{y}_2) \wedge \\ \text{pix}_1 \leq \text{pix}_2 \Leftrightarrow \text{x}_1 \leq \text{x}_2 \wedge \text{y}_1 \leq \text{y}_2 \end{array}$

We can define the offset between any two pixel positions as a pixel offset. This is defined to wrap round the edge of the pixel area and thus is a total function.

$\text{offset} : \text{Pixel} \rightarrow \text{Pixel} \rightarrow \text{Pixel}$
$\begin{array}{l} \forall \text{PixelPair} \bullet \\ \text{offset } \text{pix}_1 \text{ pix}_2 = \\ \quad ((\text{x}_1 + \text{x}_2) \bmod \text{Xsize}, (\text{y}_1 + \text{y}_2) \bmod \text{Ysize}) \end{array}$

We can also overload the '..' operator to define a rectangular area of pixels.

$- \dots - : (\text{Pixel} \times \text{Pixel}) \rightarrow \mathbf{F} \text{Pixel}$
$\begin{array}{l} \forall \text{PixelPair} \bullet \\ \text{pix}_1 \dots \text{pix}_2 = (\text{x}_1 \dots \text{x}_2) \times (\text{y}_1 \dots \text{y}_2) \end{array}$

$$pix_1, pix_2 : Pixel \vdash pix_1 .. pix_2 = \{p : Pixel \mid pix_1 \leq p \wedge p \leq pix_2\}$$

A rectangular area of pixels can be defined using any two opposing corners (e.g. returned using an attached mouse to sweep between the two). The following functions return the lower left and upper right pixel positions from two such pixel positions respectively.

$$\left| \begin{array}{l} - \underline{min} \quad - \\ - \underline{max} \quad - : (Pixel \times Pixel) \rightarrow Pixel \\ \hline \forall PixelPair \bullet \\ \quad pix_1 \underline{min} \quad pix_2 = (min \{x_1, x_2\}, min \{y_1, y_2\}) \wedge \\ \quad pix_1 \underline{max} \quad pix_2 = (max \{x_1, x_2\}, max \{y_1, y_2\}) \end{array} \right.$$

1.2 Pixel maps

A raster graphics display has a number of bit-planes. This may be considered as the Z direction of the display.

$$\mid Zsize : \mathbb{N}_1$$

Each *bit* in a bit-plane has one of two values (cleared or set).

$$\begin{aligned} Clear & == 0 \\ Set & == 1 \\ Bit & == \{Clear, Set\} \end{aligned}$$

The value of a pixel at a particular position may be modelled as a function from bit-plane number to bit value.

$$\begin{aligned} Zrange & == 0 .. Zsize - 1 \\ Value & == Zrange \rightarrow Bit \end{aligned}$$

If all the bits are 'Clear' the 'Value' is considered 'Black' and if they are all 'Set' it is considered 'White'.

$$\begin{aligned} Black & == (\mu val : Value \mid ran \, val = \{Clear\}) \\ White & == (\mu val : Value \mid ran \, val = \{Set\}) \end{aligned}$$

Note that if there is only one bit-plane (i.e. $Zsize = 1$) then pixel values can only be *Black* or *White*.

$$\vdash Zsize = 1 \Rightarrow Value = \{Black, White\}$$

A pixel map consists of a (partial) function from pixel positions to the value of the pixel contents. This can be used to describe part of a display, such as a window.

$$Pizmap == Pixel \leftrightarrow Value$$

Non-empty pixel maps may be of interest.

$$Pizmap_1 == Pizmap \setminus \{\emptyset\}$$

Pixel maps are often rectangular in area. We can define such pixel maps using their bottom left and top right pixel positions.

$$Rectangle == \{map : Pizmap_1 \mid \exists_1 p_1, p_2 : Pixel \bullet \text{dom } map = p_1 .. p_2\}$$

Sometimes it is desirable to set all the range of a pixel map to a particular value, for example when clearing a window down to the background colour. A function to set the range of a relation to a particular value is useful for this.

$\begin{array}{l} [P, V] \\ \hline \text{setval} : V \rightarrow (P \leftrightarrow V) \rightarrow (P \leftrightarrow V) \\ \hline \forall v : Value; p : P \leftrightarrow V \bullet \\ \text{setval } v \text{ } p = (\mu m : P \leftrightarrow V \mid (\text{dom } m = \text{dom } p \wedge \\ \text{ran } m = \{v\})) \end{array}$
--

The following laws apply:

$$\begin{array}{l} p : Pizmap; v : Value \vdash (\text{setval } v)^+ p = \text{setval } v \text{ } p \\ v : Value \vdash \text{setval } v \text{ } \emptyset = \emptyset \end{array}$$

Two pixel maps may overlap. For example, one window may be obscured by another. This can be captured as a relation between pixel maps:

$\begin{array}{l} [P, V] \\ \hline - \text{overlaps} - : (P \leftrightarrow V) \leftrightarrow (P \leftrightarrow V) \\ \hline \forall p_1, p_2 : P \leftrightarrow V \bullet \\ p_1 \text{ overlaps } p_2 \Leftrightarrow \text{dom } p_1 \cap \text{dom } p_2 \neq \emptyset \end{array}$

A sequence of pixel maps may be overlaid in the order given by the sequence. It is convenient to define a distributed overriding operator for this.

$[P, V]$	
$\oplus / :$	$\text{seq}(P \rightarrow V) \rightarrow (P \rightarrow V)$
$\oplus / \langle \rangle =$	\emptyset
$\forall s :$	$\text{seq}(P \rightarrow V); p : P \rightarrow V \bullet$
	$\oplus / (s \wedge \langle p \rangle) = (\oplus / s) \oplus p$

Distributed overriding is particularly useful for defining the view on a screen of a display, given a sequence of possibly overlapping pixel maps.

The following laws apply for the distributed overriding operator:

	\vdash	$\oplus / \emptyset = \emptyset$
$p : \text{Pixmap}$	\vdash	$\oplus / \langle p \rangle = p$
$s_1, s_2 : \text{seq Pixmap}$	\vdash	$\oplus / (s_1 \wedge s_2) = (\oplus / s_1) \oplus (\oplus / s_2)$
$p_1, p_2 : \text{Pixmap}$	\vdash	$\oplus / \langle p_1, p_2 \rangle = p_1 \oplus p_2$
$s : \text{seq Pixmap}$	\vdash	$\oplus / s = (\oplus / (\text{front } s)) \oplus (\text{last } s)$
		$= (\text{head } s) \oplus (\oplus / (\text{tail } s))$
$s : \text{seq Pixmap}$	\vdash	$\text{dom}(\oplus / s) = \text{dom}(\bigcup (\text{ran } s))$
$s : \text{seq Pixmap}$	\vdash	$\text{ran}(\oplus / s) \subseteq \text{ran}(\bigcup (\text{ran } s))$
$s : \text{seq Pixmap}$	\vdash	$\bigcap (\text{ran } s) = \emptyset \Rightarrow \oplus / s = \bigcup (\text{ran } s)$

Note that the first three laws above provide an alternative way of defining ' $\oplus /$ '.

1.3 Windows

A series of windows on a display screen may be viewed as a sequence in which each window is laid on the screen in the order defined by the sequence (bottom first, top last). If some of the windows are removed from a sequence, it is sometimes desirable to 'compact' the remaining windows into a sequence again.

$[W]$	
$\text{compact} :$	$(\mathbb{N} \rightarrow W) \rightarrow \text{seq } W$
$\text{compact } \emptyset =$	$\langle \rangle$
$\forall f : \mathbb{N} \rightarrow W; i : \mathbb{N} \mid i = \min(\text{dom } f) \bullet$	
	$\text{compact } f = \langle f(i) \rangle \wedge \text{compact}(\{i\} \triangleleft f)$

If the windows in a sequence overlap, it is useful to be able to move selected windows so that their contents may be viewed (or hidden). This is

analogous to shuffling a pile of sheets of paper (windows) on a desk (screen). Note that the sheets of paper may be of different sizes and in different positions on the desk.

For example, the following function may be used to move a selected window number in the sequence (if it exists) to the top of the pile (i.e. the end of the sequence).

$[W]$
$top : \mathbf{N} \rightarrow \text{seq } W \rightarrow \text{seq } W$
$\forall n : \mathbf{N}; s : \text{seq } W \bullet$
$n \in \text{dom } s \Rightarrow$
$top \ n \ s = compact(\{n\} \triangleleft s) \hat{\ } \langle s(n) \rangle \wedge$
$n \notin \text{dom } s \Rightarrow$
$top \ n \ s = s$

More generally, we can define functions to ‘select’ and ‘remove’ a set of windows from a sequence using their identifiers rather than their position in the pile.

$[W]$
$- \underline{select} \ _ : \text{seq } W \rightarrow \text{seq } W$
$- \underline{remove} \ _ : \text{seq } W \times \mathbf{P} W \rightarrow \text{seq } W$
$\forall s : \text{seq } W; w : \mathbf{P} W \bullet$
$s \underline{select} \ w = compact \ (s \triangleright w) \wedge$
$s \underline{remove} \ w = compact \ (s \triangleright w)$

We can then ‘raise’ or ‘lower’ these windows to the end or beginning of the sequence as required.

$[W]$
$- \underline{raise} \ _ : \text{seq } W \rightarrow \text{seq } W$
$- \underline{lower} \ _ : \text{seq } W \times \mathbf{P} W \rightarrow \text{seq } W$
$\forall s : \text{seq } W; w : \mathbf{P} W \bullet$
$s \underline{raise} \ w = (s \underline{remove} \ w) \hat{\ } (s \underline{select} \ w) \wedge$
$s \underline{lower} \ w = (s \underline{select} \ w) \hat{\ } (s \underline{remove} \ w)$

Every window in a system usually has an identifier, denoted ‘Window’, which allows it to be accessed uniquely.

[*Window*]

The generic functions defined in this section will normally be applied to such identifiers.

Windows often contain text. Thus, a text string is needed sometimes (e.g. for a title of a window). This is denoted as '*String*'. The string may be empty.

[*String*]

| '': *String*

This concludes the basic definitions which will be used as required in the rest of this monograph.

Chapter 2

Raster-Op Functions

Raster-op functions are useful when moving areas of pixels (e.g. parts of windows) around the screen on graphics display systems. Raster-op is now widely used in graphics systems, in particular for window systems. This chapter formally specifies raster-op functions and gives an example of its use. Readers not familiar with 'raster-op' may prefer to do some background reading first [Fole82, Newm81].

2.1 Operations on Pixel Values

Given two pixel values, we may perform a bit-wise '*NAND*' ('not and') on the two values to produce a new value. If both bits are set then the result is clear; if either bit is clear then the result is set.

$$\begin{array}{|l} \hline \text{-- } \underline{\text{NAND}} \text{ --} : (\text{Value} \times \text{Value}) \rightarrow \text{Value} \\ \hline \forall \text{val}_1, \text{val}_2, \text{val} : \text{Value} \bullet \\ \quad \text{val}_1 \underline{\text{NAND}} \text{val}_2 = \text{val} \Leftrightarrow \\ \quad (\forall n : \text{Zrange} \bullet \\ \quad \quad (\text{val}_1 n = \text{Set} \wedge \text{val}_2 n = \text{Set}) \Rightarrow \\ \quad \quad \quad \text{val } n = \text{Clear} \wedge \\ \quad \quad (\text{val}_1 n = \text{Clear} \vee \text{val}_2 n = \text{Clear}) \Rightarrow \\ \quad \quad \quad \text{val } n = \text{Set}) \\ \hline \end{array}$$

All the other binary and unary logical functions may be defined in terms of this function. For example, we can define a unary '*NOT*' function.

$$\begin{array}{|l} \hline \underline{\text{NOT}} : \text{Value} \rightarrow \text{Value} \\ \hline \forall \text{val} : \text{Value} \bullet \\ \quad \text{NOT } \text{val} = \text{val} \underline{\text{NAND}} \text{val} \\ \hline \end{array}$$

We can also define binary ‘*AND*’, ‘*OR*’ and ‘*XOR*’ functions.

$$\begin{array}{|l}
 \hline
 - \underline{\textit{AND}} \rightarrow, \\
 - \underline{\textit{OR}} \rightarrow, \\
 - \underline{\textit{XOR}} \rightarrow : (\textit{Value} \times \textit{Value}) \rightarrow \textit{Value} \\
 \hline
 \forall \textit{val}_1, \textit{val}_2 : \textit{Value} \bullet \\
 \quad \textit{val}_1 \underline{\textit{AND}} \textit{val}_2 = \textit{NOT}(\textit{val}_1 \underline{\textit{NAND}} \textit{val}_2) \wedge \\
 \quad \textit{val}_1 \underline{\textit{OR}} \textit{val}_2 = (\textit{NOT} \textit{val}_1) \underline{\textit{NAND}} (\textit{NOT} \textit{val}_2) \wedge \\
 \quad \textit{val}_1 \underline{\textit{XOR}} \textit{val}_2 = (\textit{val}_1 \underline{\textit{OR}} \textit{val}_2) \underline{\textit{AND}} (\textit{val}_1 \underline{\textit{NAND}} \textit{val}_2) \\
 \hline
 \end{array}$$

2.2 Operations on Pixel Maps

A pixel map is considered to have the same shape as another if it can be ‘moved’ using a unique one-to-one function (*offset pix₀*) in the definition below) to give it the same domain. Intuitively this implies that each of the pixel maps could be moved en masse about the domain space so that its domain is exactly the same as the other map.

$$\begin{array}{|l}
 \hline
 - \underline{\textit{sameshape}} \rightarrow : \textit{Pixmap} \leftrightarrow \textit{Pixmap} \\
 \hline
 \forall \textit{map}_1, \textit{map}_2 : \textit{Pixmap} \bullet \\
 \quad \textit{map}_1 \underline{\textit{sameshape}} \textit{map}_2 \Leftrightarrow \\
 \quad (\exists \textit{pix}_0 : \textit{Pixel} \bullet \\
 \quad \quad \textit{dom}(\textit{offset} \textit{pix}_0 ; \textit{map}_1) = \textit{dom} \textit{map}_2) \\
 \hline
 \end{array}$$

Consider a pair of (possibly overlapping) pixel maps which have the same shape.

$$\begin{array}{|l}
 \hline
 \textit{MapPair} \\
 \hline
 \textit{map}_1, \\
 \textit{map}_2 : \textit{Pixmap} \\
 \hline
 \textit{map}_1 \underline{\textit{sameshape}} \textit{map}_2 \\
 \hline
 \end{array}$$

Two pixels in a pair of pixel maps with the same shape are considered to have the same position if the same offset function which relates the two maps also relates the two pixels. This is captured in ‘*samespos*’ relation below. If the relation is true then it implies that the two pixels are in the same relative position within two pixel maps with the same shape. That is to say, if one of the pixel maps were to be moved on top of the other then one pixel would be exactly on top of the other.

$$\begin{array}{l}
 \hline
 _ \text{ samepos } _ : (\text{Pixel} \times \text{Pixmap}) \leftrightarrow (\text{Pixel} \times \text{Pixmap}) \\
 \forall \text{pix}_1, \text{pix}_2 : \text{Pixel}; \text{MapPair} \mid \\
 \text{pix}_1 \in \text{dom } \text{map}_1 \wedge \text{pix}_2 \in \text{dom } \text{map}_2 \bullet \\
 (\text{pix}_1, \text{map}_1) \text{ samepos } (\text{pix}_2, \text{map}_2) \Leftrightarrow \\
 (\exists \text{pix}_0 : \text{Pixel} \bullet \\
 \text{dom}((\text{offset } \text{pix}_0); \text{map}_1) = \text{dom } \text{map}_2 \wedge \\
 \text{pix}_1 \mapsto \text{pix}_2 \in \text{offset } \text{pix}_0)
 \end{array}$$

We may define operations similar to the bit-wise operations on values to apply to pixel maps. In this case by convention, the last operand of the function has the same domain as the result of the function. We start by defining the ‘*nand*’ function.

$$\begin{array}{l}
 \hline
 _ \text{ nand } _ : (\text{Pixmap} \times \text{Pixmap}) \leftrightarrow \text{Pixmap} \\
 \forall \text{MapPair}; \text{map} : \text{Pixmap} \bullet \\
 \text{map}_1 \text{ nand } \text{map}_2 = \text{map} \Leftrightarrow \\
 \text{dom } \text{map} = \text{dom } \text{map}_2 \wedge \\
 (\forall \text{pix}_1 : \text{dom } \text{map}_1; \text{pix}_2 : \text{dom } \text{map}_2 \mid \\
 (\text{pix}_1, \text{map}_1) \text{ samepos } (\text{pix}_2, \text{map}_2) \bullet \\
 \text{map}_1(\text{pix}_1) \text{ NAND } \text{map}_2(\text{pix}_2) = \text{map } \text{pix}_2)
 \end{array}$$

Note that this operation is not commutative, unlike ‘*NAND*’ since the second operand defines the domain of the result of the function. Specifically, if the domains of the pixel maps differ, the ordering is important. If the domains are the same then the ordering is unimportant.

$$\begin{array}{l}
 \text{MapPair} \vdash \text{dom } \text{map}_1 = \text{dom } \text{map}_2 \Rightarrow \\
 \text{map}_1 \text{ nand } \text{map}_2 = \text{map}_2 \text{ nand } \text{map}_1
 \end{array}$$

Using the basic ‘*nand*’ function, we may define fifteen further operations. Four of these degenerate to monadic functions. ‘*noop*’ leaves a pixel map unaffected, ‘*not*’ inverts all bits, ‘*clear*’ clears all the bits and ‘*set*’ sets all the bits.

$noop,$ $not,$ $clear,$ $set : Pixmap \rightarrow Pixmap$
$\forall map : Pixmap \bullet$ $noop\ map = map \wedge$ $not\ map = map \underline{and}\ map \wedge$ $clear\ map = map \underline{nand}\ (not\ map) \wedge$ $set\ map = not(clear\ map)$

There are five more important operators. These correspond to standard logical operations except 'copy' which is extremely useful for moving pixel maps.

$- \underline{and} \rightarrow,$ $- \underline{or} \rightarrow,$ $- \underline{xor} \rightarrow,$ $- \underline{nor} \rightarrow,$ $- \underline{copy} \rightarrow : (Pixmap \times Pixmap) \rightarrow Pixmap$
$\forall MapPair \bullet$ $map_1 \underline{and} map_2 =$ $not(map_1 \underline{nand} map_2) \wedge$ $map_1 \underline{or} map_2 =$ $(not\ map_1) \underline{nand}\ (not\ map_2) \wedge$ $map_1 \underline{xor} map_2 =$ $(map_1 \underline{or} map_2) \underline{and}\ (map_1 \underline{nand} map_2) \wedge$ $map_1 \underline{nor} map_2 =$ $not(map_1 \underline{or} map_2) \wedge$ $map_1 \underline{copy} map_2 =$ $map_1 \underline{or}\ (clear\ map_2)$

The rest of the operations are used less often but are detailed here for completeness.

- copyInverted \rightarrow ,
 - andReverse \rightarrow ,
 - andInverted \rightarrow ,
 - orReverse \rightarrow ,
 - orInverted \rightarrow ,
 - equiv \rightarrow : $(Pizmap \times Pizmap) \rightarrow Pizmap$

$\forall MapPair \bullet$

map_1 copyInverted $map_2 = (not\ map_1) \text{ or } (clear\ map_2) \wedge$

map_1 andReverse $map_2 = map_1 \text{ and } (not\ map_2) \wedge$

map_1 andInverted $map_2 = (not\ map_1) \text{ and } map_2 \wedge$

map_1 orReverse $map_2 = map_1 \text{ or } (not\ map_2) \wedge$

map_1 orInverted $map_2 = (not\ map_1) \text{ or } map_2 \wedge$

map_1 equiv $map_2 = (not\ map_1) \text{ xor } map_2$

This covers the sixteen possible raster-op boolean functions on two values.

2.3 Display Operations

A display screen consists of a pixel map.

$Display$ $screen : Pizmap$

During changes to the screen, its size does not change although the pixel values displayed on the screen may be updated.

$\Delta Display$ $Display$ $Display'$
$dom\ screen' = dom\ screen$

The screen may be updated using one of the raster-op functions previously defined. Some functions require a source and destination area whilst others degenerate into a single area.

$RasterOp_1$ $\Delta Display$ $area? : P Pixel$ $op? : Pixmap \rightarrow Pixmap$ $screen' = screen \oplus op?(area? \triangleleft screen)$

$RasterOp_2$ $\Delta Display$ $area? : P Pixel$ $from? : Pixel$ $op? : (Pixmap \times Pixmap) \rightarrow Pixmap$ $screen' = screen \oplus$ $op? (((offset from?) \setminus area?) \triangleleft screen , area? \triangleleft screen)$
--

2.4 An Example – Swapping Pixel Maps

Consider two separate non-overlapping pixel maps with the same shape.

$SepPair$ $MapPair$ $\neg map_1 \underline{overlaps} map_2$

$$\Delta SepPair \cong SepPair \wedge SepPair'$$

We may swap a pair of pixel maps with the same shape using the ‘copy’ operation.

$CopySwap$ $\Delta SepPair$ $map'_1 = map_2 \underline{copy} map_1$ $map'_2 = map_1 \underline{copy} map_2$
--

In practice, these two ‘copy’ operations cannot be carried out simultaneously. A third copy is necessary and additionally a temporary pixel map area is required. This can easily be expressed by using three schemas, one for each operation, and then combining them using the schema composition operator (;). This is left as an exercise for the reader.

Alternatively, the ‘xor’ raster-op function may be used. Three sequential operations are still necessary, but the use of a temporary buffer area is eliminated. The following two schemas ‘xor’ one or other of a pair of pixel maps with its opposite number.

$$\begin{array}{|l}
 \hline
 \text{Xor}_1 \\
 \hline
 \Delta \text{SepPair} \\
 \hline
 \text{map}'_1 = \text{map}_2 \text{ xor } \text{map}_1 \\
 \hline
 \text{map}'_2 = \text{map}_2 \\
 \hline
 \end{array}$$

$$\begin{array}{|l}
 \hline
 \text{Xor}_2 \\
 \hline
 \Delta \text{SepPair} \\
 \hline
 \text{map}'_1 = \text{map}_1 \\
 \hline
 \text{map}'_2 = \text{map}_1 \text{ xor } \text{map}_2 \\
 \hline
 \end{array}$$

A swap may be achieved between the two pixel maps by applying to-kenxor three times sequentially as follows.

$$\text{XorSwap} \hat{=} \text{Xor}_1 ; \text{Xor}_2 ; \text{Xor}_1$$

By symmetry, the following is also true.

$$\Delta \text{SepPair} \vdash \text{XorSwap} \Leftrightarrow (\text{Xor}_2 ; \text{Xor}_1 ; \text{Xor}_2)$$

The *XorSwap* operation has exactly the same effect as using the ‘copy’ operator twice simultaneously, as in the *CopySwap* operation.

$$\Delta \text{SepPair} \vdash \text{CopySwap} \Leftrightarrow \text{XorSwap}$$

2.5 Conclusion

We have formally specified the raster-op function and how this may be applied to a graphics display. We have given an example of its use for swapping

areas of a display. By defining operations performed on rectangular areas, we have specified some of the underlying operations necessary for a window system.

Chapter 3

An Example Window System

Window management systems are now used extensively for user interfaces to computer systems. This chapter outlines an example window system using formal specification techniques. An abstract view of the system is developed and some basic operations are given.

3.1 State of the Window System

The window display may be modelled as a sequence of windows against a background 'window' which is the size of the display itself. The order of the sequence defines which windows are on top in the case of overlapping windows, in ascending order. Only parts of windows which are contained within the background area are displayed.

$$\begin{array}{l} \text{SYS} \\ \hline \text{windows} : \text{seq Pixmap} \\ \text{screen, background} : \text{Pixmap} \\ \hline \text{screen} = \text{background} \oplus (\text{dom background} \triangleleft \oplus / \text{windows}) \end{array}$$

In this simple example, we shall ignore the complication of window identifiers. We shall simply use the position of the window in the sequence to identify it, assuming that the user of the system keeps track of which window is which.

Note that the user can only see the display screen.

$$\text{View} \hat{=} \text{SYS} \setminus (\text{windows, background})$$

Initially there are no windows and thus only the background is displayed.

$$\text{InitSYS} \equiv [\text{SYS}' \mid \text{windows}' = \langle \rangle]$$

$$\text{InitSYS} \vdash \text{screen}' = \text{background}'$$

During changes to the display, the background and hence the size of the display always remain the same.

$$\Delta\text{SYS} \equiv [\text{SYS}; \text{SYS}' \mid \text{background}' = \text{background}]$$

Some operations may leave the state of the system unchanged, for example during a status operation or if an error in the input is detected.

$$\exists\text{SYS} \equiv [\Delta\text{SYS} \mid \theta\text{SYS}' = \theta\text{SYS}]$$

3.2 Operations on Windows

In order to use the system, we must have the ability to create windows. These are created on top of all the existing windows. We specify that they must fit within the display background.

AddWindow_0
ΔSYS
$\text{window?} : \text{Pixmap}$
$\text{dom window?} \subseteq \text{dom background}$
$\text{windows}' = \text{windows} \cup \{ \text{window?} \}$

The ability to update windows is very useful. This may involve changing the size of the window or its contents or moving it about the screen. Again, the updated window must still fit within the display area.

UpdateWindow_0
ΔSYS
$\text{which?} : \mathbf{N}$
$\text{window?} : \text{Pixmap}$
$\text{which?} \in \text{dom windows}$
$\text{dom window?} \subseteq \text{dom background}$
$\text{windows}' = \text{windows} \oplus \{ \text{which?} \mapsto \text{window?} \}$

It is desirable to be able to uncover a window which is partially or even totally obscured by other windows. This can be done by moving the window to the end of the sequence of displayed windows.

<i>ExposeWindow</i> ₀
ΔSYS
$which? : \mathbf{N}$
$which? \in \text{dom windows}$
$windows' = \text{top } which? \text{ windows}$

Sometimes it is useful to simply rotate the order of the displayed windows, one at a time, moving the bottommost window to the top.

<i>RotateWindows</i> ₀
ΔSYS
$windows \neq \langle \rangle$
$windows' = \text{top } 1 \text{ windows}$

We also wish to be able to delete windows. For instance, we could delete the topmost window.

<i>RemoveTop</i> ₀
ΔSYS
$windows \neq \langle \rangle$
$windows' = \text{front windows}$

Alternatively, we may wish to specify which window is to be removed.

<i>RemoveWindow</i> ₀
ΔSYS
$which? : \mathbf{N}$
$which? \in \text{dom windows}$
$windows' = \text{compact}(\{which?\} \triangleleft windows)$

3.3 Error Conditions

The operations covered so far detail what should happen in the event of no errors. We wish to be able to report the status of an operation.

$$\begin{array}{l} \textit{Report} ::= \text{'OK'} \\ \quad | \text{'Not a window'} \\ \quad | \text{'No windows'} \\ \quad | \text{'Invalid window'} \end{array}$$

We must report the fact that the operation was successful if this is the case.

<i>Success</i>
$\textit{rep!} : \textit{Report}$
$\textit{rep!} = \text{'OK'}$

If errors do occur, then these need to be reported as well. For example, an invalid window may be specified.

<i>NotAWindow</i>
$\exists \textit{SYS}$
$\textit{which?} : \mathbb{N}$
$\textit{rep!} : \textit{Report}$
$\textit{which?} \notin \text{dom windows}$
$\textit{rep!} = \text{'Not a window'}$

It is possible that there are no windows displayed when one is required.

<i>NoWindows</i>
$\exists \textit{SYS}$
$\textit{rep!} : \textit{Report}$
$\textit{windows} = \langle \rangle$
$\textit{rep!} = \text{'No windows'}$

A specified window may not be within the background area.

$BadWindow$ $\exists SYS$ $window? : Pixmap$ $rep! : Report$
$\neg (\text{dom } window? \subseteq \text{dom } background)$ $rep! = \text{'Invalid window'}$

We may include these errors with the previously defined operations which ignored error conditions, to produce total operations.

$$\begin{aligned}
 AddWindow_1 &\cong (AddWindow_0 \wedge Success) \\
 &\quad \vee BadWindow \\
 UpdateWindow_1 &\cong (UpdateWindow_0 \wedge Success) \\
 &\quad \vee BadWindow \vee NotAWindow \\
 ExposeWindow_1 &\cong (ExposeWindow_0 \wedge Success) \\
 &\quad \vee NotAWindow \\
 RotateWindows_1 &\cong (RotateWindows_0 \wedge Success) \\
 &\quad \vee NoWindows \\
 RemoveTop_1 &\cong (RemoveTop_0 \wedge Success) \\
 &\quad \vee NoWindows \\
 RemoveWindow_1 &\cong (RemoveWindow_0 \wedge Success) \\
 &\quad \vee NotAWindow
 \end{aligned}$$

3.4 Status Operations

The contents of an existing window may be of interest.

$GetWindow_0$ SYS $which? : N$ $window! : Pixmap$
$which? \in \text{dom } windows$ $window! = windows\ which?$

We can make this operation total as well.

$$\text{GetWindow} \hat{=} (\text{GetWindow}_0 \wedge \text{Success}) \vee \text{NotAWindow}$$

It is useful to know which window a given pixel position is in, for example when pointing at a position on the screen with a mouse driven cursor. Given the position, the sequence number of the window is required. There may be several windows at the specified position but we are only interested in the topmost one.

<i>WhichWindow</i> ₀
SYS <i>position?</i> : Pixel <i>which!</i> : \mathbf{N}
<i>position?</i> $\in \text{dom}(\oplus/\text{windows})$ <i>which!</i> = $\max \{n : \text{dom windows} \mid$ <i>position?</i> $\in \text{dom}(\text{windows } n)\}$

However, it is possible that the cursor position may not be in a valid window. In this case we simply record this fact by making '*which!*' invalid. This is because this schema will be combined with others to produce an appropriate error message later.

$$\text{Invalid} == 0$$

<i>NotInWindow</i>
SYS <i>position?</i> : Pixel <i>which!</i> : \mathbf{N}
<i>position?</i> $\notin \text{dom}(\oplus/\text{windows})$ <i>which!</i> = <i>Invalid</i>

To make the operation total, we combine these two operations.

$$\text{WhichWindow} \hat{=} \text{WhichWindow}_0 \vee \text{NotInWindow}$$

3.5 Upgraded Operations

Often when a window is updated, we wish to modify the original window rather than completely replacing it. Sometimes we wish to change the shape of the window by either enlarging it or reducing it, or a combination of the two. Any new area exposed may be given any desired background pixel values.

<i>Reshape</i> <i>area?</i> , <i>window?</i> , <i>window!</i> : <i>Pixmap</i>
$window! = area? \oplus (\text{dom } area?) \triangleleft window?$

This schema is a general operation concerning the input and output of a window and hence the change in state of the window system need not be recorded.

Often we want to move a window. We can define a one-to-one mapping to do this. If the window cannot be moved in its entirety using the mapping given, then it is left unaffected.

<i>Move</i> <i>mapping?</i> : <i>Pixel</i> \leftrightarrow <i>Pixel</i> <i>window?</i> , <i>window!</i> : <i>Pixmap</i>
$\text{dom } window? \subseteq \text{ran } mapping? \Rightarrow$ $window! = mapping? ; window?$
$\neg (\text{dom } window? \subseteq \text{ran } mapping?) \Rightarrow$ $window! = window?$

At other times we simply wish to modify the contents of the window. We ignore any modifications outside the area of the window.

<i>Modify</i> <i>mods?</i> , <i>window?</i> , <i>window!</i> : <i>Pixmap</i>
$window! = window? \oplus (\text{dom } window?) \triangleleft mods?$

Again the two previous schemas have nothing to do with the state of the window system. However by combining the last three schemas with a record of the change of state, we can then combine them with previous operations to produce some new window operations.

$$\text{Reshape Window} \cong \text{GetWindow} \gg \text{Reshape} \gg \text{UpdateWindow}_0$$

$$\text{Move Window} \cong \text{GetWindow} \gg \text{Move} \gg \text{UpdateWindow}_0$$

$$\text{Modify Window} \cong \text{GetWindow} \gg \text{Modify} \gg \text{UpdateWindow}_0$$

Each of these operations selects a particular window from the system, transforms the window in some way, and then updates the system with the new window.

To bring a window to the top, we point at it using the mouse cursor and click a button to execute the operation. This can be modelled using a combination of previously defined operations.

$$\text{Expose Window} \cong \text{Which Window} \gg \text{Expose Window}_1$$

We may want a process associated with a particular window to decide whether a given operation may be executed, using some unknown protection criterion.

Decide

allowed? : P N

which?,

which! : N

which? \in *allowed?* \Rightarrow *which!* = *which?*

which? \notin *allowed?* \Rightarrow *which!* = *Invalid*

Thus a process may decide whether a user is allowed to remove a window.

$$\text{Remove Window} \cong \text{Which Window} \gg \text{Decide} \gg \text{Remove Window}_1$$

3.6 Conclusion

We have introduced a model for a simple window system together with some basic operations which may be performed. We have shown how some of these operations may be combined to produce a set of useful windowing

operations. We have used the schema piping operator (' \gg ') to combine some of the operations rather than the more normal method of using schema inclusion since this gives a better idea of how the operations might be joined together in practice.

Nothing has been said about the nature of pixels in this chapter. Little has been said about the contents of windows. However these issues could be expanded using the model given to produce a complete specification for a windowing system.

Part II

A Comparison of Three Real Window Systems

Introduction

In Part II of this monograph, **Z** [Haye87, King88, Loom88, Spiv88a, Spiv89] is used to describe parts of three window systems, **WM** [Neuw85, Rose85], the **Blit** [Pike84, Unix85] and **X** [Gett85, Gett86, Sche86]. These systems have been developed independently at Carnegie-Mellon University, AT&T Bell Laboratories at Murray Hill and Massachusetts Institute of Technology. All the systems use bit-mapped raster displays with a keyboard and mouse for user input.

A high level description of the state of each system and a selection of window operations is presented to give a flavour of each and to allow them to be compared. The operations covered are those available to applications programs via library procedure calls. In general, similar operations are available to the user under mouse control. Some background reading may be helpful for readers not familiar with windowing techniques [Fole82, Newm81].

Only operations directly concerned with windowing are covered. The systems described also include many other operations such as graphics, raster-op, mouse and cursor routines, etc. Each description covers the following:

- State of the system and initialisation.
- Windowing operations, ignoring error conditions.
- Error conditions and status reports.
- Simplifications and assumptions in the specification.
- General comments on the system.

Menu facilities are not covered in detail for the sake of brevity. Additionally, the parameters supplied to some operations have been simplified to aid clarity in the brief descriptions given.

Graphics facilities within windows are not included. However **Z** is suitable for this. For example, it has been to specify parts of the graphics standard GKS [Arno87].

Chapter 4

The ITC 'WM' Window Manager (CMU)

WM [Neuw85, Rose85], part of the 'Andrew' distributed system, is a window manager developed at the Information Technology Center (ITC) at Carnegie-Mellon University [Saty84, West85]. This runs on UNIX* based workstations designed eventually to be networked on a very large scale (c5,000-10,000 nodes). Because of the distributed file system, any authorised person may also use any other workstation on the network, and indeed create windows on other workstations remotely.

4.1 State

The state of the system is introduced in stages. In this model we consider a single machine for simplicity since we are concerned with how the window manager works rather than how the network operates. Operations over the network will be detailed later.

Each window has a number of pieces of information associated with it. These include a header area for titles and other information, and a separate body area to hold the actual contents of the window. These do not overlap and together they make up the pixel map of the displayed window. In practice, the header is a thin rectangular area just above its associated body.

*UNIX is a trademark of AT&T Bell Laboratories.

*Map**header, body, map : Pixmap**area : P Pixel**(header, body) partition map**area = dom map*

The user can request a window to lie within a specified range of dimensions and can also explicitly ask for a window body to be hidden from view or exposed on the screen. Each window has a title which can be set by the user. This information is used by the window manager to lay out the window on the screen, although there is no guarantee that what the user asks for is what the user gets!

$$\text{HideExpose} ::= \text{Hide} \mid \text{Expose}$$
*Control**title : String**control : HideExpose**xylimits : Pixel \times Pixel**first xylimits \leq second xylimits*

Together these make up the information describing a particular window.

$$\text{Info} \cong \text{Map} \wedge \text{Control}$$

There are a finite number of windows on a particular screen. One of these is considered to be the currently selected window. This may be 'undefined' sometimes. Most *WM* library functions take effect on the currently selected window. Each window has information, including a pixel map, associated with it.

$$\mid \text{Undefined} : \text{Window}$$

WM $windows : F Window$ $current : Window$ $contents : Window \rightarrow Info$ <hr/> $Undefined \notin windows$ $windows = \text{dom } contents$ $current \in windows \cup \{Undefined\}$
--

The display screen consists of the background overlaid with windows. The window pixel maps do not overlap. All windows are contained within the background area.

WM WM $maps : Window \rightarrow Pixmap$ $areas : Window \rightarrow (P Pixel)$ $screen, background : Pixmap$ <hr/> $maps = contents ; (\lambda Info \bullet map)$ $areas = contents ; (\lambda Info \bullet area)$ $\text{disjoint } areas$ $\bigcup(\text{ran } areas) \subseteq \text{dom } background$ $screen = background \oplus \bigcup(\text{ran } maps)$
--

You can have as many windows as you need, subject to the restriction that the *WM* process can handle at most 20 windows, including hidden windows and windows requested by other programs, at one time.

$MaxWindows : N$ <hr/> $MaxWindows = 20$

We can include this limitation in our model of the state.

$$WM \cong [WM \mid \#windows \leq MaxWindows]$$

The size of a window on the user's display is one of the resources that the *Window Manager* allocates. A program can request a given size, and

WM will take the requested size into account when making decisions, but it does not guarantee a particular size. This process is modelled as a function of the system. The number of windows is not changed by this function. Additionally, control information supplied by the user is left unchanged.

$$WINDOWS == Window \leftrightarrow Info$$

<i>WM</i>
<i>WM</i>
$adjust : WINDOWS \rightarrow WINDOWS$
$\forall w, w' : WINDOWS \mid w' = adjust\ w \bullet$
$\#w' = \#w \wedge$
$w' ; (\lambda Info \bullet \theta Control) = w ; (\lambda Info \bullet \theta Control)$

Initially there are no windows and the current window is undefined.

<i>InitWM</i>
<i>WM'</i>
$windows' = \emptyset$
$current' = Undefined$

Operations change the state of the system. However the background and hence the size of the screen remains constant. Additionally the algorithm to adjust the size of windows does not change.

ΔWM
<i>WM</i>
<i>WM'</i>
$background' = background$
$adjust' = adjust$

Sometimes the state of the system is unaffected during an operation.

$$\exists WM \cong [\Delta WM \mid \theta WM' = \theta WM]$$

Many operations are concerned with the current window. Hence we define a schema giving a partial specification covering all common aspects of such operations. This can be used to shorten subsequence specifications of these operations and reduce repetition. The names of such schemas are prepended with ' Φ ' to distinguish these from actual operations.

Φ <i>Current</i> Δ <i>WM</i> <i>Info</i> <i>Info'</i>
<i>current</i> \in <i>windows</i> <i>current'</i> = <i>current</i> θ <i>Info</i> = <i>contents current</i> <i>contents'</i> = <i>adjust (contents \oplus {current \mapsto θInfo'})</i>

This leaves a valid current window the same, but updates the information associated with it in some (as yet unspecified) way.

4.2 Window Creation and Deletion

When a window is created, the system adjusts all the windows in the system appropriately. The window body is exposed when it is created. In practice, the operation also takes the name of a host as input since a window may be created anywhere on the network of workstations. However this is detailed later.

<i>NewWindow</i> ₀ Δ <i>WM</i> <i>w!</i> : <i>Window</i> <i>Info</i>
$\#$ <i>windows</i> < <i>MaxWindows</i> <i>w!</i> \notin <i>windows</i> \cup { <i>Undefined</i> } <i>current'</i> = <i>w!</i> <i>control</i> = <i>Expose</i> <i>contents'</i> = <i>adjust (contents \cup {w! \mapsto θInfo})</i>

The currently selected window can be deleted.

<i>Delete Window</i> ₀
ΔWM
<i>current</i> \in <i>windows</i>
<i>current'</i> = <i>Undefined</i>
<i>contents'</i> = <i>adjust</i> (<i>{current}</i> \triangleleft <i>contents</i>)

4.3 Window Size

A program can request a given size range, and *WM* will take the requested size into account when making decisions, but it does not guarantee a particular size. The rest of the window information is unaffected. The windows will be adjusted by the system as necessary (i.e. any or all of the displayed windows may change shape as a result of setting the size of one particular window).

<i>SetDimension</i> ₀
Φ <i>Current</i>
<i>miny?</i> , <i>maxy?</i> : <i>Pixel</i>
<i>map'</i> = <i>map</i>
<i>title'</i> = <i>title</i>
<i>control'</i> = <i>control</i>
<i>xylimits'</i> = (<i>miny?</i> , <i>maxy?</i>)

The size of the body of the currently selected window can be returned. If the window is actually hidden (i.e. *WM* has adjusted the window to display the header only), then the returned size is empty.

<i>GetDimensions</i> ₀
Φ <i>Current</i>
<i>wh!</i> : <i>Pixel</i>
<i>xy₁, xy₂</i> : <i>Pixel</i>
θ <i>Info'</i> = θ <i>Info</i>
$\text{dom body} = xy_1 .. xy_2$
$wh! = xy_2 - xy_1$

4.4 Windows Visibility

A window is considered 'visible' when both its header and its body are displayed and 'hidden' when only its header is displayed. Windows are visible when they are first created and remain so unless the user hides them. A program can also control window visibility. A visible window may be hidden.

<i>HideMe</i> ₀
Φ <i>Current</i>
$map' = map$
$title' = title$
$control' = Hide$
$xylimits' = xylimits$

Similarly, a hidden window may be exposed.

<i>ExposeMe</i> ₀
Φ <i>Current</i>
$map' = map$
$title' = title$
$control' = Expose$
$xylimits' = xylimits$

Note that the state of 'control' before the operation has not been checked above, so *HideMe*₀ will leave a hidden window hidden and *ExposeMe*₀ will leave a visible window exposed.

4.5 Other Window Operations

A window may be explicitly selected as the current window, until another window is selected or created. All output will be sent to the selected window.

<i>Select Window</i> ₀
ΔWM
$w? : Window$
$w? \in windows$
$current' = w?$
$contents' = contents$

The title of a window may be set. This involves placing a text string in the header section of the window contents.

<i>SetTitle</i> ₀
$\Phi Current$
$s? : String$
$map' = map$
$title' = s?$
$control' = control$
$xylimits' = xylimits$

The body of the currently selected window may be set to white.

<i>ClearWindow</i> Φ <i>Current</i> <i>header'</i> = <i>header</i> <i>body'</i> = <i>setval White body</i> <i>title'</i> = <i>title</i> <i>control'</i> = <i>control</i> <i>xylimits'</i> = <i>xylimits</i>
--

Other operations supplied by the *WM* library include line, text and string drawing, raster operations, operations to save and restore parts of the picture, input handling, menus, mouse input, etc. In addition, new operations may be added; at the time that this specification was formulated *WM* are still under development.

4.6 Errors

There is a *Null* window identifier which is never a valid window.

| *Null* : *Window*

$WM \cong [WM \mid Null \notin windows]$

ΔWM and $\exists WM$ are redefined appropriately.

Some operations return a window identifier. If this is non-null then the operation is successful.

<i>Success</i> ΔWM <i>w!</i> : <i>Window</i> <i>w!</i> \neq <i>Null</i>
--

Alternatively an error may occur. There is a limit on the number of windows which *WM* can handle. This could cause an error when creating a new window.

<i>TooManyWindows</i>
$\equiv WM$
$w! : Window$
$\#windows \geq MaxWindows$
$w! = Null$

We can now make the operation to create a new window total.

$$NewWindow_1 \equiv (NewWindow_0 \wedge Success) \vee TooManyWindows$$

The current window may be undefined when one is required.

<i>NoCurrentWindow</i>
$\equiv WM$
$current = Undefined$

$$DeleteWindow_1 \equiv DeleteWindow_0 \vee NoCurrentWindow$$

$$SetDimensions_1 \equiv SetDimensions_0 \vee NoCurrentWindow$$

$$GetDimensions_1 \equiv GetDimensions_0 \vee NoCurrentWindow$$

$$SetTitle_1 \equiv SetTitle_0 \vee NoCurrentWindow$$

$$ClearWindow_1 \equiv ClearWindow_0 \vee NoCurrentWindow$$

A window may always be hidden or exposed, even if this does not affect its state.

An invalid window may be selected.

<i>InvalidWindow</i>
$\equiv WM$
$w? : Window$
$w? \notin windows$

$$SelectWindow_1 \equiv SelectWindow_0 \vee InvalidWindow$$

4.7 The ITC Network

In practice, as mentioned previously, there are many window managers, each running on a host workstation on a large network. Some hosts are running *WM*. All workstations have unique host names and all windows have unique IDs across the network.

$ \begin{array}{l} \text{ITC} \\ \hline \text{hosts} : \text{P String} \\ \text{wms} : \text{String} \leftrightarrow \text{WM} \\ \hline \text{dom wms} \subseteq \text{hosts} \\ \text{disjoint} (\text{wms} ; (\lambda \text{WM} \bullet \text{windows})) \end{array} $

To start with there are no hosts (and hence no window managers) on the network.

$$\text{InitITC} \cong [\text{ITC}' \mid \text{hosts}' = \emptyset]$$

Operations cause changes on the network.

$$\Delta \text{ITC} \cong \text{ITC} \wedge \text{ITC}'$$

Hosts can be added to the system (e.g. booting up) and removed (e.g. crashing or powering down).

$ \begin{array}{l} \text{AddHost} \\ \hline \Delta \text{ITC} \\ \text{host?} : \text{String} \\ \hline \text{host?} \notin \text{hosts} \cup \{\prime\} \\ \text{hosts}' = \text{hosts} \cup \{\text{host?}\} \\ \text{wms}' = \text{wms} \end{array} $

$ \begin{array}{l} \text{RemoveHost} \\ \hline \Delta \text{ITC} \\ \text{host?} : \text{String} \\ \hline \text{host?} \in \text{hosts} \\ \text{hosts}' = \text{hosts} \setminus \{\text{host?}\} \\ \text{wms}' = \{\text{host?}\} \triangleleft \text{wms} \end{array} $

Operations can be initiated on a particular 'local' host. These do not affect the host names on the network.

$\Phi Host$
ΔITC
$localhost : String$
$hosts' = hosts$
$localhost \in hosts$

For example, *WM* may be executed on a host, and may be subsequently killed.

<i>Exec WM</i>
$\Phi Host$
$initwm : InitWM$
$localhost \notin \text{dom } wms$
$wms' = wms \cup \{localhost \mapsto initwm\}$

<i>KillWM</i>
$\Phi Host$
$localhost \in \text{dom } wms$
$wms' = \{localhost\} \triangleleft wms$

WM operations can be modelled in the global context of the network by updating the state of *WM* on a particular local host which is already running *WM*.

ΦWM
$\Phi Host$
ΔWM
$host : String$
$\theta WM = wms \text{ host}$
$\theta WM' = wms' \text{ host}$

The *Window Manager* on the local host can be requested to create new windows on any machine on the ITC network that is running a *WM* process by supplying the appropriate host name. Alternatively, specifying a null host parameter results in a request for a window on the local machine. This is the normal mode of operation.

<i>NewWindow_{ITC}</i>
<i>NewWindow₁</i>
ΦWM
<i>host?</i> : <i>String</i>
<i>host?</i> = "" \Rightarrow <i>host</i> = <i>localhost</i>
<i>host?</i> \neq "" \Rightarrow <i>host</i> = <i>host?</i>

Other window operations discussed previously may be described in the global network context. For example:

$$DeleteWindow_{ITC} \hat{=} DeleteWindow_1 \wedge \Phi WM$$

If the current window is on the local machine, then the operation is executed locally, otherwise it is carried out over the network.

4.8 Simplifications and Assumptions

For the purposes of brevity, the pop-up menus supplied by *WM* have been ignored in the description given. These could easily be added to the state specification by including them as extra window information in a schema called *Menu*.

$$Info \hat{=} Info \wedge Menu$$

In practice, windows are not adjusted immediately, but when the *Window Manager* next makes a size decision (e.g. when the user requests *WM* to proportion the windows). This is not modelled here. In addition, the way in which the windows are proportioned is not specified since this is not covered in the documentation used to formulate this specification [Neuw85].

4.9 Comments

The *WM* window manager provides a simple system with non-overlapping windows. Hence no notion of window ordering is required. The idea of a

'current' window for each process using *WM* means that this information is held as part of the state of the system and need not be specified as input to many window operations. Windows are automatically reduced in size by the system when there is not enough space on the screen. This simplifies the task of organising windows for the user.

Chapter 5

Blit Windows (AT&T)

The Blit [Pike84, Unix85], developed at Bell Labs, Murray Hill, is more like an intelligent terminal than a workstation. It is diskless and interacts with a remote host running the Bell Labs Eighth Edition UNIX via a 9600 baud (slow) RS-232 serial line. It has its own simple process scheduler and a bit-mapped display. Programs can be run on the Blit (downloaded from the remote host), on the remote host itself using a standard window terminal emulator process on the Blit, or on both using two special purpose programs which interact with each other over the serial line. Deciding how to split a program between the Blit and remote host is a tricky but interesting problem.

5.1 State

The Blit contains 'layers' which are analogous to windows on most other systems. However there is no protection between layers. Each layer has an rectangular region on the screen associated with it. Here we model this simply as a partial function from pixel points to values.

Layer == *Window*

Point == *Pixel*

Each pixel point is two-valued - i.e. each window is a simple bit map.

Zsize = 1

Several layers (with associated rectangular windows) may exist simultaneously. The layers are ordered as a sequence for reasons that will be seen in a moment. There is an invalid null layer for error returns (see later).

| *NullLayer* : *Layer*

Blit

layers : seq *Layer*
windows : *Layer* → *Rectangle*
rects : seq *Rectangle*

ran *layers* = dom *windows*

NullLayer ∉ ran *layers*

rects = *layers* ; *windows*

Several processes may also exist simultaneously in the Blit. Each process has an associated program and state. A process may be disabled or enabled. However the rest of this specification is not concerned with the state of processes, but it is included here for completeness. Each process is normally associated with a layer. Creating a process without a layer or vice versa is dangerous.

{*Program*, *State*}

Proc

prog : *Program*
state : *State*
l : *Layer*

The Blit includes no window 'manager' as such since any process has access to the entire screen. There are a series of processes in the system each identified uniquely by a process id. There is a null invalid id which is returned by operations to indicate an error. One of the processes may be assigned to receive mouse and keyboard events.

[*Id*]

| *NullId* : *Id*

<i>Blit</i>
<i>Blit</i>
$procs : Id \leftrightarrow Proc$
$receiver : Id$
$NullId \notin \text{dom } procs$
$receiver \in \text{dom } procs \cup \{NullId\}$
$\forall proc : \text{ran } procs \bullet$
$proc.l \in \text{ran } layers \cup \{NullLayer\}$

A process may be associated with a default terminal emulation program if desired.

$$Blit \cong [Blit; default : Program]$$

The background may be considered as another layer which defines the size of the screen. The display consists of all the layers overlaid on top of the background. All layers are contained within the background. The order is determined from the position in the sequence (first at the bottom, last on top).

<i>Blit</i>
<i>Blit</i>
$screen, background : Rectangle$
$\text{dom}(\oplus/rects) \subseteq \text{dom } background$
$screen = background \oplus (\oplus/rects)$

Initially there are no layers or processes in the system.

<i>InitBlit</i>
<i>Blit'</i>
$layers' = \langle \rangle$
$procs' = \emptyset$
$receiver' = NullId$

Operations do not change the default terminal program or the background of the display.

$\Delta Blit$
<i>Blit</i>
<i>Blit'</i>
<i>default' = default</i>
<i>background' = background</i>

Some operations do not affect the state of the Blit.

$$\exists Blit \cong [\Delta Blit \mid \theta Blit' = \theta Blit]$$

Often operations only affect layers and all the processes in the system are left unaffected.

$\Phi Layer$
$\Delta Blit$
<i>procs' = procs</i>
<i>receiver' = receiver</i>

Similarly, processes are often changed whilst leaving all the layers in the system unaffected.

$\Phi Proc$
$\Delta Blit$
<i>layers' = layers</i>
<i>windows' = windows</i>

5.2 New Layers

A layer may be created in a specified rectangle in the physical display bitmap. The address of the layer is returned.

NewLayer_0 <hr/> ΦLayer $r? : \text{Rectangle}$ $!l : \text{Layer}$ <hr/> $\text{dom } r? \subseteq \text{dom background}$ $!l \notin \text{ran layers}$ $\text{layers}' = \text{layers} \hat{\ } \langle !l \rangle$ $\text{windows}' = \text{windows} \cup \{!l \mapsto r?\}$

A layer may also be de-allocated. The associated process should also be freed for safety, but this is a separate operation.

DelLayer_0 <hr/> ΦLayer $!l? : \text{Layer}$ <hr/> $!l? \in \text{ran layers}$ $\text{layers}' = \text{layers} \underline{\text{remove}} \{!l?\}$ $\text{windows}' = \{!l?\} \triangleleft \text{windows}$
--

5.3 New Processes

A new process can be allocated. A handle on the process is returned. Note that the associated layer is undefined. The process program is often the default terminal emulation program and in practice this is specified using a null argument.

NewProc_0 <hr/> ΦProc $f? : \text{Program}$ $id! : \text{Id}$ $proc : \text{Proc}$ <hr/> $\text{proc.prog} = f?$ $\text{procs}' = \text{procs} \cup \{id! \mapsto \text{proc}\}$ $\text{receiver}' = \text{receiver}$

A process may be created using the standard user interface to select the rectangle for the process's layer. This associates the process with the layer.

$$NewWindow_0 \cong NewProc_0 ; NewLayer_0$$

$$NewWindow_0 \cong [NewWindow_0 \mid l! = proc.l]$$

A layer may be selected so that the process in that layer becomes the receiver of mouse and keyboard events.

$ToLayer_0$ $\Phi Proc$ $l? : Layer$ $proc : Proc$
$proc \in \text{ran } procs$ $l? = proc.l$ $procs' = procs$ $receiver' = procs \sim proc$

The process whose layer is indicated by the mouse may be returned.

$GetProc_0$ $\exists Blit$ $l? : Layer$ $proc! : Proc$
$proc! \in \text{ran } procs$ $proc!.l = l?$

Alternatively, a handle on all the processes in the system may be returned.

$GetProcTab_0$ $\exists Blit$ $procs! : Id \leftrightarrow Proc$
$procs! \simeq procs$

5.4 Mouse Operations

The Blit includes a mouse. This controls the position of a cursor on the screen. Additionally, any combination of the three buttons on the mouse may be pressed at any time.

$$\textit{Button} ::= \textit{Button1} \mid \textit{Button2} \mid \textit{Button3}$$

Mouse

xy : *Point*

buttons : P *Button*

Some pixel positions on the display screen may be associated with a particular layer. This is the layer which is visible at that particular pixel. Otherwise the background is visible at that point. The gun-sight cursor is used to find a particular layer.

Gunsight

\exists *Blit*

pos? : *Mouse*

l! : *Layer*

$\textit{pos?.xy} \in \text{dom}(\oplus/\textit{rects}) \Rightarrow$

$\textit{l!} = \textit{layers}(\textit{max} \{ \textit{n} : \text{dom} \textit{rects} \mid$
 $\textit{pos?.xy} \in \text{dom}(\textit{rects} \textit{n}) \})$

$\textit{pos?.xy} \notin \text{dom}(\oplus/\textit{rects}) \Rightarrow$

$\textit{l!} = \textit{NullLayer}$

We can now give a more complete definition for *GetProc*.

$$\textit{GetProc}_1 \cong \textit{Gunsight} \gg \textit{GetProc}_0$$

The box cursor is used to pick out a rectangular area. This is done by sweeping out a rectangle whilst button 3 is depressed.

Box $\cong Blit$ $pos1?, pos2? : Mouse$ $r! : Rectangle$
$pos1?.buttons = \{Button3\}$ $pos2?.buttons = \emptyset$ $dom r! = dom background \cap$ $((pos1?.xy \underline{min} pos2?.xy) .. (pos1?.xy \underline{max} pos2?.xy))$

5.5 The 'mux' Multiplexer

The underlying library routines available for the Blit do not include all the basic operations necessary for a complete window manager. However a program called *mux* may be downloaded from the host system. This manages asynchronous windows, or layers, on the Blit terminal. Each layer is essentially a separate terminal. Layers are created, deleted, and rearranged using the mouse. Depressing mouse button 3 activates a menu of layer operations and releasing the button selects an operation. Some of these operations are covered here.

A new layer containing a terminal emulator process may be created by sweeping out a rectangle with the mouse whilst button 3 is depressed.

$$New_0 \cong Box \gg NewWindow_0$$

The size and location of a layer on the screen may be changed. A gun-sight cursor to select the layer and a box cursor to select the new position are presented to the user. The domain of the layer's rectangular area is updated.

$$Reshape_0 \cong (Gunsight \gg DelLayer_0); (Box \gg NewLayer_0)$$

A non-current layer may be selected using button 1. The layer is pulled to the front of the screen and made the current layer for keyboard and mouse input.

<i>Top₀</i>
Φ <i>Layer</i>
<i>l?</i> : <i>Layer</i>
<i>layers'</i> = <i>layers raise</i> { <i>l?</i> }
<i>windows'</i> = <i>windows</i>

$\text{Current} \& \text{Top}_0 \hat{=} \text{Gunsight} \gg (\text{ToLayer}_0 ; \text{Top}_0)$

5.6 Errors

Successful operations can be reported.

<i>Success</i>
Δ <i>Blit</i>
<i>l!</i> : <i>Layer</i>
<i>l!</i> \neq <i>NullLayer</i>

Similarly, failures can also be reported. This could be because there is not enough memory for example.

<i>Failure</i>
\exists <i>Blit</i>
<i>l!</i> : <i>Layer</i>
<i>l!</i> = <i>NullLayer</i>

A rectangle not within the background area could be given in error.

<i>InvalidRect</i>
\exists <i>Blit</i>
<i>r?</i> : <i>Rectangle</i>
$\neg (\text{dom } r? \subseteq \text{dom background})$

For example, the operations to create a new layer or process may fail because of an invalid rectangle or lack of memory.

$$NewLayer_1 \cong (NewLayer_0 \wedge Success) \vee (InvalidRect \wedge Failure) \vee Failure$$

$$NewProc_1 \cong (NewProc_0 \wedge Success) \vee Failure$$

$$NewWindow_1 \cong (NewWindow_0 \wedge Success) \vee Failure$$

Sometimes an invalid layer may be specified as input.

<i>InvalidLayer</i>
$\exists Blit$
$!?: Layer$
$! ? \notin \text{ran layers}$

A layer must exist to delete it or make it the current receiver.

$$DelLayer_1 \cong DelLayer_0 \vee InvalidLayer$$

$$ToLayer_1 \cong ToLayer_0 \vee InvalidLayer$$

Some operations return no errors.

$$GetProcTab_1 \cong GetProcTab_0$$

5.7 Simplifications, Assumptions and Comments

Many cursor and mouse operations and other graphics operations have been ignored for brevity.

The documentation [Unix85] states that the associated process must be freed when a layer is de-allocated. However it does not make it clear how to do this so this has not been specified.

The Blit is different from most other window systems in that it is a diskless intelligent terminal which interacts with a remote host in normal operation. In addition there is no protection between processes and layers within the Blit. Hence care must be exercised when programming it, but in return this allows greater flexibility and versatility.

Chapter 6

X Window System (MIT)

X [Gett85, Gett86, Sche86] is a network transparent windowing system developed at MIT and designed to run under UNIX. The X display server distributes user input to, and accepts output requests from various client programs either on the same machine or over a network.

6.1 State

The state of the X system is introduced in simple stages in order to build up the concepts involved. This is done by redefining a state schema called *X* in terms of itself and a series of manageably sized state definition fragments.

All the windows in an X server are arranged in a strict hierarchy. At the top of the hierarchy is the 'root' window. Each window has a parent except the root window. Child windows may in turn have their own children. Each window, including the root window, may be considered to consist of a pixel map in this simple description.

X
$root : Window$
$children : P Window$
$parents : Window \rightarrow Window$
$subwindows : Window \leftrightarrow Window$
$windows : Window \rightarrow Pixmap$
$root \notin children$
$children = dom\ parents$
$subwindows = parents^{\sim}$
$dom\ windows = children \cup \{root\}$

Subwindows are displayed in a particular order within their parent window. This may be modelled as a sequence of windows in ascending order of display. These consist of all the child windows.

X
X
$order : Window \rightarrow iseq\ Window$
$dom\ order = dom\ windows$
$(\forall w_1, w_2 : dom\ order \mid w_1 \neq w_2 \bullet$ $\quad ran(order\ w_1) \cap ran(order\ w_2) = \emptyset)$
$\bigcup\{w : dom\ order \bullet ran(order\ w)\} = children$
$(\forall w : dom\ windows \bullet ran(order\ w) = subwindows(\{w\}))$

Windows must be ‘mapped’ before they can be displayed. The root window is always mapped. All of a window’s ancestors (if any) must also be mapped for it to be viewable on the display. Unviewable windows are mapped but have some ancestor which is unmapped.

X
X
$mapped,$ $viewable,$ $unviewable : P\ Window$
$mapped \subseteq dom\ windows$
$root \in mapped$
$viewable = \{c : children \mid parents^*(\{c\}) \subseteq mapped\} \cup \{root\}$
$unviewable = \{c : children \mid c \in (mapped \setminus viewable)\}$

Each viewable window has an associated visible pixel map which consists of the pixel map of the window overlaid with its subwindows (in order) if any. These are ‘clipped’ to the size of the parent window.

The root window covers the entire background of the display screen. The screen displays the pixel map visible from the root window.

X
X
$visible : Window \mapsto Pixmap$
$screen, background : Pixmap$
$dom\ visible = viewable$
$(\forall w : viewable \bullet$
$\quad visible\ w = (windows\ w) \oplus (dom(windows\ w) \triangleleft$
$\quad \quad \oplus / (compact((order\ w); visible)))$
$background = windows\ root$
$screen = visible\ root$

Initially there are no children and only the root window is mapped. Hence only the background is displayed.

$InitX$
X'
$windows' = \{root' \mapsto background'\}$
$order' = \{root' \mapsto \langle \rangle\}$
$mapped' = \{root'\}$

Consider changes in the window system. The root window identifier and the background of the screen do not change.

ΔX
X
X'
$root' = root$
$background' = background$

Sometimes the state of the system is unaffected during an operation.

$$\exists X \hat{=} [\Delta X \mid \theta X' = \theta X]$$

We can now consider operations on the state of the system; initially, error-free operations will be presented for simplicity. Error conditions are covered later.

6.2 Creating and Destroying Windows

Firstly, we wish to be able to create windows. For these operations we have to supply the parent window under which the new window is to reside in the window hierarchy. The position, size and background of the window must also be specified. Here these are defined by '*bgnd?*' for simplicity. Note that the created window will not actually be displayed until it is 'mapped' (see later).

<i>CreateWindow</i> ₀
ΔX
<i>parent?</i> : <i>Window</i>
<i>bgnd?</i> : <i>Pixmap</i>
<i>w!</i> : <i>Window</i>
<i>parent?</i> \in dom <i>windows</i>
<i>w!</i> \notin dom <i>windows</i>
<i>windows'</i> = <i>windows</i> \cup { <i>w!</i> \mapsto <i>bgnd?</i> }
<i>order'</i> = <i>order</i> \oplus { <i>parent?</i> \mapsto (<i>order parent?</i>) \wedge { <i>w!</i> }}
<i>mapped'</i> = <i>mapped</i>

Note that the predicates in the schema above fully define the state after the operation since all the other state components may be derived from those given above. The other components are included in the state definition to allow us to have different views of the system, depending on the manner in which we wish to access the state.

Sometimes it is convenient to create several windows at once under a single parent window. Note that not all the windows requested may be created, but this is indicated by the information returned. This consists of a partial injection obtained from the sequence numbers of the windows which are actually created to the window identifiers which they are allocated.

<p><i>CreateWindows</i>₀</p> <p>ΔX</p> <p><i>parent?</i> : <i>Window</i></p> <p><i>defs?</i> : seq <i>Pixmap</i></p> <p><i>defs!</i> : $\mathbf{N} \mapsto \mathbf{Window}$</p> <hr/> <p><i>parent?</i> \in dom <i>windows</i></p> <p>dom <i>defs!</i> \subseteq dom <i>defs?</i></p> <p>ran <i>defs!</i> \cap dom <i>windows</i> = \emptyset</p> <p><i>windows'</i> = <i>windows</i> \cup (<i>defs!</i>\sim ; <i>defs?</i>)</p> <p><i>order'</i> =</p> <p style="padding-left: 2em;"><i>order</i> \oplus { <i>parent?</i> \mapsto (<i>order parent?</i>)\wedge (<i>compact defs!</i>) }</p> <p><i>mapped'</i> = <i>mapped</i></p>
--

We also wish to destroy windows. Given a particular window, we may wish to destroy a set of windows which are associated with it. We can define a partial specification to do this as a schema. Exactly which windows are to be destroyed is not specified for the present.

<p>Φ <i>Destroy</i></p> <p>ΔX</p> <p><i>w?</i> : <i>Window</i></p> <p><i>destroy</i> : \mathbf{P} <i>Window</i></p> <hr/> <p><i>w?</i> \in <i>children</i></p> <p><i>windows'</i> = <i>destroy</i> \triangleleft <i>windows</i></p> <p>($\forall w$: dom <i>windows</i> \bullet</p> <p style="padding-left: 2em;"><i>order'</i> <i>w</i> = (<i>order w</i>) <u>remove</u> <i>destroy</i>)</p> <p><i>mapped'</i> = <i>mapped</i> \setminus <i>destroy</i></p>
--

We may wish all subwindows, as well as the window itself, to be destroyed.

$$\text{DestroyWindows}_0 \hat{=} [\Phi \text{ Destroy} \mid \text{destroy} = \text{subwindows}^* \{ \{ w? \} \}]$$

Alternatively, we may wish to just destroy the subwindows under the specified window.

$$\text{DestroySubwindows}_0 \equiv [\Phi \text{Destroy} \mid \text{destroy} = \text{subwindows}^+ (\{w?\})]$$

Note that the 'root' background window cannot be destroyed using these operations. Only child windows may be destroyed.

6.3 Manipulating Windows

A window, and all its ancestors, must be 'mapped' to be visible on the screen. However a mapped window may still be invisible if it is obscured by a sibling window.

Mapping operations require a child window to be specified. The hierarchical relationships between windows and the contents of the windows are left unaffected.

ΦMap ΔX $w? : \text{Window}$ $w? \in \text{children}$ $\text{windows}' = \text{windows}$

Mapping a window raises the window and all its subwindows which have had map requests. Mapping a window which is already mapped has no effect on the screen - it does *not* raise it.

MapWindow_0 ΦMap $\text{parent} : \text{Window}$ $\text{parent} = \text{parents } w?$ $w? \notin \text{mapped} \Rightarrow$ $\text{order}' = \text{order} \oplus \{\text{parent} \mapsto ((\text{order } \text{parent}) \text{raise } \{w?\})\}$ $w? \in \text{mapped} \Rightarrow$ $\text{order}' = \text{order}$ $\text{mapped}' = \text{mapped} \cup \{w?\}$

All the unmapped subwindows of a given window can be mapped together. The order in which they are mapped is chosen by the system rather than the caller.

MapSubwindows_0 ΦMap $\text{neworder} : \text{iseq Window}$ $\text{newmapped} : \text{P Window}$
$\text{newmapped} = \text{subwindows}(\{w?\}) \setminus \text{mapped}$ $\text{ran neworder} = \text{newmapped}$ $\text{order}' = \text{order} \oplus$ $\quad \{w? \mapsto ((\text{order } w?) \text{ remove newmapped}) \wedge \text{neworder}\}$ $\text{mapped}' = \text{mapped} \cup \text{newmapped}$

A window can be unmapped. The window will disappear from view if it was visible.

UnmapWindow_0 ΦMap
$\text{order}' = \text{order}$ $\text{mapped}' = \text{mapped} \setminus \{w?\}$

All subwindows of a specified window can be unmapped.

UnmapSubwindows_0 ΦMap
$\text{order}' = \text{order}$ $\text{mapped}' = \text{mapped} \setminus \text{subwindows}(\{w?\})$

Windows may be manipulated in various ways. Given a window, its pixel map may be updated. It is also raised to the top of the display. We can define a general schema to simplify the definition of such operations.

Φ <i>Window</i> ΔX $w? : \textit{Window}$ $\textit{map} : \textit{Pixmap}$ $\textit{parent} : \textit{Window}$
$\textit{parent} = \textit{parents } w?$ $\textit{windows}' = \textit{windows} \oplus \{w? \mapsto \textit{map}\}$ $\textit{order}' = \textit{order} \oplus \{\textit{parent} \mapsto ((\textit{order } \textit{parent}) \underline{\textit{raise}} \{w?\})\}$ $\textit{mapped}' = \textit{mapped}$

A window may be moved and raised without changing its size. Moving a mapped window may or may not lose its contents, depending on various circumstances.

$\textit{MoveWindow}_0$ Φ <i>Window</i> $xy? : \textit{Pixel}$
$\text{dom } \textit{map} = \text{dom}((\textit{offset } xy?); (\textit{windows } w?))$

The size of a window may be changed without changing its upper left coordinate. A new width and height are given. The window is always raised. Changing the size of a mapped window loses its contents.

$\textit{ChangeWindow}_0$ Φ <i>Window</i> $\textit{wdht}? : \textit{Pixel}$ $\textit{pix}_1, \textit{pix}_2 : \textit{Pixel}$
$\text{dom}(\textit{windows } w?) = \textit{pix}_1 \dots \textit{pix}_2$ $\text{dom } \textit{map} = \textit{pix}_1 \dots (\textit{pix}_1 + \textit{wdht}?)$

The size and location of a window may be configured together by combining the last two operations. The window is raised and the contents are lost.

$$\textit{ConfigureWindow}_0 \hat{=} (\textit{MoveWindow}_0 \upharpoonright \Delta X); (\textit{ChangeWindow}_0 \upharpoonright \Delta X)$$

Some operations explicitly affect the order in which the windows are displayed. A child window is specified, and window relationships, the windows themselves and the set of mapped windows remain unchanged.

$\Phi Order$ ΔX $w' : Window$ $parent : Window$ $suborder, suborder' : seq Window$ <hr/> $parent = parents w?$ $windows' = windows$ $suborder = order parent$ $order' = order \oplus \{parent \mapsto suborder'\}$ $mapped' = mapped$
--

A window may be 'raised' so that no sibling window obscures it. If the windows are regarded as overlapping sheets of paper stacked on a desk, then raising a window is analogous to moving the sheet to the top of the stack, whilst leaving its position on the desk the same.

$$RaiseWindow_0 \hat{=} \{ \Phi Order \mid suborder' = suborder \underline{raise} \{w?\} \}$$

A window may also be 'lowered' in a complementary fashion. If the windows are regarded as overlapping sheets of paper stacked on a desk, then lowering a window is analogous to moving the sheet to the bottom of the stack, whilst leaving its position on the desk the same.

$$LowerWindow_0 \hat{=} \{ \Phi Order \mid suborder' = suborder \underline{lower} \{w?\} \}$$

Overlapping mapped subwindows of a particular window may be raised or lowered in a circular manner. The set of these windows is identified. If it is non-empty, the ordering of the window's children is updated; otherwise it is left unchanged.

ΦCirc
ΔX
$w? : \text{Window}$
$\text{submapped}, \text{circ} : \text{P Window}$
$\text{suborder}, \text{suborder}' : \text{seq Window}$
$w? \in \text{children}$
$\text{submapped} = \text{subwindows}(\{w?\}) \cap \text{mapped}$
$\text{circ} = \{w : \text{submapped} \mid (\exists w_2 : \text{submapped} \bullet w_2 \neq w \wedge (\text{visible } w_2) \underline{\text{overlaps}} (\text{visible } w))\}$
$\text{windows}' = \text{windows}$
$\text{suborder} = \text{order } w?$
$\text{circ} \neq \emptyset \Rightarrow \text{order}' = \text{order} \oplus \{w? \mapsto \text{suborder}'\}$
$\text{circ} = \emptyset \Rightarrow \text{order}' = \text{order}$
$\text{mapped}' = \text{mapped}$

For a particular window, the lowest mapped child that is partially obscured by another child may be raised. Repeated executions lead to round robin raising.

Circ WindowUp_0
ΦCirc
$\text{circ} \neq \emptyset \Rightarrow$
$\text{suborder}' = \text{suborder } \underline{\text{raise}}$
$\{\text{suborder}(\min(\text{dom}(\text{suborder} \upharpoonright \text{circ}))\}$

Similarly, the highest mapped child of a particular window that (partially) obscures another child may be lowered. Repeated executions lead to round robin lowering.

Circ WindowDown_0
ΦCirc
$\text{circ} \neq \emptyset \Rightarrow$
$\text{suborder}' = \text{suborder } \underline{\text{lower}}$
$\{\text{suborder}(\max(\text{dom}(\text{suborder} \upharpoonright \text{circ}))\}$

6.4 Other Window Operations

We can ask for information about a particular window. As well as the size, position, etc. of the window, details about the mapped state of the window are returned. ‘*IsUnmapped*’ indicates that the window is unmapped, ‘*IsMapped*’ indicates that it is mapped and displayed (i.e. all of its ancestors are also mapped), and ‘*IsInvisible*’ implies that it is mapped but some ancestor is not mapped.

$$\text{MappedState} ::= \text{IsUnmapped} \mid \text{IsMapped} \mid \text{IsInvisible}$$

<i>QueryWindow</i>
$\exists X$
$w? : \text{Window}$
$\text{info!} : \text{Pixmap}$
$\text{mapped!} : \text{MappedState}$
$w? \in \text{children}$
$\text{info!} = \text{windows } w?$
$w? \notin \text{mapped} \Rightarrow$ $\text{mapped!} = \text{IsUnmapped}$
$\text{parents}^+(\{w?\}) \subseteq \text{mapped} \Rightarrow$ $\text{mapped!} = \text{IsMapped}$
$w? \in \text{mapped} \wedge \neg (\text{parents}^+(\{w?\}) \subseteq \text{mapped}) \Rightarrow$ $\text{mapped!} = \text{IsInvisible}$

We can also find out the window identifiers of the parent and all the children (and hence the number of children) for a particular window. The children are listed in current stacking order, from bottommost (first) to topmost (last).

<i>QueryTree₀</i>
$\exists X$
$w? : \text{Window}$
$\text{parent!} : \text{Window}$
$\text{children!} : \text{seq Window}$
$\text{parent!} = \text{parents } w?$
$\text{children!} = \text{order } w?$

The X system includes many other operations. These include more detailed window operations, mouse operations, graphics for line drawing and fill operations, screen raster operations, moving bits and pixels to and from the screen, storing and freeing bit maps and pixel maps, cursor definition, colour operations, font manipulation routines, text output to a window, and so on. However the operations covered give an indication of the basic windowing facilities available under the X system.

6.5 Errors

Many operations return a status report signalling success or failure of the operation. Let this be denoted '*Status*'. Often a '*NULL*' status indicates success and a non-*NULL* status indicates failure.

[*Status*]

| *NULL* : *Status*

The operations covered so far detail what should happen in the event of no errors. In this case we also wish to report the fact that the operation was successful.

<i>Success</i>
<i>status!</i> : <i>Status</i>
<i>status!</i> = <i>NULL</i>

If errors do occur, then these need to be reported as well. For example, an invalid parent window may be specified.

<i>InvalidParent</i>
$\exists X$
<i>parent?</i> : <i>Window</i>
<i>status!</i> : <i>Status</i>
<i>parent?</i> \notin dom <i>windows</i>
<i>status!</i> \neq <i>NULL</i>

Alternatively, an invalid child window could be given as input.

InvalidWindow $\exists X$ $w? : \text{Window}$ $\text{status!} : \text{Status}$
$w? \notin \text{children}$ $\text{status!} \neq \text{NULL}$

We may include these errors with the previously defined operations which ignored error conditions, to produce total operations.

$$\text{CreateWindow}_1 \hat{=} (\text{CreateWindow}_0 \wedge \text{Success}) \vee \text{InvalidParent}$$

All the other operations covered take the following form.

$$\text{DestroyWindow}_1 \hat{=} (\text{DestroyWindow}_0 \wedge \text{Success}) \vee \text{InvalidWindow}$$

6.6 Simplifications and Assumptions

In the description given, only 'opaque' windows have been considered. The actual X system includes 'transparent' windows, mainly used for menus, and 'icon' windows which may be associated with opaque windows, but these have been ignored in this description for simplicity. These could be included in the state of the system. The operation specifications would need to be updated appropriately.

X X $\text{transparent, opaque} : \text{P Window}$ $\text{icon} : \text{Window} \leftrightarrow \text{Window}$
$(\text{transparent, opaque, ran icon})$ partition children $\text{dom icon} \subseteq \text{opaque}$

Windows have other information associated with them besides their pixel maps and their mapping status, such as border information. However this is not covered here. Exposure events that result from window operations are also ignored.

The informal description used to formulate this specification [Gett85] was not completely clear on a number of points. For example, the exact ordering of windows and their subwindows is not made explicitly clear after operations which affect this. In particular, it has been assumed here that raising and lowering a window implies that all its subwindows are also raised or lowered. Where necessary, an educated guess has been made as to the behaviour of the system.

6.7 Comments and Inconsistencies

The X window system is relatively complicated. It includes a number of basic concepts, several of which could not be included here fully because of lack of space. The hierarchical structure makes it very versatile.

Perhaps surprisingly, X has no notion of a 'current' window. Hence a large number of the library routines need a window identifier as input (including all those covered here). This is rather cumbersome and could introduce some unnecessary overhead in application programs using the system. However this is advantageous if a number of windows are being updated simultaneously since then there are effectively several current windows.

An earlier version of this specification was sent to MIT with annotations, raising questions about areas which were not well understood from the original documentation [Gett85]. A number of inconsistencies in the formal specification (compared to the implementation of X) were discovered from the feedback obtained. The major errors were as follows:

- Children are always on top of their parent, and the hierarchies of two siblings never interleave. In the original specification, an overall order (*order* : seq *Window*) was included as part of the state; it did not preclude the above. Here the ordering is defined on a per window basis, for just the immediate children.
- The contents of unmapped and invisible parts of windows are lost. For example, in the schema ΦMap , the predicate '*windows*' = *windows*' is actually incorrect since the contents of the window *w*? will be lost if it is unmapped. However the specification has not been changed in this respect since exposure events are ignored here, and these would typically restore the contents of re-exposed windows. If exposure events were added to this specification then this should be changed.

These points were missed from the original documentation. They would probably have been discovered if an implementation of X had been available

for 'testing' purposes. The documentation could be improved in these areas to avoid misunderstanding.

The original version of X specified here was Version 9 [Gett85]. Documentation for Version 10 [Gett86] and Version 11 [Sche86] were also available subsequently. Any further work on formalising X should use Version 11 documentation since this is now becoming an industry standard.

Chapter 7

Conclusions

A high level description of three window systems has been presented. Only a few operations for each system have been covered. A complete description would require a manual for each of the systems; a formal specification does not necessarily reduce the size of a description using informal methods. However it does make it much more precise. Because of this, it is possible to reason about a system and detect inconsistencies in it far more easily than the case where only an informal specification is available. Even if formal specification is not used in the final documentation, its use will clarify points which can then be described informally to the user.

7.1 Comparison of Window Systems

Of the three window systems investigated, X provides the most comprehensive features. *WM* is a much simpler system with no overlapping windows or hierarchical structure. However it does automatically adjust the size of windows when necessary. The Blit is a 'raw' machine onto which window management functions can be loaded if desired. The following table gives a comparison of the features available on each system.

Window system	Overlapping windows	Hierarchical structure	Automatic sizing	Current window
<i>WM</i>	x	x	✓	✓
Blit	✓	x	x	✓
X	✓	✓	x	x

Most of the original specification in the monograph was undertaken in 1986. Of the three window systems, X was investigated first. It turned out to be the most complicated system and took a significant amount of time to

formalise. Subsequently, the specification of *WM* and the Blit system were comparatively easy.

Since the original specification, Version 11 of X (or X11 as it is normally known) has become an industry standard and is available on many workstations. The other two systems are not so widely used. X now includes a library interface built on top of the main X interface that implements almost all of *WM*. Hence most *WM* applications will run under X without source modification.

7.2 Formal Specification of Existing Systems

This work was undertaken as part of the Distributed Computing Software (DCS) Project at the Programming Research Group. As well as designing and documenting network services using a formal notation, part of the brief of the DCS project was to undertake case studies of existing systems and to formally specify parts of them in **Z** to gain a greater understanding of their operation.

This monograph is the result of one case study on the DCS project. Originally it had been hoped to compare parts of a number of distributed systems using **Z**. However, the authors of potential systems for investigation could only supply academic papers (not enough information) or the source code (too much information). What was required was some form of informal documentation for the system. Because window systems are used directly by users, there seems to be more readable documentation for such systems.

In each case, omissions and ambiguities in the documentation were discovered by attempting to formalise the system. Where necessary, intelligent guesses were made about the actual operation. These were usually correct, but not always.

Subsequently, the formal specifications could be used to update the existing documentation, or even rewrite it from scratch. Although **Z** has been developed as a design tool, it is also well suited for *post hoc* specifications of existing systems, and for detecting and correcting errors and anomalies in the documentation of such systems [Bowe88].

The most important stage of formalising a system is selecting the right level of abstraction for modelling its state. This is normally an iterative process. On attempting to specify an operation one often needs to backtrack to change the abstract state of the system. In particular, extra state components can be convenient to provide different views of the system depending on the operation involved.

There are likely to be some inconsistencies between the specifications

given here and the actual operation of the systems described. This is due to impreciseness and misunderstanding of the informal documentation used to formulate these specifications. This illustrates one of the reasons for using formal specification techniques – to avoid ambiguity or vagueness and to aid precise communication of ideas. Because of this, formal notation forces issues to the surface much earlier in the design process than when using informal description techniques such as natural language and diagrams. Difficult areas of the design cannot be ignored until the implementation stage. This reduces the number of backtracks necessary round the design/implementation cycle.

Additionally, using formal specification techniques should reduce maintenance costs since more of the errors in a system will be discovered before it is released into the field. Although specification and design costs will be increased, implementation and maintenance costs should be lower, reducing overall costs.

Formally specifying an existing system could be particularly useful if it is to be re-engineered to comply with modern software engineering standards. In such cases there could be costs benefits by taking such an approach [Nix88].

7.3 General Conclusions

Z can be used to succinctly specify real systems. The examples given here and other case studies undertaken at the Programming Research Group and elsewhere lend support to this assertion.

Z may be used to produce readable specifications. It has been designed to be read by humans rather than computers. Thus it can form the basis for documentation.

Large specifications are manageable in **Z**, using the schema notation for structuring. It is possible to produce hierarchical specifications. A part of a system may be specified in isolation, and then this may be put into a global context.

Acknowledgements

This work was prompted by a trip to Carnegie-Mellon University, AT&T Bell Laboratories at Murray Hill and Massachusetts Institute of Technology where each of the window systems described were viewed in operation during October and November of 1985. The author is indebted to the developers of each of the systems who made the relevant documentation and other references freely available; in particular, Mahadev Satyanarayanan [Neuw85, Rose85, Saty84, West85], Rob Pike [Unix85], Dave Presotto [Pike84], Bob Scheiffer [Gett85], Bill Wehl [Gett86] and Jim Gettys [Sche86]. Additionally I would like to thank Mahadev Satyanarayanan (CMU), Dave Presotto (AT&T Bell Labs) and Bill Wehl (MIT) for acting as hosts on the trip.

Bernard Sufrin, Roger Gimson and Tim Gleeson at the PRG provided helpful suggestions and corrections on early drafts. Bob Scheiffer and Bill Wehl of MIT answered questions raised by the formal specification of the X window manager. Jim Gettys noted a clash in nomenclature with Version 11 of X on a subsequent trip to DEC Western Research Laboratory and the new nomenclature has been incorporated here.

This monograph has been produced using the \LaTeX document preparation system [Lamp86]. The **Z** specifications have been formatted and type-checked using the *fuzz* package [Spiv88b].

The Distributed Computing Software project was funded by a grant from the Science and Engineering Research Council (SERC). The author is currently working on the Software Engineering project at the PRG, also funded by the SERC. The author is grateful for their support.

References

- [Arno87] Arnold, D. J., Duce, D. A., and Reynolds, G. J., *An Approach to the Formal Specification of Configurable Models of Graphics Systems*, Proc. Eurographics 87, North Holland, 1987.
- [Bowe88] Bowen, J. P., *Formal Specification in Z as a Design and Documentation Tool*, Proc. 2nd IEE/BCS Conference, *Software Engineering 88*, Liverpool, UK, pp. 164–168, July 1988.
- [Gett85] Gettys, J. and Newman, R., *Xlib – C Language X Interface*, (Version 9), MIT Project Athena, Massachusetts Institute of Technology, USA, 1985.
- [Gett86] Gettys, J., Newman, R. and Della Fera, A., *Xlib – C Language X Interface*, (Version 10), MIT Project Athena, Massachusetts Institute of Technology, USA, 1986.
- [Fole82] Foley, J. D. and van Dam, A., *Fundamentals of Computer Graphics*, Addison Wesley, 1982.
- [Haye87] Hayes, I. J. (Editor), *Specification Case Studies*, Prentice-Hall International Series in Computer Science, 1987.
- [King88] King, S., Sørensen, I. and Woodcock, J., *Z: Grammar and Concrete and Abstract Syntazes*, Technical Monograph PRG-68, Programming Research Group, Oxford University, UK, 1988.
- [Lamp86] Lamport, L., *LaTeX: A Document Preparation System*, Addison-Wesley Publishing Company, 1986.
- [Loom88] Loomes, M. and Woodcock, J. C. P., *Software Engineering Mathematics: Formal Methods Demystified*, Pitman Publishing Ltd, London, UK, 1988.

- [Neuw85] Neuwirth, C., *Programmer's Guide to the Window Manager: An Introduction for the Uninitiated*, Information Technology Center (ITC), Carnegie-Mellon University, USA, 1985.
- [Newm81] Newman, W. M. and Sproull, R. F., *Principles of Interactive Computer Graphics*, 2nd edition, McGraw Hill, 1981.
- [Nix88] Nix, C. J. and Collins, B. P., *The Use of Software Engineering, including the Z Notation, in the Development of CICS*, Quality Assurance, 14(3), pp. 103-110, 1988.
- [Pike84] Pike, R., *The Blit: a Multiplexed Graphics Terminal*, AT&T Bell Laboratories Technical Journal, 63(8), part 2, pp. 1607-1631, October 1984.
- [Rose85] Rosenthal, D. S. H., *A Programmer's Guide to "WM" or How do I get into this Window Stuff*, Information Technology Center (ITC), Carnegie-Mellon University, USA, October 1985.
- [Saty84] Satyanarayanan, M., *The ITC Project: An Experiment in Large-Scale Distributed Personal Computing*, CMU Report CMU-ITC-035, Information Technology Center (ITC), Carnegie-Mellon University, USA, October 1984.
- [Sche86] Scheifler, R., *X Window System Protocol, Version 11*, (Draft 4), Massachusetts Institute of Technology, USA, 1986.
- [Spiv86a] Spivey, J. M., *Understanding Z: A Specification Language and its Formal Semantics*, Cambridge University Press, 1988.
- [Spiv86b] Spivey, J. M., *The fuzz Manual*, Computing Science Consultancy, 2 Willow Close, Garsington, Oxford, UK, 1988.
- [Spiv89] Spivey, J. M., *The Z Notation - A Reference Manual*, Prentice-Hall International Series in Computer Science, 1989.
- [Unix85] *UNIXTM Time-sharing System, Programmer's Manual*, Eighth Edition, Volume 1, AT&T Bell Laboratories, Murray Hill, New Jersey, USA, 1985.
- [West85] West, M., Nichols, D., Howard, J., Satyanarayanan, M. and Sidebotham, R., *The ITC Distributed File System: Principles and Design* CMU Report CMU-ITC-039, Information Technology Center (ITC), Carnegie-Mellon University, USA, March 1985.

Appendices

Appendix A

Glossary of Z notation

A glossary of the Z mathematical and schema notation used in this monograph is included here for easy reference. The complete notation is not covered. For more information, the grammar and concrete and abstract syntax for Z and a reference manual are available elsewhere [King88, Spiv89].

A.1 Abbreviated names

i, i_1, i_2	Identifiers
E, E_1, E_2	Expressions
d	Declaration
P, P_1, P_2	Predicates
D	Declaration and optional predicate (d or $d P$)
S, S_1, S_2	Schema expressions
N	Schema name with optional decoration (e.g. N')
N'	Decorated schema name (all components dashed)
f, f_1, f_2	Functions
R, R_1, R_2	Relations
x, x_1, x_2	Elements of sets
X, X_1, X_2	Sets
A	Set of sets ($\mathbf{P} X$)
n, n_1, n_2	Integer expressions
s, s_1, s_2	Sequences ($\text{seq } X$)
q	Sequence of sequences ($\text{seq}(\text{seq } X)$)
Z	Set of integers
I	Indexed family of sets ($\mathbf{P}(I \rightarrow \mathbf{P} X)$)
B, B_1, B_2	Bags

A.2 Horizontal paragraphs

$[i, \dots]$	Introduction of given set(s)
$i == E$	Abbreviation definition
$i \cong S$	Horizontal schema definition
$i ::= i_1 i_2 \dots$	Data type definition
P	Predicate (extra constraint)

A.3 Vertical paragraphs

New lines denote ‘;’. Predicates are conjoined by default.

i	Vertical schema definition. The schema name and predicate part are optional. The schema may subsequently be referenced by name with optional decoration in the document.
d	
P	

d	Axiomatic definition. The definitions may be non-unique. The predicate part is optional. The definitions apply globally in the document.
P	

$i, \dots d$	Generic definition. The definitions must be unique. The generic parameters are optional. The definitions apply globally in the document.
P	

A.4 Declarations and operators

$i : E$	Basic declaration
$i_1 : E_1; i_2 : E_2$	Multiple declarations
$i_1, i_2, \dots : E$	Declarations of same type ($i_1 : E; i_2 : E; \dots$)
$_ i _ : E$	Infix operator declaration
$i _ : E$	Prefix operator declaration
$_ i : E$	Postfix operator declaration
N	Schema reference as a declaration

A.5 Expressions

(E_1, E_2, \dots)	Ordered tuple
---------------------	---------------

$P E$	Power set (set of subsets)
$E_1 \times E_2 \times \dots$	Cartesian product
$\{E_1, E_2, \dots\}$	Set display
$\{D \bullet E\}$	Set comprehension (or $\{D\}$)
$\lambda D \bullet E$	Lambda-expression (function, given D returns E)
$\mu D \bullet E$	Definite description (value E if D unique)
$E_1 E_2$	Function application
$E.i$	Selection ($(\lambda D \bullet i)(E)$)
θS	Binding formation
N	Schema reference as an expression ($\{N \bullet \theta N\}$)
$E_1 f E_2$	Infix function expression
$E f$	Postfix function expression
(E)	Expression grouping
n	Number

A.6 Predicates

$E_1 = E_2$	Equality (or $E_1 = E_2 = \dots$)
$E_1 \in E_2$	Membership
$\neg P$	Logical negation
$P_1 \wedge P_2$	Logical conjunction
$P_1 \vee P_2$	Logical disjunction
$P_1 \Rightarrow P_2$	Logical implication ($\neg P_1 \vee P_2$)
$P_1 \Leftrightarrow P_2$	Logical equivalence ($P_1 \Rightarrow P_2 \wedge P_2 \Rightarrow P_1$)
$\forall D \bullet P$	Universal quantification
$\exists D \bullet P$	Existential quantification
$\exists_1 D \bullet P$	Unique existential quantification
N	Schema reference as a predicate
$E_1 R E_2$	Infix relation
$R E$	Prefix relation
<i>true</i>	True predicate
<i>false</i>	False predicate ($\neg \textit{true}$)
(P)	Predicate grouping

A.7 Schema expressions

$[D]$	Horizontal schema text
N	Schema reference

$\neg S$	Schema negation
$S_1 \wedge S_2$	Schema conjunction
$S_1 \vee S_2$	Schema disjunction
$S_1 \Rightarrow S_2$	Schema implication
$S_1 \Leftrightarrow S_2$	Schema equivalence
$\forall D \bullet S$	Universal schema quantification
$\exists D \bullet S$	Existential schema quantification
$S \setminus (i_1, i_2, \dots)$	Schema hiding of component(s)
$S_1 \upharpoonright S_2$	Schema projection (hiding of components not in S_2)
pre S	Precondition of schema
$S_1 ; S_2$	Schema composition (first S_1 , then S_2)
$S_1 \gg S_2$	Schema piping (S_1 outputs combine with S_2 inputs)
(S)	Schema grouping

A.8 Sets

$X_1 \neq X_2$	Inequality ($\neg (X_1 = X_2)$)
$x \notin X$	Non-membership ($\neg (x \in X)$)
\emptyset	Empty set ($\{e : X \mid \text{false}\}$)
$X_1 \subseteq X_2$	Subset relation
$X_1 \subset X_2$	Proper subset relation ($X_1 \subseteq X_2 \wedge X_1 \neq X_2$)
$X_1 \cup X_2$	Set union
$X_1 \cap X_2$	Set intersection
$X_1 \setminus X_2$	Set difference
$\mathbf{P} X$	Non-empty subsets ($\mathbf{P} X \setminus \{\emptyset\}$)
$\bigcup \mathcal{A}$	Generalised union
$\bigcap \mathcal{A}$	Generalised intersection
$\text{first}(x_1, x_2)$	Projection function of first co-ordinate (x_1)
$\text{second}(x_1, x_2)$	Projection function of second co-ordinate (x_2)

A.9 Relations

$X_1 \leftrightarrow X_2$	Binary relation ($\mathbf{P}(X_1 \times X_2)$)
$x_1 \mapsto x_2$	Maplet ((x_1, x_2))
dom R	Domain of relation
ran R	Range of relation
id X	Identity relation ($\{e : X \bullet e \mapsto e\}$)
$R_1 ; R_2$	Relational composition

$X \triangleleft R$	Domain restriction
$X \triangleright R$	Range restriction
$X \triangleleft\!\!\triangleleft R$	Domain anti-restriction
$X \triangleright\!\!\triangleright R$	Range anti-restriction
R^{\sim}	Relational inversion
$R(A)$	Relational image
R^+	Transitive closure
R^*	Reflexive-transitive closure

A.10 Functions

$X_1 \mapsto X_2$	Partial function
$X_1 \rightarrow X_2$	Total function
$X_1 \mapsto\!\!\mapsto X_2$	Partial injection
$X_1 \mapsto X_2$	Total injection
$X_1 \mapsto\!\!\mapsto X_2$	Partial surjection
$X_1 \rightarrow X_2$	Total surjection
$X_1 \mapsto X_2$	Bijection
$f_1 \oplus f_2$	Function overriding ($(\text{dom } f_2 \triangleleft f_1) \cup f_2$)

A.11 Numbers and finiteness

\mathbb{Z}	Set of integers ($\{\dots, -2, -1, 0, 1, 2, \dots\}$)
\mathbb{N}	Set of natural numbers ($\{0, 1, 2, \dots\}$)
$n_1 + n_2$	Integer addition infix total function
$n_1 - n_2$	Integer subtraction infix total function
$n_1 * n_2$	Integer multiplication infix total function
$n_1 \text{ div } n_2$	Integer division infix partial function
$n_1 \text{ mod } n_2$	Integer modulo infix partial function
$-n$	Integer negation prefix total function
$n_1 < n_2$	Less than relation
$n_1 \leq n_2$	Less than or equal relation
$n_1 \geq n_2$	Greater than or equal relation
$n_1 > n_2$	Greater than relation
\mathbb{N}_1	Strictly positive numbers ($\mathbb{N} \setminus \{0\}$)
$\text{succ } n$	Successor function ($n + 1$)
$\text{iter } n R$	Iteration (R composed n times)
R^n	Short form for iteration ($\text{iter } n R$)

$n_1 \dots n_2$	Number range ($\{i : Z \mid n_1 \leq i \leq n_2\}$)
$F X$	Set of finite subsets
$F_1 X$	Non-empty finite subsets ($F X \cap P_1 X$)
$\#X$	Number of members of a set
$X_1 \mapsto X_2$	Finite partial function
$X_1 \mapsto\!\!\!\rightarrow X_2$	Finite partial injection
$\min Z$	Minimum of a set of numbers
$\max Z$	Maximum of a set of numbers
$\#X$	Size of a finite set

A.12 Sequences

$\text{seq } X$	Set of finite sequences
$\text{seq}_1 X$	Set of non-empty finite sequences
$\text{iseq } X$	Set of finite injective sequences
$\langle E_1, E_2, \dots \rangle$	Notation for a sequence ($\{1 \mapsto E_1, 2 \mapsto E_2, \dots\}$)
$\langle \rangle$	Empty sequence (\emptyset)
$s_1 \hat{\ } s_2$	Sequence concatenation
$\text{head } s$	First element of sequence ($s(1)$)
$\text{last } s$	Last element of sequence ($s(\#s)$)
$\text{tail } s$	All but head of a sequence
$\text{front } s$	All but last element of a sequence
$\text{rev } s$	Reverse a sequence
$s \upharpoonright X$	Sequence filtering
\wedge / q	Distributed sequence concatenation
$\text{disjoint } I$	Disjoint ($\forall i_1, i_2 : \text{dom } I \mid i_1 \neq i_2 \bullet I(i_1) \cap I(i_2) = \emptyset$)
I partition A	Partitions ($\text{disjoint } I \wedge \bigcup \{i : \text{dom } I \bullet I(i)\} = A$)

A.13 Bags

$\text{bag } X$	Bag ($X \rightarrow \mathbf{N}_1$)
$[E_1, E_2, \dots]$	Notation for a bag ($\{E_1 \mapsto n_1, E_2 \mapsto n_2, \dots\}$)
$[]$	Empty bag (\emptyset)
$\text{count } B$	Multiplicity ($\lambda x : X \bullet 0$) $\oplus B$)
x in B	Bag membership ($x \in \text{dom } B$)
$B_1 \uplus B_2$	Bag union
$\text{items } s$	Bag of elements of a sequence

A.14 Conventions

$E_1 \underline{f} E_2$	Infix function underlined for clarity
$E_1 \underline{R} E_2$	Infix relation underlined ($(E_1, E_2) \in (-\underline{R}-)$)
$i?$	Input to an operation
$i!$	Output from an operation
i'	State component after an operation (i before)
N'	State schema after an operation (N before)
ΔN	Change of state (normally $N \wedge N'$)
ΞN	No change of state (normally $[\Delta N \mid \theta N = \theta N']$)
ΦN	Partial specification of an operation
$\vdash P$	Theorem
$d \vdash P$	Theorem ($\vdash \forall d \bullet P$)

Appendix B

Index of formal definitions

This index gives the pages on which **Z** identifiers and symbols are defined. Schema names are indicated by bold page numbers, state declarations are indicated by italic page numbers, and **Z** symbols are indicated by underlined page numbers.

- +, 4
- , 4
- >, 80
- 'Invalid window', 21
- 'No windows', 21
- 'Not a window', 21
- 'OK', 21
- '', 9
- ≤, 4
- .., 4
- ⊕/, 7
- ⊢, 83
- θ, 79
- λ, 79
- μ, 79
- Δ, 83
- ΔBlit, 46
- ΔDisplay, 14
- ΔITC, 40
- ΔSepPair, 15
- ΔSYS, 19
- ΔWM, 33
- ΔX, 56
- Φ, 83
- ΦCirc, 62
- ΦCurrent, 34
- ΦDestroy, 58
- ΦHost, 41
- ΦLayer, 47
- ΦMap, 59
- ΦOrder, 62
- ΦProc, 47
- ΦWindow, 60
- ΦWM, 41
- Ξ, 83
- ΞBlit, 47
- ΞSYS, 19
- ΞWM, 33
- ΞX, 56
- AddHost, 40
- AddWindow₀, 19
- AddWindow₁, 22
- adjust, 33
- AND, 11
- and, 13
- andInverted, 14
- andReverse, 14
- area, 31
- areas, 32
- background, 18, 32, 46, 56
- BadWindow, 21
- Bit, 5
- Black, 5
- Blit, 45, 48
- body, 31
- Box, 50
- Button, 50
- Button1, 50
- Button2, 50
- Button3, 50
- buttons, 50
- ChangeWindow₀, 61
- children, 54
- circ, 63
- CircWindowDown₀, 63
- CircWindowUp₀, 63
- Clear, 5
- clear, 13
- ClearWindow₀, 37
- ClearWindow₁, 39
- compact, 7
- ConfigureWindow₀, 61
- contents, 32
- Control, 31
- control, 31
- copy, 13
- copyInverted, 14
- CopySwap, 15
- count, 82

- CreateWindow*₀, 57
*CreateWindow*₁, 66
*CreateWindows*₀, 57
current, 32
*Current&Top*₀, 52

Decide, 25
default, 46
*DeleteWindow*₀, 35
*DeleteWindow*₁, 39
DeleteWindowITC, 42
*DelLayer*₀, 48
*DelLayer*₁, 53
destroy, 58
*DestroySubwindows*₀, 59
*DestroyWindow*₀, 58
*DestroyWindow*₁, 66
Display, 14

equiv, 14
EzecWM, 41
Ezpose, 31
*EzposeMe*₀, 36
EzposeWindow, 25
*EzposeWindow*₀, 20
*EzposeWindow*₁, 22

Failure, 52
false, 79
first, 80
front, 82

*GetDimensions*₀, 35
*GetDimensions*₁, 39
*GetProc*₀, 49
*GetProc*₁, 50
*GetProcTab*₀, 49
*GetProcTab*₁, 53
GetWindow, 23
*GetWindow*₀, 22
Gunsight, 50

head, 82

header, 31
Hide, 31
HideEzpose, 31
*HideMe*₀, 36
host, 41
hosts, 40

icon, 66
Id, 45
Info, 31, 42
InitBlit, 46
InitITC, 40
InitSYS, 19
InitWM, 33
InitX, 56
Invalid, 23
InvalidLayer, 53
InvalidParent, 65
InvalidRect, 52
InvalidWindow, 39, 66
IsInvisible, 64
IsMapped, 64
IsUnmapped, 64
ITC, 40
items, 82
iter, 81

KillWM, 41

l, 45
last, 82
Layer, 44
layers, 45
localhost, 41
lower, 8
*LowerWindow*₀, 62

Map, 31
map, 31, 61
*map*₁, 11
*map*₂, 11
MapPair, 11

- mapped*, 55
- MappedState*, 64
- maps*, 32
- MapSubwindows₀*, 59
- MapWindow₀*, 59
- max*, 5, 82
- MaxWindows*, 32
- Menu*, 42
- min*, 5, 82
- Modify*, 24
- ModifyWindow*, 25
- Mouse*, 50
- Movc*, 24
- MoveWindow*, 25
- MoveWindow₀*, 61

- NAND*, 10
- nand*, 12
- New₀*, 51
- NewLayer₀*, 47
- NewLayer₁*, 53
- NewProc₀*, 48
- NewProc₁*, 53
- NewWindow₀*, 34, 49
- NewWindow₁*, 39, 53
- NewWindowITC*, 42
- NoCurrentWindow*, 39
- noop*, 13
- nor*, 13
- NOT*, 10
- not*, 13
- NotAWindow*, 21
- NotInWindow*, 23
- NoWindows*, 21
- NULL*, 65
- Null*, 38
- NullId*, 45
- NullLayer*, 45

- offset*, 4
- opaque*, 66

- OR*, 11
- or*, 13
- order*, 55
- orInverted*, 14
- orReverse*, 14
- overlaps*, 6

- parent*, 61, 62
- parents*, 54
- pix₁*, 4
- pix₂*, 4
- Pixel*, 3
- PixelPair*, 4
- Pixmap*, 6
- Pixmap₁*, 6
- Point*, 44
- Proc*, 45
- procs*, 46
- prog*, 45
- Program*, 45

- QueryTree₀*, 64
- QueryWindow₀*, 64

- raise*, 8
- RaiseWindow₀*, 62
- RasterOp₁*, 14
- RasterOp₂*, 15
- receiver*, 46
- Rectangle*, 6
- rects*, 45
- remove*, 8
- RemoveHost*, 40
- RemoveTop₀*, 20
- RemoveTop₁*, 22
- RemoveWindow*, 25
- RemoveWindow₀*, 20
- RemoveWindow₁*, 22
- rep!*, 21
- Report*, 21
- Reshape*, 24
- Reshape₀*, 51

- ReshapeWindow*, 25
rev, 82
root, 54
RotateWindows₀, 20
RotateWindows₁, 22

samepos, 12
sameshape, 11
screen, 14, 18, 32, 46, 56
second, 80
select, 8
SelectWindow₀, 37
SelectWindow₁, 39
SepPair, 15
Set, 5
set, 13
SetDimensions₀, 35
SetDimensions₁, 39
SetTitle₀, 37
SetTitle₁, 39
setval, 6
State, 45
state, 45
Status, 65
status!, 65
String, 9
submapped, 63
suborder, 62, 63
subwindows, 54
succ, 81
Success, 21, 38, 52, 65
SYS, 18

tail, 82
title, 31
ToLayer₀, 49
ToLayer₁, 53
TooManyWindows, 38
top, 8
Top₀, 51
transparent, 66

true, 79

Undefined, 31
UnmapSubwindows₀, 60
UnmapWindow₀, 60
unviewable, 55
UpdateWindow₀, 19
UpdateWindow₁, 22

Value, 5
View, 18
viewable, 55
visible, 56

w?, 58, 59, 61–63
WhichWindow, 23
WhichWindow₀, 23
White, 5
Window, 9
WINDOWS, 33
windows, 18, 32, 45, 54
WM, 31–33, 38
wms, 40

X, 54, 55, 66
x₁, 4
x₂, 4
XOR, 11
zor, 13
Xor₁, 16
Xor₂, 16
XorSwap, 16
Xrange, 3
Xsize, 3
xy, 50
zylimits, 31

y₁, 4
y₂, 4
Yrange, 3
Ysize, 3

Zrange, 5