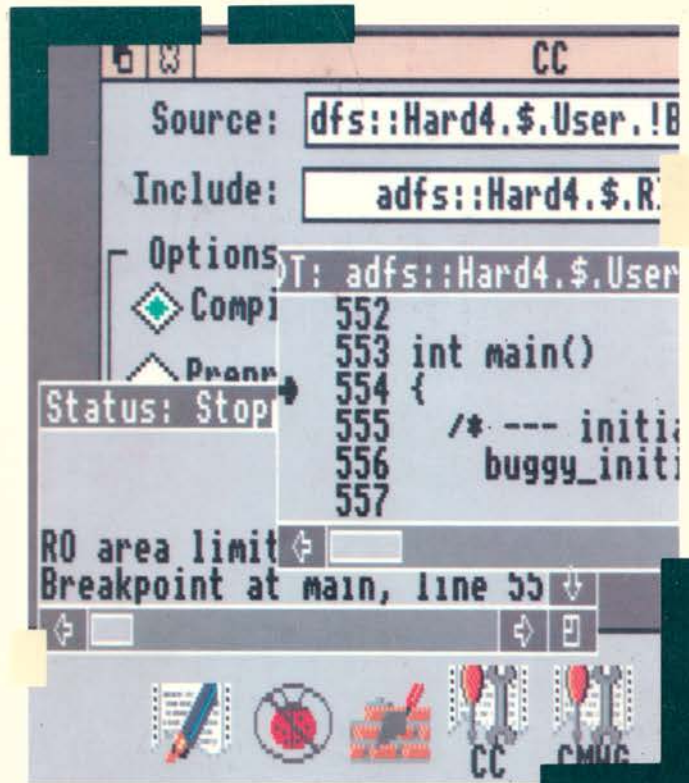


ACORN ANSI C RELEASE 4



ACORN ANSI C RELEASE 4



Copyright © Acorn Computers Limited 1991

Published by Acorn Computers Technical Publications Department

Neither the whole nor any part of the information contained in, nor the product described in, this manual may be adapted or reproduced in any material form except with the prior written approval of Acorn Computers Limited.

The product described in this manual and products for use with it are subject to continuous development and improvement. All information of a technical nature and particulars of the product and its use (including the information and particulars in this manual) are given by Acorn Computers Limited in good faith. However, Acorn Computers Limited cannot accept any liability for any loss or damage arising from the use of any information or particulars in this manual.

If you have any comments on this manual, please complete the form at the back of the manual, and send it to the address given there.

Acorn supplies its products through an international dealer network. These outlets are trained in the use and support of Acorn products and are available to help resolve any queries you may have.

Within this publication, the term 'BBC' is used as an abbreviation for 'British Broadcasting Corporation':

ACORN, ARCHIMEDES and ECONET are trademarks of Acorn Computers Limited.

UNIX is a trademark of AT&T.

Published by Acorn Computers Limited

Release 4

ISBN 1 85250 095 6

Part number 0470,529

Product number SKD35

Issue 1, May 1991

Contents

Contents iii

Introduction 1

- Installation of Acorn Desktop C 1
- The C compiler 2
- This user guide 2
- Useful references 5

Part 1 - Using the C tools 7

C tools and the DDE 9

- Using C tools through Make 9
- Editor throwback 10
- DDT debugging 13
- Using FrontEnd on your programs 18
- Making your own linkable libraries 19

CC 21

- Getting started with CC 21
- C libraries 24
- File naming and placing conventions 25
- Include file searching 29
- The SetUp dialogue box 33
- The SetUp menu 35
- The Application menu 49
- CC output messages 50
- Command line interface 51
- Worked examples 53

CMHG 59

- Starting CMHG 59
- The Application menu 60
- Example output 61
- Command line interface 61

ToANSI 63

- ToANSI C translation 63
- Starting ToANSI 64
- The Application menu 65
- Example output 65
- Command line interface 66

ToPCC 67

- ToPCC C translation 67
- Starting ToPCC 68
- The Application menu 69
- Example output 69
- Command line interface 70

Part 2 - Language issues 71

Implementation details 73

- Identifiers 73
- Data elements 73
- Structured data types 76
- Pointers 77
- Arithmetic operations 77
- Expression evaluation 78
- Implementation limits 78

Standard implementation definition 81

- Translation (A.6.3.1) 81
- Environment (A.6.3.2) 81
- Identifiers (A.6.3.3) 82
- Characters (A.6.3.4) 83
- Integers (A.6.3.5) 84
- Floating point (A.6.3.6) 84
- Arrays and pointers (A.6.3.7) 84
- Registers (A.6.3.8) 84
- Structures, unions, enumerations and bitfields (A.6.3.9) 85
- Qualifiers (A.6.3.10) 85
- Declarators (A.6.3.11) 85
- Statements (A.6.3.12) 85
- Preprocessing directives (A.6.3.13) 86
- Library functions (A.6.3.14) 86

Portability 91

- General portability considerations 91
- ANSI C vs K&R C 94
- The ToPCC and ToANSI tools 98
- pcc compatibility mode 98
- Environmental aspects 102

ANSI library reference section 105

- assert.h 105
- ctype.h 105
- errno.h 106
- float.h 107
- limits.h 107
- locale.h 107
- math.h 108
- setjmp.h 109
- signal.h 110
- stdarg.h 111
- stddef.h 113
- stdio.h 113
- stdlib.h 127
- string.h 136
- time.h 141

Part 3 - Developing software for RISC OS 147**How to write desktop applications in C 149**

- Some general principles 149
- Developing an application from scratch 151
- More RISC_OSLib facilities 159
- Using Draw files 167
- Common application features 168
- Displaying and editing text 172
- Tracing desktop applications 173
- Where do you go from here? 174

RISC OS library reference section 175

- akbd 175
- alarm 176
- baricon 178
- bbc 180
- bbc: text output functions 180
- txt: graphics output functions 181
- bbc: other calls 184
- colourmenu 185
- colourtran 186
- coords 191
- dbox 195
- dbox: creation and deletion functions 195
- dbox fields 197
- dbox: events from dialogue boxes 199
- dbox: pending operations 201
- drawdiag 205
- draw: memory allocation functions 207
- draw: unknown object handling 209
- drawferror 210
- drawfobj 212
- drawftypes 216
- drawmod 216
- event 220
- event: masking off events 222
- fileicon 222
- flex 223
- font 225
- fontlist 232
- fontselect 233
- heap 235
- help 235
- magnify 238
- menu 239
- msgs 242
- os 244
- pointer 245
- print 246
- res 252
- resspr 253

- saveas 253
- sprite 254
- sprite: simple operations 255
- sprite: operations on system/user area 255
- sprite: operations on system/user area, name/sprite pointer 257
- template 261
- trace 263
- txt 264
- txt: interface functions 264
- txt: general control operations 265
- txt: operations on the array of characters 267
- txt: layout-dependent operations 270
- txt: operations on markers 271
- txt: operations on a selection 272
- txt: input from the user 273
- txt: direct access to the array of characters 276
- txt: system hook 276
- txtedit 277
- txtopt 282
- txtscrap 283
- txtwin 283
- visdelay 284
- werr 285
- wimp 286
- Function prototypes 304
- wimpt 312
- wimpt: control of graphics environment 314
- win 316
- win: menus 320
- win: event processing 320
- win: termination 321
- xferrecv 323
- xfersend 325
- xfersend: caller-supplied function types 325
- xfersend: library functions 326

Assembly language interface 331

- Register names 332
- Register usage 332
- Control arrival 333
- Passing arguments 333
- Return link 334
- Structure results 334
- Storage of variables 335
- Function workspace 335
- Examples 335

How to write relocatable modules in C 337

- Getting started 337
- Constraints on modules written in C 338
- Overview of modules written in C 338
- Functional components of modules written in C 338

Overlays 351

- Paging vs overlays 351
- When to use overlays 352

Using memory efficiently 355

- Guidelines 355
- Recovery from lack of memory 355
- Avoiding permanent loss of memory 356
- Avoiding memory wastage 357
- Using `heap_alloc` and `heap_free` 365

Machine-specific features 367

- How to use the C library kernel 367
- Calling other programs from C 374
- The shared C library 376
- #pragma** directives 378
- Storage management (`malloc`, `calloc`, `free`) 379
- Handling host errors 379

Appendices 381

Appendix A: New features of Desktop C 383

- Acorn Desktop C tools 383
- New technical features 385
- User guides 385

Appendix B: Errors and warnings 387

Interpreting CC errors and warnings 387

Warnings 389

Non-serious errors 397

Serious errors 407

Fatal errors 422

System errors 423

Appendix C: kernel.h 425

Appendix D: The floating point emulator 433

FPE280 433

Using the compiler 434

Floating point requirements of applications 434

Subject Index 437

Function Index 447

Introduction

Acorn Desktop C is a development environment for producing RISC OS desktop applications and relocatable modules written in ANSI C. It consists of a number of programming tools which are RISC OS desktop applications. These tools interact in ways designed to help your productivity, forming an extendable environment integrated by the RISC OS desktop. Acorn Desktop C may be used with its sister product, Acorn Desktop Assembler, to provide an environment for mixed C and assembler development.

Acorn Desktop C includes tools to:

- edit program source and other text files
- search and examine text files
- convert C source and header text between ANSI and Unix dialects
- examine some binary files
- compile and link C programs
- construct relocatable modules entirely from C
- compile and construct programs under the control of makefiles, these being set up from a simple desktop interface
- squeeze finished program images to occupy less disk space
- construct linkable libraries
- debug RISC OS desktop applications interactively
- construct template files for RISC OS desktop applications.

Most of the tools in Acorn Desktop C are also of general use for constructing applications in other programming languages, and are, for example, supplied with Acorn Desktop Assembler. These non-language-specific tools are described in the accompanying *Acorn Desktop Development Environment* user guide.

Installation of Acorn Desktop C

Installation of Acorn Desktop C is described in the accompanying *Acorn Desktop Development Environment* user guide.

The C compiler

The Acorn C compiler for the Archimedes (the tool CC supplied as part of Acorn Desktop C) is a full implementation of C as defined by the 1989 ANSI language standard. To obtain this standard document, see the section entitled *Useful references* on page 5. It is tested with the Plum-Hall C Validation Suite version 2.00 available at time of publication, and passes all sections, except for failing to produce two required diagnostic messages, as described in the release note accompanying this user guide in Acorn Desktop C.

This user guide

This guide is a reference manual for the C tools CC, CMHG, ToANSI and ToPCC working as part of the development environment of Acorn Desktop C. These are the only tools in the Acorn Desktop C product which are not used for programming in other languages and described in the accompanying *Acorn Desktop Development Environment* user guide. This manual also documents the C library support provided and other aspects that are particular to this C product:

- special features of this implementation of the C language
- operating the Acorn Desktop C tools specific to the C language
- portability issues, including the portable C compiler (pcc) facility
- developing programs for the RISC OS environment:
 - Desktop applications
 - Relocatable modules
 - Overlays
 - Calling other programs and languages from C.

This guide is not intended as an introduction to C and does not teach C programming, nor is it a reference manual for the ANSI C standard. Both these needs are addressed by publications listed in the section entitled *Useful references* on page 5.

This guide is organised into four parts:

Part 1: *Using the C tools*

Part 2: *Language issues*

Part 3: *Developing software for RISC OS*

Part 4: *Appendices*

Part 1: Using the C tools

This part of the guide describes the operation of the four C specific programming tools. It consists of five chapters, the first describing the interaction of the C tools with the rest of the Desktop Development Environment, the others being each devoted to individual tools. Examples in the text and on disc are used to illustrate several points.

The chapters are:

- *C tools and the DDE*
- CC
- CMHG
- ToANSI
- ToPCC

Part 2: Language issues

This covers issues to do with the C programming language itself, in particular those parts of the ANSI standard that are necessarily machine- or operating system-specific. It also includes a chapter on portability to help with porting applications in C to and from RISC OS.

The chapters are:

- *Implementation details*
How Acorn C implements those aspects of the language which ANSI leaves to the discretion of the implementor.
- *Standard implementation definition*
How Acorn C behaves in those areas covered by Appendix A.6 of the draft standard (which lists those aspects which the standard requires each implementation to define).
- *Portability*
The chapter covers:
 - portability considerations in general
 - the major differences between ANSI and 'K&R' C
 - using the pcc compatibility mode of the Acorn compiler
 - standard headers and libraries
 - environmental aspects of portability.
- *ANSI library reference section*
This chapter works through the headers of the ANSI standard library, (assert.h to time.h), outlining the contents of each one:

- function prototypes
- macro, type and structure definitions
- constant declarations.

Part 3: Developing software for RISC OS

This part of the Guide tells you how to write software in C for the RISC OS environment. Examples in the text and on disc are used to illustrate each type of program development.

The chapters are:

- *How to write desktop applications in C*
This covers the principles of designing an application to be integrated into the RISC OS desktop environment.
- *RISC OS library reference section*
A list of fully commented headers for the RISC OS library. This library provides the high-level interface to RISC OS, with all the calls needed to program for the Wimp environment.
- *Assembly language interface*
How to handle procedure entry and exit in assembly language, so that you can write programs which interface correctly with the code produced by the C compiler.
- *How to write relocatable modules in C*
Relocatable modules – the building blocks of the RISC OS operating system – are needed for device drivers and similar low-level software.
- *Overlays*
This chapter explains how to write an application using overlays, with a worked example as an illustration.
- *Using memory efficiently*
This chapter gives hints and describes how the C library memory allocation works, to aid you in writing memory efficient RISC OS applications and relocatable modules.
- *Machine-specific features*
This chapter contains the following sections:
 - The C library kernel
 - Calling other programs from C
 - The shared C library
 - #pragma directives
 - Storage management

- Handling host errors.

Part 4: Appendices

Appendix A: New features of Desktop C

This is the fourth release of the C compiler product for Acorn computers running the RISC OS operating system. The appendix highlights all those features that are new since the previous release (release 3).

Appendix B: Errors and warnings

Messages produced by the compiler, of varying degrees of severity.

Appendix C: kernel.h

Fully-commented headers for the C library kernel. This provides the technical details needed to support the explanatory section on the kernel in the chapter *Machine-specific features*.

Appendix D: The floating point emulator

This covers what you need to know about the floating point emulator in order to use the C compiler system and write applications using it.

Conventions used

Throughout this Guide, a fixed-width font is used for text that the user should type, with an italic version representing classes of item that would be replaced in the command by actual objects of the appropriate type. For example:

`cc options filenames`

This means that you type `cc` exactly as shown, and replace *options* and *filenames* by specific examples.

A bold version of the same font is used for text that the computer responds with.

The abbreviation DDE is used in later chapters to mean Desktop Development Environment.

Useful references

C programming

- Harbison, S P and Steele, G L, (1984) *A C Reference Manual*, (second edition). Prentice-Hall, Englewood Cliffs, NJ, USA. ISBN 0-13-109802-0.

This is a very thorough reference guide to C, including a useful amount of information on the ANSI C standard.

Since the Acorn C compiler is an ANSI compiler, this book is particularly relevant, but you must get the second edition for coverage of the ANSI standard.

- Kernighan, B W and Ritchie, D M, (1988) *The C Programming Language* (second edition). Prentice-Hall, Englewood Cliffs, NJ, USA. ISBN 0-13-110362-8.

This is the original C 'bible', updated to cover the essentials of ANSI C too.

- Koenig, A, (1989) *C Traps and Pitfalls*, Addison-Wesley, Reading, Mass. ISBN 0-201-17928-8.

This book explains how to avoid the most common traps and pitfalls that ensnare even the most experienced C programmers. It provides informative reading at all levels.

RISC OS

- The *User Guide* supplied with your computer, which describes how to use the RISC OS operating system and the applications Edit, Paint and Draw.
- The *RISC OS Programmer's Reference manual*.
- The *RISC OS Style Guide*.

Reference cards

Included with this Guide are three reference cards, summarising:

- the contents of the release discs
- an overview of the installed Acorn Desktop C directory structure
- an overview of the structure of the RISC OS library.

The ANSI standard

The American National Standard for Information Systems – Programming Language C is available with the reference number ANSI X3.159-1989 for £45.00 from:

British Standards Institution
Foreign Sales Department
Linford Wood
Milton Keynes
MK14 6LE

Members of the BSI can order copies by telephone; non-members should send a cheque payable to BSI.

However, you should find the coverage of ANSI C in this manual and the books listed above adequate for all but the most demanding requirements.

Part 1 - Using the C tools

The C compiler (CC), C module header generator (CMHG) and the C dialect conversion programs ToANSI and ToPCC are the only tools included in Acorn Desktop C which are specific to programming in the C language, and hence described in this volume. All the other tools, such as the editor and debugger are described in detail in the accompanying *Acorn Desktop Development Environment* user guide.

CC, CMHG, ToANSI and ToPCC all fit into the non-interactive class of DDE tools – once you have started one of these tools with selected options, you cannot interact with it to modify its behaviour, except to view output and pause or stop it. The DDT debugger is an example of an interactive DDE tool. The non-interactive DDE tools all have many features in common, and these are described in more detail in the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide.

To load CC, CMHG, ToANSI or ToPCC into the desktop, open a directory display on the DDE directory of your work disk and double click on !CC, !CMHG, !ToANSI or !ToPCC. The tool icon then appears on the icon bar. Each icon has the standard screwdriver-and-spanner appearance of all the non-interactive DDE tools:



From these icons you have access to the interface to set options and start tasks unmanaged by Make. Each of these interfaces has its own chapter later in this guide.

Using C tools through Make

The DDE Make tool is designed to manage efficient construction of programs and libraries, usually from several source files. It is designed to avoid needless re-processing of unaltered source files and ensure consistent construction by a method specified in a Makefile. For more details, see the chapter entitled *Make* in the accompanying *Acorn Desktop Development Environment* user guide.

The C tools, like the other non-interactive DDE tools, can be used by Make to process files. When managed by Make, CC, CMHG, ToANSI and ToPCC are controlled by command lines issued by Make, and their icons need not be present on the icon bar – you don't need to double click on !CC, !CMHG, !ToANSI or !ToPCC before starting a Make job using these tools. The command lines issued by Make to the C tools are calculated from the contents of the Makefile controlling the job in progress. The command lines understood by each C tool are described in the chapters on individual tools later in this volume. You do not need to understand the details of the command lines contained in your Makefiles. Instead, you can adjust them using the same desktop interface as that available from each tool's icon bar icon. To do this you follow the Make **Tool options** menu item and click on the name of the tool concerned.

Editor throwback

During development of a program you may well find that you spend a high proportion of your time repeatedly editing, compiling and testing. Throwback to the SrcEdit editor is designed to make this development cycle very quick and easy when removing compilation errors from your sources.


If you enable the CC **Throwback** option and compile a file which causes an compiler error, then a Browser window is presented by the editor (assuming SrcEdit and the DDEUtils module are loaded). Double clicking Select on an error line in this browser window makes the editor open an edit window displaying the source file causing the error, with the offending line in view and highlighted, ready for correction. This facility can be used whether compilation is being managed by Make or performed using the CC icon bar interfaces.

Example throwback session


First double click on !SrcEdit and !CC in a directory display to load them as applications with icons on the icon bar. Next open a directory display on the subdirectory `User.CompErr.c` to show the text file `CompErr` containing the source of the program example of that name.

`CompErr` is a trivial C program which when run prints `Hello World` on the screen. It is written to be compiled with integral link step by CC to form an executable image file. Its source contains a simple error which will be detected by CC when you try to compile it.

Drag the source file `CompErr` to the CC icon to make the CC SetUp dialogue box appear with the **Source** writable icon initialised to the absolute file name. Ensure that the **Throwback** option is enabled. The correct dialogue box appearance is as follows:



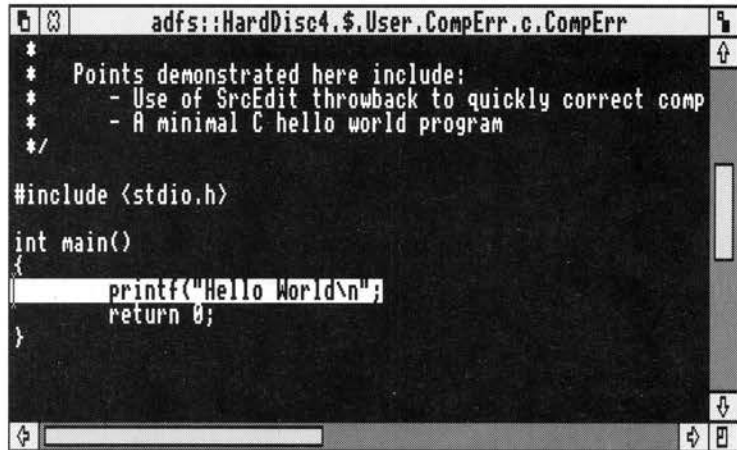
Click Menu on the setup box and ensure that the **Work directory** item on the menu displayed has the default setting of '^'. Click on the **Run** button on the SetUp box to start compilation. This has the normal effect of removing the setup box and putting the CC output display on the screen, but almost immediately afterwards the compiler produces an error and requests SrcEdit to display a Throwback error browser:



Double click Select on the compiler error message in the browser:

`expected ')' or ',' - inserted ')' before ';''`


SrcEdit displays the source file with the offending line that caused the error clearly highlighted:



```
adfs::HardDisc4.$User.CompErr.c.CompErr
*
* Points demonstrated here include:
* - Use of SrcEdit throwback to quickly correct comp
* - A minimal C hello world program
*/
#include <stdio.h>
int main()
{
printf("Hello World\n");
return 0;
}
```

Examining this line closely shows that a closing bracket is missing before the ending semicolon. Insert this bracket in SrcEdit and save the file. Click Select on the CC icon bar icon and click the Run button to repeat the last compilation. If you have changed the `CompErr` source correctly, the compilation should now complete with no errors, hence without bringing back the SrcEdit browser.

When the CC save dialogue box appears, click on the OK button to save the executable file produced in the directory `User.CompErr`. Now double click Select on the newly created executable image file in a directory display. The image file should run, printing the `Hello World` message in a RISC OS run window:



```
Run adfs::HardDisc4.$User.CompErr.CompErr
Hello World

Press SPACE or click mouse to continue
```

DDT debugging

The DDT debugger is designed to debug RISC OS desktop applications or command line programs written in C (but not relocatable modules). You can debug your C program at machine level (ie displaying current execution position on a disassembly of memory) or more conveniently, at source level (ie displaying current execution position as a line of C indicated in your source file).

To debug your program at source level you must construct it with the **Debug** option of the tools CC and Link enabled. Executing a binary linked in this way or dragging it from a directory display to the DDT icon bar icon starts a debugging session on it. See the chapter entitled *Desktop debugging tool* in the accompanying *Acorn Desktop Development Environment* user guide for more details of DDT.

Example DDT session

This session demonstrates source level debugging of a RISC OS desktop application written in C.

First double click on !Boot, then !DDT and !CC in the DDE directory display to load them as applications with icons on the icon bar. Next open a directory display on the subdirectory `User . !Buggy . c` to show the text file `Buggy` containing the source of the program example of that name.

`Buggy` is a complete RISC OS desktop application with its !RunImage executable image file constructed from the one C file mentioned above. The C makes use of calls to the RISC_OSLib library, and implements an application which when working behaves almost identically to `Ball64` – clicking Select on its icon bar icon displays a RISC OS window full of coloured balls. A bug has been deliberately added to `Buggy` for this session, resulting in the window not appearing when you click on its icon.

Drag the text source file of `Buggy` to the CC icon on the icon bar to bring up the CC setup dialogue box. Ensure that the **Compile only** option is disabled then click on the **Run** button to start compilation. Save the executable image file as `User . !Buggy . !RunImage`, then double click Select on !Buggy. The `Buggy` icon should appear on the icon bar, with an appearance similar to that of `Ball64`, but with a black cross over it:



Try clicking Select on this icon. (Nothing happens, but the balls window is supposed to appear...)

Click Select on the CC icon, and repeat the previous compilation, this time ensuring that the **Debug** option is enabled. Save !RunImage as before. You may notice that the icon of the !RunImage file now has a bug crawling on it.

Quit the previously loaded Buggy application from the icon bar, then double click on !Buggy again to start a DDT session of its !RunImage program. The two main DDT windows appear, with the application in its initialisation code before reaching your source:

```
DDT: adfs::HardDisc4.$.User.!Buggy.!RunImage
00007ff4: 00000000 dcd      0
00007ff8: 00000002 andeq    r0,r0,r2
00007ffc: 0000003b dcb      ",",0,0,0
+00008000: fb000000 blnv     &00008008
00008004: fb000000 blnv     &0000800c
00008008: eb00000c bl       &00008040
0000800c: eb0026e9 bl       Stub$$Code ; &00011bb8
00008010: ef000011 swi     OS_Exit
00008014: 0000a0ec andeq    sl,r0,ip.ror #&01
```

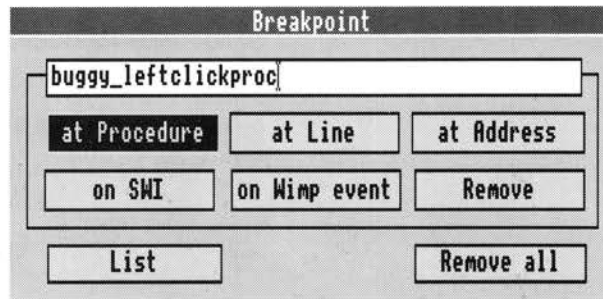
Status: Initialisation

RO area limit not on page boundary, last page not protected

If you had inspected or written the source code of Buggy, you would have known that the procedure `buggy_leftclickproc` is registered during initialisation as the handler for clicking Select on the Buggy icon. A sensible first step is therefore to set a breakpoint on entry to this procedure to see if it is ever called.

Clicking Menu on either DDT window displays the DDT menu, from which you gain access to its many features. For a detailed description of these see the chapter entitled *Desktop debugging tool* in the accompanying *Acorn Desktop Development Environment* user guide.

Select on the **Breakpoint** item to bring up the dialogue box for setting breakpoints. Type the name `buggy_leftclickproc` in the writable icon, then click on the **at Procedure** action button to set the desired breakpoint:



Now run your application by clicking Select on the DDT menu **Continue** item. If the **Stop at entry** DDT option is enabled, execution stops on entry to the C code in `main`. This is of no interest now, so merely click on **Continue** again to proceed. DDT displays disappear for the minute, and Buggy executes, putting its icon on the icon bar as before. Next click Select on the Buggy icon. The DDT windows reappear, indicating that the breakpoint has been reached:

```

DDT: adfs::HardDisc4.$User.!Buggy.!RunImage
222
223
224 static void buggy_leftclickproc(wimp_i i)
225 {
226     wimp_wstate state;
227     wimp_redrawstr r;
228     i=i;
229
230     if(displaying == FALSE)

```

```

Status: Stopped at Breakpoint


```

```

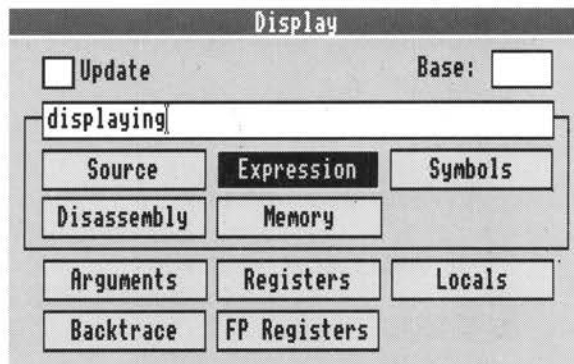
RO area limit not on page boundary, last page not protected
Breakpoint set at Procedure buggy_leftclickproc
Breakpoint at main, line 554 of adfs::HardDisc4.$User.!Buggy.c.buggy
Breakpoint at buggy_leftclickproc, line 225 of adfs::HardDisc4.$User.!B

```

We know that `buggy_leftclickproc` is being reached, but it isn't working properly, so a sensible thing to do is to single step through its code. To do this, click on the **Single step** item on the DDT menu to bring up the single step dialogue box:




Use this to single step by one source line at a time by clicking Select on the **OK** button, and watch where execution is going. Looking in the DDT Context window, you can see that the main body of this procedure is being stepped straight past without being executed. Enlarging and scrolling this window lets you examine the source of this procedure and see that the condition not being met is that the variable `displaying` is not set to `FALSE`. Use the **Display** dialogue box of DDT to examine its contents:



Rather than having a value of 0 or 1 as you would expect of a flag variable, it contains a strange large number. Perhaps it is being corrupted.

To see if the variable is being corrupted, we use the **Watchpoint** facility of DDT to see where `displaying` is being changed. First you must restart Buggy, as the variable already has its strange value at this stage. Select **Quit** on the DDT menu to


leave the session, then double click on !Buggy to restart it. The DDT windows appear with Buggy in its initialisation state again. Use the **Watchpoint** dialogue box of DDT to set a watchpoint on the variable `displaying`:



Now use **Continue** to set Buggy running again. The DDT windows reappear, showing that `displaying` has changed to its silly value. Now use the **Display** dialogue box to show a stack backtrace. This shows that the line of your source executing when `displaying` changed was line 542 of the C source. You now have all the clues you need from DDT.

Quit DDT and load the C source of Buggy into SrcEdit. Inspecting the offending line (using SrcEdit GOTO) shows that it is nothing to do with `displaying`, the only variable used being called `trans`. Look at the definition of `trans` (the first occurrence of the string `trans` in the source file). It is a static array, followed by the static variable `displaying`. If you look at other arrays of the same type as `trans`, you will see that they all have 256 elements rather than 252. So, the solution is to change the number of elements of `trans` from 252 to 256 and save the source file.

Reconstruct Buggy with CC, this time with **Debug** disabled, and double click on !Buggy. The icon appears as before, so click Select on it and the Display window now appears:



The problem should now be cured. For the best effect, set the screen mode to a 256 colour mode such as 15.

Using FrontEnd on your programs

FrontEnd is a relocatable module supplied as part of the Acorn Desktop C product which provides RISC OS desktop interfaces for non-interactive command line programs. The DDE non-interactive tools (such as CC and CMHG) are each command line programs supported in RISC OS by FrontEnd which gives you the effect that each is a fully multitasking windowed RISC OS application. See the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide for more details of non-interactive DDE tools.

You can use the power of FrontEnd to produce your own RISC OS applications. To do this you need to construct:

- A suitable command line program
- A Templates file (constructed with FormEd)
- A Sprites file (constructed with Paint)
- A !Run file
- A !Help text file containing a short description of your program
- A Messages text file
- A Desc front end description file.

To be suitable, your command line program has to be non-interactive. This means it should start with a command line, then run to error or completion without any further user interaction, outputting reports as screen text. A compiler such as CC fits this description, but an editor such as SrcEdit does not.

The Desc front end description file contains a specification of the appearance and function of the desktop interface to be provided for your program by FrontEnd. It is written in a special description language understood by FrontEnd. For more details of how to produce this file see the chapter entitled *Extending the DDE* in the accompanying *Acorn Desktop Development Environment* user guide. You may find it easier to make this file by altering a description belonging to one of the non-interactive tools rather than writing your own from scratch.

The tool ToANSI is a simple example of the non-interactive DDE tools. You may find it instructive to open a viewer on the directory DDE.!ToANSI and examine the file Desc and others.

Making your own linkable libraries

Linkable libraries, which are usually filed in o subdirectories like object files, are collections of many object files stored in one file. When presented to Link as an input file, the referenced object files within a library are linked into the output file, but those not needed are left out. A linkable library is therefore the recommended way of storing a selection of useful procedures for re-use in a number of programs. You may well find that this facility can save you a lot of time.

Use the LibFile tool to construct and modify linkable libraries. Use the DecCF tool to decode information about an existing library.

The library RISC_OSLib.o.RISC_OSLib supplied with Acorn Desktop C is a linkable library constructed with LibFile. Try taking a copy of it and using DecCF and LibFile to examine it and extract and replace object files from it.



CC is the C compiler of Acorn Desktop C. It is a full implementation of ANSI C as described in the chapter entitled *Introduction* on page 1. It processes text files containing the source and headers of programs into linkable object files. It can also drive Link directly (without double clicking on !Link to put its icon on the icon bar) to produce executable image files.

Like other non-interactive DDE tools, CC can be managed by Make, with its compiling options specified by the makefile passed to Make. For such managed use, CC is started automatically by Make – you don't have to load CC onto the icon bar.

The characteristics of CC as a language implementation are defined in parts 2 and 3 of this volume. Most of the rest of this chapter covers the CC options, and gives some programming examples.

If you are new to RISC OS and the Acorn C compiler, read the whole of this chapter before starting to use the C compiler system. If you are an experienced C programmer, you will find this chapter essential for reference, and may choose to tackle the section entitled *Worked examples* on page 53 first.


Getting started with CC

Like other non-interactive DDE tools, CC can be managed by Make, with its compiling options specified by the makefile passed to Make. For such managed use, CC is started automatically by Make – you don't have to load CC onto the icon bar.

To use CC directly, unmanaged by Make, first open a directory display on the DDE directory, then double click Select on !CC. There is no file type to double click on to start CC – CC owns no file type unlike, for example, Draw. The CC main icon appears on the icon bar:




Clicking Select on this icon, or dragging a C source file from a directory display to this icon, brings up the CC SetUp dialogue box. To see this work, try opening a directory display on the directory `User.HelloW.c` and drag the source text file `HelloW` to the CC icon. The CC SetUp dialogue box appears:



As you have dragged a source file to bring up this dialogue box, its name appears in the writable **Source** icon, otherwise this icon would have appeared containing the name of the last filename entered there. It would be empty if there were no previous filename.

Clicking Menu on the SetUp dialogue box brings up the CC SetUp menu:



The SetUp dialogue box and menu specify the next compilation to be done. You start the next job by clicking Select on the **Run** action button on the dialogue box (or Command line menu dialogue box). Clicking Select on the **Cancel** action button removes the SetUp dialogue box and clears any changes you have just made to the options settings back to the state before you brought up the SetUp box. The options last until you adjust them again or !CC is reloaded, or you can save them for future use with an item from the main icon menu.

Ensure that the option settings are the defaults, as in the above pictures. Click on the **Run** button to compile the HelloW example with an integral link step. Save the executable image file produced as User.HelloW.HelloW, then double click Select on its name in the directory display to run it. The program runs, putting a Hello World message in the standard RISC OS command line window:

```

Run adfs::HardDisc4.$User.HelloW.HelloW
Hello World

Press SPACE or click mouse to continue

```

In processing HelloW in this example, the source file was first compiled, then linked. CC operates in a work directory. This is controlled by the **Work directory** option on the CC SetUp menu. In the example, the source was in the directory User.HelloW.c, and the option had its default setting, so the work directory was User.HelloW.c.^, ie User.HelloW. The source file was first compiled to form an object file in the o subdirectory of the work directory – the file User.HelloW.o.HelloW. Since the CC SetUp dialogue box option **Compile only** was not enabled, this object file was then linked with the object file CLib.o.Stubs which gives access to the shared C library ANSI functions such as printf, to produce the executable image file you saved.

If you request assembly language or listing output from a compilation by enabling the CC SetUp menu options **Assembler** or **Listing**, these are placed in s or l subdirectories of the work directory respectively. Like an object file produced by

CC, these have the same filename as the source. The `h` subdirectory of the work directory is the place to put your program headers when you refer to them with a C source line such as:

```
#include "myheader.h"
```

For more details of file names and directories used, see the section entitled *File naming and placing conventions* on page 25.

C libraries

There are two types of library provided to support the C compiler:

- The standard ANSI library (also referred to as the C library)
This provides all the standard facilities of the language, as defined by the ANSI standard document. Code using calls to the ANSI library will be portable to other environments if an ANSI compiler and library are available for that environment.
The ANSI library used with the Acorn C compiler system is called the shared C library. It is a relocatable module, `$.!System.modules.Clib`, and must be loaded in the relocatable module area before executing C programs using it. C programs are linked with a small piece of code and data, called `Stubs`, which interfaces with the shared C library. `Stubs` is found in the directory `CLib.o`. The idea behind the shared C library is that a number of applications which are resident in memory at the same time can use it, thus economising on RAM space. It also saves space on disc, benefiting users with single floppy disc drives.
- The operating system library
This provides you with routines to harness the special facilities of the operating system, in particular the Wimp environment. Code using calls to this library will not port to other environments.
The operating system library used with the Acorn C compiler system is the RISC OS library, `RISC_OSLib`. It provides all the calls you need to program the Wimp environment and write desktop applications. The linkable library file `RISC_OSLib` is in the directory `RISC_OSLib.o`, and its headers are in `RISC_OSLib.h`.

Shared C library

The shared C library provided with Acorn Desktop C is a new version, replacing those provided with previous products such as ANSI C Release 3. It is backwards compatible so that existing software will run with it. However, software compiled with Acorn Desktop C will not work with the old shared C libraries.

C programs are linked not with the C library but with a small piece of code and data called stubs. The stubs contain your program's copy of the library's data and an *entry vector* which allows your program to locate library routines in the C library module.

Use of the shared C library:

- saves space on disc
- makes programs load faster
- costs practically nothing at run time (for example, the Dhrystone benchmark runs just as quickly using the shared C library as when linked stand-alone with ANSILib)
- typically costs less than 30KB of memory (the shared C library plus stubs occupy about 65Kb, whereas most C programs include about 40KB of ANSILib).

Without the shared C library, it would not be possible to pack so much into Acorn Desktop C.

Updating shared C libraries in your machine or those of people using the programs you have constructed has to be performed with care to avoid crashing and possible loss of data.

When an application such as Edit or SrcEdit is run which uses the shared C library, the application needs to know where the library module is in memory, so that it can locate the library routines when required. If an application such as Edit is installed when you replace the shared C library in memory with another version (with *RMLoad from the command line or by double clicking on the CLib module in a directory display) the application will crash the machine. This is because it is left pointing at the old addresses of routines which have moved.

The solution to this problem is simple – ensure that at boot up your machine loads CLib from !System.modules and that the latest CLib is in this directory, then reboot your machine. The installation of the DDE product on your machine ensures that this takes place if you have a hard disc, or when the DDE is booted if you have only floppy discs. You have to ensure that this takes place on the machines of anybody using the software you have produced with Acorn Desktop C.

File naming and placing conventions

This section explains the concept of a work directory, and describes the naming conventions used to identify the different classes of file you will come across when using Acorn Desktop C.

Work directory

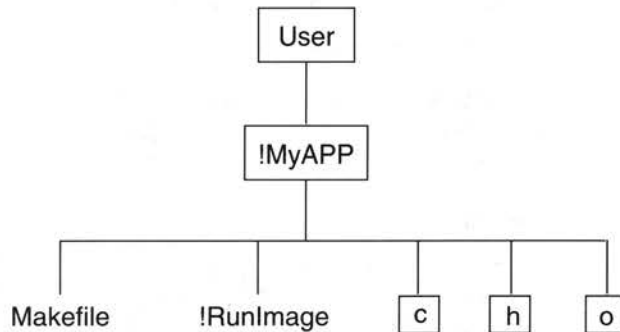
CC operates in a work directory. The work directory is where CC places all output files produced by it which are not explicitly placed by you by dragging from a Save dialogue box to a directory display. This includes object files to be linked by an integral link step, assembly language output and listing output. The work directory is also a place where some input source and header files are looked for – see the next sections for more details.

For previous Acorn C compilers, such as that provided in ANSI C Release 3, the compiler work directory was the currently-selected directory of the currently-selected filing system. Since several CC tasks and other DDE tasks depending on work directories can be multitasking, all operating at the same time, the single RISC OS current directory can no longer be used as the work directory by CC.

When managed by Make, the CC work directory is simply the directory containing the makefile controlling the job.

When not managed by Make, the CC work directory is formed from the directory containing the source file, modified by the relative path name specified by the **Work directory** SetUp menu option. The default **Work directory** SetUp menu value is ^. Thus when, for example, compiling `User.HelloW.c.HelloW` with default settings, the work directory is `User.HelloW.c.^`, ie `User.HelloW`.

A typical directory arrangement is:



The resource files (such as !Run, Templates) normally found in an application directory are not shown above. With directories arranged as above and default option settings, the managed and unmanaged work directories are the same – `User.!MyApp`.

Filename conventions

The Acorn C system, in common with many other C systems, uses naming conventions to identify the classes of file involved in the compilation and linking process. Many systems use conventional suffixes for this. For example, the suffix `.c` denotes C source files on UNIX and MS-DOS systems. This convention clashes with Acorn's use of the full-stop character in pathnames. It is more natural under Acorn filing systems to use a prefix convention, eg `c.foo`, where `c` is the directory containing C source files, and `foo` is the filename.

However, portability is an increasingly important issue in the C world. To this end, CC recognises the standard file naming conventions and performs the appropriate transformations to construct valid RISC OS pathnames. The following sections summarise the conventions for referring to source, include, object and program files.

Rooted filenames

A filename is rooted if it is

- a RISC OS filename beginning with a `$` or an `&`. For example:


```
$.RISC_Oslib.h.baricon &.h.myheader
```
- a UNIX filename beginning with a `/`. For example:


```
/RISC_Oslib/baricon.h
```
- an MS-DOS filename beginning with a `\`. For example:


```
\library\baricon.h
```

Rooted filenames are used by CC as absolute specifications of filenames, independent of work directories, search paths, etc. Rooted UNIX or MS-DOS filenames are converted into the Acorn syntax and prefix forms.

Source files

Source files to be compiled are specified on the CC command line produced from your settings of the CC SetUp dialogue box and menu for use unmanaged by Make, or produced by Make for managed development. The action of dragging a source file to the CC SetUp dialogue box specifies the file as an absolute rooted filename. Make normally specifies source files relative to the work directory.

C source files will be looked for relative to the work directory in the subdirectory `c`. To aid portability, a file specified as `foo.c` in a makefile will be looked for in `@.c.foo`, where `@` means the work directory.

Include files

The way in which the compiler searches for included files is dealt with in detail in the section entitled *Include file searching* on page 29. Here we describe the issues of naming header files and how to name them in `#include` lines in your C program source.

Include files are often headers for libraries, and are incorporated by issuing the `#include` directive – dealt with by the preprocessor – at the start of a source file. For instance, in the `HelloW` example:

```
#include <stdio.h>
```

By convention, header files are placed in subdirectory `h`. This convention is followed here. You can use subdirectory `h` of the work directory for your own header files, which can be incorporated with a source line like:

```
#include "myfile.h"
```

Note that both the example filenames `stdio.h` and `myfile.h` are in suffix form rather than Acorn prefix form. This is because you can make use of the filename processing of CC to interpret these, leaving program lines which do not need altering to port them to machines expecting suffixes.

To facilitate the porting of code from UNIX and MS-DOS to RISC OS, UNIX-style and MS-DOS-style filenames are translated to equivalent RISC OS-style filenames.

For example:

<code>../include/defs.h</code>	is translated to	<code>^.include.h.defs</code>
<code>..\cls\hash.h</code>	is translated to	<code>^.cls.h.hash</code>
<code>includes.h</code>	is translated to	<code>h.includes</code>

but

<code>system.defs</code>	is translated to	<code>system.defs</code>
--------------------------	------------------	--------------------------

In the same way, the lists of directory names given as arguments to the compiler's **Include** and **Default path** `SetUp` options (see below) are translated to RISC OS format before being used, in the rare event that this is necessary.

Object files

If you are making unmanaged use of CC, and compile one file with the `SetUp` dialogue box option **Compile only** enabled, you control the place the output object file is saved to using a standard `Save` dialogue box. With **Compile only** disabled the object files created by the compiler are stored in the directory `o` within the work directory. Thus the result of compiling `c.sieve` will be found in `o.sieve`.

Program files

If the CC SetUp dialogue box option **Compile only** is not enabled, after compiling sources to object files, CC links them with the ANSI library stubs to produce an executable program file. You may find it convenient to save this program file in the work directory itself – there is no conventional suffix for these.

Compilation list files

If the CC SetUp menu option **Listing** is enabled, a file containing a compilation listing for each compiled source file is created in the `l` subdirectory of the work directory. Thus compiling `c.sieve` with **Listing** enabled will by default result in the list file `l.sieve` being created.

Assembly list files

If the CC SetUp menu option **Assembler** is enabled, no object code is generated. Instead, an assembly listing of the code is created. If only one assembly listing file is produced, you save it from a standard save dialogue box. If more than one is produced these are placed in the subdirectory `s` of the work directory. Thus compiling `c.sieve` with **Assembler** enabled can result in the assembly language file `s.sieve` being created.

Filename validity

The compiler does not check whether the filenames you give are acceptable – whether they contain only valid characters and are of acceptable length – this is done by the filing system.

Include file searching

The process of converting text C source to linkable object files of binary code can be seen as a pipeline of several processes. The first stage is preprocessing the source. It is at this stage that the text of header files is brought in at the position of `#include` directives in the C source text.

The CC preprocessor handles `#include` directives of two forms:

```
#include <filename>
```

or

```
#include "filename"
```

You will normally include three types of header file:

- headers for the ANSI C library
- headers for the RISC OS library

- your own include files.

A special feature of the Acorn C system is that the standard ANSI headers are built into the C compiler, and are used by default. By writing the filename in the angle bracket form, you indicate that the include file is a system file, and thus ensure that the compiler looks first in its built-in filing system. Writing the filename in the double quote form indicates that the include file is a user file. Of the three common types of header above, only the ANSI headers should be referred to as system files in angle brackets.

The headers for the non-ANSI parts of the main C library – `kernel`, `pragmas`, `SWIs` and `varargs` – are not built in to the compiler. By default, they will be found by the compiler in `Clib.h`.

Headers for the RISC OS library are located in `RISC_OSLib.h`. You can make use of these by dragging the name of the directory `RISC_OSLib` from a directory display to the **Include** icon on the CC Setup dialog box. For example, in the `!Balls64` example in `User. !Balls64`:

```
in the source      #include "wimp.h" etc
```

on compilation drag the directory name `RISC_OSLib` to the **Include** Setup icon.

This is illustrated in the section entitled *Worked examples* on page 53. Placing the filename in double quotes in the `#include` directive indicates a user file.

As mentioned before, you can use the subdirectory `h` of the work directory for the third common type of header file – your own header files, which you refer to as user files with directives such as:

```
#include "myfile.h"
```

This is all you need to know for basic use of CC with largely default options. The rest of this section provides a level of detail useful for reference or studying if you wish to use CC in a non-standard way.

The way in which the compiler looks for included files depends on three factors:

- whether the filename is rooted
- whether the filename in the `#include` directive is between angle brackets `<>` or double quotes `" "`
- use of the CC **Include** and **Default path** Setup options (including the special filename `:mem`).

If a filename is not rooted (as defined earlier) CC looks for it in a sequence of directories called the search path.

Search path

The order of directories in the search path is as follows:

- 1 the compiler's own in-memory filing system (only for <filename>)
- 2 the 'current place' (see the section entitled *Nested includes* on page 31) (only for "filename")
- 3 arguments to the **Include** option, if used
- 4 the system search path:
 - the path given as an argument to the **Default path** CC SetUp menu option (see below), or
 - the value of the system variable `C$Libroot`, if this is set and the **Default path** CC SetUp menu option is not enabled; otherwise
 - `$.Clib`.

Unless the CC SetUp menu option **Default path** is enabled:

- for <filename> the search path (in order) is 1, 3, 4.
- for "filename" the search path is 2, 3, 4.

If **Default path** is used, place 1 (the in-memory filing system) is omitted for `#include <filename>`. It can be reinstated by giving the pseudo-filename `:mem` to the **Include** or **Default path** SetUp options.

Nested includes

The current place is the directory containing the source file (C source or #included header) currently being processed by the compiler. Often, this will be the work directory.

When a file is found relative to an element of the search path, the name of the directory containing that file becomes the new current place. When the compiler has finished processing that file it restores the old current place. So at any given instant, there is a stack of current places corresponding to the stack of nested #includes.

For example, suppose the current place is `$.include` and the compiler is seeking the #included file `"sys.defs.h"` (or `"sys.h.defs"`, `"sys/defs.h"`, etc). Now suppose this is found as:

```
$.include.sys.h.defs.
```

Then the new current place becomes `$.include.sys`, and files #included by `h.defs`, whose names are not rooted, will be sought relative to `$.include.sys`.

This is the search rule used by BSD UNIX systems. If you wish, you can disable the stacking of current places using the CC SetUp menu option **Feature** with the argument `K`, to get the search rule described originally by Kernighan and Ritchie in The C Programming Language. Then all non-rooted user includes are sought relative to the directory containing the source file being compiled.

In all this, the penultimate `.c` and `.h` components of the path are omitted. These are logically part of the filename – a filename extension – not logically part of the directory structure. However, directory names other than `c`, `h`, `o` and `s` are not so recognised (as filename extensions) and are used 'as is'. For example, the name `sys.new.defs` is exactly that: it is not translated to `sys.defs.new` and, if it is found, the new part of the name does become part of the new current path.

Use of `:mem`

You can use the CC SetUp menu option **Default path** to provide your own system search path, as mentioned in item 4 of Search path, above. The compiler will then use the argument you give to the **Default path** option as the system search path. You will only require this feature if you use implementations of the C library other than those provided with the Acorn C system.

Use of the **Default path** option also removes the in-memory filing system from the front of the path searched for in `#include <filename>`. It can be reinstated by using the pseudo-filename `:mem` as an argument to the **Default path** or **Include** options. If `:mem` is included in the search path in this way, its position in the path is as specified – not necessarily first – so you can take complete control over where the compiler looks for #included files.

Use of `C$Libroot`

`C$Libroot` is an environment variable that you can use to provide your own system search path, as shown in the Search path section above. It is not needed for normal use of the compiler.

The compiler will use the value of `C$Libroot`, if set, as the system search path. By default, `C$Libroot` is not set.

To set the value of `C$Libroot` to, for example, `$.mylib`, at the `*` prompt type:

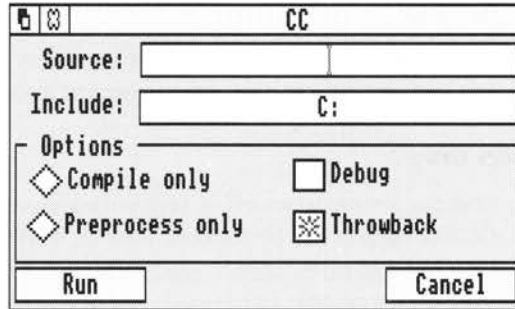
```
*set C$Libroot $.mylib
```

This variable is also used by the C compiler system as the library search path, if set. With the example given, the compiler will now look for include files in `$.mylib.h`, and for libraries in `$.mylib.o`.

The !Boot application in the DDE directory sets `C$Libroot`.

The SetUp dialogue box

Clicking Select on the CC icon bar icon or dragging a C source file from a directory display to this icon brings up the CC SetUp dialogue box:



Source

This writable icon in the SetUp dialogue box contains the names of the source files to be compiled.

When the SetUp box is obtained by clicking on the main CC icon, it comes up with this icon containing its previous setting. This helps you repeat a previous compilation, as clicking on the **Run** button repeats the last job if there was one.

If the SetUp box appears as a result of dragging a source file containing C text to the main icon, the writable **Source** icon appears containing the source file name.

When the SetUp box appears the **Source** icon has input focus, and can be edited in the normal RISC OS fashion. If you select a further source file in a directory display and drag it to this writable icon, its name is added to a list of those already there.

If you drag pre-compiled or pre-assembled object files to the **Source** icon, they are included in the set of object files linked together in an integral link step after the C source files have been compiled.

Include

This SetUp dialogue box icon adds specified directories to the list of places which are searched for `#include` files (after the in-memory or source file directory, according to the type of include file). The directories are searched in the order in which they are given in the **Include** icon. The path should end with the name of a directory, with no `.h`, which is added automatically.

For example, to make use of RISC_OSLib headers from program lines such as:

```
#include "wimp.h"
```

drag the name of the directory RISC_OSLib to this icon, so that the file RISC_OSLib.h.wimp is found correctly.

The default setting of **Include** is to C:. This allows the C compiler to search for headers in the directories listed in the RISC OS environment variable C\$Path, set by !Boot.

For more details of how to use #include lines and places searched for headers, see the section entitled *File naming and placing conventions* on page 25.

Compile only

This option switches off or on the linking of object files by CC. When enabled, the CC link step is not performed, and CC outputs object files. If only one C source file is being compiled unmanaged by Make, you drag the object file produced from a save dialogue box. Otherwise, multiple files are saved in the o subdirectory of the work directory.

If not enabled, CC performs an integral link step, linking any object files produced by compilation to any additional ones dragged to the **Source** icon, and library files, producing an executable program file. You control the saving of this from a save dialogue box.

Compile only is not enabled by default.

Preprocess only

If this option is enabled, only the preprocessor phase of the compiler is executed. The output from the preprocessor is sent to the standard output window. The standard non-interactive tool output window save facility is useful here to save this output to a file or SrcEdit window. By default, comments are stripped from the output, but see the SetUp menu option **Keep comments**.

Preprocess only is not enabled by default.

Debug

This option switches on or off production of debugging tables. When enabled, extra information is included in object files produced which enables source level debugging of the linked image (as long as the **Link Debug** option is also enabled) by the DDT debugger. If this option is disabled, any image file finally produced can only be debugged at machine level.

If the linking is done as an integral link step by CC, if the CC **Debug** option is enabled, so is that of the link step.

Debug is not enabled by default.

Throwback

This option switches editor throwback on or off. When enabled, if the DDEUtils module and SrcEdit are loaded, any compilation errors cause the editor to display an error browser. Double clicking Select on an error line in this browser makes the editor display the source file containing the error, with the offending line highlighted. See the chapter entitled *SrcEdit* in the accompanying *Acorn Desktop Development Environment* user guide for more details.

Throwback is on by default.

The SetUp menu

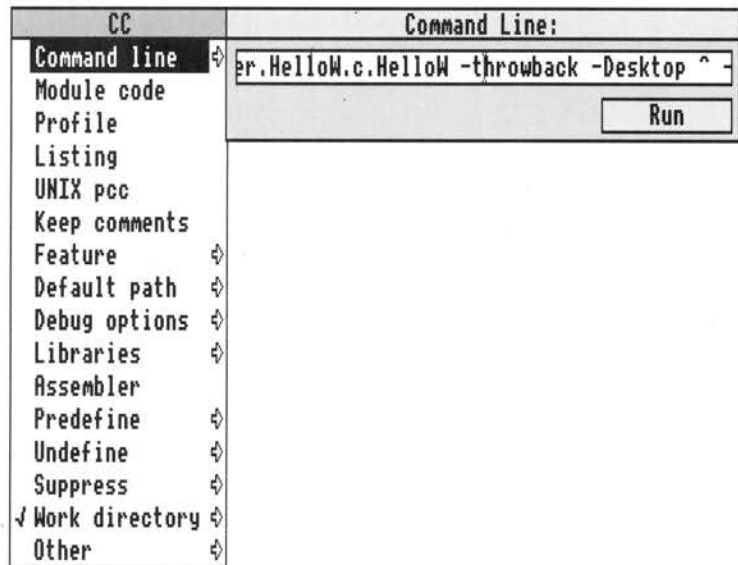
Clicking Menu on the SetUp dialogue box brings up the CC SetUp menu:

CC	
Command line	↕
Module code	
Profile	
Listing	
UNIX pcc	
Keep comments	
Feature	↕
Default path	↕
Debug options	↕
Libraries	↕
Assembler	
Predefine	↕
Undefine	↕
Suppress	↕
√ Work directory	↕
Other	↕

The options on this menu are described in the following subsections.

Command line

The CC RISC OS desktop interface works by driving a CC tool underneath with a command line constructed from your SetUp options. The **Command line** item at the top of the SetUp menu leads to a small dialogue box in which the command line equivalent of the current SetUp options is displayed:



Clicking on the **Run** action button in this dialogue box starts compilation in the same way as that in the main SetUp box. Pressing Return in the writable icon in this box has the same effect. Before starting compilation from the command line box, you can edit the command line textually, although this is not normally useful.

Module code

This SetUp menu option must be enabled when compiling code for linking into a RISC OS relocatable module, otherwise it should not be enabled. When enabled, code is produced which allows the module's static data to be separated from its code, hence be multiply instantiated.

Module code is not enabled by default.

Profile

Enabling this SetUp menu option causes the compiler to generate code to count the number of times each function is executed. This is called profiling.

The counts can be printed by calling `_mapstore()` to print them to `stderr` or by calling `_fmapstore("filename")` to print them to a named file of your choice. You should do this just before the final statement of your program.

Profiling is not supported by the shared C library, so you must link programs to be profiled with ANSILib (from `CLib.o`). If you wish, you can link with both Stubs and ANSILib, in which case only the code for `_mapstore()` and `_fmapstore()` will be included from ANSILib; your program will continue to use the shared C library, and will be much smaller than if linked with ANSILib alone.

The printed counts are lists of `lineno: count` pairs. The `lineno` value is the number of a line in your source code and the count value is the number of times it was executed. Note that `lineno` is ambiguous: it may refer to a line in a `#include` file. However, this is rare and usually causes no confusion.

Provided you didn't compile your program with the **Feature** option with `f` as an argument, blocks of counts will be interspersed with function names. In the simple cases, the output reduces to a list of line-pairs like:

```
function
lineno: count where count is the number of times function was executed.
```

If you use the Setup menu option **Other** to add the text `-px` to the command line, profiling of basic blocks within functions is performed in addition to profiling the functions. If you do this, the `lineno` values within each function relate to the start of each basic block. Sometimes, a statement (such as a For statement) may generate more than one basic block, so there can be two different counts for the same line.

Profiled programs run slowly. For example, when compiled with **Profile** enabled, Dhrystone 1.1 runs at about 5/8 speed; when compiled `-px` it runs at only about 3/8 speed.

There is no way, in this release of C, to relate execution counts to the proportion of time spent in each section of code. Nor is there any tool for annotating a source listing with profile counts. Future releases of C may address these issues.

Profile is not enabled by default.

Listing

Enabling this Setup menu option causes a listing file to be created. This consists of lines of source interleaved with error and warning messages. You can get finer control over the contents of this file using the **Feature** option (see below).

Listing is not enabled by default.

UNIX pcc

Enabling this SetUp menu option switches to compiling 'portable C compiler' C rather than ANSI C. This is based on the original Kernighan and Ritchie (K&R) definition of C, and is the dialect used on UNIX systems such as Acorn's RISC iX product. This option changes the syntax that is acceptable to the compiler, but the default header and library files are still used. See the section on this option in the chapter entitled *Portability* on page 91 for more details.

UNIX pcc is not enabled by default.

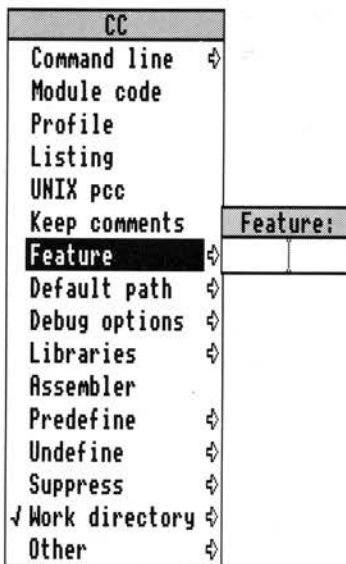
Keep comments

When enabled in conjunction with **Preprocess only**, this option retains comments in preprocessor output.

Keep comments is not enabled by default.

Feature

The **Feature** option on the SetUp menu leads to a small writable icon in which you can specify additional compiler features with single modifier letters:



This entry controls a variety of compiler features, including certain checks on your code more rigorous than usual. At least one of the following modifier letters must be entered if **Feature** is enabled:

- a Check for certain types of data flow anomalies. The compiler performs data flow analysis as part of code generation. The checks enabled by this option can sometimes indicate when an automatic variable has been used before it has been assigned a value.
- c Enable the **Limited pcc** option. This allows characters after `#else` and `#endif` preprocessor directives (treated as comments), and explicit casts of integers to function as pointers (forbidden by ANSI). These features are often required in order to use pcc-style include files in ANSI mode.
- e Check that external names used within the file are still unique when reduced to six case-insensitive characters. Some linkers only provide six significant characters in their symbol tables. This can cause problems with clashes if a system uses two names such as `getExpr1` and `getExpr2`, which are only unique in the eighth character. The check can only be made within one compilation unit (source file) so cannot catch all such problems. Acorn C allows external names of up to 256 characters, so this is a portability aid.
- f Do not embed function names in the code area. The compiler does this to make the output produced by the stack backtrace function (which is the default signal handler) and `_mapstore()` more readable. Removing the names from the compiler makes the code slightly smaller (typically 5%) at the expense of less meaningful backtraces and `_mapstore()` outputs.
- h Check that all external objects are declared in some included header file, and that all static objects are used within the compilation unit in which they are defined. These checks support good modular programming practices.
- i In the listing file (see the **Listing** option) include the lines from any files included with directives of the form:


```
#include "file"
```
- j As above, but for files included by lines of the form:


```
#include <file>
```
- k Use K&R search rules for nested `#include` directives (the 'current place' is defined by the original source file and is not stacked; see the section entitled *File naming and placing conventions* on page 25 for details).
- m Give a warning for preprocessor symbols that are defined but not used during the compilation.
- p Report on explicit casts of integers into pointers, eg:


```
char *cp = (char *) anInteger;
```

Implicit casts are reported anyway, unless suppressed by the **Suppress** option.

- u By default, the source text as 'seen' by the compiler after preprocessing (expansion) is listed. If this feature is specified then the unexpanded source text, as written by the user, is listed. Consider the line

```
p = NULL;
```

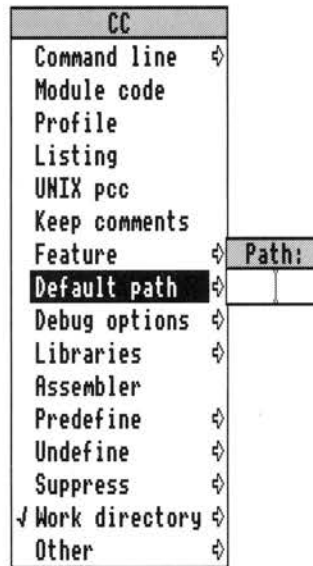
By default, this will be listed as `p= (0) ;` ; With **Feature** u specified, it will be listed as `p=NULL;` .
- v Report on all unused declarations, including those from standard headers.
- w Allow string literals to be writable, as expected by some UNIX code, by allocating them in the program's data area rather than the notionally read-only code area.
- x Turn on additional warnings about:
 - use of short integers. Shorts are slower than longs on the ARM and cause more code to be generated. They should only be used to save space in large arrays of data.
 - use of enums. ANSI defines enum values to be integers so the use of enums is not strictly type-checked. In some dialects of C, enums are more strictly type-checked than this.

When writing high-quality production software, you are encouraged to use at least the `fan` **Feature** options in the later stages of program development (the extra diagnostics produced can be annoying in the earlier stages).

Feature is not enabled by default.

Default path

The **Default path** entry on the SetUp menu leads to a writable icon in which you specify a comma-separated list of directories to be searched for included files:



This overrides the system include path with the list of directories. You can specify the memory file system in the list by using the name `:mem` (in any case). An example is:

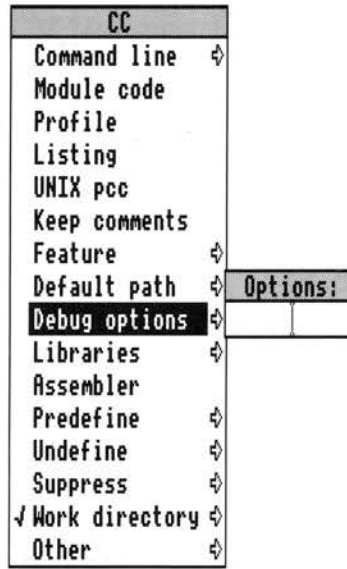
```
myhdrs, :mem, $.proj.public.hdrs
```

For more details of the system include path and searching for include files in general, see the section entitled *File naming and placing conventions* on page 25.

Default path is not enabled by default.

Debug options

The **Debug options** option on the SetUp menu leads to a writable item in which you enter a set of modifier letters:



The modifier letters limit the debugging tables generated in response to enabling the **Debug option** on the SetUp dialogue box. The letters recognised are:


- f generate information on functions and top-level variables (outside functions) only
- l generate information only describing each line in the file
- v Generate information only describing all variables

You can use these letters in any combination.

Debug options is not enabled by default.

Libraries

The **Libraries** option on the SetUp menu leads to a writable icon in which you specify a comma-separated list of filenames of libraries to be used in an integral CC link step:



The libraries specified with this option are used instead of the standard one (CLib.o.Stubs) not in addition to it.

This option can be very useful if you wish to compile and link many applications making use of RISC_OSLib using CC without the Link application. Setting **Libraries** to a list of Stubs and RISC_OSLib, and **Include** to the RISC_OSLib directory allows you to drag application sources in and save completed !RunImage files.

Libraries is not enabled by default.

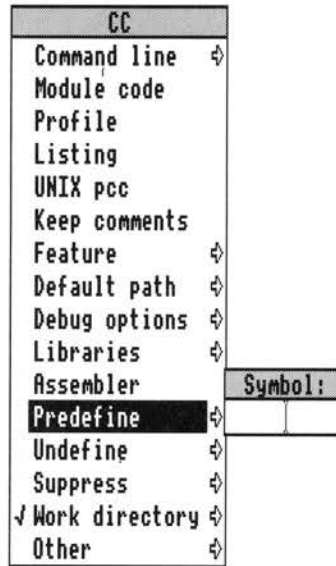
Assembler

If this SetUp menu option is enabled, no object code is generated and, naturally, no attempt is made to link it. If only one assembly listing file is produced, you save it from a standard save dialogue box. If more than one is produced these are placed in the subdirectory s of the work directory.

Assembler is not enabled by default.

Predefine

The **Predefine** option on the SetUp menu leads to a writable icon in which you can predefine preprocessor macros:



You can enter two forms of macro predefinition:

```
sym=value
```

```
sym
```

These both define sym as a preprocessor macro for the compilation. The two forms are equivalent to the lines:

```
#define sym value
```

and

```
#define sym 1
```

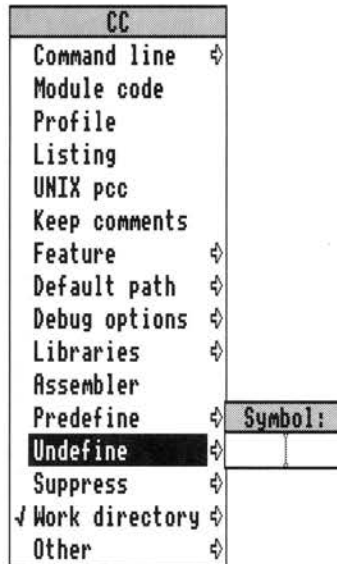
at the head of the source file.

Multiple symbols can be entered as a space-separated list.

Predefine is not enabled by default.

Undefine

The **Undefine** option on the SetUp menu leads to a writable icon in which you can undefine preprocessor macros:



You enter the name of the macro concerned, eg:

```
sym
```

Use of this option is then equivalent to the line:

```
#undef sym
```

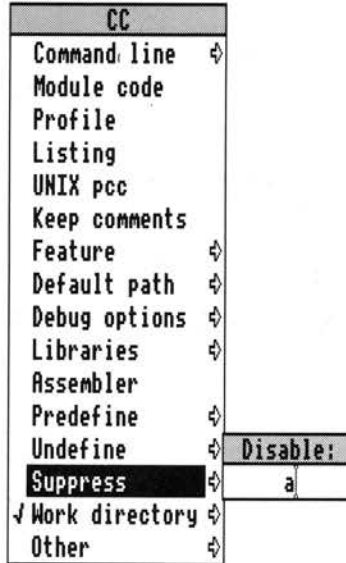
at the head of the source file.

Multiple symbols can be entered as a space-separated list.

Undefine is not enabled by default.

Suppress

The **Suppress** option on the SetUp menu leads to a writable icon in which you can enter a set of modifier letters:



The modifier letters specify various kinds of warning message to be suppressed by CC. Usually the compiler is very free with its warnings, as this tends to indicate potential portability or other problems. However, too many such messages can be a nuisance in the early stages of porting a program from old-style C, so you can disable them.

The modifier letters are:

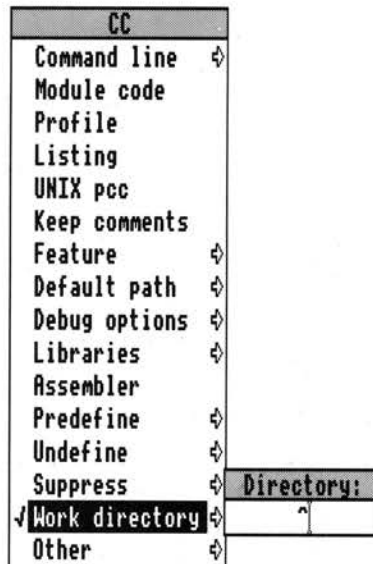
- a Give no Use of = in a condition context warning. This is given when the compiler encounters statements such as `if (a=b) {...` where it is quite possible that `==` was intended.
- d Give no Deprecated declaration `foo()` – give arg types warning. Use of old-style function declarations is deprecated in ANSI C, and in a future version of the standard this feature may be removed. However, it is useful sometimes to suppress this warning when porting old code.
- f Give no Inventing "extern int `foo()`" message. This may be useful when compiling old-style C as if it were ANSI C.

- n Give no `Implicit narrowing cast` warning. This warning is issued when the compiler detects an assignment of an expression to an object of narrower width (eg `long` to `int`, `float` to `int`). This can cause problems with loss of precision for certain values.
- v Give no `Implicit return in non-void context` warning. This is most often caused by a return from a function which was assumed to return `int` (because no other type was specified) but is in fact being used as a void function.

If you enter a space in the writable icon, then `Select` or `Return`, all warning messages are suppressed.

Work directory

The **Work directory** entry on the `SetUp` menu leads to a writable icon in which you specify the work directory:



The effect of this option is described in the section entitled *File naming and placing conventions* on page 25.

The default **Work directory** setting is `^`.

Other

The **Other** option on the SetUp menu leads to a writable icon in which you can add an arbitrary extra section of text to the command line to be passed to CC:

CC	
Command line	↔
Module code	
Profile	
Listing	
UNIX pcc	
Keep comments	
Feature	↔
Default path	↔
Debug options	↔
Libraries	↔
Assembler	
Predefine	↔
Undefine	↔
Suppress	↔
√ Work directory	↔
Other	↔
	Others:
	-px

This facility is useful if you wish to use any feature which is not supported by any of the other entries on the SetUp dialogue box and menu. This may be because the feature is used very little, or because it may not be supported in the future.

Use of `-px` in the **Other** facility is described in the **Profile** description above.

`#pragma` directives can be emulated through **Other** by adding text with the syntax `-zp<mod>`. `<mod>` is the same sequence of characters that would follow the `#pragma` directive in program text. For more details of `#pragma` directives, see the section entitled *#pragma directives* on page 378.

Preprocessor options

The first stage in processing C source is preprocessing. CC options are described above in the order they appear on the CC SetUp dialogue box and menu. The options above which control the preprocessor are:

- **Include**
- **Default path**
- **Preprocess only**

- **Keep comments**
- **Predefine**
- **Undefine**


Link step options

As described before, CC can be set up so that after generating object files, these are linked to form an executable program file. CC options are described above in the order that they appear on the SetUp dialogue box and menu. The CC options for driving the link step are:

- **Compile only**
- **Source**
- **Libraries**

The Application menu

Clicking Menu on the CC application icon on the icon bar gives access to the following menu box:



Save options saves all the current CC options, including both those set from the SetUp dialogue box and from the **Options** item on this menu. When CC is restarted it is initialised with these options rather than the defaults. If you often want to construct RISC OS applications using RISC_OSLib and the CC integral link step, you may find it useful to:

- set up the **Include** slot on the SetUp dialogue box to find RISC_OSLib headers
- set up **Libraries** on the SetUp menu to link with Stubs and RISC_OSLib
- save these options as the new default using **Save options**.

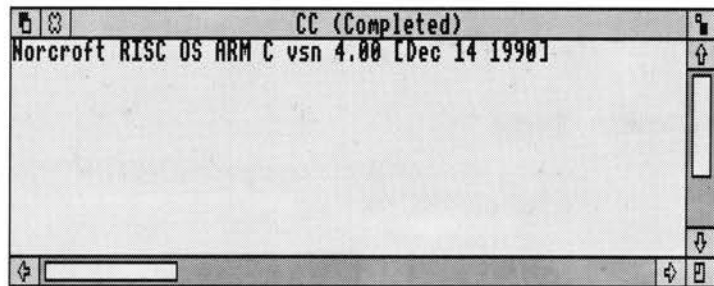
The **Options** item on the main menu allows you to enable **Auto run**, **Auto save** or start the output display as either a text window (default) or summary box. When **Auto run** is enabled, dragging a source file to the CC main icon starts a compilation immediately with the current options rather than displaying the SetUp

box first. When **Auto save** is enabled, output object files are saved to suitable places automatically without producing a save dialogue box for you to drag the file from. Both **Auto run** and **Auto save** are off by default.

For a description of each option in the application menu see the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide.

CC output messages


CC outputs text messages as it proceeds. These include preprocessed source (see **Preprocess only**) warning and error messages. By default any such text is directed into a scrollable output window:



This window is read-only; you can scroll up and down to view progress, but you cannot edit the text without first saving it. Clicking Select on the scrollable part of this window has no effect, to indicate this.

The contents of the window illustrated above are typical of those you see from a successful compilation - the title line of the compiler with version number, followed by no error messages.

Clicking Adjust on the close icon of the output window switches to the output summary dialogue box. This presents a reminder of the tool running (CC), the status of the task (Running, Paused, Completed or Aborted), the time when the task was started and the number of lines of output that have been generated (ie those that are displayed by the output window):



Clicking Adjust on the close icon of the summary box returns to the output window.

Both the above CC output displays follow the standard pattern of those of all the non-interactive DDE tools. The common features of the non-interactive DDE tools are covered in more detail in the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide. Both CC output displays and the menus brought up by clicking Menu on them offer the standard features allowing you to abort, pause or continue execution (if the execution hasn't completed) and to save output text to a file or repeat execution.

CC error messages appear in the output viewer, with copies in the editor error browser when throwback is working. *Appendix B: Errors and warnings* contains more details for interpreting error messages.

Preprocessed source appearing in the output window is often very large for compilation of complex source files. The scrolling of the output window is useful to view it, and to investigate it with the full facilities of the source editor, you can save the output text straight into the editor by dragging the output file icon to the SrcEdit main icon on the icon bar (providing Wimp\$Scrap is properly set on your machine).

Command line interface

For normal use, either managed by Make or not, you do not need to understand the syntax of the CC command line, as it is generated automatically for you from the SetUp dialogue box and menu settings before it is used.

The syntax of the CC command line is:

```
*cc [options] filenames
```

where *filenames* is the list of source and object files you set up in the **Source** icon of the SetUp dialogue box.

The options in *options* are either keywords or flags. Keyword options have the syntax:

```
-keyword
```

where *keyword* is a word of more than one alphabetic character.

Flag options have the syntax:

```
-letter[arg]
```

where *letter* is a single alphabetic character, and *arg* is only present for some flag letters. When present, *arg* is a sequence of characters not including spaces – lists as arguments must be comma-separated. CC accepts white space between *letter* and *arg*, but compilers on other systems do not, so it is not recommended to use this syntax in makefiles.

Since CC options are described earlier in detail in relation to SetUp dialogue box and menu options, here the command line options are listed with their equivalent SetUp options rather than full explanations. Note that unlike in RISC OS filenames, case of letters is important in CC command line options.

Keyword options:

-pcc	UNIX pcc
-fussy	No direct equivalent on SetUp windows – insert in Other if necessary. Enables extra strictness about enforcing conformance to the ANSI standard or to pcc conventions. For example, -fussy turns off the pre-definition of ARM and arm by the preprocessor in ANSI mode.
-list	Listing

Preprocessor flag options:

-I	Include
-J	Default path
-E	Preprocess only
-C	Keep comments
-z<mod>	No direct SetUp window equivalent – insert in Other if necessary to emulate #pragma directives.
-D<sym>=<value>	Predefine
-D<sym>	Predefine
-U<sym>	Undefine

Code generation flag options:

-g	Debug
-g<mod>	Debug options
-p	Profile
-S	Assembler
-zM	Module code

Warning feature flag options:

-w<mod> **Suppress**
 -f<mod> **Feature**

Linker flag options:

-c **Compile only**
 -l<libs> **Libraries**

Worked examples

Several examples of C program on the discs of Acorn Desktop C are worked through in other chapters of this volume such as the chapter entitled *C tools and the DDE* on page 9. A collection of examples are listed here illustrating various points and styles of working.

The following example programs are to be found in the directory `User`, each in a subdirectory with the name of the example. For each program, we give a 'recipe' for how to compile, link and run the program. Filenames are given relative to the subdirectory containing each example unless otherwise stated. If you have a machine with a single floppy disc drive, and 1Mb of RAM, you will need to clean up after running each example. It is assumed that you have read the preceding parts of this chapter. For more details of the tool `Make`, see the chapter entitled *Make* in the accompanying *Acorn Desktop Development Environment* user guide. When you enter any command lines given below, you must first ensure that the currently-selected directory is the subdirectory containing the example being tried.

HelloW

Purpose: The standard most trivial C program. Try it as an exercise.
 Source: `c.HelloW`
 Compile using: default CC SetUp options
 Run by: double clicking on `HelloW`
 Clean up by: deleting `HelloW` and `o.HelloW`

Sieve

Purpose: The Sieve of Eratosthenes is often presented as a standard benchmark, though it is not very meaningful in this context.
 Source: `c.sieve`
 Compile using: default CC SetUp options

Run by: double clicking on sieve
Clean up by: deleting sieve and o.sieve

Dhrystone 2.1

Purpose: Dhrystone 2.1 is the standard integer benchmark. Its results require careful interpretation (it often overstates the real performance of machines). Try as a first exercise in using the Make utility (!Make).

Sources: h.dhry
c.dhry_1
c.dhry2

Makefile: Makefile

Build by: double clicking on Makefile, with default Make options

Run by: double clicking on Dhrystone

Reply with any number in the range 20000 to 250000 to the prompt for number of iterations. Try a big number such as 200000 and time the execution with a stopwatch or sweep second hand to confirm the claimed performance. Note how performance depends on screen mode.

Rebuild by: double clicking on Makefile again (try altering some of the options in Makefile with Make between rebuilds: eg compile in UNIX pcc mode or link with ANSILib instead of Stubs).

Clean up by: deleting Dhrystone, o.dhry_1 and o.dhry2.

SWI_list

Purpose: To illustrate use of the SWI facilities in <kernel.h>. You can also try it as an exercise in getting going; later, you can use it to check that CLib.h.swis contains a complete list of the SWI names and numbers relevant to your machine.

Source: c.SWI_list

Compile using: default CC SetUp options

Run from: the command line, using SWI_list > h.myswis

Test using: see instructions embedded in the comment at the head of c.SWI_list

Clean up by: deleting SWI_list, o.SWI_list and h.myswis

HowToCall

Purpose:	To illustrate how to call other programs from C. Read the source, then experiment with the binary. You can also use it as another exercise in getting going. Try making your own makefile for it as an exercise in using Make.
Source:	<code>c.HowToCall</code>
Compile using:	default CC options
Run from:	the command line using: <pre>HowToCall 3 HowToCall HowToCall 2 HowToCall 3 * HowToCall 3 * etc</pre>
Clean up by:	deleting <code>HowToCall</code> and <code>o.HowToCall</code>

CModule

Purpose:	To illustrate how to implement a module in C. You can also use it as another exercise in using Make. For more details on constructing relocatable modules in C see the later chapter <i>How to write relocatable modules in C</i> .
Sources:	<code>c.CModule</code> <code>CModuleHdr</code>
Build using:	CC of <code>c.CModule</code> with options Compile only and Module code enabled, saving output object file as <code>o.CModule</code> . CMHG of <code>CModuleHdr</code> to <code>o.CModuleHdr</code> . Link of <code>o.CModule</code> , <code>o.CModuleHdr</code> and <code>\$.CLib.o.Stubs</code> with Module enabled to the output file <code>CModule</code> .
or by:	double clicking on <code>Makefile</code> , with default Make options.
Run from:	the command line using <code>CModule</code>
Test from:	the command line using: <pre>help tm1 help tm2 tm1 hello tm2 1 2 3 4 5 tm1 1 2 3 tm2 hello</pre> <p>(try other combinations too)</p>

```
*BASIC
> SYS &88000 : REM should give an error
> SYS &88001 : REM should give divide by 0 error
> SYS &88002 : REM no error, just a message
> SYS &88003 : REM no error, just a message
> SYS &88004 : REM same as &88000...
```

(now repeat some of these after issuing some invalid
* commands...)

```
>*foo
> SYS &88002
```

etc.

```
>QUIT
```

Clean up by: from the command line typing: `RMKill TestCModule`
deleting `cmodule`, `o.CModule` and `o.CModuleHdr` or
running `Make` on `Makefile` with target `clean` selected.

Desktop application examples

Some examples follow which illustrate how to write applications to run under the desktop. The following instructions assume that you have installed the DDE onto a hard disc, hence have all the tools, libraries, headers, etc in the directory tree of one disc. If you are working on a system that does not have a hard disc, your DDE is split between several floppy discs. This results in you having to change discs both to open directory displays and drag names from them to tool `SetUp` boxes, and also when requested by the execution of the tools when the files are read in.

Each example, like the others, is in a subdirectory of `User` with the name of the example. Since example applications have names beginning with `!`, hold `Shift` while double clicking on the directory names to open directory displays on the subdirectories.

When you have explored these examples, refer to the chapter entitled *How to write desktop applications in C* on page 149 if you want to go further.

Each example is built in the same way. First click on `!Boot` and `!CC` in the DDE directory. The name of the C source file is dragged to the `CC` icon bar icon to bring up the `CC Setup` box. On this dialogue box, ensure that **Compile only** is not enabled and that the **Include** icon is set to `C:`. The program is compiled, and the executable file produced saved as `!RunImage` in the example subdirectory. You can then use the `Squeeze` tool to compact the `!RunImage` file.

WExample

Purpose:	To illustrate installing an icon on the icon bar, and creating/ displaying a simple window.
Source:	<code>c.WExample</code>
Build using:	standard procedure above
Run using:	double click on <code>!WExample</code> in the User directory display. Click Select on the EG icon to get a window
Clean up by:	deleting <code>!RunImage</code> and <code>o.WExample</code>

Life

Purpose:	To illustrate use of multiple windows in an application, using the alarm facilities of RISC_OSLib and creating icons in a window.
Source:	<code>c.life</code>
Build using:	standard procedure above.
Run using:	double click on <code>!Life</code> in the User directory display
Clean up by:	deleting <code>!RunImage</code> and <code>o.life</code>

DrawEx

Purpose:	To illustrate loading files by icon dragging, and rendering draw files in a window.
Source:	<code>c.DrawEx</code>
Build using:	standard procedure above.
Run using:	double click on <code>!DrawEx</code> in the User directory display
Clean up by:	deleting <code>!RunImage</code> and <code>o.DrawEx</code>

Balls64

Purpose:	To illustrate use of a sprite as a virtual display , saving files by icon dragging, and responding to help requests. This application requires at least 320Kb of RAM to run, so you may need to quit some applications to make room for it.
Source:	<code>c.Balls64</code>
Build using:	standard procedure above.

Run using: double click on !Balls64 in the User directory display,
 click on the Balls64 icon bar icon

Clean up by: deleting !RunImage and o.Balls64

Recompiling ToANSI and ToPCC

In the directories containing the tools ToANSI and ToPCC, DDE. !ToANSI and DDE. !ToPCC, you will find c subdirectories containing the source file of each tool. If you wish to recompile these as a further exercise in using the Acorn C Compiler, you can drag each source file to the CC icon in turn and compile and link it with default options.



The C Module Header Generator (CMHG) is a tool to allow you to write a RISC OS relocatable module entirely in C without having to purchase a product such as Acorn Desktop Assembler containing an ARM assembler or understand assembly language.

Every relocatable module has at its start (ie the part that loads into memory at its lowest address) a header table pointing to various items of data and program. Most of the items pointed to are optional, the pointers being zero if not needed. When writing a relocatable module in assembly language you lay this table out yourself, but when writing in C, you use CMHG to generate this for you. In addition to generating a module header, CMHG also inserts small standard routines to, for example, initialise the C language library support and make service call handling efficient.


To construct a relocatable module you write a number of routines in C with standard prototypes, some of these routines to be called with the processor in supervisor (SVC) mode. These are accompanied by a text description file written in a special syntax which CMHG understands. For details of this language and the specifications of the C routines, see the chapter entitled *How to write relocatable modules in C* on page 337. For more details of relocatable module headers, see the chapter entitled *Modules* in the *RISC OS Programmer's Reference manual*. For some hints about memory usage from relocatable module code, see the chapter entitled *Using memory efficiently* on page 353.

The rest of this chapter explains the (simple) controls of the CMHG tool. CMHG is one of the non-interactive DDE tools, its desktop user interface being provided by the FrontEnd module. It shares many common features with the other non-interactive tools. These common features are described in the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide.

Starting CMHG


CMHG can be used as a tool 'managed' by Make, or controlled directly by you from its desktop interface for 'unmanaged' development. To start unmanaged use, first double click on !CMHG in a directory display to put its icon on the icon bar.

Clicking Select on this icon or dragging the name of a CMHG description file from a directory display to the icon brings up the CMHG SetUp dialogue box, from which you control the running of CMHG:




CMHG has hardly any options for its use, so this interface is simpler than those of most of the non-interactive tools. The **Source** writable icon is for the name of the description file to be processed. If the SetUp dialogue box was brought up by clicking on the icon bar icon, you will want to fill this icon in by dragging the name of your description file from a directory display to this icon before running CMHG.

Clicking Menu on the SetUp dialogue box brings up the CMHG SetUp menu, which owing to the simplicity of CMHG has no options:



The Application menu

Clicking Menu on the CMHG application icon on the icon bar gives access to the following options:




For a description of each option in the application menu see the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide.

Example output

The following is an example CMHG description file, similar to that used within Acorn to construct the FrontEnd module, which is itself a relocatable module written in C:

```
; Purpose: module header for the generalised front end module ;
module-is-runnable:                                ; module start code
initialisation-code:      FrontEnd_init
service-call-handler:     FrontEnd_services 0x11      ; service-memory
title-string:             FrontEnd
help-string:              FrontEnd 1.00
command-keyword-table:    FrontEnd_commands
                          FrontEnd_Start(min-args: 4, max-args: 5,
                          help-text: "Help text\n"),
                          FrontEnd_Setup(min-args: 8, max-args: 8,
                          help-text: "Help text\n")
swi-chunk-base-number:    0x081400
```

Running CMHG displays any error messages in the standard text output window for non-interactive tools. If all goes well, as it should do if you try CMHG with the above description file, this window is empty:



The output file produced is an object file. You link this with the object files compiled from your C code to produce your relocatable module.

Command line interface

For normal use, either managed by Make or not, you do not need to understand the syntax of the CMHG command line, as it is generated automatically for you from the SetUp dialogue box setting before it is used.

The syntax of the CMHG command line is simple:

```
cmhg descfile [objfile]
```

descfile Filename of the CMHG description file

objfile Filename of the output object file to link with your objects to form a relocatable module



ToANSI is a tool that helps convert program source written in the PCC style of C to program source in the ANSI style of C. PCC is the Unix Portable C Compiler, and closely follows K&R C, as defined by B Kernighan and D Ritchie in their book *The C Programming Language*.

The aim of ToANSI is to enable you to write (with care) programs that can be automatically converted between the PCC and ANSI dialects of C, hence assisting you in constructing easily portable programs. The associated tool ToPCC makes approximately the reverse translations to ToANSI. For more details of portability issues, see the chapter entitled *Portability* on page 91. The changes that ToANSI makes to C source are listed in the section entitled *ToANSI C translation* below.

ToANSI is one of the non-interactive DDE tools, its desktop user interface being provided by the FrontEnd module. It shares many common features with the other non-interactive tools. These common features are described in the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide.

ToANSI C translation

ToANSI makes the following transformations to C source code or header text:

- Function declarations with embedded comments are rewritten without the comment tokens. This reverses the action of ToPCC with regard to function declarations, rewriting:

```
type foo(/* args */);
```

as:

```
type foo(args);
```

This transformation is one which requires care in the use of ToANSI, as it can result in invalid C being uncommented.

- Function definitions of the form

```
type foo(a1, a2)
type a1;
type a2;
{...}
```

are rewritten as

```
type foo(type a1, type a2)
```


- A `va_alist` in the function definition is translated to
...
- `type foo()` is rewritten as `type foo(void)`.
- `VoidStar` (what ToPCC replaces `void *` with) is left untouched, as if it is correctly typedef'd to something suitable, thereafter its use is correct in both PCC and ANSI C.
- ToPCC rewrites unsigned and unsigned long constants using the typecasts `(unsigned)` and `(unsigned long)`. ToANSI does not reverse this change, as this is not required for correct ANSI C.

Note that ToANSI performs only simple textual translations and is not able to reliably diagnose C syntax errors, which may produce surprising results, so it is best to use ToANSI only on code you already know compiles.

Starting ToANSI

ToANSI can only be used controlled directly by you from its desktop interface for 'unmanaged' development. Since porting programs is usually a one-off process involving some experimentation, only use of ToANSI unmanaged by Make is sensible. To start unmanaged use, first double click on !ToANSI in a directory display to put its icon on the icon bar.

Clicking Select on this icon or dragging the name of a C source or header file (text) from a directory display to the icon brings up the ToANSI SetUp dialogue box, from which you control the running of ToANSI:




ToANSI has only one option for its use, so this interface is simpler than those of most of the non-interactive tools. The **Input** writable icon is for the name of the source or header file to be processed. If the SetUp dialogue box was brought up by clicking on the icon bar icon, you will want to fill this icon in by dragging the name of your source or header file from a directory display to this icon before running ToANSI.

Clicking Menu on the SetUp dialogue box brings up the ToANSI SetUp menu, which owing to the simplicity of ToANSI has no options:



The Application menu


Clicking Menu on the ToANSI application icon on the icon bar gives access to the following options:



For a description of each option in the application menu see the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide.

Example output

Running ToANSI displays any error messages in the standard text output window for non-interactive tools. If all goes well this window is empty:



As an example of using the tool ToANSI, open an empty SrcEdit text window and type the following example C source lines in it:

```
int foo(a, b)
float a;
double b;
{}
```

Check that your `Wimp$Scrap` environment variable is set to a sensible file name, then save your new text file straight onto the ToANSI icon bar icon. Run ToANSI, then save the output text file straight onto the SrcEdit icon bar icon. The translated file looks like:

```
int foo(float a, double b)
{ }
```

Command line interface

For normal use you do not need to understand the syntax of the ToANSI command line, as it is generated automatically for you from the SetUp dialogue box setting before it is used.

The syntax of the ToANSI command line is:

```
toansi infile [outfile]
```

infile Filename of the input C source or header text file.

outfile Filename of the output C source or header text file.



ToPCC is a tool that helps convert program source written in the ANSI style of C to program source in the PCC style of C. PCC is the Unix Portable C Compiler, and closely follows K&R C, as defined by B Kernighan and D Ritchie in their book *The C Programming Language*.

The aim of ToPCC is to enable you to write (with care) programs that can be automatically converted between the ANSI and PCC dialects of C, hence assist you in constructing easily portable programs. The associated tool ToANSI makes approximately the reverse translations to ToPCC. For more details of portability issues, see the chapter entitled *Portability* on page 91. The changes that ToPCC makes to C source are listed in the section entitled *ToPCC C translation* below.

ToPCC is one of the non-interactive DDE tools, its desktop user interface being provided by the FrontEnd module. It shares many common features with the other non-interactive tools. These common features are described in the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide.

ToPCC C translation

ToPCC makes the following transformations to C source code or header text:

- Function declarations of the form

```
type foo(args);
```

are rewritten as

```
type foo(/* args */);
```

Any comment tokens `/*` or `*/` in args are removed.

- Function definitions of the form

```
type foo(type a1, type a2) {...}
```

are rewritten as

```
type foo(a1, a2)
type a1;
type a2;
```

- A `...` in the function definition is interpreted as `int va_alist`. Full translation of *variadic* functions is not performed.

- `type foo(void)`
is rewritten as
`type foo()`
- Type `void *` is converted to `VoidStar` which can be typedef'd to something suitable (eg `char *`).
- Unsigned and unsigned long constants are rewritten using the typecasts `(unsigned)` and `(unsigned long)`.


For example, `300u1` becomes `(unsigned long)300L`.

Note that ToPCC performs only simple textual translations and is not able to reliably diagnose C syntax errors, which may produce surprising results, so it is best to use ToPCC only on code you already know compiles.

Starting ToPCC

ToPCC can only be used controlled directly by you from its desktop interface for 'unmanaged' development. Since porting programs is usually a one-off process involving some experimentation, only use of ToPCC unmanaged by Make is sensible. To start unmanaged use, first double click on !ToPCC in a directory display to put its icon on the icon bar.

Clicking Select on this icon or dragging the name of a C source or header file (text) from a directory display to the icon brings up the ToPCC SetUp dialogue box, from which you control the running of ToPCC:




ToPCC has only one option for its use, so this interface is simpler than those of most of the non-interactive tools. The **Input** writable icon is for the name of the source or header file to be processed. If the SetUp dialogue box was brought up by clicking on the icon bar icon, you will want to fill this icon in by dragging the name of your source or header file from a directory display to this icon before running ToPCC.

Clicking Menu on the SetUp dialogue box brings up the ToPCC SetUp menu, which owing to the simplicity of ToPCC has no options:



The Application menu


Clicking Menu on the ToPCC application icon on the icon bar gives access to the following options::



For a description of each option in the application menu see the chapter entitled *General features* in the accompanying *Acorn Desktop Development Environment* user guide.

Example output

Running ToPCC displays any error messages in the standard text output window for non-interactive tools. If all goes well this window is empty:



To try ToPCC, open an empty SrcEdit text window and type the following example C source line in it:

```
int foo(float a);
```

Check that your `Wimp$Scrap` environment variable is set to a sensible file name, then save your new text file straight onto the ToPCC icon bar icon. Run ToPCC, then save the output text file straight onto the SrcEdit icon bar icon. The translated file looks like:

```
int foo(/* float a */);
```


Command line interface

For normal use you do not need to understand the syntax of the ToPCC command line, as it is generated automatically for you from the SetUp dialogue box setting before it is used.

The syntax of the ToPCC command line is:

```
topcc infile [outfile]
```

infile Filename of the input C source or header text file.

outfile Filename of the output C source or header text file.

Part 2 - Language issues

7 Implementation details

This chapter gives details of those aspects of the compiler which the ANSI standard identifies as implementation-defined, and some other points of interest to programmers. They are grouped here by subject; the section entitled *Implementation limits* on page 78 lists the points required to be documented as set out in appendix A.6 of the standard.

Identifiers

Identifiers can be of any length. They are truncated by the compiler to 256 characters, all of which are significant (the standard requires a minimum of 31).

The source character set expected by the compiler is 7-bit ASCII, except that within comments, string literals, and character constants, the full ISO 8859-1 8-bit character set is recognised. At run time, the C library processes the full ISO 8859-1 8-bit character set, except that the default locale is the C locale (see the chapter entitled *Standard implementation definition* on page 81). The `ctype` functions therefore all return 0 when applied to codes in the range 160–255. By calling `setlocale(LC_CTYPE, "ISO8859-1")` you can cause the `ctype` functions such as `isupper()` and `islower()` to behave as expected over the full 8-bit Latin alphabet, rather than just over the 7-bit ASCII subset.

Upper and lower case characters are distinct in all identifiers, both internal and external.

Data elements

The sizes of data elements are as follows:

Type	Size in bits
char	8
short	16
int	32
long	32
float	32
double	64
long double	64 (subject to future change)
all pointers	32

Integers are represented in two's complement form.

Data items of type `char` are unsigned by default, though they may be explicitly declared as `signed char` or `unsigned char` (in `-pcc` mode `chars` are signed by default). Single-character constants are thus always positive.

Floating point quantities are stored in the IEEE format. In double and long double quantities, the word containing the sign, the exponent and the most significant part of the mantissa is stored at the lower machine address.

Limits: `limits.h` and `float.h`

The standard defines two headers, `limits.h` and `float.h`, which contain constant declarations describing the ranges of values which can be represented by the arithmetic types. The standard also defines minimum values for many of these constants.

The following table sets out the values in these two headers on the ARM, and a brief description of their significance. See the standard for a full definition of their meanings.

Number of bits in smallest object that is not a bit field (ie a byte):

`CHAR_BIT` 8

Maximum number of bytes in a multibyte character, for any supported locale:

`MB_LEN_MAX` 1

Numeric ranges of integer types: The column on the left gives the numerical values. The column on the right gives the bit patterns (in hexadecimal) that would be interpreted as these values in C. When entering constants you must be careful about the size and signed-ness of the quantity. Furthermore, constants are interpreted differently in decimal and hexadecimal/octal. See the ANSI standard or Harbison and Steele for more details.

CHAR_MAX	255	0xff
CHAR_MIN	0	0x00
SCHAR_MAX	127	0x7f
SCHAR_MIN	-128	0x80
UCHAR_MAX	255	0xff
SHRT_MAX	32767	0x7fff
SHRT_MIN	-32768	0x8000
USHRT_MAX	65535	0xffff
INT_MAX	2147483647	0x7fffffff
INT_MIN	-2147483648	0x80000000
UINT_MAX	4294967295	0xffffffff
LONG_MAX	2147483647	0x7fffffff
LONG_MIN	-2147483648	0x80000000
ULONG_MAX	4294967295	0xffffffff

Characteristics of floating point:

FLT_RADIX	2
FLT_ROUNDS	1

Ranges of floating types:

FLT_MAX	3.40282347e+38F
DBL_MAX	1.79769313486231571e+308
LDBL_MAX	1.79769313486231571e+308
FLT_MIN	1.17549435e-38F
DBL_MIN	2.22507385850720138e-308
LDBL_MIN	2.22507385850720138e-308

Ranges of base two exponents:

FLT_MAX_EXP	128
DBL_MAX_EXP	1024
LDBL_MAX_EXP	1024
FLT_MIN_EXP	(-125)
DBL_MIN_EXP	(-1021)
LDBL_MIN_EXP	(-1021)

Ranges of base ten exponents:

FLT_MAX_10_EXP	38
DBL_MAX_10_EXP	308
LDBL_MAX_10_EXP	308
FLT_MIN_10_EXP	(-37)
DBL_MIN_10_EX	(-307)
LDBL_MIN_10_EXP	(-307)

Decimal digits of precision:

FLT_DIG	6
DBL_DIG	15
LDBL_DIG	15

Digits (base two) in mantissa:

FLT_MANT_DIG	24
DBL_MANT_DIG	53
LDBL_MANT_DIG	53

Smallest positive values such that $(1.0 + x \neq 1.0)$:

FLT_EPSILON	1.19209290e-7F
DBL_EPSILON	2.2204460492503131e-16
LDBL_EPSILON	2.2204460492503131e-16L

Structured data types

The standard leaves details of the layout of the components of structured data types up to each implementation. The following points apply to the Acorn C compiler:

- Structures are aligned on word boundaries.
- Structures are arranged with the first-named component at the lowest address.
- `char` components are placed in adjacent bytes.
- `short` components are aligned at even-addressed bytes.
- All other arithmetic type components are word-aligned, as are pointers and `ints` containing bitfields.
- The only valid type for bitfields is `int`, either signed or unsigned.
- A bitfield of type `int` is treated as unsigned by default (signed by default in `-pcc` mode).
- Bitfields must be contained within the 32 bits of an `int`.
- Bitfields are allocated within `ints` so that the first field specified occupies the least significant bits of the word.

Pointers

The following remarks apply to pointer types:

- Adjacent bytes have addresses which differ by one.
- The macro `NULL` expands to the value 0.
- Casting between integers and pointers results in no change of representation.
- The compiler faults casts between pointers to functions and pointers to data (but not in `-pcc` mode).

Pointer subtraction

When two pointers are subtracted, the difference is obtained as if by the expression:

```
((int)a - (int)b) / (int)sizeof(type pointed to)
```

If the pointers point to objects whose size is no greater than four bytes, word alignment of data ensures that the division will be exact in all cases. For longer types, such as doubles and structures, the division may not be exact unless both pointers are to elements of the same array. Moreover the quotient may be rounded up or down at different times, leading to potential inconsistencies.

Arithmetic operations

The compiler performs all of the 'usual arithmetic conversions' set out in the standard.

The following points apply to operations on the integral types:

- All signed integer arithmetic uses a two's complement representation.
- Bitwise operations on signed integral types follow the rules which arise naturally from two's complement representation.
- Right shifts on signed quantities are arithmetic.
- Any quantity which specifies the amount of a shift is treated as an unsigned 8-bit value.
- Any value to be shifted is treated as a 32-bit value.
- Left shifts of more than 31 give a result of zero.
- Right shifts of more than 31 give a result of zero from an unsigned or positive signed value, -1 from a negative signed value.
- The remainder on integer division has the same sign as the divisor.

- If a value of integral type is truncated to a shorter signed integral type, the result is obtained by masking the original value to the length of the destination and then sign extending.
- Conversions between integral types never cause exceptions to be raised.
- Integer overflow does not cause an exception to be raised.
- Integer division by zero causes an exception to be raised.

The following points apply to operations on floating types:

- The ARM's floating point registers are wider than stored floating point numbers, so that some values may be computed to a slightly higher precision than the stated limits imply.
- When a `double` or `long double` is converted to a `float`, rounding is to the nearest representable value.
- Conversions from floating to integral types cause exceptions to be raised only if the value cannot be represented in a `long int` (or `unsigned long int` in the case of conversion to an `unsigned int`).
- Floating point underflow is not detected; any operation which underflows returns zero.
- Floating point overflow causes an exception to be raised.
- Floating point divide by zero causes an exception to be raised.

Expression evaluation

The compiler performs the 'usual arithmetic conversions' (promotions) set out in the standard before evaluating any expression.

- The compiler may re-order expressions involving only associative and commutative operators, even in the presence of parentheses.
- Between sequence points, the compiler may evaluate expressions in any order, regardless of parentheses. Thus the side effects of expressions between sequence points may occur in any order.
- Similarly, the compiler may evaluate function arguments in any order; moreover, this order may change from release to release.

Implementation limits

The standard sets out certain minimum translation limits which a conforming compiler must cope with; you should be aware of these if you are porting applications to other compilers. A summary is given here. The 'mem' limit indicates that no limit is imposed other than that of available memory.

Description	Requirement	Acorn C
Nesting levels of compound statements and iteration/selection control structures	15	mem
Nesting levels of conditional compilation	6	mem
Declarators modifying a basic type	12	mem
Expressions nested by parentheses	127	mem
Significant characters		
in internal identifiers and macro names	31	256
in external identifiers	6	256
External identifiers in one source file	511	mem
Identifiers with block scope in one block	127	mem
Macro identifiers in one source file	1024	mem
Parameters in one function definition/call	31	mem
Parameters in one macro definition/invocation	31	mem
Characters in one logical source line	509	no limit
Characters in a string literal	509	mem
Bytes in a single object	32767	mem
Nesting levels for #included files	8	mem
Case labels in a switch statement	255	mem
atexit-registered functions	32	33

8

Standard implementation definition

This chapter discusses aspects of the compiler which are not defined by the ANSI standard, but are implementation-defined and must be documented.

Appendix A.6 of the standard X3.159-1989 collects together information about portability issues; section A.6.3 lists those points which are implementation defined, and directs that each implementation shall document its behaviour in each of the areas listed. This chapter corresponds to appendix A.6.3, answering the points listed in the appendix, under the same headings and in the same order.

Translation (A.6.3.1)

- Diagnostic messages produced by the compiler are of the form "source-file", line #: *severity*: *explanation* where *severity* is one of
 - *warning*: not a diagnostic in the ANSI sense, but an attempt by the compiler to be helpful to you.
 - *error*: a violation of the ANSI specification from which the compiler was able to recover by guessing your intentions.
 - *serious error*: a violation of the ANSI specification from which no recovery was possible because the compiler could not reliably guess what you intended.
 - *too many errors/fatal error*: (for example, 'not enough memory') these are not really diagnostics but indicates that the compiler limits have been exceeded.

Environment (A.6.3.2)

- The arguments given to `main()` are the words of the Command Line (not including I/O redirections, covered in the next point), delimited by white spaces, except where the white space characters are contained in double

quotes. A white space character is any one of: space, form-feed, newline, carriage return, tab or vertical tab (note that the RISC OS Command Line interpreter filters out some of these).

A double quote or backslash character (\) inside double quotes must be preceded by a backslash character. An I/O redirection will not be recognised inside double quotes.

- The term *interactive device* denotes either the keyboard or the screen (:tt). No buffering is done on any stream connected to :tt unless I/O redirection has taken place. If I/O redirection other than to :tt has taken place, full buffering is used except where both `stdout` and `stderr` have been redirected to the same file, in which case line buffering is used.
- The standard input, output and error streams, `stdin`, `stdout`, and `stderr` can be redirected at runtime in the following way. For example, if `copy` is a compiled and linked program which simply copies the standard input to the standard output, the following line:

```
*copy <infile >outfile 2>errfile
```

runs the program, redirecting `stdin` to the file `infile`, `stdout` to the file `outfile` and `stderr` to the file `errfile`.

The following table shows all allowed redirections:

0< <i>filename</i>	read <code>stdin</code> from <i>filename</i>
< <i>filename</i>	read <code>stdin</code> from <i>filename</i>
1> <i>filename</i>	write <code>stdout</code> to <i>filename</i>
> <i>filename</i>	write <code>stdout</code> to <i>filename</i>
2> <i>filename</i>	write <code>stderr</code> to <i>filename</i>
>& <i>filename</i>	write both <code>stdout</code> and <code>stderr</code> to <i>filename</i>
1>&2	write <code>stdout</code> to wherever <code>stderr</code> is currently going
2>&1	write <code>stderr</code> to wherever <code>stdout</code> is currently going

Identifiers (A.6.3.3)

- 256 characters are significant in identifiers without external linkage. (Allowed characters are letters, digits, and underscores.)
- 256 characters are significant in identifiers with external linkage. (Allowed characters are letters, digits, and underscores.)
- Case distinctions are significant in identifiers with external linkage.

Characters (A.6.3.4)

- The characters in the source character set are ISO 8859-1 (Latin Alphabet), a superset of the ASCII character set. The printable characters are those in the range 32 to 126 and 160 to 255. All printable characters may appear in string or character constants, and in comments.
- There are no locales implemented for which a multibyte character shift state exists.
- The execution character set is identical to the source character set.
- There are four `chars` in an `int`. The bytes are ordered from least significant at the lowest address to most significant at the highest address.
- There are eight bits in a character in the execution character set.
- All integer character constants that contain a character or escape sequence are represented in the source and execution character set.
- Characters of the source character set in string literals and character constants map identically into characters in the execution character set.
- No locale is used to convert multibyte characters into the corresponding wide characters (codes) for a wide character constant.
- A character constant containing more than one character has the type `int`. Up to four characters of the constant are represented in the integer value. The first character contained in the constant occupies the lowest-addressed byte of the integer value; up to three following characters are placed at ascending addresses. Unused bytes are filled with the `NULL` (or `/0`) character. This is not portable.
- A plain `char` is treated as unsigned (signed in `-pcc` mode).
- Escape codes are:

Escape sequence	Char value	Description
<code>\a</code>	7	Attention (bell)
<code>\b</code>	8	Backspace
<code>\f</code>	12	Form feed
<code>\n</code>	10	Newline
<code>\r</code>	13	Carriage return
<code>\t</code>	9	Tab
<code>\v</code>	11	Vertical tab
<code>\xnn</code>	<code>nn</code>	ASCII code in hexadecimal
<code>\nnn</code>	<code>nnn</code>	ASCII code in octal

Integers (A.6.3.5)

The representations and sets of values of the integral types are set out in the section entitled *Data elements* on page 73. Note also that:

- The result of converting an integer to a shorter signed integer, if the value cannot be represented, is as if the bits in the original value which cannot be represented in the final value were masked out, and the resulting integer sign-extended. The same applies when you convert an unsigned integer to a signed integer of equal length.
- Bitwise operations on signed integers yield the expected result given two's complement representation. No sign extension takes place.
- The sign of the remainder on integer division is the same as defined for the function `div()`.
- Right shift operations on signed integral types are arithmetic.

Floating point (A.6.3.6)

The representations and ranges of values of the floating point types have been given in the section entitled *Data elements* on page 73. Note also that:

- When a floating point number is converted to a shorter floating point one, it is rounded to the nearest representable number.
- The properties of floating point arithmetic accord with IEEE 754.

Arrays and pointers (A.6.3.7)

The ANSI standard specifies three areas in which the behaviour of arrays and pointers must be documented. The points to note are:

- The type `size_t` is defined as `unsigned int`.
- Casting pointers to integers and vice versa involves no change of representation. Thus any integer obtained by casting from a pointer will be positive.
- The type `ptrdiff_t` is defined as `(signed) int`.

Registers (A.6.3.8)

In the Acorn C compiler, you can declare up to six objects as having the storage class `register`. There are six available registers, so declaring more than six objects with register storage class will result in at least one of them not being held in a register. It is advisable to declare no more than four. The valid types are:

- any integer type

- any pointer type
- any structure type which contains only bitfields and which is no more than one word long.

Note that other variables, not declared as `register`, may be held in registers for extended periods, and that `register` variables may be held in memory for some periods.

Structures, unions, enumerations and bitfields (A.6.3.9)

The Acorn C compiler handles structures in the following way:

- When a member of a union is accessed using a member of a different type, the resulting value can be predicted from the representation of the original type. No error is given.
- Structures are aligned on word boundaries. Characters are aligned in bytes, shorts on even numbered byte boundaries and all other types, except bitfields, are aligned on word boundaries. Bitfields are parts of `ints`, themselves aligned on word boundaries.
- A 'plain' bitfield (declared as `int`) is treated as `unsigned int` (`signed int` in `-pcc` mode).
- A bitfield which does not fit into the space remaining in an `int` is placed in the next `int`.
- The order of allocation of bitfields within `ints` is such that the first field specified occupies the least significant bits of the word.
- Bitfields do not straddle storage unit (`int`) boundaries.
- The integer type chosen to represent the values of an enumeration type is `int` (`signed int`).

Qualifiers (A.6.3.10)

A read or write constitutes an access to an object that has volatile-qualified type.

Declarators (A.6.3.11)

The number of declarators that may modify an arithmetic, structure or union type is limited only by available memory.

Statements (A.6.3.12)

The number of `case` values in a `switch` statement is limited only by memory.

Preprocessing directives (A.6.3.13)

- A single-character constant in a preprocessor directive cannot have a negative value.
- The standard header files are contained within the compiler itself. The mechanism for translating the standard suffix notation to an Acorn filename is described in the chapter CC.
- Quoted names for includable source files are supported. The rules for directory searching are given in the chapter CC.
- The recognized `#pragma` directives and their meaning are described in the section entitled *#pragma directives* on page 378.
- The date and time of translation are always available, so `__DATE__` and `__TIME__` always give respectively the date and time.

Library functions (A.6.3.14)

When using library functions in the Acorn C compiler, note the following points:

- The macro `NULL` expands to the integer constant 0.
- If a program redefines a reserved external identifier, then an error may occur when the program is linked with the standard libraries. If it is not linked with standard libraries, no error will be detected.
- The `assert()` function prints the following message:

```
*** assertion failed: expression, file filename, line,  
line-number
```

and then calls the function `abort()`.

- The functions:

```
isalnum()
isalpha()
iscntrl()
islower()
isprint()
isupper()
ispunct()
```

usually test only for characters whose values are in the range 0 to 127 (inclusive). Characters with values greater than 127 return a result of 0 for all of these functions, except `iscntrl()` which returns non-zero for 0 to 31, and 128 to 255.

After the call `setlocale(LC_CTYPE, "ISO8859-1")` the following statements also apply for characters:

```
0 to 31 are control characters
128 to 159 are control characters
192 to 223 except 215 are upper case
224 to 255 except 247 are lower case
160 to 191, and 215 and 247 are punctuation
```

The results returned by the functions reflect this.

- The mathematical functions return the following values on domain errors:

Function	Condition	Returned value
<code>log(x)</code>	$x \leq 0$	-HUGE_VAL
<code>log10(x)</code>	$x \leq 0$	-HUGE_VAL
<code>sqrt(x)</code>	$x < 0$	-HUGE_VAL
<code>atan2(x, y)</code>	$x = y = 0$	-HUGE_VAL
<code>asin(x)</code>	$\text{abs}(x) > 1$	-HUGE_VAL
<code>acos(x)</code>	$\text{abs}(x) > 1$	-HUGE_VAL

Where -HUGE_VAL is written above, a number is returned which is defined in the header `h.math`. Consult the `errno` variable for the error number.

- The mathematical functions set `errno` to `ERANGE` on underflow range errors.
- A domain error occurs if the second argument of `fmod` is zero, and -HUGE_VAL returned.
- The set of signals for the `signal()` function is as follows:

SIGABRT	Abort
SIGFPE	Arithmetic exception
SIGILL	Illegal instruction
SIGINT	Attention request from user

SIGSEGV	Bad memory access
SIGTERM	Termination request
SIGSTAK	Stack overflow

- The default handling of all the signals recognised is the printing of a suitable message followed by a stack backtrace. This default behaviour applies at program start-up.
- When a signal occurs, if `func` points to a function, the equivalent of `signal(sig, SIG_DFL);` is first executed.
- If the `SIGILL` signal is received by a handler specified to the signal function, the default handling is reset.
- The last line of a text stream does not require a terminating newline character.
- Space characters written out to a text stream immediately before a newline character do appear when read in.
- No null characters are appended to a binary output stream.
- The file position indicator of an append mode stream is initially placed at the end of the file.
- A write to a text stream does not cause the associated file to be truncated beyond that point.
- The characteristics of file buffering are as intended in the standard (section 4.9.3).
- A zero-length file (on which no characters have been written by an output stream) does exist.
- The validity of filenames is defined by the host computer's filing system.
- The same file can be opened many times for reading, and once for writing or updating. A file cannot however be open for reading on one stream and for writing or updating on another.
- Local time zones and Daylight Saving Time are not implemented. The values returned will always indicate that the information is not available.
- Note also the following points about library functions:
 - `remove()` Cannot remove an open file.
 - `rename()` The effect of calling the `rename()` function when the new name already exists is dependent on the host filing system. Not all renames are valid: examples of invalid renames

```

include
("net:file1", "net:$.file2") and
("net:file1", "adfs:file2").
fprintf() Prints %p arguments in hexadecimal format (lower case) as if
a precision of 8 had been specified. If the variant form (%#p)
is selected, the number is preceded by the character @.
fscanf() Treats %p arguments identically to %x arguments.
fscanf() Always treats the character - in a %[ argument as a literal
character.
ftell() and
fgetpos() Set errno to the value of EDOM on failure.
perror() Generates the following messages:
Error: Message:
0 No error (errno = 0)
EDOM EDOM - function argument out of range
ERANGE ERANGE - function result not representable
ESIGNUM ESIGNUM - illegal signal number to signal() or raise()
others Error code number has no associated message
calloc(),
malloc() and
realloc() If size of area requested is zero, NULL is returned.
abort() Closes all open files, and deletes all temporary files.
exit() The status returned by exit is the same value that was
passed to it. For a definition of EXIT_SUCCESS and
EXIT_FAILURE refer to the header file stdlib.h.
getenv() Returns the value of the named RISC OS Environmental
variable, or NULL if the variable had no value.
eg root = getenv ("C$libroot");
if (root == NULL) root = "$.arm.clib";
system() Used either to CHAIN to another application or built-in
command or to CALL one as a sub-program. When a program
is chained, all trace of the original program is removed from
memory and the chained program invoked. If a program is
called (which is the default if no CHAIN: or CALL: precedes
the program name - a change from Release 2), the calling
program and data are moved in memory to somewhere safe
and the callee loaded and started up. The return value from

```

the `system()` call is `-2` (indicating a failure to invoke the program) or the value of `Sys$ReturnCode` set by the called program (`0` indicates success).

`strerror()` The error messages given by this function are identical to those given by the `perror()` function.

`clock()` Returns the time taken by the program since its invocation, as indicated by the host's operating system.

The C programming language has gained a reputation for being portable across machines, while still providing capabilities at a machine-specific level. The fact that a program is written in C by no means indicates the effort required to port software from one machine to another, or indeed from one compiler to another. Obviously the most time-consuming task is porting between two entirely different hardware environments, running different operating systems with different compilers. Since many users of the Acorn C compiler will find themselves in this situation, this chapter deals with a number of issues you should be aware of when porting software to or from our environment. The chapter covers the following:

- general portability considerations
- major differences between ANSI C and the well-known 'K&R' C as defined in the book *The C Programming Language*, (first edition) by Kernighan and Ritchie
- using the Acorn C compiler in 'pcc' compatibility mode
- environmental aspects of portability.

General portability considerations

If you intend your code to be used on a variety of different systems, there are certain aspects which you should bear in mind in order to make porting an easy and relatively error-free process. It is essential to single out items which may make software system-specific, and to employ techniques to avoid non-portable use of such items. In this section, we describe general portability issues for C programs.

Fundamental data types

The size of fundamental data types such as `char`, `int`, `long int`, `short int` and `float` will depend mainly on the underlying architecture of the machine on which the C program is to run. Compiler writers usually implement these types in a manner which best fits the architectures of machines for which their compilers are targetted. For example, Release 5 of the Microsoft C Compiler has `int`, `short int` and `long int` occupying 2, 2 and 4 bytes respectively, where the Acorn C Compiler uses 4, 2 and 4 bytes. Certain relations are guaranteed by the ANSI C Standard (such as the fact that the size of `long int` is at least that of `short int`), but code which makes any assumptions regarding implementation-defined issues such as whether `int` and `long int` are the same size will not be maximally portable.

A common non-portable assumption is embedded in the use of hexadecimal constant values. For example:

```
int i;
i = i & 0xffffffff8; /* set bottom 3 bits to zero, assuming 32-bit int */
```

Such non-portability can be avoided by using:

```
int i;
i = i & ~0x07; /* set bottom 3 bits to zero, whatever sizeof(int) */
```

If you find that some size assumptions are inevitable, then at least use a series of `assert` calls when the program starts up, to indicate any conditions under which successful operation is not guaranteed. Alternatively, write macros for frequently-used operations so that size assumptions are localised and can be altered locally.

Byte ordering

A highly non-portable feature of many C programs is the implicit or explicit exploitation of byte ordering within a word of store. Such assumptions tend to arise when copying objects word by word (rather than byte by byte), when inputting and outputting binary values, and when extracting bytes from or inserting bytes into words using a mix of shift-and-mask and byte addressing. A contrived example is the following code which copies individual bytes from an `int` variable `w` into an `int` variable pointed to by `p`, until a null byte is encountered. The code assumes that `w` does contain a null byte.

```
int a;
char *p = (char *)&a;
int w = AN_ARBITRARY_VALUE;

for (;;)
{
    if ((*p++ = w) == 0) break;
    w >>= 8;
}
```

This code will only work on a machine with even (or little-endian) byte-sex, and so is not portable. The best solution to such problems is either to write code which does not rely on byte-sex, or to have different code to deal appropriately with different byte-sex and to compile the correct variant conditionally, depending on your target machine architecture.

Store alignment

The only guarantee given in the ANSI C Standard regarding alignment of members of a `struct`, is that a 'hole' (caused by padding) cannot exist at the beginning of the `struct`. The values of 'holes' created by alignment restrictions are undefined, and you should not make assumptions about these values. In particular, two structures with identical members, each having identical values, will only be considered equal if field-by-field comparison is used; a byte-by-byte, or word-by-word comparison may not indicate equality.

This may also have implications on the size requirements of large arrays of `structs`. Given the following declarations:

```
#define ARRSIZE 10000
typedef struct
{
    int i;
    short s;
} ELEM;
ELEM arr[ARRSIZE];
```

this may require significantly different amounts of store under, say, a compiler which aligns `ints` on even boundaries, as opposed to one which aligns them on word boundaries.

Pointers and pointer arithmetic

A deficiency of the original definition of C, and of its subsequent use, has been the relatively unrestrained interchanging between pointers to different data types and integers or longs. Much existing code makes the assumption that a pointer can safely be held in either a `long int` or `int` variable. While such an assumption may indeed be true in many implementations on many machines, it is a highly non-portable feature on which to rely.

This problem is further compounded when taking the difference of two pointers by performing a subtraction. When the difference is large, this approach is full of possible errors. For this purpose, ANSI C defines a type `ptrdiff_t`, which is capable of reliably storing the result of subtracting two pointer values of the same type; a typical use of this mechanism would be to apply it to pointers into the same array.

Function argument evaluation

Whilst the evaluation of operands to such operators as `&&` and `||` is defined to be strictly left-to-right (including all side-effects), the same does not apply to function argument evaluation. For example, in the function call `f(i, i++)`; the issue of

whether the post-increment of `i` is performed after the first use of `i` is implementation-dependent. In any case, this is an unwise form of statement, since it may be decided later to implement `f` as a macro, instead of a function.

System-specific code

The direct use of operating system calls is, as you would expect, non-portable. If you use code which is obviously targetted for a particular environment, then it should be clearly documented as such, and should preferably be isolated into a system-specific module, which needs to be modified when porting to a new machine or operating system. Pathnames of system files should be `#defined` and not hard-coded into the program, and, as far as possible, all processing of filenames should be made easy to modify. Many file operations can be written in terms of the ANSI input/output library functions, which will make an application more portable. Obviously, binary data files are inherently non-portable, and the only solution to this problem may be the use of some portable external representation.

ANSI C vs K&R C

The ANSI C Standard has succeeded in tightening up many of the vague areas of K&R C. This results in a much clearer definition of a correct C program. However, if programs have been written to exploit particular vague features of K&R C, then their authors may find surprises when porting to an ANSI C environment. In the following sections, we present a list of what we consider to be the major differences between ANSI and K&R C. These differences are at the language level, and we defer discussion of library differences until a later section. The order in which this list is presented follows approximately relevant parts of the ANSI C Standard Document.

Lexical elements

The ordering of phases of translation is well-defined. Of special note is the preprocessor which is conceptually token-based (which does not yield the same results as might naively be expected from pure text manipulation).

A number of new keywords have been introduced with the following meanings:

- The type qualifier `volatile` which means that the object may be modified in ways unknown to the implementation, or have other unknown side effects. Examples of objects correctly described as `volatile` include device registers, semaphores and flags shared with asynchronous signal handlers. In general, expressions involving `volatile` objects cannot be optimised by the compiler.

- The type qualifier `const` which indicates that a variable's value should not be changed.
- The type specifier `void` to indicate a non-existent value for an expression.
- The type specifier `void *`, which is a generic pointer to or from which pointer variables can be assigned, without loss of information.
- The `signed` type qualifier, to sign any integral types explicitly.
- `structs` and `unions` have their own distinct name spaces.
- There is a new floating-point type `long double`.
- The K&R C practice of using `long float` to denote `double` is now outlawed in ANSI C.
- Suffixes `U` and `L` (or `u` and `l`), can be used to explicitly denote `unsigned` and `long` constants (eg. `32L`, `64U`, `1024UL` etc).
- The use of 'octal' constants `8` and `9` (previously defined to be octal `10` and `11` respectively) is no longer supported.
- Literal strings are to be considered as read-only, and identical strings may be stored as one shared version (as indeed they are, in the Acorn C Compiler). For example, given:

```
char *p1 = "hello";
char *p2 = "hello";
```

`p1` and `p2` will point at the same store location, where the string `hello` is held. Programs should not therefore modify literal strings.

- Variadic functions (ie those which take a variable number of arguments) are declared explicitly using an ellipsis (...). For example, `int printf(const char *fmt, ...);`
- Empty comments `/**/` are replaced by a single space (use the preprocessor directive `##` to do token-pasting if you previously used `/**/` to do this).

Conversions

ANSI C uses value-preserving rules for arithmetic conversions (whereas K&R C implementations tend to use unsigned-preserving rules). Thus, for example:

```
int f(int x, unsigned char y)
{
    return (x+y)/2;
}
```

does signed division, where unsigned-preserving implementations would do unsigned division.

Aside from value-preserving rules, arithmetic conversions follow those of K&R C, with additional rules for `long double` and `unsigned long int`. It is now also possible to perform `float` arithmetic without widening to `double`. Floating-point values truncate towards zero when they are converted to integral types.

It is illegal to attempt to assign function pointers to data pointers and vice versa (even using explicit casts). The only exception to this is the value 0, as in:

```
int (*pfi)();
pfi = 0;
```

Assignment compatibility between `structs` and `unions` is now stricter. For example, consider the following:

```
struct {char a; int b;} v1;
struct {char a; int b;} v2;
v1 = v2; /* illegal because v1 and v2
          strictly have different types*/
```

Expressions

- `structs` and `unions` may be passed by value as arguments to functions.
- Given a pointer to function declared as, say, `int (*pfi)();`, then the function to which it points can be called either by `pfi();` or `(*pfi)();`.
- Due to the use of distinct name spaces for `struct` and `union` members absolute machine addresses must be explicitly cast before being used as `struct` and `union` pointers. For example:

```
((struct io_space *)0x00ff)->io_buf;
```

Declarations

Perhaps the greatest impact on C of the ANSI Standard has been the adoption of function prototypes. A function prototype declares the return type and argument types of a function. For example, `int f(int, float);` declares a function returning `int` with one `int` and one `float` argument. This means that a function's argument types are part of the type of that function, thus giving the advantage of stricter argument type-checking, especially across source files. A function definition (which is also a prototype) is similar except that identifiers must be given for the arguments. For example, `int f(int i, float f);`. It is still possible to use 'old style' function declarations and definitions, but you are advised to convert to the 'new style'. It is also possible to mix old and new styles of function declaration. If the function declaration which is in scope is an old style

one, normal integral promotions are performed for integral arguments, and floats are converted to `double`. If the function declaration which is in scope is a new style one, arguments are converted as in normal assignment statements.

Empty declarations are now illegal.

Arrays cannot be defined to have zero or negative size.

Statements

- ANSI has defined the minimum attributes of control statements (eg the minimum number of case limbs which must be supported by a compiler). These values are almost invariably greater than those supported by PCCs, and so should not present a problem.
- A value returned from `main()` is guaranteed to be used as the program's exit code.
- Values used in the controlling statement and labels of a `switch` can be of any integral type.

Preprocessor

- Preprocessor directives cannot be redefined.
- There is a new `##` directive for token-pasting.
- There is a directive `#` which produces a string literal from its following characters. This is useful for cases where you want replacement of macro arguments in strings.
- The order of phases of translation is well defined and is as follows for the preprocessing phases:
 - 1 Map source file characters to the source character set (this includes replacing trigraphs).
 - 2 Delete all newline characters which are immediately preceded by `\`.
 - 3 Divide the source file into preprocessing tokens and sequences of white space characters (comments are replaced by a single space).
 - 4 Execute preprocessing directives and expand macros.

Any `#include` files are passed through steps 1-4 recursively.

The macro `__STDC__` is `#defined` to 1 in ANSI-conforming compilers.

The `ToPCC` and `ToANSI` tools

The DDE tools `ToPCC` and `ToANSI` help you to translate C programs and headers between the ANSI and PCC dialects of C. For more details of their use and capabilities see the earlier chapters `ToANSI` and `ToPCC`.

`pcc` compatibility mode

This section discusses the differences apparent when the compiler is used in 'PCC' mode. When the **UNIX `pcc`** setup option is enabled, the C compiler will accept (Berkeley) UNIX-compatible C, as defined by the implementation of the Portable C Compiler and subject to the restrictions which are noted below.

In essence, PCC-style C is K&R C, as defined by B Kernighan and D Ritchie in their book *The C Programming Language*, with a small number of extensions and clarifications of language features that the book leaves undefined.

Language and preprocessor compatibility

In `Unix pcc` mode, the Acorn C compiler accepts K&R C, but it does not accept many of the old-style compatibility features, the use of which has been deprecated and warned against for many years. Differences are listed briefly below:

- Compound assignment operators where the = sign comes first are accepted (with a warning) by some PCCs. An example is `=+` instead of `+=`. Acorn C does not allow this ordering of the characters in the token.
- The = sign before a `static` initialiser was not required by some very old C compilers. Acorn C does not support this syntax.
- The following very peculiar usage is found in some UNIX tools pre-dating UNIX Version 7:

```
    struct {int a, b;};
    double d;

    d.a = 0;
    d.b = 0x....;
```

This is accepted by some UNIX PCCs and may cause problems when porting old (and badly written) code.

- `enums` are less strongly typed than is usual under PCCs. `enum` is a non-K&R extension to C which has been standardised by ANSI somewhat differently from the usual PCC implementation.
- `chars` are signed by default in `Unix pcc` mode.
- In `Unix pcc` mode, the compiler permits the use of the ANSI `'...'` notation which signifies that a variable number of formal arguments follow.

- In order to cater for PCC-style use of variadic functions, a version of the PCC header file `varargs.h` is supplied with the release.
- With the exception of enums, the compiler's type checking is generally stricter than PCC's – much more akin to lint's, in fact. In writing the Acorn C compiler, we have attempted to strike a balance between generating too many warnings when compiling known, working code, and warning of poor or non-portable programming practices. Many PCCs silently compile code which has no chance of executing in just a slightly different environment. We have tried to be helpful to those who need to port C among machines in which the following varies:
 - the order of bytes within a word (eg little-endian ARM, VAX, Intel versus big-endian Motorola, IBM370)
 - the default size of `int` (four bytes versus two bytes in many PC implementations)
 - the default size of pointers (not always the same as `int`)
 - whether values of type `char` default to signed or unsigned `char`
 - the default handling of undefined and implementation-defined aspects of the C language.

If the verbosity of CC in UNIX `pcc` mode is found undesirable, all warnings can be turned off using the **Suppress** setup option.

- The compiler's preprocessor is believed to be equivalent to UNIX's `cpp`, except for the points listed below. Unfortunately, `cpp` is only defined by its implementation, and although equivalence has been tested over a large body of UNIX source code, completely identical behaviour cannot be guaranteed. Some of the points listed below only apply when the **Preprocess only** option is used with the CC tool.
 - There is a different treatment of whitespace sequences (benign).
 - `n1` is processed by CC with **Preprocess only** enabled, but passed by `cpp` (making lines longer than expected).
 - `Cpp` breaks long lines at a token boundary; CC with **Preprocess only** enabled doesn't (this may break line-size constraints when the source is later consumed by another program).
 - The handling of unrecognised `#` directives is different (this is mostly benign).

Standard headers and libraries

Use of the compiler in UNIX `pcc` mode precludes neither the use of the standard ANSI headers built in to the compiler nor the use of the run-time library supplied with the C compiler. Of course, the ANSI library does not contain the whole of the

UNIX C library, but it does contain almost all the commonly used functions. However, look out for functions with different names, or a slightly different definition, or those in different 'standard' places. Unless the user directs otherwise using `Default_path`, the C compiler will attempt to satisfy references to, say, `<stdio.h>` from its in-store filing system.

Listed below are a number of differences between the ANSI C Library, and the BSD UNIX library. They are placed under headings corresponding to the ANSI header files:

ctype.h

There are no `isascii()` and `toascii()` functions, since ANSI C is not character-set specific.

errno.h

On BSD systems there are `sys_nerr` and `sys_errlist()` defined to give error messages corresponding to error numbers. ANSI C does not have these, but provides similar functionality via `perror(const char *s)`, which displays the string pointed to by `s` followed by a system error message corresponding to the current value of `errno`.

There is also `char *strerror(int errnum)` which, when given a purported value of `errno`, returns its textual equivalent.

math.h

The `#defined` value `HUGE`, found in BSD libraries, is called `HUGE_VAL` in ANSI C. ANSI C does not have `asinh()`, `acosh()`, `atanh()`.

signal.h

In ANSI C the `signal()` function's prototype is:

```
extern void (*signal(int, void(*func)(int)))(int);
```

`signal()` therefore expects its second argument to be a pointer to a function returning `void` with one `int` argument. In BSD-style programs it is common to use a function returning `int` as a signal handler. The PCC-style function definitions shown below will therefore produce a compiler warning about an implicit cast between different function pointers (since `f()` defaults to `int f()`) This is just a warning, and correct code will be generated anyway.

```
f(signo)
int signo;
{
.....
}

main()
{
extern f();
signal(SIGINT, f);
}
```

stdio.h

`sprintf()` now returns the number of characters 'printed' (following UNIX System V), whereas the BSD `sprintf()` returns a pointer to the start of the character buffer.

The BSD functions `ecvt()`, `fcvt()` and `gcvt()` are not included in ANSI C, since their functionality is provided by `sprintf()`.

string.h

On BSD systems, string manipulation functions are found in `strings.h`, whereas ANSI C places them in `<string.h>`. The Acorn C Compiler also has `strings.h` for PCC-compatibility.

The BSD functions `index()` and `rindex()` are replaced by the ANSI functions `strchr()` and `strrchr()` respectively.

Functions which refer to string lengths (and other sizes) now use the ANSI type `size_t`, which in our implementation is `unsigned int`.

stdlib.h

`malloc()` returns `void *`, rather than the `char *` of the BSD `malloc()`.

float.h

A new header added by ANSI giving details of floating point precision etc.

limits.h

A new header added by ANSI to give maximum and minimum limit values for data types.

locale.h

A new header added by ANSI to provide local environment-specific features.

Environmental aspects

When porting an application, the most extensive changes will probably need to be made at the operating system interface level. The following is a brief description of aspects of RISC OS and Acorn C which differ from systems such as UNIX and MS-DOS.

The most apparent interface between a C program and its environment is via the arguments to `main()`. The ANSI Standard declares that `main()` is a function defined as the program entry point with either no arguments or two arguments (one giving a count of command line arguments, commonly called `int argc`, the other an array of pointers to the text of the arguments themselves, after removal of input/output redirection, commonly called `char *argv[]`). As discussed in the section entitled *Environment* (A.6.3.2) on page 81, Acorn C supports the style of input/output redirection used by UNIX BSD4.3, but does not support filename wildcarding. Further parameters to `main()` are not supported.

Under UNIX and MS-DOS, it is common to use a third parameter, normally called `char *environ[]` under UNIX and `char *envp[]` under Microsoft C for MS-DOS, to give access to environment variables. The same effect can be achieved in our system by using `getenv()` to request system variable values explicitly; the names of these variables are as they appear from a RISC OS `*Show` command. The string pointed at by `argv[0]` is the program name (similar to UNIX and MS-DOS, except the name is exactly that typed on invocation, so if a full pathname is used to invoke the program, this is what appears in `argv[0]`).

File naming is one of the least portable aspects in any programming environment. RISC OS uses a full stop (`.`) as a separator in pathnames and does not support filename extensions (nor does UNIX, but existing UNIX tools make assumptions about file naming conventions). The best way to simulate extensions is to create a directory whose name corresponds to the required extension (in a manner similar to the use of `c` and `h` directories for C source and header files). RISC OS filename components are limited to 10 characters.

The Acorn C compiler has support for making Software Interrupt (SWI) calls to RISC OS routines, which can be used to replace any system calls which you make under UNIX or MS-DOS. The include file `kernel.h` has function prototypes and appropriate `typedefs` for issuing SWIs. Briefly, the type `_kernel_swi_regs` allows values to be placed in registers R0-R9, and `_kernel_swi()` can then be used to issue the SWI; a list of SWI numbers can be found in the include file `swis.h`. File information, for example, can be obtained in a way similar to `stat()` under UNIX, by making an `OS_GBPB` SWI with R0 set to the reason code

11 (full file information). Most of the UNIX/MS-DOS low-level I/O can be simulated in this way, but the ANSI C run-time library provides sufficient support for most applications to be written in a portable style. If the application is running under the desktop, then limited piping facilities can be achieved by using the calls `wimp_transferblock` and `wimp_sendmessage` to synchronise the data transfer.

RISC OS does not support different memory models as in MS-DOS, so programs which have been written to exploit this will need modification; this should only require the removal of Microsoft C keywords such as `near`, `far` and `huge`, if the program has otherwise been written with portability in mind.

10 ANSI library reference section

assert.h

The `assert` macro puts diagnostics into programs. When it is executed, if its argument expression is false, it writes information about the call that failed (including the text of the argument, the name of the source file, and the source line number, the last two of these being, respectively, the values of the preprocessing macros `__FILE__` and `__LINE__`) on the standard error stream. It then calls the `abort` function. If its argument expression is true, the `assert` macro returns no value.

If `NDEBUG` is `#defined` prior to inclusion of `assert.h`, calls to `assert` expand to null statements. This provides a simple way to turn off the generation of diagnostics selectively.

Note that `assert.h` may be included more than once in a program with different settings of `NDEBUG`.

ctype.h

`ctype.h` declares several functions useful for testing and mapping characters. In all cases the argument is an `int`, the value of which is representable as an unsigned `char` or equal to the value of the macro `EOF`. If the argument has any other value, the behaviour is undefined.

<code>int isalnum(int c)</code>	Returns true if <code>c</code> is alphabetic or numeric
<code>int isalph(int c)</code>	Returns true if <code>c</code> is alphabetic
<code>int iscntrl(int c)</code>	Returns true if <code>c</code> is a control character (in the ASCII locale)
<code>int isdigit(int c)</code>	Returns true if <code>c</code> is a decimal digit
<code>int isgraph(int c)</code>	Returns true if <code>c</code> is any printable character other than space
<code>int islower(int c)</code>	Returns true if <code>c</code> is a lower-case letter
<code>int isprint(int c)</code>	Returns true if <code>c</code> is a printable character (in the ASCII locale this means 0x20 (space) → 0x7E (tilde) inclusive).

<code>int ispunct(int c)</code>	Returns true if <code>c</code> is a printable character other than a space or alphanumeric character
<code>int isspace(int c)</code>	Returns true if <code>c</code> is a white space character viz: space, newline, return, linefeed, tab or vertical tab
<code>int isupper(int c)</code>	Returns true if <code>c</code> is an upper-case letter
<code>int isxdigit(int c)</code>	Returns true if <code>c</code> is a hexadecimal digit, ie in 0...9, a...f, or A...F
<code>int tolower(int c)</code>	Forces <code>c</code> to lower case if it is an upper-case letter, otherwise returns the original value
<code>int toupper(int c)</code>	Forces <code>c</code> to upper case if it is a lower-case letter, otherwise returns the original value

errno.h

This file contains the definition of the macro `errno`, which is of type `volatile int`. It contains three macro constants defining the error conditions listed below.

EDOM

If a domain error occurs (an input argument is outside the domain over which the mathematical function is defined) the integer expression `errno` acquires the value of the macro `EDOM` and `HUGE_VAL` is returned. `EDOM` may be used by non-mathematical functions.

ERANGE

A range error occurs if the result of a function cannot be represented as a double value. If the result overflows (the magnitude of the result is so large that it cannot be represented in an object of the specified type), the function returns the value of the macro `HUGE_VAL`, with the same sign as the correct value of the function; the integer expression `errno` acquires the value of the macro `ERANGE`. If the result underflows (the magnitude of the result is so small that it cannot be represented in an object of the specified type), the function returns zero; the integer expression `errno` acquires the value of the macro `ERANGE`. `ERANGE` may be used by non-mathematical functions.

ESIGNUM

If an unrecognised signal is caught by the default signal handler, `errno` is set to `ESIGNUM`.

float.h

This file contains a set of macro constants which define the limits of computation on floating point numbers. These are discussed in the chapter entitled *Implementation details* on page 73.

limits.h

This set of macro constants determines the upper and lower value limits for integral objects of various types, as follows:

Object type	Minimum value	Maximum value
Byte (number of bits)	0	8
Signed char	-128	127
Unsigned char	0	255
Char	0	255
Multibyte character (number of bytes)	0	1
Short int	-0x8000	0x7fff
Unsigned short int	0	65535
Int	(~0x7fffffff)	0x7fffffff
Unsigned int	0	0xffffffff
Long int	(~0x7fffffff)	0x7fffffff
Unsigned long int	0	0xffffffff

See also the chapter entitled *Implementation details* on page 73.

locale.h

This file handles national characteristics, such as the different orderings month-day-year (USA) and day-month-year (UK).

```
char *setlocale(int category, const char *locale)
```

Selects the appropriate part of the program's locale as specified by the `category` and `locale` arguments. The `setlocale` function may be used to change or query the program's entire current locale or portions thereof. Locale information is divided into the following types:

LC_COLLATE	string collation
LC_CTYPE	character type
LC_MONETARY	monetary formatting
LC_NUMERIC	numeric string formatting
LC_TIME	time formatting
LC_ALL	entire locale

The locale string specifies which locale set of information is to be used. For example,

setlocale

```
setlocale(LC_MONETARY, "uk")
```

would insert monetary information into the `lconv` structure. To query the current locale information, set the `locale` string to null and read the string returned.

lconv

```
struct lconv *localeconv(void)
```

Sets the components of an object with type `struct lconv` with values appropriate for the formatting of numeric quantities (monetary and otherwise) according to the rules of the current locale. The members of the structure with type `char *` are strings, any of which (except `decimal_point`) can point to "", to indicate that the value is not available in the current locale or is of zero length. The members with type `char` are non-negative numbers, any of which can be `CHAR_MAX` to indicate that the value is not available in the current locale. The members included are described above.

`localeconv` returns a pointer to the filled in object. The structure pointed to by the return value will not be modified by the program, but may be overwritten by a subsequent call to the `localeconv` function. In addition, calls to the `setlocale` function with categories `LC_ALL`, `LC_MONETARY`, or `LC_NUMERIC` may overwrite the contents of the structure.

math.h

This file contains the prototypes for 22 mathematical functions. All return the type `double`.

Function

```
double acos(double x)
```

```
double asin(double x)
```

```
double atan(double x)
```

```
double atan2(double x, double y)
```

```
double cos(double x)
```

```
double sin(double x)
```

```
double tan(double x)
```

```
double cosh(double x)
```

Returns

arc cosine of x . A domain error occurs for arguments not in the range -1 to 1

arc sine of x . A domain error occurs for arguments not in the range -1 to 1

arc tangent of x

arc tangent of y/x

cosine of x (measured in radians)

sine of x (measured in radians)

tangent of x (measured in radians)

hyperbolic cosine of x

<code>double sinh(double x)</code>	hyperbolic sine of x
<code>double tanh(double x)</code>	hyperbolic tangent of x
<code>double exp(double x)</code>	exponential function of x
<code>double frexp(double x, int *exp)</code>	the value x , such that x is a double with magnitude in the interval 0.5 to 1.0 or zero, and value equals x times 2 raised to the power $*exp$
<code>double ldexp(double x, int exp)</code>	x times 2 raised to the power of exp
<code>double log(double x)</code>	natural logarithm of x
<code>double log10(double x)</code>	log to the base 10 of x
<code>double modf(double x, double *iptr)</code>	signed fractional part of x . Stores integer part of x in object pointed to by $iptr$.
<code>double pow(double x, double y)</code>	x raised to the power of y
<code>double sqrt(double x)</code>	positive square root of x
<code>double ceil(double x)</code>	smallest integer not less than x (ie rounding up)
<code>double fabs(double x)</code>	absolute value of x
<code>double floor(double x)</code>	largest integer not greater than x (ie rounding down)
<code>double fmod(double x, double y)</code>	floating-point remainder of x/y

setjmp.h

This file declares two functions, and one type, for bypassing the normal function call and return discipline (useful for dealing with unusual conditions encountered in a low-level function of a program). It also defines the `jmp_buf` structure type required by these routines.

setjmp

```
int setjmp(jmp_buf env)
```

The calling environment is saved in `env`, for later use by the `longjmp` function. If the return is from a direct invocation, the `setjmp` function returns the value zero. If the return is from a call to the `longjmp` function, the `setjmp` function returns a non-zero value.

longjmp

```
void longjmp(jmp_buf env, int val)
```


The environment saved in `env` by the most recent call to `setjmp` is restored. If there has been no such call, or if the function containing the call to `setjmp` has terminated execution (eg with a return statement) in the interim, the behaviour is undefined. All accessible objects have values as at the time `longjmp` was called, except that the values of objects of automatic storage duration that do not have volatile type and that have been changed between the `setjmp` and `longjmp` calls are indeterminate.

As it bypasses the usual function call and return mechanism, the `longjmp` function executes correctly in contexts of interrupts, signals and any of their associated functions. However, if the `longjmp` function is invoked from a nested signal handler (that is, from a function invoked as a result of a signal raised during the handling of another signal), the behaviour is undefined.

After `longjmp` is completed, program execution continues as if the corresponding call to `setjmp` had just returned the value specified by `val`. The `longjmp` function cannot cause `setjmp` to return the value 0; if `val` is 0, `setjmp` returns the value 1.

signal.h

Signal declares a type (`sig_atomic_t`) and two functions.

It also defines several macros for handling various signals (conditions that may be reported during program execution). These are `SIG_DFL` (default routine), `SIG_IGN` (ignore signal routine) and `SIG_ERR` (dummy routine used to flag error return from signal).

```
void (*signal (int sig, void (*func)(int)))(int)
```

Think of this as

```
typedef void Handler(int);  
Handler *signal(int, Handler *);
```

Chooses one of three ways in which receipt of the signal number `sig` is to be subsequently handled. If the value of `func` is `SIG_DFL`, default handling for that signal will occur. If the value of `func` is `SIG_IGN`, the signal will be ignored. Otherwise `func` points to a function to be called when that signal occurs.

When a signal occurs, if `func` points to a function, first the equivalent of `signal(sig, SIG_DFL)` is executed. (If the value of `sig` is `SIGILL`, whether the reset to `SIG_DFL` occurs is implementation-defined (under RISC OS the reset does occur)). Next, the equivalent of `(*func)(sig)` is executed. The function may terminate by calling the `abort`, `exit` or `longjmp` function. If `func` executes a return statement and the value of `sig` was `SIGFPE` or any other

implementation-defined value corresponding to a computational exception, the behaviour is undefined. Otherwise, the program will resume execution at the point it was interrupted.

If the signal occurs other than as a result of calling the `abort` or `raise` function, the behaviour is undefined if the signal handler calls any function in the standard library other than the signal function itself or refers to any object with static storage duration other than by assigning a value to a volatile static variable of type `sig_atomic_t`. At program start-up, the equivalent of `signal(sig, SIG_IGN)` may be executed for some signals selected in an implementation-defined manner (under RISC OS this does not occur); the equivalent of `signal(sig, SIG_DFL)` is executed for all other signals defined by the implementation.

If the request can be honoured, the `signal` function returns the value of `func` for most recent call to `signal` for the specified signal `sig`. Otherwise, a value of `SIG_ERR` is returned and the integer expression `errno` is set to indicate the error.

raise

```
int raise(int /*sig*/)
```

Sends the signal `sig` to the executing program. Returns zero if successful, non-zero if unsuccessful.

stdarg.h

This file declares a type and defines three macros, for advancing through a list of arguments whose number and types are not known to the called function when it is translated. A function may be called with a variable number of arguments of differing types. Its parameter list contains one or more parameters, the rightmost of which plays a special role in the access mechanism, and will be called `parmN` in this description.

`stdio.h` is required to declare `vfprintf()` without defining `va_list`. Clearly the type `__va_list` there must keep in step.

va_list

```
char *va_list[1]
```

An array type suitable for holding information needed by the macro `va_arg` and the function `va_end`. The called function declares a variable (referred to as `ap`) having type `va_list`. The variable `ap` may be passed as an argument to another

function. `va_list` is an array type so that when an object of that type is passed as an argument it gets passed by reference, but this is not required by the ANSI specification and cannot be relied on.

va_start

The `va_start` macro will be executed before any access to the unnamed arguments. The parameter `ap` points to an object that has type `va_list`. The `va_start` macro initialises `ap` for subsequent use by `va_arg` and `va_end`. The parameter `parmN` is the identifier of the rightmost parameter in the variable parameter list in the function definition (the one just before the `,` . . .). If the parameter `parmN` is declared with the register storage class the behaviour is undefined.

Returns: no value.

va_arg

The `va_arg` macro expands to an expression that has the type and value of the next argument in the call. The parameter `ap` is the same as the `va_list ap` initialised by `va_start`. Each invocation of `va_arg` modifies `ap` so that successive arguments are returned in turn. The parameter `type` is a type name such that the type of a pointer to an object that has the specified type can be obtained simply by postfixing a `*` to `type`. If `type` disagrees with the type of the actual next argument (as promoted according to the default argument promotions), the behaviour is undefined.

Returns: The first invocation of the `va_arg` macro after that of the `va_start` macro returns the value of the argument after that specified by `parmN`. Successive invocations return the values of the remaining arguments in succession. Care is taken in `va_arg` so that illegal things like `va_arg(ap, char)` – which may seem natural but are in fact illegal – are caught. `va_arg(ap, float)` is wrong but cannot be patched up at the C macro level.

va_end

```
#define va_end(ap) ((void)* (ap) = (char *)-256)
```

The `va_end` macro facilitates a normal return from the function whose variable argument list was referenced by the expansion of `va_start` that initialised the `va_list ap`. If the `va_end` macro is not invoked before the return, the behaviour is undefined.

stddef.h

This file contains a macro for calculating the offset of fields within a structure. It also defines the pointer constant NULL and three types.

<code>ptrdiff_t</code> (here int)	the signed integral type of the result of subtracting two pointers
<code>size_t</code> (here unsigned int)	the unsigned integral type of the result of the <code>sizeof</code> operator
<code>wchar_t</code> (here int)	also in <code>stdlib.h</code> . An integral type whose range of values can represent distinct codes for all members of the largest extended character set specified among the supported locales; the null character has the code value zero and each member of the basic character set has a code value when used as the lone character in an integer character constant.
<code>size_t</code> <code>offsetof</code> (type, member)	Expands to an integral constant expression that has type <code>size_t</code> , the value of which is the offset in bytes from the beginning of a structure designated by <code>type</code> , of the member designated by <code>member</code> (if the specified member is a bit-field, the behaviour is undefined).

stdio.h

`stdio` declares two types, several macros, and many functions for performing input and output. For a discussion on Streams and Files refer to sections 4.9.2 and 4.9.3 in the ANSI standard or to one of the other references given in the *Introduction* to this Guide.

<code>fpos_t</code>	<code>fpos_t</code> is an object capable of recording all information needed to specify uniquely every position within a file.
<code>FILE</code>	is an object capable of recording all information needed to control a stream, such as its file position indicator, a pointer to its associated buffer, an error indicator that records whether a read/write error has occurred and an end-of-file indicator that records whether the end-of-file has been reached. The objects contained in the <code>#ifdef __system_io</code> clause are for system use only, and cannot be relied on between releases of C.

remove

```
int remove(const char * filename)
```

Causes the file whose name is the string pointed to by *filename* to be removed. Subsequent attempts to open the file will fail, unless it is created anew. If the file is open, the behaviour of the `remove` function is implementation-defined (under RISC OS the operation fails).

Returns: zero if the operation succeeds, non-zero if it fails.

rename

```
int rename(const char * old, const char * new)
```

Causes the file whose name is the string pointed to by *old* to be henceforth known by the name given by the string pointed to by *new*. The file named *old* is effectively removed. If a file named by the string pointed to by *new* exists prior to the call of the `rename` function, the behaviour is implementation-defined (under RISC OS, the operation fails).

Returns: zero if the operation succeeds, non-zero if it fails, in which case if the file existed previously it is still known by its original name.

tmpfile

```
FILE *tmpfile(void)
```

Creates a temporary binary file that will be automatically removed when it is closed or at program termination. The file is opened for update.

Returns: a pointer to the stream of the file that it created. If the file cannot be created, a null pointer is returned.

tmpnam

```
char *tmpnam(char * s)
```

Generates a string that is not the same as the name of an existing file. The `tmpnam` function generates a different string each time it is called, up to `TMP_MAX` times. If it is called more than `TMP_MAX` times, the behaviour is implementation-defined (under RISC OS the algorithm for the name generation works just as well after `tmpnam` has been called more than `TMP_MAX` times as before; a name clash is impossible in any single half year period).

Returns: If the argument is a null pointer, the `tmpnam` function leaves its result in an internal static object and returns a pointer to that object. Subsequent calls to the `tmpnam` function may modify the same object. If the argument is not a null

pointer, it is assumed to point to an array of at least `L_tmpnam` characters; the `tmpnam` function writes its result in that array and returns the argument as its value.

fclose

```
int fclose(FILE * stream)
```

Causes the stream pointed to by *stream* to be flushed and the associated file to be closed. Any unwritten buffered data for the stream are delivered to the host environment to be written to the file; any unread buffered data are discarded. The stream is disassociated from the file. If the associated buffer was automatically allocated, it is deallocated.

Returns: zero if the stream was successfully closed, or EOF if any errors were detected or if the stream was already closed.

fflush

```
int fflush(FILE * stream)
```

If the stream points to an output or update stream in which the most recent operation was output, the `fflush` function causes any unwritten data for that stream to be delivered to the host environment to be written to the file. If the stream points to an input or update stream, the `fflush` function undoes the effect of any preceding `ungetc` operation on the stream.

Returns: EOF if a write error occurs.

fopen

```
FILE *fopen(const char * filename, const char * mode)
```

Opens the file whose name is the string pointed to by *filename*, and associates a stream with it. The argument *mode* points to a string beginning with one of the following sequences:

<code>r</code>	open text file for reading
<code>w</code>	create text file for writing, or truncate to zero length
<code>a</code>	append; open text file or create for writing at eof
<code>rb</code>	open binary file for reading
<code>wb</code>	create binary file for writing, or truncate to zero length
<code>ab</code>	append; open binary file or create for writing at eof
<code>r+</code>	open text file for update (reading and writing)
<code>w+</code>	create text file for update, or truncate to zero length
<code>a+</code>	append; open text file or create for update, writing at eof
<code>r+b</code> or <code>rb+</code>	open binary file for update (reading and writing)

w+b or wb+	create binary file for update, or truncate to zero length
a+b or ab+	append; open binary file or create for update, writing at eof

- Opening a file with read mode (*r* as the first character in the *mode* argument) fails if the file does not exist or cannot be read.
- Opening a file with append mode (*a* as the first character in the *mode* argument) causes all subsequent writes to be forced to the current end of file, regardless of intervening calls to the *fseek* function.
- In some implementations, opening a binary file with append mode (*b* as the second or third character in the *mode* argument) may initially position the file position indicator beyond the last data written, because of null padding (but not under RISC OS).
- When a file is opened with update mode (*+* as the second or third character in the *mode* argument), both input and output may be performed on the associated stream. However, output may not be directly followed by input without an intervening call to the *fflush* function or to a file positioning function (*fseek*, *fsetpos*, or *rewind*), nor may input be directly followed by output without an intervening call to the *fflush* function or to a file positioning function, unless the input operation encounters end-of-file.
- Opening a file with update mode may open or create a binary stream in some implementations (but not under RISC OS). When opened, a stream is fully buffered if and only if it does not refer to an interactive device. The error and end-of-file indicators for the stream are cleared.

Returns: a pointer to the object controlling the stream. If the open operation fails, *fopen* returns a null pointer.

freopen

```
FILE *freopen(const char * filename, const char * mode,
             FILE * stream)
```

Opens the file whose name is the string pointed to by *filename* and associates the stream pointed to by *stream* with it. The *mode* argument is used just as in the *fopen* function. The *freopen* function first attempts to close any file that is associated with the specified stream. Failure to close the file successfully is ignored. The error and end-of-file indicators for the stream are cleared.

Returns: a null pointer if the operation fails. Otherwise, *freopen* returns the value of the stream.

setbuf

```
void setbuf(FILE * stream, char * buf)
```

Except that it returns no value, the `setbuf` function is equivalent to the `setvbuf` function invoked with the values `_IOFBF` for `mode` and `BUFSIZ` for `size`, or if `buf` is a null pointer, with the value `_IONBF` for `mode`.

Returns: no value.

setvbuf

```
int setvbuf(FILE * stream, char * buf, int mode, size_t
            size)
```

This may be used after the stream pointed to by `stream` has been associated with an open file but before it is read or written. The argument `mode` determines how `stream` will be buffered, as follows:

- `_IOFBF` causes input/output to be fully buffered.
- `_IOLBF` causes output to be line buffered (the buffer will be flushed when a newline character is written, when the buffer is full, or when interactive input is requested).
- `_IONBF` causes input/output to be completely unbuffered.

If `buf` is not the null pointer, the array it points to may be used instead of an automatically allocated buffer (the buffer must have a lifetime at least as great as the open stream, so the stream should be closed before a buffer that has automatic storage duration is deallocated upon block exit). The argument `size` specifies the size of the array. The contents of the array at any time are indeterminate.

Returns: zero on success, or non-zero if an invalid value is given for `mode` or `size`, or if the request cannot be honoured.

fprintf

```
int fprintf(FILE * stream, const char * format, ...)
```

writes output to the stream pointed to by `stream`, under control of the string pointed to by `format` that specifies how subsequent arguments are converted for output. If there are insufficient arguments for the format, the behaviour is undefined. If the format is exhausted while arguments remain, the excess arguments are evaluated but otherwise ignored. The `fprintf` function returns when the end of the format string is reached. The format must be a multibyte character sequence, beginning and ending in its initial shift state (in all locales supported under RISC OS this is the same as a plain character string). The format is composed of zero or more directives: ordinary multibyte characters (not `%`), which are copied unchanged to the output stream; and conversion specifiers, each of which results in fetching zero or more subsequent arguments. Each conversion specification is introduced by the character `%`. For a complete description of the

available conversion specifiers refer to section 4.9.6.1 in the ANSI standard or to one of the other references in the *Introduction* to this Guide. The minimum value for the maximum number of characters that can be produced by any single conversion is at least 509.

A brief and incomplete description of conversion specifications is:

[flags][field width][.precision]specifier-char

flags is most commonly -, indicating left justification of the output item within the field. If omitted, the item will be right justified.

field width is the minimum width of field to use. If the formatted item is longer, a bigger field will be used; otherwise, the item will be right (left) justified in the field.

precision is the minimum number of digits to print for a d, i, o, u, x or X conversion, the number of digits to appear after the decimal digit for e, E and f conversions, the maximum number of significant digits for g and G conversions, or the maximum number of characters to be written from strings in an s conversion.

Either of both of *field width* and *precision* may be *, indicating that the value is an argument to printf.

The *specifier chars* are:

d, i	int printed as signed decimal
o, u, x, X	unsigned int value printed as unsigned octal, decimal or hexadecimal
f	double value printed in the style [-]ddd.ddd
e, E	double value printed in the style [-]d.ddd...e dd
g, G	double printed in f or e format, whichever is more appropriate
c	int value printed as unsigned char
s	char * value printed as a string of characters
p	void * argument printed as a hexadecimal address
%	write a literal %

Returns: the number of characters transmitted, or a negative value if an output error occurred.

printf

```
int printf(const char * format, ...)
```

Equivalent to fprintf with the argument stdout interposed before the arguments to printf.

Returns: the number of characters transmitted, or a negative value if an output error occurred.

sprintf

```
int sprintf(char * s, const char * format, ...)
```

Equivalent to `fprintf`, except that the argument *s* specifies an array into which the generated output is to be written, rather than to a stream. A null character is written at the end of the characters written; it is not counted as part of the returned sum.

Returns: the number of characters written to the array, not counting the terminating null character.

fscanf

```
int fscanf(FILE * stream, const char * format, ...)
```

Reads input from the stream pointed to by *stream*, under control of the string pointed to by *format* that specifies the admissible input sequences and how they are to be converted for assignment, using subsequent arguments as pointers to the objects to receive the converted input. If there are insufficient arguments for the format, the behaviour is undefined. If the format is exhausted while arguments remain, the excess arguments are evaluated but otherwise ignored. The format is composed of zero or more directives, one or more white-space characters, an ordinary character (not %), or a conversion specification. Each conversion specification is introduced by the character %. For a description of the available conversion specifiers refer to section 4.9.6.2 in the ANSI standard, or to any of the references listed in the chapter entitled *Introduction* on page 1. A brief list is given above, under the entry for `fprintf`.

If end-of-file is encountered during input, conversion is terminated. If end-of-file occurs before any characters matching the current directive have been read (other than leading white space, where permitted), execution of the current directive terminates with an input failure; otherwise, unless execution of the current directive is terminated with a matching failure, execution of the following directive (if any) is terminated with an input failure.

If conversions terminate on a conflicting input character, the offending input character is left unread in the input stream. Trailing white space (including newline characters) is left unread unless matched by a directive. The success of literal matches and suppressed assignments is not directly determinable other than via the `%n` directive.

Returns: the value of the macro EOF if an input failure occurs before any conversion. Otherwise, the `fscanf` function returns the number of input items assigned, which can be fewer than provided for, or even zero, in the event of an early conflict between an input character and the format.

scanf

```
int scanf(const char * format, ...)
```

Equivalent to `fscanf` with the argument `stdin` interposed before the arguments to `scanf`.

Returns: the value of the macro EOF if an input failure occurs before any conversion. Otherwise, the `scanf` function returns the number of input items assigned, which can be fewer than provided for, or even zero, in the event of an early matching failure.

sscanf

```
int sscanf(const char * s, const char * format, ...)
```

Equivalent to `fscanf` except that the argument `s` specifies a string from which the input is to be obtained, rather than from a stream. Reaching the end of the string is equivalent to encountering end-of-file for the `fscanf` function.

Returns: the value of the macro EOF if an input failure occurs before any conversion. Otherwise, the `scanf` function returns the number of input items assigned, which can be fewer than provided for, or even zero, in the event of an early matching failure.

vprintf

```
int vprintf(const char * format, va_list arg)
```

Equivalent to `printf`, with the variable argument list replaced by `arg`, which has been initialised by the `va_start` macro (and possibly subsequent `va_arg` calls). The `vprintf` function does not invoke the `va_end` function.

Returns: the number of characters transmitted, or a negative value if an output error occurred.

vfprintf

```
int vfprintf(FILE * stream, const char * format, va_list arg)
```

Equivalent to `fprintf`, with the variable argument list replaced by `arg`, which has been initialised by the `va_start` macro (and possibly subsequent `va_arg` calls). The `vfprintf` function does not invoke the `va_end` function.

Returns: the number of characters transmitted, or a negative value if an output error occurred.

vsprintf

```
int vsprintf(char * s, const char * format, va_list arg)
```

Equivalent to `sprintf`, with the variable argument list replaced by `arg`, which has been initialised by the `va_start` macro (and possibly subsequent `va_arg` calls). The `vsprintf` function does not invoke the `va_end` function.

Returns: the number of characters written in the array, not counting the terminating null character.

fgetc

```
int fgetc(FILE * stream)
```

Obtains the next character (if present) as an unsigned char converted to an int, from the input stream pointed to by `stream`, and advances the associated file position indicator (if defined).

Returns: the next character from the input stream pointed to by `stream`. If the stream is at end-of-file, the end-of-file indicator is set and `fgetc` returns EOF. If a read error occurs, the error indicator is set and `fgetc` returns EOF.

fgets

```
char *fgets(char * s, int n, FILE * stream)
```

Reads at most one less than the number of characters specified by `n` from the stream pointed to by `stream` into the array pointed to by `s`. No additional characters are read after a newline character (which is retained) or after end-of-file. A null character is written immediately after the last character read into the array.

Returns: `s` if successful. If end-of-file is encountered and no characters have been read into the array, the contents of the array remain unchanged and a null pointer is returned. If a read error occurs during the operation, the array contents are indeterminate and a null pointer is returned.

fputc

```
int fputc(int c, FILE * stream)
```

Writes the character specified by `c` (converted to an unsigned char) to the output stream pointed to by `stream`, at the position indicated by the associated file position indicator (if defined), and advances the indicator appropriately. If the file cannot support positioning requests, or if the stream was opened with append mode, the character is appended to the output stream.

Returns: the character written. If a write error occurs, the error indicator is set and `fputc` returns EOF.

fputs

```
int fputs(const char * s, FILE * stream)
```

Writes the string pointed to by `s` to the stream pointed to by `stream`. The terminating null character is not written.

Returns: EOF if a write error occurs; otherwise it returns a non-negative value.

getc

```
int getc(FILE * stream)
```

Equivalent to `fgetc` except that it may be (and is under RISC OS) implemented as a macro. `stream` may be evaluated more than once, so the argument should never be an expression with side effects.

Returns: the next character from the input stream pointed to by `stream`. If the stream is at end-of-file, the end-of-file indicator is set and `getc` returns EOF. If a read error occurs, the error indicator is set and `getc` returns EOF.

getchar

```
int getchar(void)
```

Equivalent to `getc` with the argument `stdin`.

Returns: the next character from the input stream pointed to by `stdin`. If the stream is at end-of-file, the end-of-file indicator is set and `getchar` returns EOF. If a read error occurs, the error indicator is set and `getchar` returns EOF.

gets

```
char *gets(char * s)
```

Reads characters from the input stream pointed to by `stdin` into the array pointed to by `s`, until end-of-file is encountered or a newline character is read. Any newline character is discarded, and a null character is written immediately after the last character read into the array.

Returns: `s` if successful. If end-of-file is encountered and no characters have been read into the array, the contents of the array remain unchanged and a null pointer is returned. If a read error occurs during the operation, the array contents are indeterminate and a null pointer is returned.

putc

```
int putc(int c, FILE * stream)
```

Equivalent to `fputc` except that it may be (and is under RISC OS) implemented as a macro. `stream` may be evaluated more than once, so the argument should never be an expression with side effects.

Returns: the character written. If a write error occurs, the error indicator is set and `putc` returns EOF.

putchar

```
int putchar(int c)
```

Equivalent to `putc` with the second argument `stdout`.

Returns: the character written. If a write error occurs, the error indicator is set and `putc` returns EOF.

puts

```
int puts(const char * s)
```

Writes the string pointed to by `s` to the stream pointed to by `stdout`, and appends a newline character to the output. The terminating null character is not written.

Returns: EOF if a write error occurs; otherwise it returns a non-negative value.

ungetc

```
int ungetc(int c, FILE * stream)
```

Pushes the character specified by `c` (converted to an unsigned char) back onto the input stream pointed to by `stream`. The character will be returned by the next read on that stream. An intervening call to the `fflush` function or to a file positioning function (`fseek`, `fsetpos`, `rewind`) discards any pushed-back characters. The external storage corresponding to the stream is unchanged. One character pushback is guaranteed. If the `ungetc` function is called too many times on the same stream without an intervening read or file positioning operation on that stream, the operation may fail. If the value of `c` equals that of the macro `EOF`, the operation fails and the input stream is unchanged.

A successful call to the `ungetc` function clears the end-of-file indicator. The value of the file position indicator after reading or discarding all pushed-back characters will be the same as it was before the characters were pushed back. For a text stream, the value of the file position indicator after a successful call to the `ungetc` function is unspecified until all pushed-back characters are read or discarded. For

a binary stream, the file position indicator is decremented by each successful call to the `ungetc` function; if its value was zero before a call, it is indeterminate after the call.

Returns: the character pushed back after conversion, or EOF if the operation fails.

fread

```
size_t fread(void * ptr, size_t size,
             size_t nmemb, FILE * stream)
```

Reads into the array pointed to by `ptr`, up to `nmemb` members whose size is specified by `size`, from the stream pointed to by `stream`. The file position indicator (if defined) is advanced by the number of characters successfully read. If an error occurs, the resulting value of the file position indicator is indeterminate. If a partial member is read, its value is indeterminate. The `ferror` or `feof` function shall be used to distinguish between a read error and end-of-file.

Returns: the number of members successfully read, which may be less than `nmemb` if a read error or end-of-file is encountered. If `size` or `nmemb` is zero, `fread` returns zero and the contents of the array and the state of the stream remain unchanged.

fwrite

```
size_t fwrite(const void * ptr,
             size_t size, size_t nmemb, FILE * stream)
```

Writes, from the array pointed to by `ptr` up to `nmemb` members whose size is specified by `size`, to the stream pointed to by `stream`. The file position indicator (if defined) is advanced by the number of characters successfully written. If an error occurs, the resulting value of the file position indicator is indeterminate.

Returns: the number of members successfully written, which will be less than `nmemb` only if a write error is encountered.

fgetpos

```
int fgetpos(FILE * stream, fpos_t * pos)
```

Stores the current value of the file position indicator for the stream pointed to by `stream` in the object pointed to by `pos`. The value stored contains unspecified information usable by the `fsetpos` function for repositioning the stream to its position at the time of the call to the `fgetpos` function.

Returns: zero, if successful. Otherwise non-zero is returned and the integer expression `errno` is set to an implementation-defined non-zero value (under RISC OS `fgetpos` cannot fail).

fseek

```
int fseek(FILE * stream, long int offset, int whence)
```

Sets the file position indicator for the stream pointed to by *stream*. For a binary stream, the new position is at the signed number of characters specified by *offset* away from the point specified by *whence*. The specified point is the beginning of the file for `SEEK_SET`, the current position in the file for `SEEK_CUR`, or end-of-file for `SEEK_END`. A binary stream need not meaningfully support *fseek* calls with a *whence* value of `SEEK_END`, though the Acorn implementation does. For a text stream, *offset* is either zero or a value returned by an earlier call to the `ftell` function on the same stream; *whence* is then `SEEK_SET`. The Acorn implementation also allows a text stream to be positioned in exactly the same manner as a binary stream, but this is not portable. The *fseek* function clears the end-of-file indicator and undoes any effects of the `ungetc` function on the same stream. After an *fseek* call, the next operation on an update stream may be either input or output.

Returns: non-zero only for a request that cannot be satisfied.

fsetpos

```
int fsetpos(FILE * stream, const fpos_t * pos)
```

Sets the file position indicator for the stream pointed to by *stream* according to the value of the object pointed to by *pos*, which is a value returned by an earlier call to the `fgetpos` function on the same stream. The *fsetpos* function clears the end-of-file indicator and undoes any effects of the `ungetc` function on the same stream. After an *fsetpos* call, the next operation on an update stream may be either input or output.

Returns: zero, if successful. Otherwise non-zero is returned and the integer expression `errno` is set to an implementation-defined non-zero value (under RISC OS the value is that of `EDOM` in `math.h`).

ftell

```
long int ftell(FILE * stream)
```

Obtains the current value of the file position indicator for the stream pointed to by *stream*. For a binary stream, the value is the number of characters from the beginning of the file. For a text stream, the file position indicator contains unspecified information, usable by the *fseek* function for returning the file position indicator to its position at the time of the `ftell` call; the difference between two such return values is not necessarily a meaningful measure of the

number of characters written or read. However, for the Acorn implementation, the value returned is merely the byte offset into the file, whether the stream is text or binary.

Returns: if successful, the current value of the file position indicator. On failure, the `ftell` function returns `-1L` and sets the integer expression `errno` to an implementation-defined non-zero value (under RISC OS `ftell` cannot fail).

rewind

```
void rewind(FILE * stream)
```

Sets the file position indicator for the stream pointed to by *stream* to the beginning of the file. It is equivalent to `(void)fseek(stream, 0L, SEEK_SET)` except that the error indicator for the stream is also cleared.

Returns: no value.

clearerr

```
void clearerr(FILE * stream)
```

Clears the end-of-file and error indicators for the stream pointed to by *stream*. These indicators are cleared only when the file is opened or by an explicit call to the `clearerr` function or to the `rewind` function.

Returns: no value.

feof

```
int feof(FILE * stream)
```

Tests the end-of-file indicator for the stream pointed to by *stream*.

Returns: non-zero if the end-of-file indicator is set for *stream*.

ferror

```
int ferror(FILE * stream)
```

Tests the error indicator for the stream pointed to by *stream*.

Returns: non-zero if the error indicator is set for *stream*.

perror

```
void perror(const char * s)
```

Maps the error number in the integer expression `errno` to an error message. It writes a sequence of characters to the standard error stream thus: first (if `s` is not a null pointer and the character pointed to by `s` is not the null character), the string pointed to by `s` followed by a colon and a space; then an appropriate error message string followed by a newline character. The contents of the error message strings are the same as those returned by the `strerror` function with argument `errno`, which are implementation-defined.

Returns: no value.

stdlib.h

`stdlib.h` declares four types, several general purpose functions, and defines several macros.

atof

```
double atof(const char * nptr)
```

Converts the initial part of the string pointed to by `nptr` to double representation.

Returns: the converted value.

atoi

```
int atoi(const char * nptr)
```

Converts the initial part of the string pointed to by `nptr` to int representation.

Returns: the converted value.

atol

```
long int atol(const char * nptr)
```

Converts the initial part of the string pointed to by `nptr` to long int representation.

Returns: the converted value.

strtod

```
double strtod(const char * nptr, char ** endptr)
```

Converts the initial part of the string pointed to by `nptr` to double representation. First it decomposes the input string into three parts: an initial, possibly empty, sequence of white-space characters (as specified by the `isspace` function), a subject sequence resembling a floating point constant, and a final string of one or

more unrecognised characters, including the terminating null character of the input string. It then attempts to convert the subject sequence to a floating point number, and returns the result. A pointer to the final string is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

Returns: the converted value if any. If no conversion could be performed, zero is returned. If the correct value is outside the range of representable values, plus or minus `HUGE_VAL` is returned (according to the sign of the value), and the value of the macro `ERANGE` is stored in `errno`. If the correct value would cause underflow, zero is returned and the value of the macro `ERANGE` is stored in `errno`.

strtoul

```
long int strtoul(const char * nptr, char **endptr, int
                base)
```

Converts the initial part of the string pointed to by *nptr* to long int representation. First it decomposes the input string into three parts: an initial, possibly empty, sequence of white-space characters (as specified by the `isspace` function), a subject sequence resembling an integer represented in some radix determined by the value of *base*, and a final string of one or more unrecognised characters, including the terminating null character of the input string.

It then attempts to convert the subject sequence to an integer, and returns the result. If the value of *base* is 0, the expected form of the subject sequence is that of an integer constant (described precisely in the ANSI standard, section 3.1.3.2), optionally preceded by a + or - sign, but not including an integer suffix. If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence of letters and digits representing an integer with the radix specified by *base*, optionally preceded by a plus or minus sign, but not including an integer suffix. The letters from a (or A) through z (or Z) are ascribed the values 10 to 35; only letters whose ascribed values are less than that of the base are permitted. If the value of *base* is 16, the characters 0x or 0X may optionally precede the sequence of letters and digits following the sign if present. A pointer to the final string is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

Returns: the converted value if any. If no conversion could be performed, zero is returned. If the correct value is outside the range of representable values, `LONG_MAX` or `LONG_MIN` is returned (according to the sign of the value), and the value of the macro `ERANGE` is stored in `errno`.

strtoul

```
unsigned long int strtoul(const char * nptr, char **
                        endptr, int base)
```

Converts the initial part of the string pointed to by `nptr` to unsigned long int representation. First it decomposes the input string into three parts: an initial, possibly empty, sequence of white space characters (as determined by the `isspace` function), a subject sequence resembling an unsigned integer represented in some radix determined by the value of `base`, and a final string of one or more unrecognised characters, including the terminating null character of the input string.

It then attempts to convert the subject sequence to an unsigned integer, and returns the result. If the value of `base` is zero, the expected form of the subject sequence is that of an integer constant (described precisely in the ANSI Draft, section 3.1.3.2), optionally preceded by a + or - sign, but not including an integer suffix. If the value of `base` is between 2 and 36, the expected form of the subject sequence is a sequence of letters and digits representing an integer with the radix specified by `base`, optionally preceded by a + or - sign, but not including an integer suffix. The letters from a (or A) through z (or Z) stand for the values 10 to 35; only letters whose ascribed values are less than that of the base are permitted. If the value of `base` is 16, the characters 0x or 0X may optionally precede the sequence of letters and digits following the sign, if present. A pointer to the final string is stored in the object pointed to by `endptr`, provided that `endptr` is not a null pointer.

Returns: the converted value if any. If no conversion could be performed, zero is returned. If the correct value is outside the range of representable values, `ULONG_MAX` is returned, and the value of the * macro `ERANGE` is stored in `errno`.

rand

```
int rand(void)
```

Computes a sequence of pseudo-random integers in the range 0 to `RAND_MAX`, where `RAND_MAX = 0x7fffffff`.

Returns: a pseudo-random integer.

srand

```
void srand(unsigned int seed)
```

Uses its argument as a seed for a new sequence of pseudo-random numbers to be returned by subsequent calls to `rand`. If `srand` is then called with the same seed value, the sequence of pseudo-random numbers will be repeated. If `rand` is called before any calls to `srand` have been made, the same sequence is generated as when `srand` is first called with a seed value of 1.

calloc

```
void *calloc(size_t nmemb, size_t size)
```

Allocates space for an array of *nmemb* objects, each of whose size is *size*. The space is initialised to all bits zero.

Returns: either a null pointer or a pointer to the allocated space.

free

```
void free(void * ptr)
```

Causes the space pointed to by *ptr* to be deallocated (made available for further allocation). If *ptr* is a null pointer, no action occurs. Otherwise, if *ptr* does not match a pointer earlier returned by `calloc`, `malloc` or `realloc` or if the space has been deallocated by a call to `free` or `realloc`, the behaviour is undefined.

malloc

```
void *malloc(size_t size)
```

Allocates space for an object whose size is specified by *size* and whose value is indeterminate.

Returns: either a null pointer or a pointer to the allocated space.

realloc

```
void *realloc(void * ptr, size_t size)
```

Changes the size of the object pointed to by *ptr* to the size specified by *size*. The contents of the object is unchanged up to the lesser of the new and old sizes. If the new size is larger, the value of the newly allocated portion of the object is indeterminate. If *ptr* is a null pointer, the `realloc` function behaves like a call to `malloc` for the specified size. Otherwise, if *ptr* does not match a pointer earlier returned by `calloc`, `malloc` or `realloc`, or if the space has been deallocated by a call to `free` or `realloc`, the behaviour is undefined. If the space cannot be allocated, the object pointed to by *ptr* is unchanged. If *size* is zero and *ptr* is not a null pointer, the object it points to is freed.

Returns: either a null pointer or a pointer to the possibly moved allocated space.

abort

```
void abort(void)
```

Causes abnormal program termination to occur, unless the signal `SIGABRT` is being caught and the signal handler does not return. Whether open output streams are flushed or open streams are closed or temporary files removed is implementation-defined (under RISC OS all these occur). An implementation-defined form of the status 'unsuccessful termination' (1 under RISC OS) is returned to the host environment by means of a call to `raise(SIGABRT)`.

atexit

```
int atexit(void (* func)(void))
```

Registers the function pointed to by `func`, to be called without its arguments at normal program termination. It is possible to register at least 32 functions.

Returns: zero if the registration succeeds, non-zero if it fails.

exit

```
void exit(int status)
```

Causes normal program termination to occur. If more than one call to the `exit` function is executed by a program (for example, by a function registered with `atexit`), the behaviour is undefined. First, all functions registered by the `atexit` function are called, in the reverse order of their registration. Next, all open output streams are flushed, all open streams are closed, and all files created by the `tmpfile` function are removed. Finally, control is returned to the host environment. If the value of `status` is zero or `EXIT_SUCCESS`, an implementation-defined form of the status 'successful termination' (0 under RISC OS) is returned. If the value of `status` is `EXIT_FAILURE`, an implementation-defined form of the status 'unsuccessful termination' (1 under RISC OS) is returned. Otherwise the status returned is implementation-defined (the value of `status` is returned under RISC OS).

getenv

```
char *getenv(const char * name)
```

Searches the environment list, provided by the host environment, for a string that matches the string pointed to by `name`. The set of environment names and the method for altering the environment list are implementation-defined.

Returns: a pointer to a string associated with the matched list member. The array pointed to is not modified by the program, but may be overwritten by a subsequent call to the `getenv` function. If the specified name cannot be found, a null pointer is returned.

system

```
int system(const char * string)
```

Passes the string pointed to by *string* to the host environment to be executed by a command processor in an implementation-defined manner. A null pointer may be used for *string*, to inquire whether a command processor exists. Under RISC OS, care must be taken, when executing a command, that the command does not overwrite the calling program. To control this, the string `chain:` or `call:` may immediately precede the actual command. The effect of `call:` is the same as if `call:` were not present. When a command is called, the caller is first moved to a safe place in application workspace. When the callee terminates, the caller is restored. This requires enough memory to hold caller and callee simultaneously. When a command is chained, the caller may be overwritten. If the caller is not overwritten, the caller exits when the caller terminates. Thus a transfer of control is effected and memory requirements are minimised.

Returns: If the argument is a null pointer, the `system` function returns non-zero only if a command processor is available. If the argument is not a null pointer, it returns an implementation-defined value (under RISC OS 0 is returned for success and -2 for failure to invoke the command; any other value is the return code from the executed command).

bsearch

```
void *bsearch(const void *key, const void * base,  
             size_t nmemb, size_t size, int (* compar)  
             (const void *, const void *))
```

Searches an array of *nmemb* objects, the initial member of which is pointed to by *base*, for a member that matches the object pointed to by *key*. The size of each member of the array is specified by *size*. The contents of the array must be in ascending sorted order according to a comparison function pointed to by *compar*, which is called with two arguments that point to the key object and to an array member, in that order. The function returns an integer less than, equal to, or greater than zero if the key object is considered, respectively, to be less than, to match, or to be greater than the array member.

Returns: a pointer to a matching member of the array, or a null pointer if no match is found. If two members compare as equal, which member is matched is unspecified.

qsort

```
void qsort(void * base, size_t nmemb, size_t size,  
          int (* compar)(const void *, const void *))
```

Sorts an array of *nmemb* objects, the initial member of which is pointed to by *base*. The size of each object is specified by *size*. The contents of the array are sorted in ascending order according to a comparison function pointed to by *compar*, which is called with two arguments that point to the objects being compared. The function returns an integer less than, equal to, or greater than zero if the first argument is considered to be respectively less than, equal to, or greater than the second. If two members compare as equal, their order in the sorted array is unspecified.

abs

```
int abs(int j)
```

Computes the absolute value of an integer *j*. If the result cannot be represented, the behaviour is undefined.

Returns: the absolute value.

div

```
div_t div(int numer, int denom)
```

Computes the quotient and remainder of the division of the numerator *numer* by the denominator *denom*. If the division is inexact, the resulting quotient is the integer of lesser magnitude that is the nearest to the algebraic quotient. If the result cannot be represented, the behaviour is undefined; otherwise, $\text{quot} * \text{denom} + \text{rem}$ equals *numer*.

Returns: a structure of type `div_t`, comprising both the quotient and the remainder. The structure contains the following members: `int quot`; `int rem`. You may not rely on their order.

labs

```
long int labs(long int j)
```

Computes the absolute value of an long integer *j*. If the result cannot be represented, the behaviour is undefined.

Returns: the absolute value.

ldiv

```
ldiv_t ldiv(long int numer, long int denom)
```

Computes the quotient and remainder of the division of the numerator *numer* by the denominator *denom*. If the division is inexact, the sign of the resulting quotient is that of the algebraic quotient, and the magnitude of the resulting

quotient is the largest integer less than the magnitude of the algebraic quotient. If the result cannot be represented, the behaviour is undefined; otherwise, $quot * denom + rem$ equals $numer$.

Returns: a structure of type `ldiv_t`, comprising both the quotient and the remainder. The structure contains the following members: `long int quot`; `long int rem`. You may not rely on their order.

Multibyte character functions

The behaviour of the multibyte character functions is affected by the `LC_CTYPE` category of the current locale. For a state-dependent encoding, each function is placed into its initial state by a call for which its character pointer argument, *s*, is a null pointer. Subsequent calls with *s* as other than a null pointer cause the internal state of the function to be altered as necessary. A call with *s* as a null pointer causes these functions to return a non-zero value if encodings have state dependency, and a zero otherwise. After the `LC_CTYPE` category is changed, the shift state of these functions is indeterminate.

mblen

```
int mblen(const char * s, size_t n)
```

If *s* is not a null pointer, the `mblen` function determines the number of bytes comprising the multibyte character pointed to by *s*. Except that the shift state of the `mbtowc` function is not affected, it is equivalent to `mbtowc((wchar_t *)0, s, n)`.

Returns: If *s* is a null pointer, the `mblen` function returns a non-zero or zero value, if multibyte character encodings, respectively do or do not have state-dependent encodings. If *s* is not a null pointer, the `mblen` function either returns a 0 (if *s* points to a null character), or returns the number of bytes that comprise the multibyte character (if the next *n* of fewer bytes form a valid multibyte character), or returns -1 (if they do not form a valid multibyte character).

mbtowc

```
int mbtowc(wchar_t * pwc, const char * s, size_t n)
```

If *s* is not a null pointer, the `mbtowc` function determines the number of bytes that comprise the multibyte character pointed to by *s*. It then determines the code for value of type `wchar_t` that corresponds to that multibyte character. (The value of the code corresponding to the null character is zero). If the multibyte character is valid and *pwc* is not a null pointer, the `mbtowc` function stores the code in the object pointed to by *pwc*. At most *n* bytes of the array pointed to by *s* will be examined.

Returns: If *s* is a null pointer, the `mbtowc` function returns a non-zero or zero value, if multibyte character encodings, respectively do or do not have state-dependent encodings. If *s* is not a null pointer, the `mbtowc` function either returns a 0 (if *s* points to a null character), or returns the number of bytes that comprise the converted multibyte character (if the next *n* of fewer bytes form a valid multibyte character), or returns -1 (if they do not form a valid multibyte character).

wctomb

```
int wctomb(char * s, wchar_t wchar)
```

Determines the number of bytes need to represent the multibyte character corresponding to the code whose value is *wchar* (including any change in shift state). It stores the multibyte character representation in the array object pointed to by *s* (if *s* is not a null pointer). At most `MB_CUR_MAX` characters are stored. If the value of *wchar* is zero, the `wctomb` function is left in the initial shift state).

Returns: If *s* is a null pointer, the `wctomb` function returns a non-zero or zero value, if multibyte character encodings, respectively do or do not have state-dependent encodings. If *s* is not a null pointer, the `wctomb` function returns a -1 if the value of *wchar* does not correspond to a valid multibyte character, or returns the number of bytes that comprise the multibyte character corresponding to the value of *wchar*.

Multibyte string functions

The behaviour of the multibyte string functions is affected by the `LC_CTYPE` category of the current locale.

mbstowcs

```
size_t mbstowcs(wchar_t * pwcs, const char * s, size_t n)
```

Converts a sequence of multibyte characters that begins in the initial shift state from the array pointed to by *s* into a sequence of corresponding codes and stores not more than *n* codes into the array pointed to by *pwcs*. No multibyte character that follow a null character (which is converted into a code with value zero) will be examined or converted. Each multibyte character is converted as if by a call to the `mbtowc` function. If an invalid multibyte character is found, `mbstowcs` returns `(size_t)-1`. Otherwise, the `mbstowcs` function returns the number of array elements modified, not including a terminating zero code, if any.

wcstombs

```
size_t wcstombs(char * s, const wchar_t * pwcs, size_t n)
```

Converts a sequence of codes that correspond to multibyte characters from the array pointed to by *pwcs* into a sequence of multibyte characters that begins in the initial shift state and stores these multibyte characters into the array pointed to by *s*, stopping if a multibyte character would exceed the limit of *n* total bytes or if a null character is stored. Each code is converted as if by a call to the `wctomb` function, except that the shift state of the `wctomb` function is not affected. If a code is encountered which does not correspond to any valid multibyte character, the `wcstombs` function returns `(size_t) - 1`. Otherwise, the `wcstombs` function returns the number of bytes modified, not including a terminating null character, if any.

string.h

`string.h` declares one type and several functions, and defines one macro useful for manipulating character arrays and other objects treated as character arrays. Various methods are used for determining the lengths of the arrays, but in all cases a `char *` or `void *` argument points to the initial (lowest addresses) character of the array. If an array is written beyond the end of an object, the behaviour is undefined.

memcpy

```
void *memcpy(void * s1, const void * s2, size_t n)
```

Copies *n* characters from the object pointed to by *s2* into the object pointed to by *s1*. If copying takes place between objects that overlap, the behaviour is undefined.

Returns: the value of *s1*.

memmove

```
void *memmove(void * s1, const void * s2, size_t n)
```

Copies *n* characters from the object pointed to by *s2* into the object pointed to by *s1*. Copying takes place as if the *n* characters from the object pointed to by *s2* are first copied into a temporary array of *n* characters that does not overlap the objects pointed to by *s1* and *s2*, and then the *n* characters from the temporary array are copied into the object pointed to by *s1*.

Returns: the value of *s1*.

strcpy

```
char *strcpy(char * s1, const char * s2)
```

Copies the string pointed to by *s2* (including the terminating null character) into the array pointed to by *s1*. If copying takes place between objects that overlap, the behaviour is undefined.

Returns: the value of *s1*.

strncpy

```
char *strncpy(char * s1, const char * s2, size_t n)
```

Copies not more than *n* characters (characters that follow a null character are not copied) from the array pointed to by *s2* into the array pointed to by *s1*. If copying takes place between objects that overlap, the behaviour is undefined. If terminating nul has not been copied in chars, no term nul is placed in *s2*.

Returns: the value of *s1*.

strcat

```
char *strcat(char * s1, const char * s2)
```

Appends a copy of the string pointed to by *s2* (including the terminating null character) to the end of the string pointed to by *s1*. The initial character of *s2* overwrites the null character at the end of *s1*.

Returns: the value of *s1*.

strncat

```
char *strncat(char * s1, const char * s2, size_t n)
```

Appends not more than *n* characters (a null character and characters that follow it are not appended) from the array pointed to by *s2* to the end of the string pointed to by *s1*. The initial character of *s2* overwrites the null character at the end of *s1*. A terminating null character is always appended to the result.

Returns: the value of *s1*.

The sign of a non-zero value returned by the comparison functions is determined by the sign of the difference between the values of the first pair of characters (both interpreted as unsigned char) that differ in the objects being compared.

memcmp

```
int memcmp(const void * s1, const void * s2, size_t n)
```

Compares the first *n* characters of the object pointed to by *s1* to the first *n* characters of the object pointed to by *s2*.

Returns: an integer greater than, equal to, or less than zero, depending on whether the object pointed to by *s1* is greater than, equal to, or less than the object pointed to by *s2*.

strcmp

```
int strcmp(const char * s1, const char * s2)
```

Compares the string pointed to by *s1* to the string pointed to by *s2*.

Returns: an integer greater than, equal to, or less than zero, depending on whether the string pointed to by *s1* is greater than, equal to, or less than the string pointed to by *s2*.

strncmp

```
int strncmp(const char * s1, const char * s2, size_t n)
```

Compares not more than *n* characters (characters that follow a null character are not compared) from the array pointed to by *s1* to the array pointed to by *s2*.

Returns: an integer greater than, equal to, or less than zero, depending on whether the string pointed to by *s1* is greater than, equal to, or less than the string pointed to by *s2*.

strcoll

```
int strcoll(const char * s1, const char * s2)
```

Compares the string pointed to by *s1* to the string pointed to by *s2*, both interpreted as appropriate to the LC_COLLATE category of the current locale.

Returns: an integer greater than, equal to, or less than zero, depending on whether the string pointed to by *s1* is greater than, equal to, or less than the string pointed to by *s2* when both are interpreted as appropriate to the current locale.

strxfrm

```
size_t strxfrm(char * s1, const char * s2, size_t n)
```

Transforms the string pointed to by *s2* and places the resulting string into the array pointed to by *s1*. The transformation function is such that if the `strcmp` function is applied to two transformed strings, it returns a value greater than, equal to or less than zero, corresponding to the result of the `strcoll` function applied to the same two original strings. No more than *n* characters are placed into the resulting array pointed to by *s1*, including the terminating null character. If *n* is zero, *s1* is permitted to be a null pointer. If copying takes place between objects that overlap, the behaviour is undefined.

Returns: The length of the transformed string is returned (not including the terminating null character). If the value returned is n or more, the contents of the array pointed to by $s1$ are indeterminate.

memchr

```
void *memchr(const void * s, int c, size_t n)
```

Locates the first occurrence of c (converted to an unsigned char) in the initial n characters (each interpreted as unsigned char) of the object pointed to by s .

Returns: a pointer to the located character, or a null pointer if the character does not occur in the object.

strchr

```
char *strchr(const char * s, int c)
```

Locates the first occurrence of c (converted to a char) in the string pointed to by s (including the terminating null character). The BSD UNIX name for this function is `index()`.

Returns: a pointer to the located character, or a null pointer if the character does not occur in the string.

strcspn

```
size_t strcspn(const char * s1, const char * s2)
```

Computes the length of the initial segment of the string pointed to by $s1$ which consists entirely of characters not from the string pointed to by $s2$. The terminating null character is not considered part of $s2$.

Returns: the length of the segment.

strpbrk

```
char *strpbrk(const char * s1, const char * s2)
```

Locates the first occurrence in the string pointed to by $s1$ of any character from the string pointed to by $s2$.

Returns: returns a pointer to the character, or a null pointer if no character from $s2$ occurs in $s1$.

strrchr

```
char *strrchr(const char * s, int c)
```

Locates the last occurrence of *c* (converted to a char) in the string pointed to by *s*. The terminating null character is considered part of the string. The BSD UNIX name for this function is `rindex()`.

Returns: returns a pointer to the character, or a null pointer if *c* does not occur in the string.

strspn

```
size_t strspn(const char * s1, const char * s2)
```

Computes the length of the initial segment of the string pointed to by *s1* which consists entirely of characters from the string pointed to by *s2*.

Returns: the length of the segment.

strstr

```
char *strstr(const char * s1, const char * s2)
```

Locates the first occurrence in the string pointed to by *s1* of the sequence of characters (excluding the terminating null character) in the string pointed to by *s2*.

Returns: a pointer to the located string, or a null pointer if the string is not found.

strtok

```
char *strtok(char * s1, const char * s2)
```

A sequence of calls to the `strtok` function breaks the string pointed to by *s1* into a sequence of tokens, each of which is delimited by a character from the string pointed to by *s2*. The first call in the sequence has *s1* as its first argument, and is followed by calls with a null pointer as their first argument. The separator string pointed to by *s2* may be different from call to call. The first call in the sequence searches for the first character that is not contained in the current separator string *s2*. If no such character is found, then there are no tokens in *s1* and the `strtok` function returns a null pointer. If such a character is found, it is the start of the first token. The `strtok` function then searches from there for a character that is contained in the current separator string. If no such character is found, the current token extends to the end of the string pointed to by *s1*, and subsequent searches for a token will fail. If such a character is found, it is overwritten by a null character, which terminates the current token. The `strtok` function saves a pointer to the following character, from which the next search for a token will start. Each subsequent call, with a null pointer as the value for the first argument, starts searching from the saved pointer and behaves as described above.

Returns: pointer to the first character of a token, or a null pointer if there is no token.

memset

```
void *memset(void * s, int c, size_t n)
```

Copies the value of *c* (converted to an unsigned char) into each of the first *n* characters of the object pointed to by *s*.

Returns: the value of *s*.

strerror

```
char *strerror(int errnum)
```

Maps the error number in *errnum* to an error message string.

Returns: a pointer to the string, the contents of which are implementation-defined. Under RISC OS and Arthur the strings for the given *errnums* are as follows:

- 0 No error (errno = 0)
- EDOM EDOM – function argument out of range
- ERANGE ERANGE – function result not representable
- ESIGNUM ESIGNUM – illegal signal number to `signal()` or `raise()`
- others Error code (errno) has no associated message).

The array pointed to may not be modified by the program, but may be overwritten by a subsequent call to the `strerror` function.

strlen

```
size_t strlen(const char * s)
```

Computes the length of the string pointed to by *s*.

Returns: the number of characters that precede the terminating null character.

time.h

`time.h` declares two macros, four types and several functions for manipulating time. Many functions deal with a calendar time that represents the current date (according to the Gregorian calendar) and time. Some functions deal with local time, which is the calendar time expressed for some specific time zone, and with Daylight Saving Time, which is a temporary change in the algorithm for determining local time.

struct tm

struct tm holds the components of a calendar time called the broken-down time. The value of tm_isdst is positive if Daylight Saving Time is in effect, zero if Daylight Saving Time is not in effect, and negative if the information is not available (as is the case under RISC OS).

```
struct tm {
    int tm_sec;      /* seconds after the minute, 0 to 60
                     (0-60 allows for the occasional leap
                     second) */
    int tm_min      /* minutes after the hour, 0 to 59 */
    int tm_hour     /* hours since midnight, 0 to 23 */
    int tm_mday     /* day of the month, 0 to 31 */
    int tm_mon      /* months since January, 0 to 11 */
    int tm_year     /* years since 1900 */
    int tm_wday     /* days since Sunday, 0 to 6 */
    int tm_yday     /* days since January 1, 0 to 365 */
    int tm_isdst    /* Daylight Saving Time flag */
};
```

clock

```
clock_t clock(void)
```

Determines the processor time used.

Returns: the implementation's best approximation to the processor time used by the program since program invocation. The time in seconds is the value returned, divided by the value of the macro `CLOCKS_PER_SEC`. The value `(clock_t)-1` is returned if the processor time used is not available. In the desktop, `clock()` returns all processor time, not just that of the program.

difftime

```
double difftime(time_t time1, time_t time0)
```

Computes the difference between two calendar times: `time1 - time0`. Returns: the difference expressed in seconds as a double.

mktime

```
time_t mktime(struct tm * timeptr)
```

Converts the broken-down time, expressed as local time, in the structure pointed to by `timeptr` into a calendar time value with the same encoding as that of the values returned by the `time` function. The original values of the `tm_wday` and `tm_yday` components of the structure are ignored, and the original values of the

other components are not restricted to the ranges indicated above. On successful completion, the values of the `tm_wday` and `tm_yday` structure components are set appropriately, and the other components are set to represent the specified calendar time, but with their values forced to the ranges indicated above; the final value of `tm_mday` is not set until `tm_mon` and `tm_year` are determined.

Returns: the specified calendar time encoded as a value of type `time_t`. If the calendar time cannot be represented, the function returns the value `(time_t)-1`.

time

```
time_t time(time_t * timer)
```

Determines the current calendar time. The encoding of the value is unspecified.

Returns: the implementation's best approximation to the current calendar time. The value `(time_t)-1` is returned if the calendar time is not available. If `timer` is not a null pointer, the return value is also assigned to the object it points to.

asctime

```
char *asctime(const struct tm * timeptr)
```

Converts the broken-down time in the structure pointed to by `timeptr` into a string in the style `Sun Sep 16 01:03:52 1973\n\n0`.

Returns: a pointer to the string containing the date and time.

ctime

```
char *ctime(const time_t * timer)
```

Converts the calendar time pointed to by `timer` to local time in the form of a string. It is equivalent to `asctime(localtime(timer))`.

Returns: the pointer returned by the `asctime` function with that broken-down time as argument.

gmtime

```
struct tm *gmtime(const time_t * timer)
```

Converts the calendar time pointed to by `timer` into a broken-down time, expressed as Greenwich Mean Time (GMT).

Returns: a pointer to that object or a null pointer if GMT is not available (it is not available under RISC OS).

localtime

```
struct tm *localtime(const time_t * timer)
```

Converts the calendar time pointed to by *timer* into a broken-down time, expressed a local time.

Returns: a pointer to that object.

strftime

```
size_t strftime(char * s, size_t maxsize, const char *  
                format, const struct tm * timeptr)
```

Places characters into the array pointed to by *s* as controlled by the string pointed to by *format*. The format string consists of zero or more directives and ordinary characters. A directive consists of a % character followed by a character that determines the directive's behaviour. All ordinary characters (including the terminating null character) are copied unchanged into the array. No more than *maxsize* characters are placed into the array. Each directive is replaced by appropriate characters as described in the following list. The appropriate characters are determined by the LC_TIME category of the current locale and by the values contained in the structure pointed to by *timeptr*.

Directive	Replaced by
%a	the locale's abbreviated weekday name
%A	the locale's full weekday name
%b	the locale's abbreviated month name
%B	the locale's full month name
%c	the locale's appropriate date and time representation
%d	the day of the month as a decimal number (01–31)
%H	the hour (24-hour clock) as a decimal number (00–23)
%I	the hour (12-hour clock) as a decimal number (01–12)
%j	the day of the year as a decimal number (001–366)
%m	the month as a decimal number (01–12)
%M	the minute as a decimal number (00–61)
%p	the locale's equivalent of either AM or PM designation associated with a 12-hour clock
%S	the second as a decimal number (00–61)
%U	the week number of the year (Sunday as the first day of week 1) as a decimal number (00–53)
%w	the weekday as a decimal number (0(Sunday)–6)
%W	the week number of the year (Monday as the first day of week 1) as a decimal number (00–53)
%x	the locale's appropriate date representation
%X	the locale's appropriate time representation

%y	the year without century as a decimal number (00–99)
%Y	the year with century as a decimal number
%Z	the time zone name or abbreviation, or by no character if no time zone is determinable
%%	%.

If a directive is not one of the above, the behaviour is undefined.

Returns: If the total number of resulting characters including the terminating null character is not more than *maxsize*, the *strftime* function returns the number of characters placed into the array pointed to by *s* not including the terminating null character. Otherwise, zero is returned and the contents of the array are indeterminate.

Part 3 - Developing software for RISC OS

How to write desktop applications in C

In this chapter, you will learn how to construct desktop applications in C, using the facilities provided by the RISC OS library (RISC_OSLib). You will probably find it useful to scan through the contents of the library before reading the chapter. Some familiarity with the *Window Manager* part of the *RISC OS Programmer's Reference manual* is also assumed. The description of RISC_OSLib here is not exhaustive: it is intended to introduce some common programming techniques used in desktop applications.

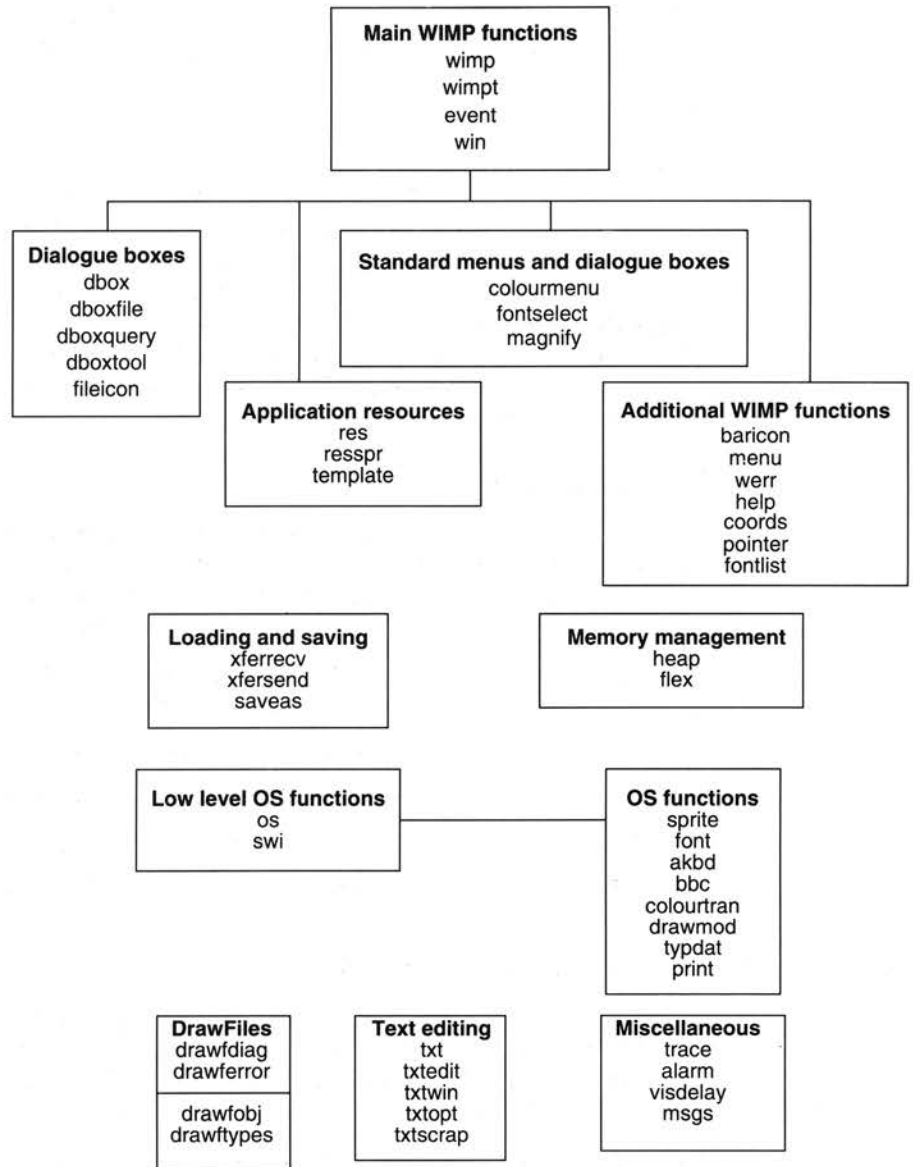
You are also advised to read the *RISC OS Style Guide*. This describes certain standards to which all desktop applications must conform in order to have an appearance which is consistent with the applications supplied by Acorn. Following these guidelines will make your own applications look and feel like part of the same environment, which makes them easier to learn and use.

The diagram over the page shows approximately how the various parts of the RISC OS library fit together. The diagram is reproduced on one of the reference cards; you may find it useful to refer to it as you work through this chapter.

Some general principles

Event handling

If you have read the *Window Manager* chapter in the *RISC OS Programmer's Reference manual*, you may be familiar with the idea of Wimp polling, as the means whereby an application finds out what the window manager requires it to do. In this method, an application uses a single polling loop, which must work out which of its windows each request from the Wimp is associated with, and take the appropriate action. RISC_OSLib makes available an alternative means of communicating with the Wimp, using functions called *event handlers*. An application may register event handlers (in the form of C functions) for windows, menus, icon bar icons, etc. It then calls a function in RISC_OSLib which processes events (ie polls the Wimp), and directs each event in turn to the relevant event handler. Event handlers may be added and removed whilst the application is running. This approach simplifies keeping track of which window, menu, or whatever, is associated with a window manager event.



When a call to register an event handler is made, an extra piece of data may be registered with it. This value (or handle) is then passed as an argument to the event handler when it is called by RISC_OSLib. It is sometimes convenient to use this as a way of allowing an event handler to retain private data. For example, you could use the same event handler for several windows, the handle being a pointer

to the data structure associated with the window. The event handler would then be able to locate the data structure for the window immediately, rather than having to work it out from the window handle passed into the event handler.

Windows and templates

In order to define the windows and dialogue boxes used by your program, you can either set up data structures which correspond to those used by the window manager SWIs, or you can use templates created by the template editor, FormEd. The template editor is an application which allows you to define windows on the screen, and save the definitions in a file ready for loading by your application. This is the approach used in Acorn's own applications, and you will find it makes the process of creating windows for your applications much easier. FormEd is described in the chapter entitled *FormEd* in the accompanying *Acorn Desktop Development Environment* user guide.

Application resources

Most desktop applications make use of a number of *resource files*. These should be considered as an integral part of your application. You can find a full list of the resource files typically used by an application in the section entitled *Application Notes* in the *RISC OS Programmer's Reference manual*; the following are usually present:

- `!Boot` *Run when the application directory is first displayed
- `!Run` *Run to start the application
- `!RunImage` the executable code for the application
- `!Sprites` used for application and file icon sprites
- `Sprites` containing other sprites used by the application
- `Templates` containing window and dialogue box templates.

Developing an application from scratch

This section contains an example of how to develop an application from scratch. You can use the description and the code given as a starting point for writing your own applications. The example application is called `!WExample`. It can be found as `User.!WExample` (the source code is there too).

The application is very simple. When started, it places an icon on the icon bar. The icon has a menu consisting of **Info**, which leads to a dialogue box containing information about the program, and **Quit** which closes the program down. Clicking **Select** on the icon opens a window, which may be resized, moved, closed, and so

on, in the normal way. If the window is already open, an error is reported. In itself, this is not a very useful program, but it illustrates the basic principles of writing a program which uses the RISC OS library.

The source code of the program is to be found as `User.!WExample.c.WExample`. Fragments of the code are given in this chapter to illustrate the points being described. You may also find it useful to have a listing of the whole program available to see how it all fits together.

The program illustrates the following:

- the general form of a desktop application
- how to initialise a desktop application
- how to create windows, icons and menus
- how to open a window, and respond to Wimp requests for it
- how to respond to user choices from a menu
- how to report errors to the user.

General form of a desktop application

A Wimp application normally consists of initialisation of both the RISC OS library and the application itself, followed by an event processing loop. `main()` in `c.WExample` is an example. You will see from this that the final step of the `main` function is to enter an infinite loop of calls to `event_process()`. The application will be closed down as a result of a call to the ANSI library function `exit()` elsewhere in the program. If you don't like this approach, an alternative is to define a global flag and test it in the event processing loop, for example:

```
/* --- global closedown flag --- */
BOOL all_done = FALSE;
...
/* --- the main event loop --- */
while ( !all_done ) event_process();
```

Note also that `event_process()` can automatically close down the application. To do this, it keeps an *active count*. The active count is initially zero; if it is zero again when `event_process()` is called, the application is closed down. You can change the active count by calling the functions `win_activeinc()` to increment it and `win_activedec()` to decrement it. Typically, you would call the first of these on opening a window, and the second on closing it. When you call `baricon()` to place an icon on the icon bar, `win_activeinc()` is called for you. If your application does not place an icon on the icon bar, you must make sure you call `win_activeinc()` yourself before entering the event processing loop. See the description of the `win` functions in the section entitled *win* on page 316 for more details.

Initialising a desktop application

The initialisation of `WExample` occurs in `example_initialise()`. The first few lines initialise various parts of the RISC OS library:

```
/* RISC_Oslib initialisation */
wimpt_init("RISC_Oslib example"); /* Main Wimp initialisation */
res_init("WExample");           /* Resources */
resspr_init();                  /* Application sprites */
template_init();                /* Templates */
dbox_init();                     /* Dialogue boxes */
```

Most applications will start with something similar, although there may be more or fewer parts of the library which need initialising. One point to note is the use of the arguments to `wimpt_init()` and `res_init()`.

`wimpt_init()` uses its argument in any circumstances where the application is to be referred to by name, for example in the task display, or in error boxes.

`res_init()` uses its argument to locate the application resources; in this case they will be expected to be in a directory whose name can be found from the system variable `WExample$Dir`. This variable must therefore be set up in the `!Run` file for the application, for example by:

```
*set WExample$Dir <Obey$Dir>
```

Creating windows, icons and menus

The remainder of `example_initialise()` creates the window which will be used in the program, sets up a menu to go with the icon, and places the icon on the icon bar. We will consider these one by one.

Creating the window is very straightforward. A pointer to a window definition read from the templates file is passed to `wimp_create_wind()`. An event handler is then registered for the window. Here is the code to do it, from `example_initialise()`:

```
/* Create the main window, and declare its event handler */
if (!example_create_window("MainWindow", &example_win_handle))
    return FALSE; /* Window creation failed */
win_register_event_handler(example_win_handle, example_event_handler, 0);
```

The code for creating the window is in a separate function, so that we could use it for creating other windows in a more complex program. It is as follows:

```
static BOOL example_create_window(char *name, wimp_w *handle)
{
    wimp_wind *window;    /* Pointer to window definition */

    /* Find template for the window */
    window = template_syshandle(name);
    if (window == 0) return FALSE;

    /* Create the window, dealing with errors */
    return (wimpt_complain(wimp_create_wind(window, handle)) == 0);
}
```

`example_create_window()` illustrates the value of using templates: all the work needed to set up the window definition in the program is avoided. The event handler will not be called unless the window is open; we will come back to this later.

The next step is to create the menu. If possible, you should create all menus during the initialisation. That way, when the user activates a menu, you can be sure that it exists and can be displayed. This may not be possible in some applications, because the menus have to change with circumstances, but you will nearly always be able to create at least part of the menu tree. In addition, clicking with Adjust when menus are created dynamically may fail, for subtle reasons connected with when the window manager calls the menu maker code.

The code to create the menu in the example is:

```
/* Create the menu tree */
if (example_menu = menu_new("Example", ">Info,Quit"), example_menu) == NULL)
    return FALSE; /* Menu create failed */
```

Example is the name which appears in the title bar of the menu. `menu.h` explains the syntax of the second argument to `menu_new()` in detail. In this case, `>Info` means that the menu entry for **Info** is to be marked as leading to a submenu consisting of a dialogue box.

If you want to check the menu before it is displayed, perhaps because you want to tick or shade items in it, then instead of calling `event_attachmenu()`, you can call `event_attachmenumaker()` with the name of a function to be called just before the menu is displayed. See the description of the file `event.h` in the chapter entitled *RISC OS library reference section* on page 175 for details.

Finally, the icon is placed on the icon bar, and event handlers registered for it. There are two event handlers here: `example_iconclick()`, which is called when Select is clicked on the icon, and `example_menuproc()` which is called when a choice is made from the menu. The work of displaying the menu is handled by `RISC_OSLib`. Here is the code:

```

/* Set up the icon on the icon bar, and register its event handlers */
baricon("!WExample", (int)resspr_area(), example_iconclick);
if (!event_attachmenu(win_ICONBAR, example_menu, example_menuproc, 0))
    return FALSE; /* Unable to attach menu */

```

There are two points to note here. First, the sprite used for the icon will be loaded from the 'Sprites' file in the application directory. The function `resspr_area()` returns a pointer to the sprite area into which this file is loaded. Second, `event_attachmenu()` is used to associate a menu with a window by specifying the window handle. The value `win_ICONBAR` is a special window handle which is used to represent the icon bar.

Opening and maintaining a window

The window is to be opened when the user clicks Select on the icon. As we saw above, clicking Select calls the function `example_iconclick()`, which is as follows:

```

static void example_iconclick(wimp_i icon)
{
    icon = icon; /* We don't need the handle: this stops compiler warning */

    /* Open the window - only one allowed */
    if (example_window_open)
        werr(FALSE, "Only one window may be opened");
    else
    {
        wimp_wstate state;

        /* Get the state of the window */
        if (wimpt_complain(wimp_get_wind_state(example_win_handle, &state)) == 0)
        {
            state.o.behind = -1; /* Make sure window is opened in front */
            wimpt_noerr(wimp_open_window(&state.o));
            example_window_open = TRUE;
        }
    }
}

```

You can ignore the lines of this up to the 'else' part for now: they just report an error if the window is already open. When we open the window, we want to make sure it is in front of any others on the screen. To do this, we read the current state of the window with `wimp_get_wind_state()`, and then ensure that our window is behind the window with handle `-1`, ie in front of all others. The window is actually opened with `wimp_open_wind()`.

Once the window is open, the event handler which we registered earlier will be called by `event_process()` when the window manager generates events for the window. The code for the event handler is:

```

static void example_event_handler(wimp_eventstr *e, void *handle)
{
    handle = handle; /* We don't need the handle: this stops compiler warning */

    /* Deal with event */
    switch (e->e)
    {
        case wimp_EREDRAW:
            example_redraw_window(e->data.o.w);
            break;

        case wimp_EOPEN:
            example_open_window(&e->data.o);
            break;

        case wimp_ECLOSE: /* Pass on close request */
            wimp_noerr(wimp_close_wind(e->data.o.w));
            example_window_open = FALSE;
            break;

        default: /* Ignore any other event */
            break;
    }
}

```

In this case, the event handler is very simple. Redraw and open requests are handled as described in the chapter entitled *Window Manager* in the *RISC OS Programmer's Reference manual*: see `c.wexample` for the full details of the functions. On a close request (generated by the user clicking the close icon of the window), we simply call `wimp_close_wind()`. After this, the event handler will not be called again, unless the window is re-opened. In an editor, some checks would normally be made before passing on the close request to the window manager: for example, ensuring that the contents of the window had been saved, and either warning the user or rejecting the close window request if they had not.

All events other than redraw, open and close requests are simply ignored. A way of improving the efficiency of the program would be to call `event_setmask()`. This indicates to the window manager that some events are never to be returned to the program. It must be used with care, since it masks the events to all windows. Thus, although the main window of the program has no menu, we could not mask out menu events, since they are used by the icon bar event handler. However, we could safely mask out 'pointer entering window' and 'pointer leaving window' events. Some suitable code for doing this would be:

```
event_setmask(wimp_EPTRENTER | wimp_EPTRLEAVE);
```

Responding to user choices from a menu

The menu is displayed when the user presses Menu over the icon: no special action is needed by the application for this. When the user makes a choice from the menu, the menu event handler we registered earlier is called:

```

/* Menu items */
#define Example_menu_info      1
#define Example_menu_quit     2

...

static void example_menuproc(void *handle, char *hit)
{
    handle = handle; /* We don't need the handle: this stops compiler warning */

    /* Find which menu item was hit and take action as appropriate */
    switch (hit[0])
    {
        case Example_menu_info:
            example_info_about_program();
            break;

        case Example_menu_quit:
            /* Exit from the program. The Wimp gets rid of the window and icon */
            exit(0);
    }
}

```

`handle` is the fourth argument which was given to `event_attachmenu()`: we make no use of it in this example. `hit` is a string in which each entry corresponds to a selection from the menu tree: the first character is the number of the selection from the top level menu, the second from the first submenu chosen, and so on. In this example, only a single, top-level menu was set up, so we are only interested in `hit[0]`.

Handling **Quit** is easy: the program just exits. A hit on **Info** occurs when either the user chooses it by clicking, or when he follows the submenu arrow leading from it. In this case, we call the following function to display a dialogue box containing information about the application:


```
static void example_info_about_program(void)
{
    dbox d; /* Dialogue box handle */

    /* Create the dialogue box */
    if (d = dbox_new("ProgInfo"), d != NULL)
    {
        /* Fill in the version number */
        dbox_setfield(d, Example_info_field, example_Version_String);

        /* Show the dialogue box */
        dbox_show(d);

        /* Keep it on the screen as long as needed */
        dbox_fillin(d);

        /* Dispose of the dialogue box */
        dbox_dispose(&d);
    }
}
```

First, the dialogue box is created from the template named `ProgInfo`; (this name is case-sensitive). Most of the fields are also taken from the template, but we want to fill in one field with the current version of the program, from the string `example_Version_String`. When this has been done, using `dbox_setfield()`, the dialogue box is displayed with `dbox_show()`.

The call to `dbox_fillin()` needs some explanation. It will not return until the window manager detects that the dialogue box has been finished with. For example, in this case a click elsewhere on the screen (and which therefore removes the menu tree) would cause `dbox_fillin()` to return. However, in the intervening time, other event handlers for the program may still be called. When the dialogue box is finished with, `dbox_fillin()` removes it from the screen and returns. The event handler then calls `dbox_dispose()` which deletes any internal data that was set up by the call to `dbox_new()`.

Reporting errors

The example application shows three different ways of dealing with errors. In each case, the error is reported in a standard error box, and the application waits until **OK** has been clicked. You will probably have seen this format from the desktop and the applications suite.

First, there are errors generated by the application itself. These are reported with, for example:

```
werr(FALSE, "Only one window may be opened");
```

(in the function `example_iconclick()`). The first parameter indicates whether this error is fatal: in this case it is not. A fatal error causes the application to exit. You can specify a number of parameters to `werr`, using the second one as a format string in the same way as for the ANSI library function `printf`.

When an error is returned by a RISC_OSLib function, we can report it in one of two further ways. The first is illustrated by the following line from `example_redraw_window()`:

```
wimpt_noerr(wimp_redraw_wind(&r, &more));
```

This reports the error in a dialogue box and halts the application.

An alternative is `wimpt_complain()`. This is similar to `wimpt_noerr()`, except that it also returns a pointer to the error, allowing the application to detect the error and take further action, as well as reporting it. A returned value of 0 indicates no error. For example (this is from `example_iconclick()`):

```
if (wimpt_complain(wimp_get_wind_state(example_win_handle, &state)) == 0)
{
... actions if there is no error ...
}
```

With one exception, it is strongly recommended that you report errors using `wimpt_noerr()` or `wimpt_complain()` on *all* calls to RISC_OSLib functions that return an error. This will help you find errors as soon as they occur. If an error does occur and you discard it, the effects of the errors may cause confusion at later stages in the program.

The one exception is reporting errors during redraw, using `wimpt_complain()`. Here you must take some care. If the error box lies over your window, when it is removed a new redraw will be issued, which can lead to the same error again. A possible solution is to keep a flag to avoid this happening, resetting the flag when the contents of the window have been mended.

More RISC_OSLib facilities

In this section, we will examine some more of the facilities provided by the RISC OS library. There is no complete example program to illustrate all of them, but fragments of code are given as illustrations.

The topics covered are:

- memory management
- responding to idle and unknown events
- loading and saving.

All of these require some practice to get right.

Memory management

RISC_OSLib includes two sets of functions for memory management: see `flex.h` and `heap.h`. The functions are an alternative to the ANSI C library functions such as `malloc()` and `free()`. For more details of how and when to use `flex` and `keep` routines, see the chapter entitled *Using memory efficiently* on page 355.

The `flex` functions can be useful for overcoming two of the problems of the standard ANSI functions, although they have different limitations themselves. Memory blocks allocated by the standard functions have fixed addresses, so when allocated cannot be shuffled to collect together free areas. This results in fragmentation of free space, wasting memory. Memory blocks allocated by `flex` functions are relocatable, so do not cause fragmentation. The standard ANSI functions allocate memory within the application's wimp slot. The size of this slot is determined initially by the use of `*WimpSlot` at application start-up, then grows if required to meet `malloc` demands. Releasing memory with `free` does not reduce the size of the application's wimp slot, so does not make it available for use by other applications. The `flex` routines allocate blocks outside the wimp slot, from memory available to all applications, and return it to this state when deallocated. The standard ANSI functions are tuned to provide good performance for a variety of applications, and operate quicker than the `flex` functions, especially for small memory blocks.

The basic difference between `flex` storage allocation and `malloc` allocation is that blocks allocated with `flex` may be moved in order to make the best use of the available memory, without the application being informed directly. This would cause problems if the application simply kept pointers to blocks of allocated memory: when RISC_OSLib moves blocks around, the pointers would cease to point to the right place. The approach that is used instead is to tell `flex` the address of the pointer to the block, which it will note in its own internal tables. If the block is moved, `flex` can then change the pointer. Thus, instead of writing code such as:

```
char *pointer;
pointer = malloc(size);
```

you would write:

```
char *pointer;
/* Allocate memory, passing in the address of 'pointer' */
flex_alloc((flex_ptr) &pointer, size);
```

The value `&pointer` is called the *anchor* of the `flex` block.

This may sound a little awkward if you are used to using `malloc` and `free`, but you will soon find that it becomes easy to use.

There is a restriction which you must be aware of if you use `flex`. The anchor of each flex block allocated must not itself move. This means that you cannot have flex pointers within blocks of memory allocated by `flex`. The following program fragments shows why this will not work. (You can skip this explanation if you like – the important point to remember is not to place flex pointers in areas of memory allocated by `flex`.)

```
#define MemSize 100
struct s_control_block
{
    char *data;
    ... other fields ...
    int size;
} *pointer;

/* Allocate a control block */
flex_alloc((flex_ptr) &pointer, sizeof(struct s_control_block));

/* Allocate the data block itself */
/* The next line is wrong!!! */
flex_alloc((flex_ptr) &(pointer->data), MemSize);
```

Suppose that `flex` moves the control block we allocated at `pointer`, and then tries to move the data block referenced by `pointer->data`. When the memory was allocated, `flex` made a note of the anchors `&pointer` and `&(pointer->data)`. Now suppose it moves the control block, and changes the value of the variable `pointer`. It then moves the data block and attempts to change its anchor. But the anchor it noted was an address within the control block at its original location, and the control block is no longer there. Consequently, `pointer->data`, with `pointer` having its new value, is not changed, rendering the pointer to the data block no longer valid. Not only that, but `flex` will have changed the location which originally contained `pointer->data` and which may by now be part of some other block.

A second restriction is that you must not make a copy of the pointer to a block allocated by `flex`. The reason here is simply that `flex` only knows of one anchor for each block. If the location of the block changes, the original pointer will be changed, but not any copies of it. A place where this can easily cause problems is in passing a pointer to a flex block as an argument to a function. Thus, the following example would not work:

```

void some_function(char *data)
{
    printf("%s\n", data);
}
...

char *pointer;
flex_alloc((flex_ptr) &pointer, 256);
...

/* The next call can go wrong!!! */
some_function(pointer);

```

A safe alternative is to introduce an extra level of indirection:

```

void some_function(char **data)
{
    printf("%s\n", *data);
}
...

char *pointer;
flex_alloc((flex_ptr) &pointer, 256);
...

/* Pass pointer to reference - this is OK */
some_function(&pointer);

```

You must call `flex_init()` before attempting to use `flex`. There are functions for allocating and freeing memory, and for changing the size of an allocated block of memory both by adding or removing memory from the end of the block, and adding or removing memory from part way through the block.

A flex block can move as the result of other calls to `flex`. The version of RISC_OSLib supplied with Desktop C by default inhibits expansion of the `malloc` area once `flex` has been initialised. For details of how to override this see the later chapter entitled *Using memory efficiently*. The default behaviour means that the first block allocated by `flex` does not move, and so pointers to data within this block can be used in the normal way. The heap functions provide facilities to manage a heap of fixed blocks within the first flex block.

Heap allocation is similar to `malloc()` in that a pointer to the block allocated is returned to the caller: the routine to do this is called `heap_alloc()`. Memory may be released with `heap_free()`. Before you use heap, you must call `heap_init()`. This must be done after `flex` has been initialised with `flex_init()`, and before any calls to flex functions.

Responding to idle and unknown events

There are two special types of event, which are only passed to your application if you specifically claim them.

Idle events

Idle events are generated by the window manager when nothing else is happening; they correspond to the `Null_Reason_Code` generated by the `SWI Wimp_Poll`. Applications should only claim idle events if they want to do some activity which continually needs processor time; an example is dragging the selection box in Draw. When you claim idle events, they are directed to the event handler for the window you specify, as an event of type `wimp_ENULL`. Idle events should be released as soon as they are no longer needed by the application. To claim and release idle events, use the function `win_claim_idle_events()`. See `win.h` for more details. Note that, by default, idle events are masked out, so to claim them you must also enable them with `event_setmask(0)`.

Unknown events

Unknown events are events which are not associated with a specific window. The following events are considered to be unknown:

- user drag events: `wimp_EUSERDRAG`
- menu events: `wimp_EMENU`
- losing and gaining the caret: `wimp_ELOSECARET` and `wimp_EGAINCARET`
- user message send events, `wimp_ESEND` and `wimp_ESENDWANTACK`, for any **except** the following message types: `wimp_MCLOSEDOWN`, `wimp_MDATASAVE`, `wimp_MDATALOAD`, `wimp_MHELPREQUEST`
- user message acknowledge events: `wimp_EACK`.

To claim unknown events, register one or more unknown event processors, and optionally an unknown event handler. When an unknown event occurs, it is offered to each of the unknown event processors in turn, until either one deals with the event, or they have all been tried. If none of them deals with the unknown event, it is then passed on to the unknown event handler, if any, or discarded if there isn't one.

To register an unknown event processor, call `win_add_unknown_event_processor()`, giving it the name of the unknown event handler, and a handle for any extra data you wish to be passed to the processor (as for window event handlers). The processors are called in the reverse of the order in which you register them, ie the most recently registered one is called first. See the type `win_unknown_event_processor()` in `win.h` for the type used for unknown

event processors. Each unknown event processor must return a value indicating whether it has handled the event or not. Unknown event processors may be cancelled by calling `win_remove_unknown_event_processor()`.

To register an unknown event handler, call `win_claim_unknown_events()`, specifying a window handle. Unknown events are then directed to the normal event handler for that window. You can cancel the unknown event handler by calling the same function with an argument of `-1`.

Loading and saving

There are functions in the RISC OS library for loading and saving data, using the same style as Acorn's own applications. The functions implement the data transfer protocol, as described under `Wimp_SendMessage` in the chapter entitled *Window Manager* in the *RISC OS Programmer's Reference manual*. They may be used for loading from and saving to files, and for transfers from and to other applications via RAM.

Loading

To load data, use the functions in `xferrecv`. Loading a file is initiated when the user drags a file to either a window opened by the application or its icon bar entry. A message is then sent to the corresponding event handler. In the event handler, you should have something like:

```
... other event cases ...
case wimp_ESEND:
case wimp_ESENDWANTACK:

switch (e->data.msg.hdr.action)
{
case wimp_MDATASAVE: /* import data */
load_ram();
break;

case wimp_MDATALOAD: /* insert data */
case wimp_MDATAOPEN:
load_file();
break;

... other message cases ...
}
```

For a load via RAM, the code is as follows (in outline):

```

static char *data;
void load_ram(void)
{
    int estsize;    /* Estimate size of file */

    /* Get the type of the file being loaded, and an estimate of its size */
    int filetype = xferrecv_checkimport(&estsize)

    if (filetype != -1)
    {
        int final_size;

        ... any necessary pre-load checks, e.g. valid filetype ...
        ... allocate a block 'estsize' long, at 'data' ...

        /* Initiate the load */
        if (final_size = xferrecv_do_import(data, estsize, buffer_processor),
            final_size >= 0)
        {
            ... load was ok ...
        }
        else
        {
            ... error during loading ...
        }
    }
    else /* Filetype of -1 indicates we should try to load via a file */
    {
        load_file();
    }
}

```

Here we check that the load really is via RAM transfer, and if not try to load from a file instead. If we decide to go ahead with the load, a call is made to `xferrecv_do_import()`. If the data being loaded fills up the buffer, then `buffer_processor()` is called. This function is not defined here: what it must do is either to empty the buffer, or to extend it. For a more precise specification, see the definition of `xferrecv_buffer_processor()` in `xferrecv.h`.

The code for loading from a file is:

```

void load_file(void)
{
    char *filename;

    /* Fetch the type and name of the file */
    int filetype = xferrecv_checkinsert(&filename);

    ... any necessary pre-load checks, e.g. valid filetype ...
    ... load file ...

    /* Indicate load is completed */
    xferrecv_insertfileok();
}

```


The work of loading the file here can be done using the standard methods for reading files, such as `os_file()`, or the ANSI C file functions. The file size is usually read first, so that the entire buffer can be allocated before loading starts.

In this function, a pointer to the name of the file being loaded is placed in `filename` by `xferrecv_checkinsert()`. This pointer does not remain valid permanently, and if you want to preserve it (for example, to use in a window title), you should copy it to a buffer of your own. The call `xferrecv_file_is_safe()` may be used to check the validity of the name.

In both cases, it is good practice to turn the hourglass on during the load. Suitable calls for turning it on and off are:

```
/* Turn hourglass on */
visdelay_begin();
...
/* Turn hourglass off */
visdelay_end();
```

Saving

There are two levels to the functions for saving. The bulk of the work is handled by the functions in `xfersend`, which are used for transferring data from the application to the destination of the save operation. The functions in `saveas` are used to display a save dialogue box and respond to dragging the icon from it. It is better to use `saveas`, since this makes the user interface for saving consistent with Acorn's applications. However, even if you are using `saveas`, you will still call some of the functions in `xfersend`, as described in this section.

A save operation is typically initiated by the user choosing something like **Save as** from a menu. In this case, you would start the operation with code such as:

```
int filetype = ...; /* Type of file */
char *name = ...; /* File name to be placed in dialogue box */
int estsize = ...; /* Estimated size of file */
char *data = ...; /* Data to be saved */
saveas(filetype, name, estsize, saver_proc, sender_proc, print_proc,
       (void *)data);
```

The three functions are used for:

- saving the file directly (`xfersend_saveproc`)
- transferring it via RAM a buffer-full at a time (`xfersend_sendproc`)
- printing the file (`xfersend_printproc`).

The last parameter to `saveas()` is an arbitrary handle which is passed to these functions. It is a convenient way of indicating the source of the data to be saved. The functions are called when the user has dragged the file icon to its destination, or specified the name of the file to be saved.

saveas() requires the presence of a template called xfer_send in the application's resources: it contains the save dialogue box.

Outlines of functions for saver_proc() and sender_proc() are:

```

/* saver_proc: type is the same as xfersend_saveproc */
BOOL saver_proc(char *filename, void *handle)
{
    ... any checks, eg valid file name ...
    ... save file, using any conventional method ...
}

/* sender_proc: type is the same as xfersend_sendproc */
BOOL sender_proc(void *handle, int *maxbuf)
{
    char *data = ...; /* Location of the data being sent, initially from
handle */
    int length = ...; /* Size of the block being sent */

    ... here there would be some sort of loop, getting chunks of data up to
        *maxbuf in length, and sending them with code something like: */

    while (...)
    {
        /* The data save itself */
        if (!xfersend_sendbuf(data, length)) return FALSE;
        else
        {
            /* Advance to next block */
            data += length; /* For example */
        }
    }
}

```

As with loading, you may want to turn on the hourglass during the save operation.

Note that you can specify the send and print functions as being NULL.

Using Draw files

You can use the RISC OS library to display files in the format used by Draw in your own applications. The format of Draw files is described in the RISC OS *Programmer's Reference manual*; Acorn intends that this should be treated as a standard for graphical data in RISC OS. There are two interfaces to the code for displaying Draw files. You can either draw entire files: the header for this is drawfdiag.h. Alternatively, you can draw files object by object; see drawfobj.h. The object-level interface also includes functions for adding and deleting objects. In both cases, it is possible to define your own object types, by specifying a function to handle unknown object types thus allowing you to extend the Draw file format.

Draw files use their own coordinate system. When rendering a Draw file, the origin of the file (ie coordinate (0,0)) is mapped to work area coordinate (0,0). The function `draw_shift_diag()`, may be used to shift all the coordinates in a Draw file. In addition, the coordinates used in Draw objects are not the same as those used for work area and screen coordinates. There are macros and functions which convert between the two systems. These just multiply the coordinates by the relevant factor: they take no account of where the Draw file origin is. Note also that the Draw file headers refer to the work area and screen coordinates collectively as *screen units*. This refers purely to their size: you are responsible for applying any further origin shifts to convert them to the coordinate system of the work area or the screen as a whole.

An application which illustrates many of the points described in the preceding sections can be found in `User.!DrawEx`. It does not set itself up on the icon bar when it is started; instead, it simply opens a window. Draw files dragged to the window are displayed in it. There is a window menu, with the usual **Info** and **Quit** entries, plus an entry to save the contents of the window: this entry is shaded when there is no file loaded. The application is closed down either by choosing **Quit**, or by closing the window.

The source code is not described here: it is left as an exercise for the reader to see how it works. You may also like to look at the `!Run` file, which shows how to make sure the necessary modules are loaded. Overlooking this is a common source of apparently serious errors in desktop applications.

The programming techniques illustrated by the application include:

- loading and saving files
- using `flex`
- using the active window count to handle closedown
- rendering Draw files.

One thing you may find it useful to look at in detail is the method used to extend `flex` blocks during a load via RAM. The code to do this is quite simple, but it is easy to get wrong. See the functions `drawex_load_ram()` and `drawex_ram_loader()`.

Common application features

There are a number of functions in the RISC OS library which are intended to help you produce applications with a similar appearance to those written by Acorn. Some of these have already been examined. This section briefly describes some of the others. As usual, for full details, look at the relevant parts of the chapter entitled RISC OS *library reference section* on page 175.

Coordinate conversion

You will often need to convert between the work area coordinates and screen. This is not difficult: the chapter entitled *Window Manager* in the *RISC OS Programmer's Reference manual* describes how to do it. However, you may find it convenient to use the functions in `coord.h` to do the conversion. Using these functions may make it clearer exactly what is happening in the source of your program. There are functions for converting x and y coordinates, points and boxes to either work area or screen coordinates, together with some extra functions used to move boxes, and determine if boxes overlap, and if a line intersects with a box. The conversion functions take a pointer to a `coords_cvtstr` object as a parameter. This consists of a box and two scroll values. You can obtain a suitable value for this parameter from the data structures returned by a number of Wimp functions. For example, the 'box', 'x' and 'y' fields of a `wimp_openstr`, or the 'box', 'scx' and 'scy' fields of a `wimp_redrawstr` are both suitable. Thus a typical fragment which might appear in a redraw loop is:

```
wimp_redrawstr r;
int screen_x, workarea_x;
...
screen_x = coords_x_toscreen(workarea_x, (coords_cvtstr *)&r.box);
```

You can always obtain the box and scroll values for the current window by finding the window state with `wimp_get_wind_state()`.

Colour translation

Some of the RISC OS graphics primitives such as the draw module and sprite plotting allow colours to be specified as full RGB (red/green/blue) values. RGB colours are usually referred to as 'true' colours. At any instant the desktop will only be able to display approximations to the true colours, specified using 'Gcol' values. The functions in `colourtran.h` are used to convert between these two ways of referring to colours. You can find further details in the chapter entitled *ColourTrans Module* in the *RISC OS Programmer's Reference manual*. One point to note about using these functions is that they require the ColourTrans module to be loaded. If you use them, the application's !Run file should include (something like) the following:

```
if "<System$Path>" = "" then Error 0 System resources cannot be found
|
RMEnsure ColourTrans 0.51 RMLoad System:Modules.Colours
RMEnsure ColourTrans 0.51 Error You need ColourTrans 0.51 or later
```

There are separate sets of functions for setting colours to be used in ordinary graphics operations, and for use with anti-aliased fonts.

Colour menus

The function `colourmenu_make()` constructs a menu of the current desktop colours. You can see an example of this kind of menu in Edit, where it is used for the **Foreground** and **Background** entries of the **Display** submenu. Menus of this form are used when you want to select one of the standard desktop colours, rather than a true colour.

If you do want the user to be able to select a true colour, you can call the function `dboxtcol`, which allows the red, green and blue levels of a colour to be set using sliders, or by specifying numerical values.

Dialogue boxes

There are a number of functions for handling dialogue boxes. Some of these have already been introduced; here we look at some more of them. You may also find it instructive to look at some dialogue boxes in the templates files of standard applications, using the template editor: this will give you some idea of how they are constructed and what button types you use for the various sorts of field. You can use dialogue boxes both as part of the menu trees, as already described, or on their own. The only difference between these is how you display the dialogue box: for a menu tree, use `dbox_show()` and for a 'standalone' one, use `dbox_showstatic()`.

As described in the RISC OS *Programmer's Reference manual*, the fields of a dialogue box consist of icons. You can change the contents of the fields using the routine `dbox_setfield()`. `dbox_setnumeric()` can be used to place a number in a field. Values from the fields may be read back with `dbox_getfield()` and `dbox_getnumeric()`. Fields may be faded, as for menu items, with `dbox_fadefield()` and `dbox_unfadefield()`, to cancel the effect.

To recognise when an action has occurred in a dialogue box, you can either call `dbox_fillin()`, which enters a Wimp polling loop until a field has been activated, or register your own event handler for the dialogue box with `dbox_eventhandler()`. The first of these is simpler and usually provides all the flexibility you need. When `dbox_fillin()` returns, you should call `dbox_persist()`. This will tell you whether the dialogue box is to be removed from the screen or not. A typical use of these functions is:

```

dbox  dialogue;
BOOL  filling = TRUE;  /* TRUE until the dbox is to be removed */

/* Create dialogue box */
if ((dialogue = dbox_new(<name of the dbox>)) == 0)
    ... error ...

... fill in initial values for fields with dbox_setfield, etc. ...

/* Display the dbox. This is for a dbox in a menu tree */
dbox_show(dialogue);

/* Fill in the dialogue box */
while (filling)
{
    switch (dbox_fillin(dialogue))
    {
        /* Clauses for each field that has an effect */
        case <field number>:
            ... get field contents with dbox_getfield, etc. ...

...
        /*
            Use the following line on (for example) OK and Cancel buttons
        */
        filling = dbox_persist();
        break;

        ... more similar clauses ...

        /* Use the next clause if the dbox has a close icon */
        case dbox_CLOSE:
            filling = FALSE;
            break;

        /* Clauses for uninteresting fields */
        default:          /* Do nothing */
            break;
    }
}

/* Get rid of the dialogue box */
dbox_dispose(&dialogue);

```

Some special properties of dialogue boxes are worth noting. If there are writable fields in the dialogue box, the dbox code interprets the up and down arrows to move the caret between them, in field order. Pressing Return advances the caret to the next writable field. Field 0 may be used in a special way here. If you press Return on the last writable field, field 0 will be activated, and `dbox_fillin()` will hence return 0 to the caller. If your dialogue box contains an **OK** button, it should normally be field 0, so that repeatedly pressing Return will eventually activate it.

Besides the functions in `dbox.h`, there are also three subsidiary dialogue box functions. `dboxtool` has already been described. `dboxfile` is a function for handling file dialogue boxes, similar to those used by `xfersend`. `dboxquery.h` is used to handle dialogue boxes that consist simply of a message to the user with **YES** and **NO** buttons, as used by Edit and Draw to ask whether unsaved data is to be discarded or not when a window is closed. See the header files for more details of these functions.

Finally, don't forget to call `template_init()` and `dbox_init()` during the initialisation of your application, in order to load the templates from the application's resources.

Magnifier

The magnifier is used for operations such as **Zoom** in Draw. The function `magnify_select()` can be used to read a magnification factor from a zoom dialogue box. To use this function, you must have a template called `magnifier` in your application's template file.

Displaying and editing text

There are a large number of functions for displaying and editing text in a window, in a similar way to Edit. See `txt.h`, `txtedit.h` and `txtwin.h` for full details. The conceptual model used is as follows.

Text is kept as a linear array of characters, known as a 'txt'. All character codes are allowed. There is a pointer into this called the 'dot', which marks the current editing position, and some other pointers known as markers, which are used (for example) for selecting blocks of text.

The characters are displayed in a window, with a newline for each '\n' character in the buffer. Screen updates happen for each text operation, but the result is only sure to be good when redraws can happen too. When a txt is displayed, the dot is constrained to be visible and the text will be scrolled in order to achieve this.

You can insert and delete characters at the dot, during which the markers will continue to point at the character that they pointed at before. There are functions for moving the dot and querying its position.

You can indicate a part of the buffer as being selected. Characters in the selection are displayed highlighted. No other special meaning is given to the selection. The selection and the dot need not coincide. There are functions to create, delete, move and query markers.

A txt is implemented using a single buffer containing the text, with a gap at the dot. Moving the dot involves a block move of the intervening text, but insertions and deletions are fast. The text buffer is expanded if necessary (it is held in a flex block).

The basic text editing functions are defined in `txt.h`. There are also higher level functions, which are intended for building complete text editors, in `txtedit.h`. `txtwin.h` adds further functions for displaying the same text in multiple windows.

The functions are based on the code in `Edit`, and you may find it useful to compare them with the way you can see `Edit` working.

Alarm functions

If your application needs to do some activity after a fixed length of time (for example, periodically updating a window), there are two ways in which it can do this. The first is to claim idle events and repeatedly examine the time. The second, and preferable, way is to use the functions defined in `alarm.h`. These allow you to set one or more alarms, specifying the time when they will occur. When the alarm is triggered an event handler is called. You may have more than one alarm set simultaneously. See `alarm.h` for details of the functions.

Tracing desktop applications

During the development of your program, you will probably want to trace what is happening. The DDT debugger provides several facilities to perform this for you, and this is the recommended tool. Some primitive facilities are also provided for tracing by `RISC_OSLib`. These were added before DDT was constructed, and are now retained largely for backwards compatibility.

One way of using `RISC_OSLib` for tracing is to use `werr`, but this is often inconvenient, since it requires acknowledgement by clicking in the **OK** box, and because it obscures part of the screen, which will cause problems if it is used in a redraw loop.

An alternative is to use the functions defined in `trace.h`. They display their results directly onto the screen using `printf`. This is rather messy, since the trace output does not appear in a window and may thus be overwritten by the output from other application, though it will never interfere with the application. One trick that is sometimes useful is to spool the output to a file, using `*Spool` so that the trace output can be examined later. In this case, all the other graphics output will also be sent to the file, and you may find it useful to include some sort of distinctive text in your trace output which you can search for using a text editor; for example:

```
tracef0(">>> This is some trace\n");
```

In order to use tracing, you will have to define `TRACE`, either using a line in your program such as

```
#define TRACE
```


or using the `-D` command line parameter to the C compiler. When trace is not set, the trace functions are treated as macros which convert into empty statements. Thus, the call to the trace function may be left in your program even when you no longer need the trace. This is often useful for generating debugging and production versions of the program from the same source. Tracing may also be turned on and off dynamically, with `trace_on()` and `trace_off()`, when trace has been compiled in.

There is a general trace function which takes an arbitrary number of parameters (like `printf`), and five functions which take a fixed number of parameters. The general trace function cannot be omitted by leaving `TRACE` undefined, because of the properties of C macro expansion. The functions with a fixed number of parameters are therefore generally preferable.

Where do you go from here?

The next step is to try writing a desktop application of your own. You might like to take one of the example programs and extend it. For example, you could add multiple windows to `DrawEx`, or allow it to display text and sprite files as well as `Draw` files, or to display an animated sequence of pictures. Don't try to use all of `RISC_OSLib` in one go! It is better to become familiar with it gradually, using the functions as you need them. You may also find it useful to glance at the `RISC_OSLib` header files which have not been mentioned here. They all correspond more or less exactly to sections in the *RISC OS Programmer's Reference manual*.

Writing desktop applications takes a little getting used to. In particular, the flow of control through the program is driven primarily by events from the window manager. This makes the programming a little harder, but it leads to applications which respond better to user actions. Using `RISC_OSLib`, you should find that programming in this style soon comes naturally.

Example programs

The following example desktop applications are supplied in the directory `User`:

- `!Wexample` and `!DrawEx`, as described above.
- `!Balls64`, which displays coloured balls in a window.
- `!Life`, which runs Conway's game of life in several windows simultaneously. This is coded as a demonstration of `RISC_OSLib`, not for speed or as a high-quality animation of `Life`.
- `!Automata`, which displays one dimensional cellular automata patterns in one or more windows. Among other points, this demonstrates C and assembler interworking, use of `flex`, and user file types.

12 RISC OS library reference section

This chapter presents brief summaries of all the functions in the RISC OS library, grouped alphabetically by header. You should also refer to the RISC OS *Programmer's Reference manual* for related information.

Additional functions are exported by RISC_OSLib but not documented here or in the library headers. These are experimental, and you should not use them.

akbd

These functions provide access to the keyboard under the Wimp.

akbd_pollsh

Checks if Shift is depressed.

Syntax: `int akbd_pollsh(void)`
Returns: 1 if Shift is depressed, 0 otherwise.

akbd_pollctl

Checks if Control is depressed.

Syntax: `int akbd_pollctl(void)`
Returns: 1 if Control is depressed, 0 otherwise.

akbd_pollkey

Checks if user has typed ahead.

Syntax: `int akbd_pollkey(int *keycode)`
Parameters: `int *keycode` – value of key pressed
Returns: 1 if user has typed ahead. Also passes value of key back through `keycode`.

Other Information: Function keys appear as values > 256 (produced by Wimp)

alarm

These functions provide alarm facilities for Wimp programs, using non-busy waiting.

alarm_init

Initialises the alarm system.

Syntax: `void alarm_init(void)`

Parameters: `void`.

Returns: `void`.

Other Information: If this call is made more than once, any pending alarms are cancelled.

alarm_timenow

Reports the current monotonic time.

Syntax: `int alarm_timenow(void)`

Parameters: `void`.

Returns: the current monotonic time.

Other Information: This timer cannot be set by programs, and can therefore be relied on to increment every centisecond. It wraps every few months.

alarm_timedifference

Returns the difference between two times.

Syntax: `int alarm_timedifference(int t1, int t2)`

Parameters: `int t1` – the earlier time

`int t2` – the later time.

Returns: difference between `t1` and `t2`.

Other Information: Times are as in SWI `OS_ReadMonotonicTime`. Deals with wrap-round of timer.

alarm_set

Sets an alarm at the given time.

Syntax: `void alarm_set(int at, alarm_handler proc, void *handle)`

Parameters: `int at` – time at which alarm should occur

`alarm_handler proc` – function to be called at alarm time

void *handle – caller-supplied handle to be passed to function.

Returns: void.

Other Information: The supplied function is called before passing the event on to any idle event claimer windows. at is in terms of the monotonic centisecond timer. The supplied function is passed the time at which it was called. If you have enabled idle events, these are still returned to you; otherwise, RISC_OSLib uses idle events internally to implement alarm calls (using non-busy waiting via wimp_pollidle()).

alarm_remove

Removes an alarm which was set for a given time with a given handle.

Syntax: void alarm_remove(int at, void *handle)

Parameters: int at – the time at which the alarm was to be made
void *handle – the given handle.

Returns: void.

Other Information: If no such alarm exists, this call has no effect.

alarm_removeall

Removes all alarms which have a given handle.

Syntax: void alarm_removeall(void *handle)

Parameters: void *handle – the handle to search for.

Returns: void.

alarm_anypending

Informs the caller whether an alarm with a given handle is pending.

Syntax: BOOL alarm_anypending(void *handle)

Parameters: void *handle – the handle.

Returns: True if there are any pending alarms for this handle.

alarm_next

Informs the caller whether an alarm is pending and, if so, for when it is.

Syntax: BOOL alarm_next(int *result)

Parameters: int *result – time for which alarm is pending

Returns: True if an alarm is pending.

Other Information: This should be used by polling loops (if you use the standard `while(TRUE) event_process();` loop, this is done for you). If a polling loop finds that an alarm is set it should use `wimp_pollidle` (with earliest time set to *result of `alarm_next()`) rather than `wimp_poll` to do its polling. If you handle idle events yourself, your handler should use `call_next` to call the next alarm function upon receiving an idle event (ie `wimp_ENULL`).

alarm_callnext

Calls the next alarm handler function.

Syntax: `void alarm_callnext(void)`

Parameters: `void`.

Returns: `void`.

Other Information: This is done for you if you use `event_process()` to do your polling (or even if you reach down as far as using `wimpt` for polling).

baricon

Installs the named sprite as an icon on the icon bar and registers a function to be called when Select is clicked.

Syntax: `wimp_i baricon(char *spritename, int spritearea, baricon_clickproc p)`

Parameters: `char *spritename` – name of sprite to be used
`int spritearea` – area where sprite is
`baricon_clickproc p` – pointer to function to be called on click of Select

Returns: the icon number of the installed icon (of type `wimp_i`). This will be passed to function `p` on click of Select.

Other Information: For details of installing a menu handler for this icon see `event_attachmenu()`.

baricon_left

Installs the named sprite as an icon on the left of the icon bar and registers a function to be called when Select is clicked.

Syntax:	<code>wimp_i baricon_left(char *spritename, int spritearea, baricon_clickproc p);</code>
Parameters:	As for baricon above.
Returns:	As for baricon above.
Other information:	As for baricon above.

baricon_newsprite

Changes the sprite used on the icon bar.

Syntax:	<code>wimp_i baricon_newsprite(char *newsprite)</code>
Parameters:	<code>char *newsprite</code> – name of new sprite to be used
Returns:	the icon number of the installed icon sprite.
Other information:	<code>newsprite</code> must be held in the same area as the sprite used in <code>baricon()</code> .

baricon_textandsprite

Installs the named sprite as an icon on the right of the icon bar with some given text below it, and registers a function to be called when Select is clicked.

Syntax:	<code>wimp_i baricon_textandsprite(char *spritename, char *text, int buflen, int spritearea, baricon_clickproc p);</code>
Parameters:	<code>char *spritename</code> – name of sprite to be used <code>char *text</code> – text to appear under sprite <code>int buflen</code> – length of text buffer <code>int spritearea</code> – area in which sprite is held <code>baricon_clickproc p</code> – pointer to function to be called on left mouse click
Returns:	the icon number of the installed icon (of type <code>wimp_i</code>). This will be passed to function <code>p</code> on left mouse click.
Other information:	For details of installing a menu handler for this icon see <code>event_attachmenu()</code> . The width of the icon is taken as the greater of <code>buflen</code> system font characters and the width of the sprite used.

baricon_textandsprite_left

Installs the named sprite as an icon on the right of the icon bar with some given text below it, and registers a function to be called when Select is clicked.

Syntax:	<code>wimp_i baricon_textandsprite_left(char *spritename, char *text, int bufflen, int spritearea, baricon_clickproc p);</code>
Parameters:	<p><code>char *spritename</code> – name of sprite to be used</p> <p><code>char *text</code> – text to appear under sprite</p> <p><code>int bufflen</code> – length of text buffer</p> <p><code>int spritearea</code> – area in which sprite is held</p> <p><code>baricon_clickproc p</code> – pointer to function to be called on left mouse click.</p>
Returns:	the icon number of the installed icon (of type <code>wimp_i</code>). This will be passed to function <code>p</code> on left mouse click.
Other information:	For details of installing a menu handler for this icon see <code>event_attachmenu()</code> . The width of the icon is taken as the greater of <code>bufflen</code> system font characters and the width of the sprite used.

bbc

These functions provide BBC-style graphics and mouse/keyboard control.

bbc: text output functions

The following functions provide BBC-style text output functions. They are retained to allow 'old-style' operations; you are recommended to use SWI calls via `kernel.h` in the C library.

bbc_vdu

Outputs a single character.

Syntax: `os_error *bbc_vdu(int)`

bbc_vduw

Outputs a double character.

Syntax: `os_error *bbc_vduw(int)`

bbc_vduq

Outputs multiple characters. Ctl is a control character. The number of further parameters is appropriate to Ctl (vduq knows how many to expect, and assumes the correct parameters have been passed).

Syntax: `os_error *bbc_vduq(int ctl,...)`

bbc_stringprint

Displays a null-terminated string.

Syntax: `os_error *bbc_stringprint(char *)`

bbc_cls

Clears text window.

Syntax: `os_error *bbc_cls(void)`

bbc_colour

Sets text colour.

Syntax: `os_error *bbc_colour(int)`

bbc_pos

Returns X coordinate of text cursor.

Syntax: `os_error *bbc_pos(void)`

bbc_vpos

Returns Y coordinate of text cursor.

Syntax: `os_error *bbc_vpos(void)`

bbc_tab

Positions text cursor at given coordinates.

Syntax: `os_error *bbc_tab(int,int)`

txt: graphics output functions**bbc_plot**

Carries out a given plot operation.

Syntax: `os_error *bbc_plot(int plotnumber, int x, int y)`

bbc_mode

Sets the screen mode.

Syntax: `os_error *bbc_mode(int)`

bbc_move

Moves the graphics cursor to the absolute coordinates given.

Syntax: `os_error *bbc_move(int, int)`

bbc_moveby

Moves the graphics cursor to a position relative to its current text cursor position.

Syntax: `os_error *bbc_moveby(int, int)`

bbc_draw

Draws a line to the given absolute coordinates.

Syntax: `os_error *bbc_draw(int, int)`

bbc_drawby

Draws a line to a position relative to the current graphics cursor.

Syntax: `os_error *bbc_drawby(int, int)`

bbc_rectangle

Plots a rectangle, given LeftX, BottomY, Width, and Height.

Syntax: `os_error *bbc_rectangle(int, int, int, int)`

bbc_rectanglefill

Plots a solid rectangle, given LeftX, BottomY, Width, and Height.

Syntax: `os_error *bbc_rectanglefill(int, int, int, int)`

bbc_circle

Draws a circle, given Xcoord, Ycoord, and Radius.

Syntax: `os_error *bbc_circle(int, int, int)`

bbc_circlefill

Draws a solid circle, given Xcoord, Ycoord, and Radius.

Syntax: `os_error *bbc_circlefill(int, int, int)`

bbc_origin

Moves the graphics origin to the given absolute coordinates.

Syntax: `os_error *bbc_origin(int, int)`

bbc_gwindow

Sets up a graphics window, given BottomLeftX, BottomLeftY, TopRightX, and TopRightY.

Syntax: `os_error *bbc_gwindow(int, int, int, int)`

bbc_clg

Clears the graphics window.

Syntax: `os_error *bbc_clg(void)`

bbc_fill

Flood-fills area X,Y, filling from X,Y until either a non-background colour or the edge of the screen is reached.

Syntax: `os_error *bbc_fill(int, int)`

bbc_gcol

Sets a graphics colour to the given value.

Syntax: `os_error *bbc_gcol(int, int)`

bbc_tint

Sets the grey level of a colour: use tint 0-3, as it gets shifted for you.

Syntax: `os_error *bbc_tint(int, int)`

bbc_palette

Physical to logical colour definition: Logical colour, Physical colour, Red level, Green level, Blue level.

Syntax: `os_error *bbc_palette(int, int, int, int, int)`

bbc_point

Finds the logical colour of a pixel at the indicated coordinates x, y.

Syntax: `int bbc_point(int, int)`

bbc_vduvar

Reads a single VDU or mode variable value, for the current screen mode.

Syntax: `int bbc_vduvar(int varno)`

bbc_vduvars

Reads several VDU or mode variable values. `vars` points to a sequence of ints, terminated by `-1`. Each is a VDU or mode variable number, and the corresponding values will be replaced by the value of that variable.

Syntax: `os_error *bbc_vduvars(int *vars /*in*/, int *values /*out*/)`

bbc_modevar

Reads a single mode variable, for the given screen mode.

Syntax: `int bbc_modevar(int mode, int varno)`

bbc: other calls

bbc_get

Returns a character from the input stream. `0x1xx` is returned if an escape condition exists.

Syntax: `int bbc_get(void)`

bbc_cursor

Alters cursor appearance. Argument takes values 0 to 3.

Syntax: `os_error *bbc_cursor(int)`

bbc_adval

Reads data from analogue ports or gives buffer data.

Syntax: `int bbc_adval(int)`

bbc_getbeat

Returns current beat value.

Syntax: `int bbc_getbeat(void)`

bbc_getbeats

Reads beat counter cycle length.

Syntax: `int bbc_getbeats(void)`

bbc_gettempo

Reads rate at which beat counter counts.

Syntax: `int bbc_gettempo(void)`

bbc_inkey

Returns character code from an input stream or the keyboard.

Syntax: `int bbc_inkey(int)`

bbc_setbeats

Sets beat counter cycle length.

Syntax: `os_error *bbc_setbeats(int)`

bbc_settempo

Sets rate at which beat counter counts.

Syntax: `os_error *bbc_settempo(int)`

bbc_sound

Makes or schedules a sound. Parameters: Channel, Amplitude, Pitch, Duration, and Future Time.

Syntax: `os_error *bbc_sound(int, int, int, int, int)`

bbc_soundoff

Deactivates the sound system.

Syntax: `os_error *bbc_soundoff(void)`

bbc_soundon

Activates the sound system.

Syntax: `os_error *bbc_soundon(void)`

bbc_stereo

Sets the stereo position for the specified channel.

Syntax: `os_error *bbc_stereo(int, int)`

bbc_voices

Sets the number of sound channels.

Syntax: `os_error *bbc_voices(int)`

colourmenu

Creates a Wimp colour setting menu.

colourmenu_make

Creates a menu containing the sixteen Wimp colours, with an optional None entry. Text in colour is written in black or white, depending on the background.

Syntax: `menu colourmenu_make(char *title, BOOL includeNone)`

Parameters: `char *title` – null-terminated string for menu title
`BOOL includeNone` – whether to include 'None' entry

Returns: On successful completion, pointer to created menu structure, otherwise null.

Other Information: Hits on this menu start from 1 as for other menus (see menu module for details).

colourtran

C interface to the ColourTrans SWIs.

colourtran_select_table

Sets up a translation table in a buffer, given a source mode and palette, and a destination mode and palette.

Syntax: `os_error *colourtran_select_table (int source_mode, wimp_paletteword *source_palette, int dest_mode, wimp_paletteword *dest_palette, void *buffer)`

Parameters: `int source_mode` – source mode
`wimp_paletteword *source_palette` – source palette
`int dest_mode` – destination mode
`wimp_paletteword *dest_palette` – destination palette
`void *buffer` – pointer to store for the table.

Returns: possible error condition.

colourtran_select_GCOLtable

Sets up a list of GCOLs in a buffer, given a source mode and palette, and a destination mode and palette.

Syntax: `os_error *colourtran_select_GCOLtable (int source_mode, wimp_paletteword *source_palette, int dest_mode, wimp_paletteword *dest_palette, void *buffer)`

Parameters: `int source_mode` – source mode
 `wimp_paletteword *source_palette` – source palette
 `int dest_mode` – destination mode
 `wimp_paletteword *dest_palette` – destination palette
 `void *buffer` – pointer to store for the list of GCOLs.

Returns: possible error condition.

colourtran_returnGCOL

Informs the caller of the closest GCOL in the current mode to a given palette entry.

Syntax: `os_error *colourtran_returnGCOL (wimp_paletteword entry, int *gcol)`

Parameters: `wimp_paletteword entry` – the palette entry
 `int *gcol` – returned GCOL value.

Returns: possible error condition.

colourtran_setGCOL

Informs the caller of the closest GCOL in the current mode to a given palette entry, and also sets the GCOL.

Syntax: `os_error *colourtran_setGCOL (wimp_paletteword entry, int fore_back, int gcol_in, int *gcol_out)`

Parameters: `wimp_paletteword entry` – the palette entry
 `int fore_back` – set to 0 for foreground, set to 128 for background
 `int gcol_in` – GCOL action
 `int *gcol_out` – returned closest GCOL.

Returns: possible error condition.

colourtran_return_colournumber

Informs the caller of the closest colour number to a given palette entry, in the current mode and palette.

Syntax: `os_error *colourtran_return_colournumber (wimp_paletteword entry, int *col)`

Parameters: `wimp_paletteword` – the palette entry
 `int *col` – returned colour number.

Returns: possible error condition.

colourtran_return_GCOLformode

Informs the caller of the closest GCOL to a given palette entry, destination mode and destination palette.

Syntax: `os_error *colourtran_return_GCOLformode (wimp_paletteword entry, int dest_mode, wimp_paletteword *dest_palette, int *gcol)`

Parameters: `wimp_paletteword entry` – the palette entry
`int dest_mode` – destination mode
`wimp_paletteword *dest_palette` – destination palette
`int *gcol` – returned closest GCOL.

Returns: possible error condition.

colourtran_return_colourformode

Informs the caller of the closest colour number to a given palette entry, destination mode and destination palette.

Syntax: `os_error *colourtran_return_colourformode (wimp_paletteword entry, int dest_mode, wimp_paletteword *dest_palette, int *col)`

Parameters: `wimp_paletteword entry` – the palette entry
`int dest_mode` – destination mode
`wimp_paletteword *dest_palette` – destination palette
`int *col` – returned closest colour number.

Returns: possible error condition.

colourtran_return_OppGCOL

Informs the caller of the furthest GCOL in the current mode from a given palette entry.

Syntax: `os_error *colourtran_return_OppGCOL (wimp_paletteword entry, int *gcol)`

Parameters: `wimp_paletteword entry` – the palette entry
`int *gcol` – returned GCOL value.

Returns: possible error condition.

colourtran_setOppGCOL

Informs the caller of the furthest GCOL in the current mode from a given palette entry, and also sets the GCOL.

Syntax: `os_error *colourtran_setOppGCOL (wimp_paletteword entry, int fore_back, int gcol_in, int *gcol_out)`

Parameters: `wimp_paletteword entry` – the palette entry
`int fore_back` – set to 0 for foreground, set to 128 for background
`int gcol_in` – GCOL action
`int *gcol_out` – returned furthest GCOL.

Returns: possible error condition.

colourtran_return_Oppcolournumber

Informs the caller of the furthest colour number from a given palette entry, in the current mode and palette.

Syntax: `os_error *colourtran_return_Oppcolournumber (wimp_paletteword entry, int *col)`

Parameters: `wimp_paletteword` – the palette entry
`int *col` – returned colour number.

Returns: possible error condition.

colourtran_return_OppGCOLformode

Informs the caller of the furthest GCOL from a given palette entry, destination mode and destination palette.

Syntax: `os_error *colourtran_return_OppGCOLformode (wimp_paletteword entry, int dest_mode, wimp_paletteword *dest_palette, int *gcol)`

Parameters: `wimp_paletteword entry` – the palette entry
`int dest_mode` – destination mode
`wimp_paletteword *dest_palette` – destination palette
`int *gcol` – returned furthest GCOL.

Returns: possible error condition.

colourtran_return_Oppcolourformode

Informs the caller of the furthest colour number from a given palette entry, destination mode and destination palette.

Syntax: `os_error *colourtran_return_Oppcolourformode (wimp_paletteword entry int dest_mode, wimp_paletteword *dest_palette, int *col)`

Parameters: `wimp_paletteword entry` – the palette entry
`int dest_mode` – destination mode
`wimp_paletteword *dest_palette` – destination palette
`int *col` – returned furthest colour number.

Returns: possible error condition.

colourtran_GCOL_tocolournumber

Translates a GCOL to a colour number (assuming 256-colour mode).

Syntax: `os_error *colourtran_GCOL_tocolournumber (int gcol, int *col)`

Parameters: `int gcol` – the GCOL
`int *col` – returned colour number.

Returns: possible error condition.

colourtran_colournumbertoGCOL

Translates a colour number to a GCOL (assuming 256-colour mode).

Syntax: `os_error *colourtran_colournumbertoGCOL (int col, int *gcol)`

Parameters: `int col` – the colour number
`int *gcol` – the returned GCOL.

Returns: possible error condition.

colourtran_returnfontcolours

Informs the caller of the font colours to match the given colours.

Syntax: `os_error *colourtran_returnfontcolours (font *handle, wimp_paletteword *backgnd, wimp_paletteword *foregnd, int *max_offset)`

Parameters: `font *handle` – the font's handle
`wimp_paletteword *backgnd` – background palette entry
`wimp_paletteword *foregnd` – foreground palette entry
`int *max_offset`

- Returns: possible error condition.
- Other Information: Closest approximations to fore/background colours will be set, and as many intermediate colours as possible (up to a maximum of `*max_offset`). Values are returned through the parameters.

colourtran_setfontcolours

Informs the caller of the font colours to match the given colours, and calls `font_setfontcolour()` to set them.

- Syntax: `os_error *colourtran_setfontcolours (font *handle, wimp_paletteword *backgnd, wimp_paletteword *foregnd, int *max_offset)`
- Parameters: `font *handle` – the font's handle
`wimp_paletteword *backgnd` – background palette entry
`wimp_paletteword *foregnd` – foreground palette entry
`int *max_offset`
- Returns: possible error condition.
- Other Information: Closest approximations to fore/background colours will be set, and as many intermediate colours as possible (up to a maximum of `*max_offset`). Values are returned through the parameters. `Font_setfontcolours()` is then called with these as parameters.

colourtran_invalidate_cache

To be called when the palette has changed since a call was last made to a function in this module, or a Draw object was rendered.

- Syntax: `os_error *colourtran_invalidate_cache (void)`
- Parameters: `void`
- Returns: possible error condition

coords

This file contains functions for working in the window coordinate system. Functions are provided to convert between screen and work area coordinates, and perform other simple operations on points, lines, or 'boxes'.

It is conventional to think of the point (0,0) as appearing at the top lefthand corner of a document.

coords_x_toscreen/coords_y_toscreen

Converts x/y work area coordinates into x/y absolute screen coordinates.

Syntax: `int coords_x_toscreen(int x, coords_cvtstr *r)`
`int coords_y_toscreen(int y, coords_cvtstr *r)`

Parameters: `int x` or `int y` – x or y coordinate in work area coordinates
`coords_cvtstr *r` – conversion box (screen coordinates and scroll offsets).

Returns: x or y screen coordinates.

coords_x_toworkarea/coords_y_toworkarea

Converts x/y screen coordinates into x/y work area coordinates.

Syntax: `int coords_x_toworkarea(int x, coords_cvtstr *r)`
`int coords_y_toworkarea(int y, coords_cvtstr *r)`

Parameters: `int x` or `int y` – x or y coordinate in screen coordinates
`coords_cvtstr *r` – conversion box (screen coordinates and scroll offsets).

Returns: x or y work area coordinates.

coords_box_toscreen

Converts a 'box' of workarea coordinates into a 'box' of screen coordinates.

Syntax: `void coords_box_toscreen(wimp_box *b, coords_cvtstr *r)`

Parameters: `wimp_box *b` – workarea box to be converted
`coords_cvtstr *r` – conversion box (screen coordinates and scroll offsets).

Returns: `void`.

Other Information: `b` is converted 'in situ' into screen coordinates (ie its contents change).

coords_box_toworkarea

Converts a 'box' of screen coordinates into a 'box' of workarea coordinates.

Syntax: `void coords_box_toworkarea(wimp_box *b, coords_cvtstr *r)`

Parameters: `wimp_box *b` – screen box to be converted
`coords_cvtstr *r` – conversion box (screen coordinates and scroll offsets).

Returns: `void`.

Other Information: `b` is converted 'in situ' into workarea coordinates (ie its contents are changed).

coords_point_toscreen

Converts a point (x,y) from workarea coordinates to screen coordinates.

Syntax: `void coords_point_toscreen(coords_pointstr *point, coords_cvtstr *r)`

Parameters: `coords_pointstr *point` – the point in workarea coordinates
`coords_cvtstr *r` – conversion box (screen coordinates and scroll offsets).

Returns: `void`.

Other Information: `point` is converted 'in situ' into screen coordinates (ie its contents are changed).

coords_point_toworkarea

Converts a point (x,y) from screen coordinates to workarea coordinates.

Syntax: `void coords_point_toworkarea(coords_pointstr *point, coords_cvtstr *r)`

Parameters: `coords_pointstr *point` – the point in screen coordinates
`coords_cvtstr *r` – conversion box (screen coordinates and scroll offsets).

Returns: `void`.

Other Information: `point` is converted 'in situ' into workarea coordinates (ie its contents are changed).

coords_withinbox

Informs the caller if a point (x,y) lies within a 'box'.

Syntax: `BOOL coords_withinbox(coords_pointstr *point, wimp_box *box)`

Parameters: `coords_pointstr *point` – the point
`wimp_box *box` – the box.

Returns: True if point lies within the box.

coords_offsetbox

Offset a 'box' by a given x and y displacement.

Syntax: `void coords_offsetbox(wimp_box *source, int byx, int byy, wimp_box *result)`

Parameters: `wimp_box *source` – the box to be moved
`int byx` – x displacement
`int byy` – y displacement
`wimp_box *result` – box when offset.

Returns: `void`.

Other Information: `source` and `result` are permitted to point at the same box.

coords_intersects

Informs the caller whether a line intersects a given 'box'.

Syntax: `BOOL coords_intersects(wimp_box *line, wimp_box *rect, int widen)`

Parameters: `wimp_box *line` – the line
`wimp_box *rect` – the box
`int widen` – width of line (same units as line and rect).

Returns: True if line intersects box.

coords_boxesoverlap

Informs the caller whether two 'boxes' cover any common area.

Syntax: `BOOL coords_boxesoverlap(wimp_box *box1, wimp_box *box2)`

Parameters: `wimp_box *box1` – one box
`wimp_box *box2` – the other box.

Returns: True if boxes overlap.

dbox

This file contains functions concerned with the creation, deletion and manipulation of dialogue boxes. It is important to note that the structure of your dialogue templates is an integral part of your program. Always use symbolic names for templates and for fields and action buttons within them. Templates for the dialogue boxes can be loaded using the template module in this library. See the chapter entitled *How to use the Template Editor* for how to use the RISC OS Template Editor in conjunction with this interface. A dbox is an abstract dialogue box handle.

dbox: creation and deletion functions

dbox_new

Builds a dialogue box from a named template. Template editor (FormEd) may have been used to create this template in the `Templates` file for the application.

Syntax: `dbox dbox_new(char *name)`

Parameters: `char *name` – template name (from templates previously read in by `template_init`), from which to construct dialogue box. `name` is as given when using FormEd to create template.

Returns: On successful completion, pointer to a dialogue box structure, otherwise null (eg when not enough space).

Other Information: This only creates a structure; it doesn't display anything! However, it does register the dialogue box as an active window with the window manager.

dbox_dispose

Disposes of a dialogue box structure.

Syntax: `void dbox_dispose(dbox*)`

Parameters: `dbox*` – pointer to pointer to a dialogue box structure

Returns: `void`.

Other Information: This also has the side-effect of hiding the dialogue box, so that it no longer appears on the screen. It also 'un-registers' it as an active window with the window manager and event processor.

dbox_show

Displays the given dialogue box on the screen.

Syntax: `void dbox_show(dbox)`

Parameters: `dbox` – dialogue box to be displayed (typically created by `dbox_new`)

Returns: `void`.

Other Information: Typically used when dialogue box is from a submenu so that it disappears when the menu is closed. If called when this dialogue box is showing, it has no effect. The show will occur near the last menu selection or the last caret setting (whichever is most recent).

dbox_showstatic

Displays the given dialogue box on the screen, and leaves it there, until explicitly closed.

Syntax: `void dbox_showstatic(dbox)`

Parameters: `dbox` – dialogue box to be displayed (typically created by `dbox_new`)

Returns: `void`.

Other Information: This is typically not used from menu selection, as it will persist longer than the menu (otherwise, it is the same as `dbox_show`).

dbox_hide

Hides a previously displayed dialogue box.

Syntax: `void dbox_hide(dbox)`

Parameters: `dbox` – dialogue box to be hidden

Returns: `void`.

Other Information: This does not release any storage; it just hides the dialogue box. If called when the dialogue box is already hidden, it has no effect.

dbox fields

A dialogue box has a number of fields, labelled from 0. There are the following distinct field types:

- **action fields**
Mouse clicks here are communicated to the client. The fields are usually labelled `go`, `quit`, etc. `Set/GetField` apply to the label on the field, although this is usually set up in the template.
- **output fields**
These display a message to the user, using `SetField`. Mouse clicks etc. have no effect.
- **input fields**
The user can type into these, and simple local editing is provided. `Set/GetField` can be used on the textual value, or `Set/GetNumeric` if the user should type in numeric values.
- **on/off fields**
The user can click on these to display their on/off status. They are always 'off' when the dialogue box is first created. The template editor can set up mutually exclusive sets of these at will. `Set/GetField` apply to the label on this field, `Set/GetNumeric` set/get 1 (on) and 0 (off) values.

The function keys can be used instead of the mouse to 'click' action and on/off fields. In addition, if a letter key is pressed, an attempt will be made to match the first capital letter found in any action field, and 'click' on that field. For example, 'y' will match Yes, and 'd' will match reDo.

dbox_field/dbox_fieldtype

type `dbox_field` values are field numbers within a dialogue box. They indicate what sort a field is (ie action, output, input, on/off).

dbox_setfield

Sets the given field, within the given dialogue box, to the given text value.

Syntax: `void dbox_setfield(dbox, dbox_field, char*)`

Parameters: `dbox` – the chosen dialogue box
`dbox_field` – chosen field number
`char*` – text to be displayed in field.

Returns: `void`.

Other Information: If the function is applied to a non-text field, it has no effect. If the field is an indirected text icon, the text length is limited by the size value used when setting up the

template in the template editor. Any longer text will be truncated to this length. Otherwise, text is truncated to 12 characters (11 text + 1 null). If the dialogue box is currently showing, the change is immediately visible. This function is really only useful with indirect icons.

dbox_getfield

Puts the current contents of the chosen text field into a buffer, whose size is given as the third parameter.

Syntax: `void dbox_getfield(dbox, dbox_field, char *buffer, int size)`

Parameters: `dbox` – the chosen dialogue box
`dbox_field` – the chosen field number
`char *buffer` – buffer to be used
`int size` – size of buffer.

Returns: `void`.

Other Information: If the function is applied to a non-text field, the null string is put in the buffer. If the length of the chosen field (plus null-terminator) is larger than the buffer, the result will be truncated.

dbox_setnumeric

Sets the given field, in the given dialogue box, to the given integer value.

Syntax: `void dbox_setnumeric(dbox, dbox_field, int)`

Parameters: `dbox` – the chosen dialogue box
`dbox_field` – the chosen field number
`int` – field's contents will be set to this value.

Returns: `void`.

Other Information: If the field is of type `input/output`, the integer value is converted to a string and displayed in the field. If the field is of type `action` or `on/off`, a non-zero integer value selects this field; zero deselects it.

dbox_getnumeric

Gets the integer value held in the chosen field of the chosen dialogue box.

Syntax: `int dbox_getnumeric(dbox, dbox_field)`

Parameters:	<code>dbox</code> – the chosen dialogue box <code>dbox_field</code> – the chosen field number.
Returns:	integer value held in chosen field.
Other Information:	If the field is of type <code>on/off</code> then return value of 0 means <code>off</code> , 1 means <code>on</code> . Otherwise, the return value is the integer equivalent of the field contents.

dbox_fadefield

Makes a field unselectable (ie faded by Wimp).

Syntax:	<code>void dbox_fadefield(dbox d, dbox_field f)</code>
Parameters:	<code>dbox d</code> – the dialogue box in which field resides <code>dbox_field f</code> – the field to be faded.
Returns:	<code>void</code> .
Other Information:	Fading an already faded field has no effect.

dbox_unfadefield

Makes a field selectable (ie 'unfades' it).

Syntax:	<code>void dbox_unfadefield(dbox d, dbox_field f)</code>
Parameters:	<code>dbox d</code> – the dialogue box in which field resides <code>dbox_field f</code> – the field to be unfaded.
Returns:	<code>void</code> .
Other Information:	Unfading an already selectable field has no effect.

dbox: events from dialogue boxes

A dialogue box acts as an input device: a stream of characters comes from it as if it were a keyboard, and an up-call can be arranged when input is waiting. Dialogue boxes may have a `close` button that is separate from their action buttons, usually in the header of the window. If this is pressed, `CLOSE` is returned: this should lead to the dialogue box being invisible. If the dialogue box represents a particular pending operation, the operation should be cancelled.

dbox_get

Tells caller which action field has been activated in the chosen dialogue box.

Syntax:	<code>dbox_field dbox_get(dbox d)</code>
---------	--

Parameters: `dbox` – the chosen dialogue box.

Returns: field number of activated field.

Other Information: This should only be called from an event handler (since this is the only situation where it makes sense).

dbox_read

Tells caller which action field has been activated in the chosen dialogue box. Does not cancel the event.

Syntax: `dbox_field dbox_read(dbox d);`

Parameters: `dbox` – the chosen dialogue box

Returns: field number of activated field

Other Information: This should only be called from an event handler (since this is the only situation where it makes sense).

dbox_eventhandler

Registers an event handler function for the given dialogue box.

Syntax: `void dbox_eventhandler(dbox, dbox_handler_proc, void* handle)`

Parameters: `dbox` – the chosen dialogue box
`dbox_handler_proc` – name of handler function
`void *handle` – user-defined handle.

Returns: `void`.

Other Information: When a field of the given dialogue box has been activated, the user-supplied handler function is called. The handler should be defined in the form: `void foo (dbox d, void *handle)`. When called, the function `foo` will be passed the relevant dialogue box, and its user-defined handle. A typical action in `foo` would be to call `dbox_get` to determine which field was activated. If `handler==0` then no function is installed as a handler (and any existing handler is 'un-registered').

dbox_raw_eventhandler

Registers a 'raw' event handler for the given dialogue box.

Syntax: `void dbox_raw_eventhandler(dbox, dbox_raw_handler_proc, void *handle)`

Parameters:	<code>dbox</code> – the given dialogue box <code>dbox_raw_handler_proc</code> – handler function for event <code>void *handle</code> – user-defined handle.
Returns:	<code>void</code> .
Other Information:	This registers a function which will be passed 'unvetted' window events. Under the window manager in RISC OS, the event will be a <code>wimp_eventstr*</code> (see Wimp module). The supplied handler function should return <code>True</code> if it processed the event; if it returns <code>False</code> , the event will be passed on to any event handler defined using <code>dbox_eventhandler()</code> as above. The form of the handler's function header is: <code>BOOL func (dbox d, void *event, void *handle)</code> .

dbox: pending operations

Dialogue boxes are often used to fill in the details of a pending operation. In this case a down-call driven interface to the entire interaction is often convenient. The following facilities aid this form of use.

dbox_fillin

Process events until a field in the given dialogue box has been activated.

Syntax: `dbox_field dbox_fillin(dbox d)`

Parameters: `dbox d` – the given dialogue box

Returns: field number of activated field.

Other Information: Handling of harmful events, like `dbox_popup` (below). On each call to `dbox_fillin`, the caret is set to the end of the lowest numbered writable icon.

dbox_fillin_fixedcaret

Process events until a field in the given dialogue box has been activated.

Syntax: `dbox_field dbox_fillin_fixedcaret(dbox d);`

Parameters: `dbox d` – the given dialogue box

Returns: field number of activated field.

Other Information: Same as `dbox_fillin`, except caret is not set to end of lowest numbered writable icon

dbox_popup

Build a dialogue box, from a named template, assign message to field 1, do a `dbox_fillin`, destroy the dialogue box, and return the number of the activated field.

Syntax: `dbox_field dbox_popup(char *name, char *message)`

Parameters: `char *name` – template name for dialogue box
`char *message` – message to be displayed in field 1.

Returns: field number of activated field.

Other Information: 'Harmful' events are those which could cause the dialogue to fail (eg keystrokes, mouse clicks). These events will cause the dialogue box to receive a `CLOSE` event.

dbox_persist

When `dbox_fillin` has returned an action event, this function returns `True` if the user wishes the action to be performed, but the dialogue box to remain.

Syntax: `BOOL dbox_persist(void)`

Parameters: `void`.

Returns: `BOOL` – does the user want the dialogue box to remain on screen?

Other Information: The current implementation returns `True` when the user has clicked `Adjust`. The caller should continue round the fill-in loop if the return value is `True` (ie don't destroy the dialogue box).

dbox_syshandle

Allows the caller to get a handle on the window associated with the given dialogue box.

Syntax: `int dbox_syshandle(dbox)`

Parameters: `dbox` – the given dialogue box

Returns: window handle of dialogue box (this is a `wimp_w` under the RISC OS window manager).

Other Information: This could be used to hang a menu off a dialogue box, or to 'customise' the dialogue box in some way. `dbox_dispose` will also dispose of any such attached menus.

dbox_init

Prepare for use of dialogue boxes from templates.

Syntax: `void dbox_init(void)`

Parameters: `void`

Returns: `void`

Other Information: This function must be called **once** before any dialogue box functions are used. You should call `template_init()` before this function.

dboxfile

Displays dialogue box with message, input field, and OK field and allows input of filename.

Syntax: `void dboxfile(char *message, unsigned filetype, char *a, int bufsize)`

Parameters: `char *message` – informative message to be displayed in dialogue box
`unsigned filetype` – OS-dependent type of file
`char *a` – default filename (initially) and also used for user-typed filename
`int bufsize` – size of `a`.

Returns: `void`.

Other Information: The template for the dialogue box must be called `dboxfile_db`. Parameters correspond to the template's icon numbers as follows:

icon #0	OK button
icon #1	message
icon #2	filename

The template should have the following icons:

icon #0	a text icon containing text OK with button type <code>menu icon</code>
icon #1	an indirected text icon (possibly with a default message) with button type <code>never</code>
icon #2	an indirected text icon with button type <code>writable</code> . See the <code>dboxfile_db</code> template used by Edit for an example.

The maximum length of `message` is 20. The char array pointed to by `a` will contain the typed-in file name of maximum length `bufsize` (if longer, truncated).

dboxquery

Displays a dialogue box, with YES and NO buttons, and a question, and gets reply.

Syntax: `dboxquery_REPLY dboxquery(char *question)`

Parameters: `char *question` – the question to be asked

Returns: reply by user.

Other Information: Question can be up to 120 chars long, 3 lines of 40 characters. Return will reply `yes`; Escape or CLOSE event will reply `cancel`. A call of `dboxquery(0)` will reserve space for the dialogue box and return with no display. This will mean that space is always available for important things like asking to quit! The template for the dialogue box should have the following attributes:

`window flags` `moveable, auto-redraw`. It is also advisable to have a title icon containing the name of your program (or other suitable text).

`icon #1` the message icon: should have indirected text flag set, with button type `never` and validation string L40.

`icon #0` the YES icon: should be text icon with text string set to YES; button type should be `menu icon`.

`icon #2` the NO icon: should be text icon with text string set to NO; button type should be `menu icon`. See the query dialogue box in Edit for an example.

dboxtcol

Displays a dialogue box to allow the editing of a true colour value.

Syntax: `BOOL dboxtcol(dboxtcol_colour *colour /*inout*/, BOOL allow_transparent, char *name, dboxtcol_colourhandler proc, void *handle)`

Parameters: `dboxtcol_colour *colour` – colour to be edited
`BOOL allow_transparent` – enables selection of a 'see-through' colour
`char *name` – title to put in dialogue box.

`dboxtcol_colourhandler` `proc` – function to act on the colour change
`void *handle` – the handle passed to `proc`.

Returns: True if colour edited, user clicks OK.

Other Information: The dialogue box to be used should be the same as that used by Paint to edit the palette. If the user clicks Select on OK, the `proc` is called and the dialogue box is closed. If the user clicks Adjust on OK, the `proc` is called and the dialogue box stays on the screen. This allows the client of this function to use `proc` to, say, change a sprite's palette to reflect the edited colour value and then to cause a redraw of the sprite.

drawdiag

This file contains functions concerned with the processing of Draw format files (diagram level interface). It defines the interface to the simplest version of the DrawFile module. It can read in files to diagrams and render them. There is no checking of whether the end of the diagram has been overrun.

To read in Draw files, it is expected that the caller will do the work of the I/O itself. To dispose of a diagram, the caller can just throw it away: the module does not keep any hidden information about what diagrams it has seen.

Some calls return an offset to the bad data on an error. This is not necessarily the start of an object: it may be bad data part way through it. The offset is relative to the start of the diagram.

The module cannot handle rectangle or ellipse objects: you should use a path instead.

Data types

Diagram: a pointer to the data and a length field. The length must be an exact number of words, and is the amount of space used in the diagram, not the size of the memory allocated to it.

Abstract handle for an object: The object handle is an offset from the start of the diagram to the object data. You may use it to set a pointer directly to an object, when using the object level interface

Error types: Where a routine can produce an error, the actual value returned is a `BOOL`, which is `True` if the routine succeeded. The error itself is returned in a block passed by the user; if `NULL`, then the details of the error are not passed back.

The error block may contain either an operating system error or an internal error. In the latter case, it consists of a code and possibly a pointer to the location in the file where the error occurred (if `NULL`, the location is not known or not specified). By convention, this should be reported by the caller in the form `message (location &xx in file)`. For a list of codes and standard errors, see `h.DrawfErrors`. The location is relative to the start of the data block in the diagram.

draw_verify_diag

Verifies a diagram which has been read in from a file.

Syntax: `BOOL draw_verify_diag(draw_diag *diag, draw_error *error)`

Parameters: `draw_diag *diag` – the diagram to be verified
`draw_error *error` – the first error encountered (if any).

Returns: True if diagram is correct.

Other Information: Each object in the file is checked and the first error encountered causes return (with error set appropriately).

draw_append_diag

Merges two diagrams into one.

Syntax: `BOOL draw_append_diag(draw_diag *diag1, draw_diag *diag2, draw_error *error)`

Parameters: `draw_diag *diag1` – diagram to which to append `diag2`
`draw_diag *diag2` – diagram to be appended to `diag1`
`draw_error *error` – possible error condition.

Returns: True if merge was successful.

Other Information: Both diagrams should have been processed by `draw_verify_diag()`. `Diag1`'s data block must be at least `diag1.length + diag2.length`. `Diag1.length` will be updated to its new appropriate value. `Diag1`'s bounding box will be set to the union of the bounding boxes of the two diagrams. Offsets of objects in `Diag1` may change due to a change in font table size (if `Diag2` has fonts). Errors referring to specific locations, refer to `Diag2`.

draw_render_diag

Renders a diagram with a given scale factor, in a given Wimp redraw rectangle.

Syntax: `BOOL draw_render_diag(draw_diag *diag, draw_redrawstr *r, double scale, draw_error *error)`

Parameters: `draw_diag *diag` – the diagram to be rendered
`draw_redrawstr *r` – the Wimp redraw rectangle
`double scale` – scale factor
`draw_error *error` – possible error condition.

Returns: True if render was successful.

Other Information: The diagram must have been processed by `draw_verify_diag()`. `draw_redrawstr` is the same as `wimp_redrawstr`, which may be cast to it. Very small and negative scale factors will result in a run-time error (safe > 0.00009). The caller should do range checking on the scale factor. Following the normal convention for coordinate mapping, the part of the diagram rendered is found by mapping the top left of the diagram, in draw coord space onto a point: (`r->box.x0 - r->scx`, `r->box.y1 - r->scy`) in screen coordinates.

draw: memory allocation functions**draw_registerMemoryFunctions**

Registers three functions to be used to allocate, extend and free memory, when rendering text objects.

Syntax: `void draw_registerMemoryFunctions(draw_allocate alloc, draw_extend extend, draw_free free)`

Parameters: `draw_allocate alloc` – pointer to function to be used for memory allocation
`draw_extend extend` – pointer to function to be used for memory extension
`draw_free free` – pointer to function to be used for memory freeing.

Returns: void.

Other Information: These three functions will be used only when rendering text area objects. Any memory allocated during rendering will be freed (using the supplied function) after rendering. If `draw_registerMemoryFunctions()` is never called, or if memory allocation fails, then an attempt to render a text area will produce no effect. The three functions should operate as follows:

- `int alloc(void **anchor, int n)`: allocate `n` bytes of store and set `*anchor` to point to them. Return 0 if store can't be allocated, otherwise non-zero.
- `int extend(void **anchor, int n)`: extend the block of memory which starts at `*anchor` to a total size of `n` bytes. `n` will always be positive, and the new memory should be appended to the existing block (which may be moved by the operation). Return 0 if the memory can't be allocated, otherwise non-zero.
- `void free(void **anchor)`: free the block of memory which starts at `*anchor`, and set `*anchor` to 0.

The specification for these three functions is the same as that for `flex_alloc`, `flex_extend` and `flex_free` (in the `flex` module), so these can be used as the three required functions.

draw_shift_diag

Shifts a diagram by a given distance.

Syntax: `void draw_shift_diag(draw_diag *diag, int xMove, int yMove)`

Parameters: `draw_diag *diag` – the diagram to be shifted
`int xMove` – distance to shift in x direction
`int yMove` – distance to shift in y direction.

Returns: `void`.

Other Information: All coordinates in the diagram are moved by the given distance.

draw_querybox

Finds the bounding box of a diagram.

Syntax: `void draw_queryBox(draw_diag *diag, draw_box *box, BOOL screenUnits)`

Parameters: `draw_diag *diag` – the diagram
 `draw_box *box` – the returned bounding box
 `screenUnits` – indication whether the box is to be
 specified in draw or screen units.

Returns: `void`.

Other Information: The bounding box of `diag` is returned in `box`. If `screenUnits` is true, `box` is in screen units, otherwise, it is in draw units.

draw_convertBox

Converts a box to/from screen coordinates.

Syntax: `void draw_convertBox(draw_box *from, draw_box *to, BOOL toScreen)`

Parameters: `draw_box *from` – box to be converted
 `draw_box *to` – converted box
 `BOOL toScreen` – should set to True if conversion is to be from draw coordinates to screen coordinates. False makes conversion from screen coordinates to draw coordinates.

Returns: `void`.

Other Information: `from` and `to` may point to the same box.

draw_rebind_diag

Force the header of a diagram's bounding box to be exactly the union of the objects in it.

Syntax: `void draw_rebind_diag(draw_diag *diag)`

Parameters: `draw_diag *diag` – the diagram.

Returns: `void`.

Other Information: The diagram should have been processed by `draw_verify_diag()` first.

draw: unknown object handling

New types of object can be added by registering an unknown object handler. The handler is called whenever an attempt is made to render an object whose tag is not one of the standard ones known to DrawFile. It is passed a pointer to the object to be rendered (cast to a `void *`), and a pointer to a block into which to write any

error status. The object pointer may be cast to one of the standard Draw types (defined in the object level interface), or to a client-defined type. If an error occurs, the handler must return False and set up the error block; otherwise it must return True. Unknown objects must conform to the standard convention for object headers, ie one-word object tag; one-word object size; four-word bounding box. The unknown object handler is only called if the object is visible, ie if there is an overlap between its bounding box and the region of the diagram being rendered. The object size field must be correct, otherwise catastrophes will probably result.

draw_set_unknown_object_handler

Registers a function to be called when an attempt is made to render an object with an object tag which is not known.

Syntax: `draw_unknown_object_handler`
`draw_set_unknown_object_handler`
`(draw_unknown_object_handler handler, void *handle)`

Parameters: `draw_unknown_object_handler handler` – the handler function
`void *handle` – arbitrary handle to pass to function.

Returns: The previous handler.

Other Information: The handler can be removed by calling with 0 as a parameter.

drawfdiag_init

Initialise the diagram level interface

Syntax: `BOOL drawfdiag_init(void);`

Parameters: `void`

Returns: `TRUE` if all went OK.

Other Information: `none`

drawferror

Definition of error codes and standard messages for the Drawfile rendering functions. For each error, a code and the standard message are listed. See `drawfdiag`, above, for how to use the errors.

BadObject 1	Bad object
BadObjectHandle 2	Bad object handle
TooManyFonts 3	Too many font definitions

BBoxWrong 101	Bounding box coordinates are in the wrong order
BadCharacter 102	Bad character in string
ObjectTooSmall 103	Object size is too small
ObjectTooLarge 104	Object size is too large
ObjectNotMult4 105	Object size is not a multiple of 4
ObjectOverrun 106	Object data is larger than specified size
ManyFontTables 107	There is more than one font table
LateFontTable 108	The font table appears after text object(s)
BadTextStyle 109	Bad text style word
MoveMissing 110	Path must start with a move
BadPathTag 111	Path contains an invalid tag
NoPathElements 112	Path does not contain any line or curve elements
PathExtraData 113	There is extra data present at the end of a path object
BadSpriteSize 114	The sprite definition size is inconsistent with the object size
BadTextColumnEnd 115	Missing end marker in text columns
ColumnsMismatch 116	Actual number of columns in a text area object does not match specified number of columns
NonZeroReserved 117	Non-zero reserved words in a text area object
NotDrawFile 118	This is not a Draw file
VersionTooHigh 119	Version number too high
BadObjectType 120	Unknown object type
CorruptTextArea 121	Corrupted text area (must start with '\!')
TextAreaVersion 121	Text area version number is wrong or missing
MissingNewline 122	Text area must end with a newline character
BadAlign 123	Text area: bad \A code (must be L, R, C or D)
BadTerminator 124	Text area: bad number or missing terminator
ManyDCommands 125	Text area: more than one \D command
BadFontNumber 126	Text area: bad font number
UnexpectedCharacter 127	Text area: unexpected character in \F command
BadFontWidth 128	Text area: bad or missing font width in \F command
BadFontSize 129	Text area: bad or missing font size in \F command
NonDigitV 130	Text area: non-digit in \V command
BadEscape 131	Text area: bad escape sequence
FewColumns 133	Text area must have at least one column
TextColMemory 134	Out of memory when building text area (location field is always 0 for this error).

drawfobj

This file handles the processing of Draw format files (object level interface), and supplements the diagram level interface with routines for dealing with individual objects.

draw_create_diag

Creates an empty diagram (ie just the file header), with a given bounding box.

Syntax: `void draw_create_diag(draw_diag *diag, char *creator, draw_box bbox)`

Parameters: `draw_diag *diag` – pointer to store to hold diagram
`char *creator` – pointer to character string holding creator's name
`draw_box bbox` – the bounding box (in Draw units).

Returns: `void`.

Other Information: `diag` must point at sufficient memory to hold the diagram. The first 12 chars of `creator` are stored in the file header. `diag.length` is set appropriately by this function.

draw_doObjects

Renders a specified range of objects from a diagram.

Syntax: `BOOL draw_doObjects(draw_diag *diag, draw_object start, draw_object end, draw_redrawstr *r, double scale, draw_error *error)`

Parameters: `draw_diag *diag` – the diagram
`draw_object start` – start of range of objects to be rendered
`draw_object end` – end of range of objects to be rendered
`draw_redrawstr *r` – Wimp-style redraw rectangle
`double scale` – the scale factor for rendering
`draw_error *error` – possible error condition.

Returns: True if render was successful.

Other Information: Parameters (except range) are used as in `draw_render_diag`, in diagram level module. The diagram must be verified before a call to this function. If the range of objects includes text with anti-aliasing fonts,

you **must** call `draw_setFontTable` first. Very small (<0.00009) or negative scale factors will cause run-time errors.

draw_setFontTable

Scans a diagram for a font table object and records it for a subsequent call of `draw_doObjects`.

Syntax: `void draw_setFontTable(draw_diag *diag)`

Parameters: `draw_diag *diag` – the diagram to be scanned.

Returns: `void`.

Other Information: This function must be called for `draw_doObjects` to work on a sequence of objects that includes text objects using anti-aliasing fonts, but no font table object. The font table remains valid until either a different one is encountered during a call to `draw_doObjects`, or until `draw_render_diag` is called, or until a different diagram is rendered.

draw_verifyObject

Verifies the data for an existing object in a diagram.

Syntax: `BOOL draw_verifyObject(draw_diag *diag, draw_object object, int *size, draw_error *error)`

Parameters: `draw_diag *diag` – the diagram
`draw_object object` – the object to be verified
`int *size` – gets set to the amount of memory occupied by the object
`draw_error *error` – possible error condition.

Returns: True if object found and verified.

Other Information: Verifying an object ensures that its bounding box is consistent with the data in it; if not, no error is reported, but the box is made consistent. On an error, the location is relative to the start of the diagram. The object's size is returned only if `size` is a non-null pointer.

draw_createObject

Creates an object after a specified object in a given diagram.

Syntax: `BOOL draw_createObject(draw_diag *diag, draw_objectType newObject, draw_object after, BOOL rebind, draw_object *object, draw_error *error)`

Parameters: `draw_diag *diag` – the diagram
`draw_objectType newObject` – the created object
`draw_object after` – the object after which the new object should be created
`BOOL rebind` – if True, the bounding box of the diagram is updated to the union of its existing value and that of the new object
`draw_object *object` – new object's handle
`draw_error *error` – possible error condition.

Returns: True if object was created OK.

Other Information: All data after the insertion point is moved down. `after` may be set to `draw_FirstObject`/`draw_LastObject` for inserting at the start/end of the diagram. The diagram must be large enough for the new data; its length field is updated. On an error, the location is not meaningful. The handle of the new object is returned in `object`. If this function is used to create a font table, `after` is ignored, and the object merged with the existing one (if such exists) or inserted at the start of the diagram otherwise. This can cause the font reference numbers to change; if a call to this function is followed by a `draw_translateText()`, the font change will be applied (this is only needed when anti-aliased fonts are used in text objects).

draw_deleteObjects

Deletes the specified range of objects from a diagram.

Syntax: `BOOL draw_deleteObjects(draw_diag *diag, draw_object start, draw_object end, BOOL rebind, draw_error *error)`

Parameters: `draw_diag *diag` – the diagram
`draw_object start` – start of range of objects to be deleted
`draw_object end` – end of range of objects to be deleted
`BOOL rebind` – if set to True, then the diagram's

bounding box will be set to the union of those remaining objects
`draw_error *error` – possible error condition.

Returns: True if objects deleted successfully.

Other Information: diagram length is updated appropriately.

draw_extractObject

Extracts an object from a diagram into a supplied buffer.

Syntax: `BOOL draw_extractObject(draw_diag *diag, draw_object object, draw_objectType result, draw_error *error)`

Parameters: `draw_diag *diag` – the diagram
`draw_object object` – the object to be extracted
`draw_objectType result` – pointer to the buffer
`draw_error *error` – possible error division

Returns: True if the object was extracted successfully.

Other Information: The buffer for the result must be large enough to hold the extracted object (an object's size can be ascertained by calling `draw_verifyObject()`).

draw_translateText

Updates all font reference numbers for text objects following creation of a font table.

Syntax: `void draw_translateText(draw_diag *diag)`

Parameters: `draw_diag *diag` – the diagram.

Returns: `void`.

Other Information: If the font table has not been changed, this function does nothing.

drawfobj_init

Initialise the object level interface

Syntax: `BOOL drawfobj_init(void);`

Parameters: `void`

Returns: TRUE if all went OK.

Other Information: none

drawftypes

This file contains declarations of all the data types needed for manipulating Draw objects at a low level, enabling you to examine or change their individual properties. For full details, refer to the header file on Disc 3:
`$.RISC_OSlib.h.drawftypes`.

drawmod

This file provides a C interface to the Draw module (not to be confused with the Draw application). It defines a number of types used for PostScript-like operations, with enhancements (for full details, refer to the header file on Disc 3:
`$.RISC_OSlib.h.drawmod`). The enhancements consist mainly of choice of fill style (fill including/excluding boundary etc), extra winding numbers, differing leading/trailing line caps and triangular line caps. It calls the Draw SWIs.

drawmod_fill

Emulates the Postscript 'fill' operator – ie closes open subpaths, flattens a path, transforms it to standard coordinates and fills the result.

Syntax: `os_error *drawmod_fill(drawmod_pathelem_ptr path_seq,
drawmod_filltype fill_style, drawmod_transmat *matrix, int
flatness)`

Parameters: `drawmod_pathelem_ptr path_seq` – sequence of path elements
`drawmod_filltype fill_style` – style of fill
`drawmod_transmat *matrix` – transformation matrix (0 for the identity matrix)
`int flatness` – flatness in user coordinates (0 means default).

Returns: possible error condition

drawmod_stroke

Emulates PostScript 'stroke' operator.

Syntax: `os_error *drawmod_stroke(drawmod_pathelem_ptr path_seq,
drawmod_filltype fill_style, drawmod_transmat *matrix,
drawmod_line *line_style)`

Parameters: `drawmod_pathelem_ptr path_seq` – sequence of path elements
`drawmod_filltype fill_style` – style of fill
`drawmod_transmat *matrix` – transformation

`matrix` (0 means identity matrix)
`drawmod_line *line_style` – (see typedef in header file for details).

Returns: possible error condition.

drawmod_do_strokepath

Puts a path through all stages of `drawmod_stroke` except the final fill. The resulting path is placed in the buffer.

Syntax: `os_error *drawmod_do_strokepath(drawmod_pathelem_ptr path_seq, drawmod_transmat *matrix, drawmod_line *line_style, drawmod_buffer *buffer)`

Parameters: `drawmod_pathelem_ptr path_seq` – sequence of path elements
`drawmod_transmat *matrix` – transformation matrix
`drawmod_line *line_style` – see typedef in header file
`drawmod_buffer *buffer` – buffer to hold stroked path.

Returns: possible error condition.

drawmod_ask_strokepath

Puts a path through all stages of `drawmod_stroke`, except the fill, and returns the size of buffer needed to hold such a path.

Syntax: `os_error *drawmod_ask_strokepath(drawmod_pathelem_ptr path_seq, drawmod_transmat *matrix, drawmod_line *line_style, int *buflen)`

Parameters: `drawmod_pathelem_ptr path_seq` – sequence of path elements
`drawmod_transmat *matrix` – transformation matrix
`drawmod_line *line_style` – (see typedef in header for details)
`int *buflen` – returned length of required buffer.

Returns: possible error condition.

drawmod_do_flattenpath

Flattens the given path, and puts it into the supplied buffer.

Syntax: `os_error *drawmod_do_flattenpath(drawmod_pathelem_ptr path_seq, drawmod_buffer *buffer, int flatness)`

Parameters: `drawmod_pathelem_ptr path_seq` – sequence of path elements
`drawmod_buffer *buffer` – buffer to hold flattened path
`int flatness` – required flatness.

Returns: possible error condition.

drawmod_ask_flattenpath

Puts the given path through the stages of `drawmod_flattenpath` and returns the size of buffer needed to hold the resulting path.

Syntax: `os_error *drawmod_ask_flattenpath(drawmod_pathelem_ptr path_seq, int flatness, int *buflen)`

Parameters: `drawmod_pathelem_ptr path_seq` – sequence of path elements
`int flatness` – required flatness
`int *buflen` – returned length of required buffer.

Returns: possible error condition.

drawmod_buf_transformpath

Puts a path through a transformation matrix and puts the result in the supplied buffer.

Syntax: `os_error *drawmod_buf_transformpath(drawmod_pathelem_ptr path_seq, drawmod_buffer *buffer, drawmod_transmat *matrix)`

Parameters: `drawmod_pathelem_ptr path_seq` – sequence of path elements
`drawmod_buffer *buffer` – buffer to hold transformed path
`drawmod_transmat *matrix` – the transformation matrix.

Returns: possible error condition.

drawmod_insitu_transformpath

Puts a path through a transformation matrix by modifying the supplied path itself.

Syntax: `os_error *drawmod_insitu_transformpath(drawmod_pathelem_ptr path_seq, drawmod_transmat *matrix)`

Parameters: `drawmod_pathelemptr path_seq` – sequence of path elements
`drawmod_transmat *matrix` – the transformation matrix.

Returns: possible error condition.

drawmod_processpath

Puts a path through a set of processes used when doing Stroke and Fill.

Syntax: `os_error *drawmod_processpath(drawmod_pathelemptr path_seq, drawmod_filltype fill_style, drawmod_transmat *matrix, drawmod_line *line_style, drawmod_options *options, int *buflen)`

Parameters: `drawmod_pathelemptr path_seq` – sequence of path elements
`drawmod_filltype fill_style` – style of fill
`drawmod_transmat *matrix` – the transformation matrix
`drawmod_line *line_style` – (see typedef in header for details)
`drawmod_options *options` – this can have the values detailed below. Note: pass in address of a `draw_options` struct
`int *buflen` – returned length of required buffer (only used when `options->tagtype == tag_fill` && `options->data.opts == option_countsiz`).

Returns: possible error condition.

Other Information: Possible values for options:

<code>drawmod_insitu</code>	output to the input path (only if path size wouldn't change)
<code>drawmod_fillnormal</code>	fill path normally
<code>drawmod_fillsubpath</code>	fill path, subpath by subpath
OR an address	output bounding box of path to the word-aligned address, and three next words, with word-order lowX, lowY, highX, highY

OR a buffer to hold the processed path.

event

This file handles system-independent central processing for window system events.

event_process

Processes one event.

Syntax: `void event_process(void)`

Parameters: `void`.

Returns: `void`.

Other Information: If the number of current active windows is 0, the program exits. One event is polled and processed (with the exception of some complex menu handling, this really means passing the event on to the `win` module). Unless an application window is claiming idle events, this function waits when the processor is idle. Typically this should be called in a loop in the main function of the application.

event_anywindows

Informs the caller if there are any windows active that can process events.

Syntax: `BOOL event_anywindows(void)`

Parameters: `void`.

Returns: True if there are any active windows.

event_attachmenu

Attaches a menu and its associated handler function to the given window.

Syntax: `BOOL event_attachmenu(event_w, menu, event_menu_proc, void *handle)`

Parameters: `event_w` – the window to which menu should be attached
`menu` – the menu structure
`event_menu_proc` – the handler for the menu
`void *handle` – caller-defined handle.

Returns: True if able to attach menu.

Other Information: The menu should have been created by a call to `menu_new` or something similar. When the user invokes a menu from the given window, this menu will be activated. The handler function will be called when the user selects a menu entry. The handler's parameter `hit` is a string containing a character for each level of nesting in a hierarchical menu structure, terminated by a 0 character. A call with `menu == 0` removes the attachment. To catch menu events on the icon bar, attach a menu to `win_ICONBAR` (defined in the `win` module).

event_attachmenumaker

Attaches to the given window a function which makes a menu when the user invokes a menu.

Syntax: `BOOL event_attachmenumaker(event_w, event_menu_maker, event_menu_proc, void *handle)`

Parameters: `event_w` – the window to which the menu maker should be attached
`event_menu_maker` – the menu maker function
`event_menu_proc` – handler for the menu
`void *handle` – caller-defined handle

Returns: True if able to attach menu maker

Other Information: This works similarly to `event_attachmenu`, except that it allows you to make last minute changes to flags in the menu (such as ticks or fades), before displaying it. A call with `event_menu_maker==0` removes the attachment.

event_clear_current_menu

Clears the current menu tree.

Syntax: `void event_clear_current_menu(void)`

Parameters: `void`.

Returns: `void`.

Other Information: To be used to force all menus to be cleared from the screen.

event_is_menu_being_recreated

Informs the caller if a menu is being recreated.

Syntax: `BOOL event_is_menu_being_recreated(void)`
Parameters: `void`.
Returns: `void`.
Other Information: Useful for when RISC_OSLib is recreating a menu in response to a click on Adjust (call it in a menu maker).

event: masking off events

event_setmask

Sets the mask used by `wimp_poll` and `wimpt_poll` when polling the Wimp.

Syntax: `void event_setmask(wimp_emask mask)`
Parameters: `wimp_emask mask` – the desired mask.
Returns: `void`.
Other Information: Bits of the mask are set if you want the corresponding events ignored (as in the `wimp_poll` SWI). For example, `event_setmask(wimp_ENULL | wimp_EPTRENTER)` will ignore nulls and pointer entering window events. The default mask is to ignore null events only.

event_getmask

Informs the caller of the current mask being used to poll the Wimp.

Syntax: `wimp_emask event_getmask(void)`
Parameters: `void`.
Returns: The mask currently used.

fileicon

Displays an icon representing a file, in a given window.

Syntax: `void fileicon(wimp_w, wimp_i, int filetype)`
Parameters: `wimp_w` – the given window's handle
`wimp_i` – an existing icon
`int filetype` – RISC OS file type (eg `0x0ffe`)
Returns: `void`.

Other Information: If you want a file icon in a dialogue box then pass that dialogue box's window handle through first parameter, eg `fileicon((wimp_w) dbox_syshandle(d), ...)`. The second parameter is the icon number of the required icon, within the template set up using FormEd. For an example see the `fileInfo` template for Edit.

flex

These functions provide memory allocation for interactive programs requiring large chunks of store.

flex_alloc

Allocates `n` bytes of store, obtained from the Wimp.

Syntax: `int flex_alloc(flex_ptr anchor, int n)`

Parameters: `flex_ptr anchor` – to be used to access allocated store
`int n` – number of bytes to be allocated.

Returns: 0 == failure, 1 == success

Other Information: You should pass the `&` of a pointer variable as the first parameter. The allocated store **must** then be accessed indirectly, through this, ie `(*anchor)[0] .. (*anchor)[n]`. This is important because the allocated store may later be moved. If there isn't enough store, returns zero leaving `anchor` unchanged.

flex_free

Frees the previously allocated store.

Syntax: `void flex_free(flex_ptr anchor)`

Parameters: `flex_ptr anchor` – pointer to allocated store.

Returns: `void`.

Other Information: `*anchor` will be set to 0.

flex_size

Informs the caller of the number of bytes allocated.

Syntax: `int flex_size(flex_ptr)`

Parameters: `flex_ptr` – pointer to allocated store
Returns: number of allocated bytes.

flex_extend

Extend or truncate the store area to have a new size.

Syntax: `int flex_extend(flex_ptr, int newsize)`
Parameters: `flex_ptr` – pointer to allocated store
`int newsize` – new size of store
Returns: 0 == failure, 1 == success.

flex_midextend

Extend or truncate store, at any offset.

Syntax: `int flex_midextend(flex_ptr, int at, int by)`
Parameters: `flex_ptr` – pointer to allocated store
`int at` – offset within the allocated store
`int by` – extent.
Returns: 0 == failure, 1 == success.
Other Information: If `by` is +ve, store is extended, and locations above `at` are copied up by `by`. If `by` is -ve, store is reduced, and any bytes beyond `at` are copied down to `at+by`.

flex_budge

Moves the flex store when the C library needs to extend the heap.

Syntax: `int flex_budge(int n, void **a)`
Parameters: `int n` – number of bytes needed by C library.
`void **a` – address of acquired store.
Returns: amount of store acquired.
Other Information: Do not call this function directly, but register it with the C library via:
`_kernel_register_slotextend(flex_budge)`
This will cause flex store to be moved up if the C library needs to extend the heap. Note that in this state, you can rely on pointers into flex blocks across function calls which do not extend the stack and do not call `malloc`.
The default state is `flex_dont_budge`, so, if required, this function should be registered after calling `flex_init()`.

flex_dont_budge

Refuses to move the flex store when the C library needs to extend the heap.

Syntax: `int flex_dont_budge(int n, void **a)`

Parameters: `int n` – number of bytes needed by C library.
`void **a` – address of acquired store.

Returns: amount of store acquired (always 0).

Other Information: Do not call this function directly, but register it with the C library via:

`_kernel_register_slotextend(flex_dont_budge)`

If the C library needs to extend the heap, flex will refuse to move. This means that you can rely on pointers into flex blocks across function calls.

This is the default state after calling `flex_init()`.

flex_init

Initialise store allocation module.

Syntax: `void flex_init(void)`

Parameters: `void`.

Returns: `void`.

Other Information: Must be called before any other functions in this module.

font

These functions provide access to RISC OS font facilities.

font_cacheaddress

Informs the caller of font cache used and font cache size.

Syntax: `os_error * font_cacheaddress(int *version, int *cacheused, int *cachesize)`

Parameters: `int *version` – version number
`int *cacheused` – amount of font cache used (in bytes)
`int *cachesize` – total size of font cache (in bytes).

Returns: Possible error condition

Other Information: Version number is *100, so v.1.07 would be returned as 107.

font_find

Gives the caller a handle to font, given its name.

Syntax: `os_error * font_find(char* name, int xsize, int ysize, int xres, int yres, font*)`

Parameters: `char *name` – the font name
`int xsize, ysize` – x/y point size (in 16ths of a point)
`int xres, yres` – x/y resolution in dots per inch
`font*` – the returned font handle

Returns: Possible error condition.

font_lose

Informs the font manager that a font is no longer needed.

Syntax: `os_error * font_lose(font f)`

Parameters: `font f` – the font.

Returns: possible error condition.

font_readdef

Gets details about a font, given its handle.

Syntax: `os_error * font_readdef(font, font_def*)`

Parameters: `font` – the font handle
`font_def*` – pointer to buffer to hold returned details.

Returns: possible error condition.

Other Information: This function fills in details about a font into the supplied buffer (a variable of type `font_def`). The fields of this buffer are as follows:

<code>name</code>	font name
<code>xsize, ysize</code>	x/y point size * 16
<code>xres, yres</code>	x/y resolution (dots per inch)
<code>usage</code>	number of times <code>Font_FindFont</code> has found the font minus number of times <code>Font_LoseFont</code> has been used on it
<code>age</code>	number of font accesses made since this one was last accessed.

font_readinfo

Informs the caller of the minimal area covering any character in the font bounding box.

Syntax: `os_error * font_readinfo(font, font_info*)`

Parameters: `font` – the font handle
`font_info*` – pointer to buffer to hold returned details.

Returns: possible error condition.

Other Information: Fills in details of the font in the supplied buffer (variable of type `font_info`). The fields of this buffer are as follows:

<code>minx</code>	min x coord in pixels (inclusive)
<code>maxx</code>	max x coord in pixels (inclusive)
<code>miny</code>	min y coord in pixels (exclusive)
<code>maxy</code>	max y coord in pixels (exclusive).

font_strwidth

Determines the width of a string.

Syntax: `os_error * font_strwidth(font_string *fs)`

Parameters: `font_string *fs` – the string, with fields:
`s` string itself
`x` max x offset before termination
`y` max y offset before termination
`split` string split character
`term` index of char to terminate by

Returns: possible error condition.

Other Information: On exit `fs` fields hold:

<code>s</code>	unchanged
<code>x</code>	x offset after printing string
<code>y</code>	y offset after printing string
<code>split</code>	number of split characters found; number of printable characters if <code>split</code> was <code>-1</code>
<code>term</code>	index into string at which terminated.

font_paint

Paints the given string at coordinates x,y.

Syntax: `os_error * font_paint(char*, int options, int x, int y)`

Parameters: `char` – the string
 `int options` – set using 'paint options' defined in the header file
 `int x, y` – coordinates (either OS or 1/72000 inch)

Returns: possible error condition.

font_caret

Sets the colour, size and position of the caret.

Syntax: `os_error *font_caret(int colour, int height, int flags, int x, int y)`

Parameters: `int colour` – EORed onto screen
 `int height` – in OS coordinates
 `int flags` – bit 4 set ==> OS coordinates, otherwise 1/72000 inch
 `int x,y` – x/y coordinates.

Returns: possible error condition.

font_converttoos

Converts coordinates in 1/72000 inch to OS units.

Syntax: `os_error *font_converttoos(int x_inch, int y_inch, int *x_os, int *y_os)`

Parameters: `int x_inch, y_inch` – x/y coordinates in 1/72000 inch
 `int *x_os, *y_os` – x/y coordinates in OS units.

Returns: possible error condition.

font_converttopoints

Converts OS units to 1/72000 inch.

Syntax: `os_error *font_converttopoints(int x_os, int y_os, int *x_inch, int *y_inch)`

Parameters: `int x_os, y_os` – x/y coordinates in OS units
 `int *x_inch, *y_inch` – x/y coordinates in 1/72000 inch.

Returns: possible error condition.

font_setfont

Sets up the font used for subsequent painting or size-requests.

Syntax: `os_error * font_setfont(font)`

Parameters: `font` – the font handle

Returns: possible error condition.

font_current

Informs the caller of the current font state.

Syntax: `os_error *font_current(font_state *f)`

Parameters: `font_state *f` – pointer to buffer to hold font state

Returns: possible error condition.

Other Information: returned buffer (into variable of type `font_state`):

<code>font f</code>	handle of current font
<code>int back_colour</code>	current background colour
<code>int fore_colour</code>	current foreground colour
<code>int offset</code>	foreground colour offset.

font_future

Informs the caller of font characteristics after a future `font_paint`.

Syntax: `os_error *font_future(font_state *f)`

Parameters: `font_state *f` – pointer to buffer to hold font state.

Returns: possible error condition.

Other Information: buffer contents:

<code>font f</code>	handle of font which would be selected
<code>int back_colour</code>	future background colour
<code>int fore_colour</code>	future foreground colour
<code>int offset</code>	foreground colour offset.

font_findcaret

Informs the caller of the nearest point in a string to the caret position.

Syntax: `os_error *font_findcaret(font_string *fs)`

Parameters: `font_string *fs` – the string
 fields: `char *s` – the string itself
`int x,y` – x/y coordinates of caret

Returns: possible error condition.
Other Information: returned fields of fs as in `font_strwidth`.

font_charbbox

Informs the caller of the bounding box of a character in a given font.

Syntax: `os_error * font_charbbox(font, char, int options, font_info*)`

Parameters: `font` – the font handle
`char` – the ASCII character
`int options` – only relevant option if `font_OSCOORDS`
`font_info*` – pointer to buffer to hold font information.

Returns: possible error condition.
Other Information: if OS coordinates are used and font has been scaled, box may be surrounded by area of blank pixels.

font_readscaletfactor

Informs the caller of the x and y scale factors used by the font manager for converting between OS coordinates and 1/72000 inch.

Syntax: `os_error *font_readscaletfactor(int *x, int *y)`

Parameters: `int *x, *y` – returned scale factors.

Returns: possible error condition.

font_setscaletfactor

Sets the scale factors used by the font manager.

Syntax: `os_error *font_setscaletfactor(int x, int y)`

Parameters: `int x, y` – the new scale factors

Returns: possible error condition.

Other Information: scale factors may have been changed by another application; well-behaved applications save and restore scale factors.

font_list

Gives the name of an available font.

Syntax: `os_error * font_list(char*, int*)`

Parameters: `char*` – pointer to buffer to hold font name
`int*` – count of fonts found (0 on first call).

Returns: possible error condition.

Other Information: count is -1 if no more names. Typically used in loop until count == -1.

font_setcolour

Sets the current font (optionally), changes foreground and background colours, and offset for that font.

Syntax: `os_error * font_setcolour(font, int background, int foreground, int offset)`

Parameters: `font` – the font handle (0 for current font)
`int background`, `foreground` – back/foreground colours
`int offset` – foreground offset colour (-14 to +14).

Returns: possible error condition.

font_setpalette

Sets the anti-alias palette.

Syntax: `os_error *font_setpalette(int background, int foreground, int offset, int physical_back, int physical_fore)`

Parameters: `int background` – logical background colour
`int foreground` – logical foreground colour
`int offset` – foreground colour offset
`int physical_back` – physical background colour
`int physical_fore` – physical foreground colour

Returns: possible error condition.

Other Information: `physical_back` and `physical_fore` are of the form 0xBBGGRR00.

font_readthresholds

Reads the list of threshold values that the font manager uses when painting characters.

Syntax: `os_error *font_readthresholds(font_threshold *th)`

Parameters: `font_threshold *th` – pointer to result buffer.

Returns: possible error condition.

font_setthresholds

Sets up threshold values for painting colours.

Syntax: `os_error *font_setthresholds(font_threshold *th)`

Parameters: `font_threshold *th` – pointer to a threshold table.

Returns: possible error condition.

font_findcaretj

Finds the nearest point where the caret can go (using justification offsets).

Syntax: `os_error *font_findcaretj(font_string *fs, int offset_x, int offset_y)`

Parameters: `font_string *fs` – the string (set up as in `font_findcaret`)
`int offset_x, offset_y` – the justification offsets.

Returns: possible error condition.

Other Information: If the offsets are both zero, the function is the same as `font_findcaret`.

font_stringbbox

Measures the size of a string (without printing it).

Syntax: `os_error *font_stringbbox(char *s, font_info *fi)`

Parameters: `char *s` – the string
`font_info *fi` – pointer to buffer to hold font information.


Returns: possible error condition.

Other Information: fields returned in `fi` are:
`minx, miny` bounding box min x/y
`maxx, maxy` bounding box min x/y.

fontlist

These functions count the fonts on the system into a tree structure.

As an example of a font tree structure, consider a system providing Corpus.Medium, Corpus.Bold, Selwyn, Trinity.Medium, Trinity.Bold, Trinity.Medium.Italic and Widget.Medium.Italic.Outline. This will be stored in the following way (#'s denote flags which are TRUE):



Brothers are connected vertically, sons to their parents right-to-left

fontlist_list_all_fonts

Read in the font list into a font tree

Syntax: `fontlist_node *fontlist_list_all_fonts(BOOL system);`
 Parameters: `BOOL system` – TRUE if System font should be included in the list
 Returns: a pointer to the start of the font tree
 Other Information: None

fontlist_free_font_tree

Free a font tree

Syntax: `void fontlist_free_font_tree(fontlist_node *font_tree);`
 Parameters: `fontlist_node *font_tree` – the font tree to free
 Returns: None
 Other Information: None

fontselect

These functions provide an interface to font choosing

fontselect_init

Read in the font list and prepare data for the font selector window

Syntax: `int fontselect_init(void);`
 Parameters: None
 Returns: TRUE if initialisation succeeded.
 Other Information: None

fontselect_closedown

Close the font selector windows aif they are open, and free the font selector data structures

Syntax: `void fontselect_closedown(void);`
 Parameters: None
 Returns: None
 Other Information: This call is provided to return the machine to the state it was in before a call of `fontselect_init()`

fontselect_closewindows

Close the font selector windows if they are open

Syntax: `void fontselect_closewindows(void);`

Parameters: None

Returns: None

Other Information: This call will close the font selector windows and unattach the handlers, if they are open.

fontselect_selector

Opens up or reopens the font chooser window.

Syntax: `int fontselect_selector(char *title, int flags, char *font_name, double width, double height, fontselect_fn unknown_icon_routine);`

Parameters: `char *title` – The title for the window (can be NULL if flags SETTITLE is clear)

`int flags` – The flags for the call

`char *font_name` – The font name to set the window contents to (only if SETFONT is set)

`double width` – The width in point size of the font
`double height` `fontselect_fn` `unknown_icon`

Returns: The window handle of the font selector main window, if the function call was successful. Otherwise it returns 0.

Other Information: The flags word allows the call to have different effects. If fontselect_SETFONT is set then the window contents will be updated to reflect the font choice passed in. If fontselect_SETTITLE is set then the title of the window will be set, otherwise title is ignored. If fontselect_REOPEN is set then the font selector will only open the window if it is already open. This lets the application update the contents of the window only if it is currently open. Note that the fontselect_init() must be called before this routine.

fontselect_attach_menu

Attaches a menu to all four font selector windows

Syntax: `BOOL fontselect_attach_menu(menu mn, event_menu_proc menu_processor, void *handle);`

Parameters: `menu mn` – menu to attach

`event_menu_proc menu_processor` – menu processor for the menu events

void *handle – handle to pass to the menu processor
 Returns: TRUE if the menus were attached, FALSE otherwise
 Other Information: None

heap

These functions provide malloc-style heap allocation within a flex block of memory. They should only be used when flex is set up so that it cannot be moved by malloc expansion (the default). See the later chapter entitled *Using memory efficiently* for more details.

heap_init

Initialises the heap allocation system.

Syntax: void heap_init(BOOL heap_shrink)
 Parameters: BOOL heap_shrink – if True, the flex block will be shrunk (when possible) after heap_free().
 Returns: void.
 Other Information: You must call flex_init before calling this routine.

heap_alloc

Allocates a block of storage from the heap.

Syntax: void *heap_alloc(unsigned int size)
 Parameters: unsigned int size – size of block to be allocated.
 Returns: pointer to allocated block (or 0 if failed).
 Other Information: None

heap_free

Frees a previously allocated block of heap storage.

Syntax: void heap_free(void *heapptr)
 Parameters: void *heapptr – pointer to block to be freed.
 Returns: possible error condition.

help

These functions provide support for interactive help, including on menu entries when this is supported by versions of the Wimp following 2.00.

help_process

Returns TRUE if the given event is a menu interactive help message, which has now been processed.

Syntax: `BOOL help_process(wimp_eventstr *e);`

Parameters: `e` – the event to be considered.

Returns: TRUE if the event has now been processed.

Other Information: This should be called by the unknown event handler of the program. For it to work, you must inform wimpt that you are aware of Wimps beyond version 2.00. The rest of this interface for the handling facilities that this call invokes.

help_register_handler

Record the handler to be used when `help_process` is next called.

Syntax: `void help_register_handler(event_menu_proc, void *handle);`

Parameters: `event_menu_proc` – the handler procedure
`handle` – a data handle for the handler procedure

Returns: void

Other Information: This should be called by the menu-maker proc of every menu in the program. When `help_process` is called, the most recently installed `event_menu_proc` handler is assumed to be the correct one. Call with NULL as a proc if no help is available on this menu.

help_genmessage

From a given menu hit and prefix, generate a message tag which is looked up in `msgs_lookup`. If a message is found then return it as the interactive help message, and return TRUE. If not, return FALSE.

Syntax: `BOOL help_genmessage(char *prefix, char *hit);`

Parameters: `prefix` – the prefix for all message tags used
`hit` – the hit string handed to the `event_menu_proc` registered using `help_register_handler`

Returns: TRUE if this was a menu help message which has now been handled

Other Information: The tag for the `msgs_lookup` call is generated by:

Start with the prefix (maximum length – 20 characters) append '0'..'9' or 'a'..'z' for each character in the hit (eg character 1 counts as '0', character 10 counts as '9', character 11 counts as 'a', etc. More than 35 seems unlikely in a menu of fixed size).

Example:

Original hit 0,1,2 gets translated to string "\x001\x002\x003", which gets turned into tag "FOO012" for prefix "FOO". A prefix consisting entirely of upper case letters is conventional. If there's only one menu tree, the prefix "HELP" is conventional. For the icon bar, "IHELP" is conventional.

help_simplehandler

A simple event_menu_proc suitable for giving to a help handler. The implementation is simply { help_genmessage((char*) handle, hit); }

Syntax:

```
void help_simplehandler(void *handle, char *hit);
```

Parameters:

handle – prefix to pass to help_genmessage
hit – the menu hit string being processed

Returns:

void

Other Information:

This will suffice for cases where there are no alternatives or additional cases to consider. For menus where the message can vary at run time, a more complex handler will be required which parses the hit string, or which calls help_genmessage more than once, or perhaps a combination of the two.

help_dboxrawevents

A routine suitable for passing to dbox_raw_eventhandler, for providing help on dialogue boxes.

Syntax:

```
BOOL help_dboxrawevents(dbox, void *event, void *handle);
```

Parameters:

dbox – the dbox for which help is being provided
event – the wimp event being processed
handle – message tag prefix, really a char*

Returns:

TRUE if this was a menu help message which has now been handled

Other Information:

The handle passed to it should be a message prefix. A single character suffix may be added to it in the style of help_genmessage, containing the icon number. If this is

not found (or if no icon is being pointed to), then the prefix alone will be used as a message tag. There is no error if the message is not found.

Typical use:

```
dbox d = dbox_new("foo");
dbox_raw_eventhandler(d, help_dboxrawevents, "FOO");
dbox_show(d);
... now fill in the dbox as normal ...
```

The test used in implementing this is:

```
if (
    (e->e == wimp_ESEND || e->e == wimp_ESENDWANTACK)
    &&
    e->data.msg.hdr.action == wimp_MHELPREQUEST
)
```

{ ... construct a reply, and send it using help_reply ... }

The Messages file should contain:

FOO: This message will appear in the dbox background.

FOO0: This message will appear for icon 0 of the dbox etc.

help_reply

Reply to the help message in `wimpt_last_event()` with the (already translated) message provided.

Syntax:

```
void help_reply(char *m);
```

Parameters:

m – help message to display in interactive help window

Returns:

```
void
```

Other Information:

This is useful when creating your own versions of `help_dboxrawevents`, which must also handle other events.

magnify

This function allows the display and entry of magnification factors.

magnify_select

Displays a dialogue box to set magnification factors.

Syntax:

```
void magnify_select (int *mul, int *div, int maxmul, int
maxdiv, void (*proc)(void *), void *phandle)
```

Parameters:	<p><code>int *mul, *div</code> – multiplication/division factors <code>int maxmul, maxdiv</code> – maximum mult/div factors <code>void(*proc)(void *)</code> – caller-supplied function <code>void *phandle</code> – handle passed to user function.</p>
Returns:	<code>void</code> .
Other Information:	Displays a template called 'magnifier' (which must be one of your loaded templates). <code>mul</code> and <code>div</code> are the initial values on the left and right of the <code>:</code> in the ratio shown in the dialogue box. They are modified according to user mouse clicks on the arrow icons. <code>proc</code> (if non-null) is called each time the magnification factor changes.

The template should have the following attributes:

- window flags – moveable, auto-redraw. It is advisable to have a title icon with the text `magnifier` or similar.
- icon #0 – the multiplication factor icon. This should have an indirected text flag set with text something like `999` and a maximum length of 4. It is also advisable to have a validation string `a0-9` (allowing numeric input). The button type should be 'writeable'.
- icon #1 – the division factor icon (same as icon #0)
- icon #2 – the increase multiplication factor icon should have its text flag set and contain the `↑` character (like the arrow used in scroll bars). The button type should be 'auto-repeat'.
- icon #3 – the decrease multiplication factor icon (same as icon #2, but using the `↓` char).
- icon #4 – the increase division factor icon (same as icon #2).
- icon #5 – the decrease division factor icon (same as icon #3).
- icon #6 – (optional but advisable) a text icon placed between icons #0 and #1 as a separator eg `:`

These icons can be arranged in the window however you wish, but a recommended layout is that of the Magnifier dialogue box in Draw or Paint.

menu

These functions deal with the creation, deletion and manipulation of menus.

A menu description string defines a sequence of entries, with the following syntax (curly brackets mean 0 or more, square brackets mean 0 or 1):

```

opt ::= ! or ~ or > or space
sep ::= , or |
l1 ::= any char but opt or sep
l2 ::= any char but sep
name ::= l1 {l2}
entry ::= {opt} name
descr ::= entry {sep entry}

```

Each entry defines a single entry in the menu. | as a separator means that there should be a gap or line between these menu components.

```

opt ! means 'put a tick by it'
opt ~ means 'make it non-selectable'
opt > means 'has a dialogue box as 'submenu''
space has no effect as an opt.

```

menu_new

Creates a new menu structure from the given textual description (arranged as above).

```

Syntax:          menu menu_new(char *name, char *description)
Parameters:     char *name - name to appear in title of menu
                char *description - textual description of menu
Returns:        pointer to menu structure created
Other Information:  Creates a menu structure, with entries as given in the
                    textual description. Entries are indexed from 1. For
                    example:

```

```
m=menu_new("Edit", ">Info Create Quit")
```

Handler needs to be attached using `event_attachmenu`.

menu_dispose

Disposes of a menu structure.

```

Syntax:          void menu_dispose(menu*, int recursive)
Parameters:     menu* - the menu to be disposed of
                int recursive - non-zero ==recursively dispose of
                    submenus.
Returns:        void.

```

menu_extend

Adds entries to the end of a menu.

```
Syntax:          void menu_extend(menu, char *description)
```

Parameters: menu – the menu to which extension is being made
char *description – textual description of extension.

Returns: void.

Other Information: extension has the format:
[sep] entry {sep entry}

A menu which is already a submenu of another menu cannot be extended.

menu_setflags

Sets or changes flags on an already existing menu entry.

Syntax: void menu_setflags(menu, int entry, int tick, int fade)

Parameters: menu – the menu
int entry – index into menu entries (from 1)
int tick – non-zero == tick this entry
int fade – non-zero == fade this entry (ie make it unselectable).

Returns: void.

menu_submenu

Attaches a menu as a submenu of another at a given entry in the parent menu.

Syntax: void menu_submenu(menu, int entry, menu submenu)

Parameters: menu – the menu
int entry – entry at which to attach submenu
menu submenu – pointer to the submenu.

Returns: void.

Other Information: This replaces any previous submenu at this entry. Use 0 for submenu to remove an existing entry. Only a strict hierarchy is allowed. When attached as a submenu, a menu can't be extended or explicitly deleted.

menu_make_writeable

Makes a menu entry writeable.

Syntax: void menu_make_writeable(menu m, int entry, char *buffer, int bufferlength, char *validstring)

Parameters: menu m – the menu
int entry – the entry to make writeable
char *buffer – pointer to buffer to hold text of entry
int bufferlength – size of buffer
char *validstring – pointer to validation string

Returns: void.
 Other Information: The lifetimes of `buffer` and `validstring` must be long enough.

menu_make_sprite

Makes a menu entry into a sprite.

Syntax: `void menu_make_sprite(menu m, int entry, char *spritename)`

Parameters: `menu m` – the menu
`int entry` – entry to be made into sprite
`char *spritename` – name of the sprite.

Returns: void.

Other Information: Entry which is initially a non-indirected text entry is changed to an indirected sprite, with sprite area given by `resspr_area()`, and name given by `spritename`.

menu_syshandle

Gives low-level handle to a menu.

Syntax: `void *menu_syshandle(menu)`

Parameters: `menu` – the menu

Returns: pointer to underlying Wimp menu structure.

Other Information: Allows the massaging of a menu by means other than those provided in this module. The returned pointer is in fact a pointer to a `wimp_menustr` (ie `wimp_menuhdr` followed by zero or more `wimp_menuitems`).

msgs

These functions provide support for message resource files. Use them to make your applications easily convertible to other natural languages. A messages file for RISC_OSLib error messages is provided; it is not needed if you just want English messages, since these are the defaults.

msgs_init

Reads in the messages file, and initialise message system.

Syntax: `void msgs_init(void)`

Parameters: `void`

Returns: `void`.

Other Information: The messages file is a resource of your application and should be named `messages`. Each line of this file is a message with the following format:

```
<tag><colon><message text><newline>
```

The tag is an alphanumeric identifier for the message, which will be used to search for the message, when using `msgs_lookup()`. It has a maximum length of 9 characters.

msgs_lookup

Finds the text message associated with a given tag.

Syntax: `char *msgs_lookup(char *tag_and_default)`

Parameters: `char *tag_and_default` – the tag of the message, and an optional default message (to be used if tagged message not found).

Returns: pointer to the message text (if all is well).

Other Information: If the caller just supplies a tag, he will receive a pointer to its associated message (if found). A default message can be given after the tag (separated by a colon). A typical use would be:

```
werr(1, msgs_lookup("error1"))
or
werr(1, msgs_lookup("error1:Not enough memory")).
```

msgs_readfile

Read in the messages file, and initialise msg system

Syntax: `void msgs_readfile(char *name);`

Parameters: `char *name` – the name of the messages file to be read.

Returns: `void`.

Other Information: the messages file is a resource of your application and should normally be named "Messages". For non-standard applications this call has been provided to enable the messages file to be read from elsewhere. Each line of this file is a message with the following format:

```
<tag><colon><message text><newline>
```

The tag is an alphanumeric identifier for the message, which will be used to search for the message, when using `msgs_lookup()`.

OS

This file is provided as an alternative to `kernel.h`. It provides low-level access to RISC OS. `os_error` functions return a pointer to an error if one has occurred, otherwise return `NULL (0)`.

os_swi

Performs the given SWI instruction, with the given registers loaded. An error results in a RISC OS error being raised. A `NULL` `regset` pointer means that no inout parameters are used.

Syntax: `void os_swi(int swicode, os_regset *regs)`

os_swix

Performs the given SWI instruction, with the given registers loaded. Calls returning `os_error*` use the X form of the relevant SWI. If an error is returned then the `os_error` should be copied before further system calls are made. If no error occurs then `NULL` is returned.

Syntax: `os_error *os_swix(int swicode, os_regset *regs)`

If `swicode` does not have the X bit set, `os_swi` is called and these functions return `NULL` (regardless of whether an error was raised). You should therefore use X bit set `swicodes` to save confusion.

SWIs with varying numbers of arguments and results:

`NULL` result pointers mean that the result from that register is not required. The `swi` codes can be of the X form if required, as specified by `swicode`.

```
os_error *os_swi0(int swicode); /* zero arguments and results */
os_error *os_swi1(int swicode, int r0)
os_error *os_swi2(int swicode, int r0, int r1)
os_error *os_swi3(int swicode, int r0, int r1, int r2)
os_error *os_swi4(int swicode, int r0, int r1, int r2, int r3)
os_error *os_swi6(int swicode, int r0, int r1, int r2, int r3, int r4, int r5)
os_error *os_swilr(int swicode, int r0in, int *r0out)
os_error *os_swi2r(int swicode, int r0in, int r1in, int *r0out, int *r1out)
os_error *os_swi3r(int swicode, int, int, int, int*, int*, int*)
os_error *os_swi4r(int swicode, int, int, int, int, int*, int*, int*, int*)
os_error *os_swi6r(int swicode,
int r0, int r1, int r2, int r3, int r4, int r5,
int *r0out, int *r1out, int *r2out, int *r3out, int *r4out, int *r5out)
```

os_byte

Performs an `OS_Byte` SWIx, with `x` and `y` passed in register `r1` and `r2` respectively.

Syntax: `os_error *os_byte(int a, int *x /*inout*/, int *y /*inout*/)`

os_word

Performs an OS_Word SWI, with operation number given in wordcode and p pointing at necessary parameters to be passed in r1.

Syntax: `os_error *os_word(int wordcode, void *p)`

os_gpbp

Performs an OS_GBPB SWI. `os_gpbpstr` should be used like an `os_regset`.

Syntax: `os_error *os_gpbp(os_gpbpstr*)`

os_file

Performs an OS_FILE SWI.

Syntax: `os_error *os_file(os_filestr*)`

os_args

Performs an OS_Args SWI.

Syntax: `os_error *os_args(os_regset*)`

os_find

Performs an OS_Find SWI.

Syntax: `os_error *os_find(os_regset*)`

os_cli

Performs an OS_CLi SWI.

Syntax: `os_error *os_cli(char *cmd)`

os_read_var_val

Reads a named environment variable into a given buffer (of size `bufsize`). If the variable doesn't exist, `buf` points at a null string.

`os_read_var_val(char *name, char *buf /*out*/, int bufsize)`

pointer

These functions deal with setting the pointer shape.

pointer_set_shape

Sets pointer shape 2, to sprite, from sprite area.

Syntax:	<code>os_error *pointer_set_shape(sprite_area *, sprite_id *, int, int)</code>
Parameters:	<code>sprite_area*</code> – area where sprite is to be found <code>sprite_id*</code> – identity of the sprite <code>int, int</code> – active point for pointer.
Returns:	possible error condition.
Other Information:	A typical use is to change pointer shape on entering or leaving application window (appropriate events are returned from <code>wimp_poll</code>).

pointer_reset_shape

Resets pointer shape to shape 1.

Syntax:	<code>void pointer_reset_shape(void)</code>
Parameters:	<code>void</code> .
Returns:	<code>void</code> .
Other Information:	Typically should be called when leaving an application window.

print

These functions provide access to printer driver facilities. The descriptions here are only brief. For more details, see the *Printer Drivers* chapter in the RISC OS *Programmer's Reference manual*.

Several enumerations and structures are defined in the `print` header. These correspond closely to the data passed to or returned by the printer driver SWIs documented in the above mentioned *Programmer's Reference Manual* chapter:

print_identity

```
typedef enum
{
    print_PostScript      = 0,
    print_FX80compatible  = 1
} print_identity;
```

print_features

```
typedef enum
{
    print_colour          = 0x0000001, - colour
```

```

print_limited      = 0x00000002, - if print_COLOUR bit set, full colour
                    range not available
print_discrete    = 0x00000004, - only a discrete colour set supported
print_NOFILL      = 0x00000100, - cannot handle filled shapes well
print_NOTHICKNESS = 0x00000200, - cannot handle thick lines well
print_NOOVERWRITE = 0x00000400, - cannot overwrite colours properly
print_SCREEN_DUMP = 0x10000000, - supports PDriver_ScreenDump
print_TRANSFORM   = 0x20000000 - supports arbitrary transformations
                    (else only axis-preserving ones).
} print_features;

```

print_infostr

```

typedef struct print_infostr
{
    short int version;      - version number *100
    short int identity;    - driver identity (eg 0=Postscript,1=FX80)
    int xres, yres;        - x, y resolution (pixels/inch)
    int features;          - see print_features
    char *description;     - printers supported, <=20chars + null
    int xhalf, yhalf;      - halftone resolution (repeats/inch)
    int number;            - configured printer number
} print_infostr;

```

print_box

```

typedef struct
{
    int x0, y0, x1, y1;
} print_box;

```

print_pagesizestr

```

typedef struct print_pagesizestr
{
    int xsize, ysize;      - size of page, including margins (1/72000 inch)
}

```

```
    print_box bbox;           – bounding box of printable portion (1/72000  
                               inch)  
} print_pagesizestr;
```

print_transmatstr

```
typedef struct print_transmatstr  
{ int xx, xy, yx, yy;  
} print_transmatstr;
```

print_positionstr

```
typedef struct print_positionstr  
{ int dx, dy;  
} print_positionstr;
```

print_info

Read details of current printer driver (version, resolution, features etc).

Syntax: `os_error * print_info(print_infostr*);`
Parameters: Pointer to `print_infostr` structure to be filled in
Returns: Any error returned from the system call

print_setinfo

Reconfigure current printer driver.

Syntax: `os_error * print_setinfo(print_infostr *i);`
Parameters: Pointer to the `print_infostr` structure to be used to
update the printer driver configuration. The version,
identity and description fields are not used. Leave bit 0
clear in the features field for monochrome, set bit 0 for
colour.
Returns: Any error returned from the system call

print_checkfeatures

Checks the features of a printer, returning an error if the current printer does not
have the specified features.

Syntax: `os_error * print_checkfeatures(int mask, int value);`

Parameters: `int mask` – set bits correspond to the features of interest (bits as `print_features`)

Returns: `int value` – required values of the bits of interest
 Error returned from system call if the printer does not have the specified features.

print_pagesize

Find how large paper and print area is.

Syntax: `os_error * print_pagesize(print_pagesizestr*)`;
 Parameters: Pointer to the `print_pagesizestr` structure to be filled in.
 Returns: Any error returned from the system call.

print_setpagesize

Set how large paper and print size is.

Syntax: `os_error * print_setpagesize(print_pagesizestr *p)`;
 Parameters: Pointer to the `print_pagesizestr` structure to be used to update the printer driver.
 Returns: Any error returned from the system call.

print_selectjob

Make a given print job the current one.

Syntax: `os_error * print_selectjob(int job, char *title, int *oldjob)`;
 Parameters: `int job` – file handle for selected job, or 0 to leave no print job selected
`char *title` – title string for job
`int *oldjob` – pointer to integer to fill in with file handle of previously active job
 Returns: Any error returned from the system call.

print_currentjob

Get the file handle of the current print job.

Syntax: `os_error * print_currentjob(int *curjob)`;
 Parameters: / Pointer to integer to be filled in with the file handle of the current print job.
 Returns: Any error returned from the system call.

print_endjob

End a print job normally.

Syntax: `os_error * print_endjob(int job);`
Parameters: File handle of print job to be ended.
Returns: Any error returned from the system call.

print_abortjob

End a print job without any further output.

Syntax: `os_error * print_abortjob(int job);`
Parameters: File handle of print job to be aborted.
Returns: Any error returned from the system call.

print_canceljob

Stops a specified print job from printing.

Syntax: `os_error * print_canceljob(int job);`
Parameters: File handle of print job to be cancelled.
Returns: Any error returned from the system call.

print_reset

Abort all print jobs.

Syntax: `os_error * print_reset(void);`
Parameters: Void
Returns: Any error returned from the system call.

print_selectillustration

Makes the specified print job the current one, and treats it as an illustration. The difference with `print_selectjob` is that an error is generated if the job does not contain one page, and certain printer drivers (such as the PostScript printer driver) generate different output for illustrations.

Syntax: `os_error * print_selectillustration(int job, char *title, int *oldjobp);`
Parameters: `int job` – file handle for selected job, or 0 to leave no print job selected
`char *title` – title string for job
`int *oldjobp` – pointer to integer to fill in with file handle of previously active job
Returns: Any error returned from the system call.

print_giverectangle

Specify a rectangle to be printed.

Syntax: `os_error *print_giverectangle(int ident, print_box*, print_transmatstr*, print_positionstr*, int bgcol);`

Parameters:

- `ident` – rectangle identification word
- Pointer to structure specifying rectangle to be plotted (OS coordinates)
- Pointer to structure specifying transformation matrix (fixed point, 16 binary places)
- Pointer to structure containing the position of bottom left of rectangle on page (1/72000 inch)
- `bgcol` – background colour for this rectangle, &BBGGRRXX

Returns: Any error returned from the system call.

print_drawpage

This should be called after specifying all rectangles to be plotted on the current page with `print_giverectangle`.

Syntax: `os_error * print_drawpage(int copies, int sequ, char *page, print_box *clip, int *more, int *ident);`

Parameters:

- `copies` – number of copies
- `sequ` – zero or pages sequence number within document
- `page` – zero or a string containing a textual page number (no spaces)
- Pointer to structure to be filled in with the rectangle to print */
- `more` – pointer to integer to be filled in with the number of copies left to print
- `ident` – pointer to integer to be filled in with the rectangle identification word

Returns: Any error returned from the system call.

print_getrectangle

Get the next print rectangle.

Syntax: `os_error * print_getrectangle(print_box *clip, int *more, int *ident);`

Parameters:

- Pointer to the structure to be filled in with the clip rectangle

`more` – pointer to integer to be filled in with the number of rectangles left to print

`ident` – pointer to integer to be filled in with the rectangle identification word

Returns: Any error returned from the system call.

print_screendump

Output a screen dump to the printer.

Syntax: `os_error * print_screendump(int job);`

Parameters: File handle of file to receive the dump.

Returns: Any error returned from the system call.

res

These functions provide access to resources.

res_init

Initialises, ready for calling other `res` functions.

Syntax: `void res_init(const char *programe)`

Parameters: `const char *a` – your program name.

Returns: `void`.

Other Information: Call this before using any `res` or `resspr` functions.

res_findname

Creates a full pathname for a `resname` file.

Syntax: `int res_findname(const char *resname, char *buf /*out*/)`

Parameters: `const char *resname` – name of one of your resource files

`char *buf` – buffer to put full pathname in.

Returns: True (always).

Other Information: the full pathname is constructed as:
`<ProgramName$Dir>.resname` where
`ProgramName` has been set using `res_init`.

res_openfile

Opens a named resource file, in a given ANSI-style mode.

Syntax: `FILE *res_openfile(const char *resname, const char *mode)`

Parameters: `const char *resname` – name of the resource file
`const char *mode` – usual ANSI open mode (r, w, etc)

Returns: ANSI FILE pointer for opened file.

Other Information: `resname` should be a 'leafname' (a call to `res_findname` is made for you).

resspr

These functions provide access to sprite resources.

resspr_init

Initialises, ready for calls to `resspr` functions.

Syntax: `void resspr_init(void)`

Parameters: `void`

Returns: `void`.

Other Information: call before using any `resspr` functions and before using `template_init()`, if your templates have sprites. This function reads in your sprites.

resspr_area

Returns a pointer to the sprite area being used.

Syntax: `sprite_area *resspr_area(void)`

Parameters: `void`

Returns: pointer to sprite area being used.

Other Information: Useful for passing parameters to functions like `baricon` which expect to be told sprite area to use.

saveas

These functions handle the export of data by dragging the icon from the dialogue box.

saveas

Displays a dialogue box to enable the user to export application data.

Syntax: `BOOL saveas(int filetype, char *name, int estsize, xfersend_saveproc, xfersend_sendproc, xfersend_printproc, void *handle)`

Parameters: `int filetype` – type of file to save to
`char *name` – suggested file name
`int estsize` – estimated size of the file

`xfersend_saveproc` – caller-supplied function for saving application data to a file
`xfersend_sendproc` – caller-supplied function for RAM data transfer (if application is able to do this)
`xfersend_printproc` – caller-supplied function for printing application data, if Save icon is dragged onto printer icon
`void *handle` – handle to be passed to handler functions.

Returns: True if data exported successfully.

Other Information: This function displays a dialogue box with the following fields:

- a sprite icon appropriate to the given file type
- the suggested filename
- an OK button.

A template called `xfer_send` must be in the application's templates file to use this function, set up as in the Edit, Draw and Paint applications). `xfer_send` deals with the complexities of message-passing protocols to achieve the data transfer. Refer to the typedefs in `xfersend.h` for an explanation of what the three caller-supplied functions should do. If you pass 0 as the `xfersend_sendproc`, no in-core data transfer will be attempted. If you pass 0 as the `xfersend_printproc`, the file format for printing is assumed to be the same as for saving. The estimated file size is not essential, but may improve performance.

saveas_read_leafname_during_send

Gets the 'leaf' of the filename in the `filename` field of the `xfer-send` dialogue box.

Syntax: `void saveas_read_leafname_during_send(char *name, int length)`

Parameters: `char *name` – buffer to put filename in
`int length` – size in bytes of supplied buffer.

Returns: `void`.

sprite

These functions provide access to RISC OS sprite facilities. Only a brief description is given for each call. More details can be found in the RISC OS *Programmer's Reference manual*, in the chapter entitled *Sprites*.

sprite: simple operations

sprite_screensave

Saves the current graphics window as a sprite file, with optional palette (equivalent to *ScreenSave).

Syntax: `os_error *sprite_screensave(const char *filename, sprite_palflag)`

sprite_screenload

Load a sprite file onto the screen (equivalent to *ScreenLoad).

Syntax: `os_error *sprite_screenload(const char *filename)`

sprite: operations on system/user area

sprite_area_initialise

Initialises an area of memory as a sprite area.

Syntax: `void sprite_area_initialise(sprite_area *, int size)`

sprite_area_readinfo

Reads information from a sprite area control block.

Syntax: `os_error *sprite_area_readinfo(sprite_area *, sprite_area *resultarea)`

sprite_area_reinit

Reinitialises a sprite area. If the sprite area is a system area, the function is equivalent to *SNew.

Syntax: `os_error *sprite_area_reinit(sprite_area *)`

sprite_area_load

Loads a sprite file into a sprite area. If the file is a system area, the function is equivalent to *SLoad.

Syntax: `os_error *sprite_area_load(sprite_area *, const char /*filename)`

sprite_area_merge

Merges a sprite file with a sprite area. If the file is a system area, the function is equivalent to *SMerge.

Syntax: `os_error *sprite_area_merge(sprite_area *, const char *filename)`

sprite_area_save

Saves a sprite area as a sprite file. If the sprite area is a system area, the function is equivalent to *SSave.

Syntax: `os_error *sprite_area_save(sprite_area *, const char *filename)`

sprite_getname

Returns the name and length of the nth sprite in a sprite area into a buffer.

Syntax: `os_error *sprite_getname(sprite_area *, void *buffer, int *length, int index)`

sprite_get

Copies a rectangle of screen delimited by the last pair of graphics cursor positions as a named sprite in a sprite area, optionally storing the palette with the sprite.

Syntax: `os_error *sprite_get(sprite_area *, char *name, sprite_palflag)`

sprite_get_rp

Copies a rectangle of screen delimited by the last pair of graphics cursor positions as a named sprite in a sprite area, optionally storing the palette with the sprite. The address of the sprite is returned in `resultaddress`.

Syntax: `os_error *sprite_get_rp(sprite_area *, char *name, sprite_palflag, sprite_ptr *resultaddress)`

sprite_get_given

Copies a rectangle of screen delimited by the given pair of graphics coordinates as a named sprite in a sprite area, optionally storing the palette with the sprite.

Syntax: `os_error *sprite_get_given(sprite_area *, char *name, sprite_palflag, int x0, int y0, int x1, int y1)`

sprite_get_given_rp

Copies a rectangle of screen delimited by the given pair of graphics coordinates as a named sprite in a sprite area, optionally storing the palette with the sprite. The address of the sprite is returned in `resultaddress`.

Syntax: `os_error *sprite_get_given_rp(sprite_area *, char *name, sprite_palflag, int x0, int y0, int x1, int y1, sprite_ptr *resultaddress)`

sprite_create

Creates a named sprite in a sprite area of specified size and screen mode, optionally reserving space for palette data with the sprite.

Syntax: `os_error *sprite_create(sprite_area *, char *name, sprite_palflag, int width, int height, int mode)`

sprite_create_rp

Creates a named sprite in a sprite area of specified size and screen mode, optionally reserving space for palette data with the sprite. The address of the sprite is returned in `resultaddress`.

Syntax: `os_error *sprite_create_rp(sprite_area *, char *name, sprite_palflag, int width, int height, int mode, sprite_ptr *resultaddress)`

sprite: operations on system/user area, name/sprite pointer**sprite_select**

Selects the specified sprite for plotting using `plot (0xed, x, y)`.

Syntax: `os_error *sprite_select(sprite_area *, sprite_id *)`

sprite_select_rp

Selects the specified sprite for plotting using `plot (0xed, x, y)`. The address of the sprite is returned in `resultaddress`.

Syntax: `os_error *sprite_select_rp(sprite_area *, sprite_id *, sprite_ptr *resultaddress)`

sprite_delete

Deletes the specified sprite.

Syntax: `os_error *sprite_delete(sprite_area *, sprite_id *)`

sprite_rename

Renames the specified sprite within the same sprite area.

Syntax: `os_error *sprite_rename(sprite_area *, sprite_id *, char *newname)`

sprite_copy

Copies the specified sprite as another named sprite in the same sprite area.

Syntax: `os_error *sprite_copy(sprite_area *, sprite_id *, char *copyname)`

sprite_put

Plots the specified sprite using the given GCOL action.

Syntax: `os_error *sprite_put(sprite_area *, sprite_id *, int gcol)`

sprite_put_given

Plots the specified sprite at (x,y) using the given GCOL action.

Syntax: `os_error *sprite_put_given(sprite_area *, sprite_id *, int gcol, int x, int y)`

sprite_put_scaled

Plots the specified sprite at (x,y) using the given GCOL action, and scaled using the given scale factors.

Syntax: `os_error *sprite_put_scaled(sprite_area *, sprite_id *, int gcol, int x, int y, sprite_factors *factors, sprite_pixtrans pixtrans[])`

sprite_put_greyscaled

Plots the specified sprite at (x,y) using the given GCOL action, greyscaled using the given scale factors.

Syntax: `os_error *sprite_put_greyscaled(sprite_area *, sprite_id *, int x, int y, sprite_factors *factors, sprite_pixtrans pixtrans[])`

sprite_put_mask

Plots the specified sprite mask in the background colour.

Syntax: `os_error *sprite_put_mask(sprite_area *, sprite_id *)`

sprite_put_mask_given

Plots the specified sprite mask at (x,y) in the background colour.

Syntax: `os_error *sprite_put_mask_given(sprite_area *, sprite_id *, int x, int y)`

sprite_put_mask_scaled

Plots the sprite mask at (x,y) scaled, using the background colour/action.

Syntax: `os_error *sprite_put_mask_scaled(sprite_area *, sprite_id *, int x, int y, sprite_factors *factors)`

sprite_put_char_scaled

Paints char scaled at (x,y).

Syntax: `os_error *sprite_put_char_scaled(char ch, int x, int y, sprite_factors *factors)`

sprite_create_mask

Creates a mask definition for the specified sprite.

Syntax: `os_error *sprite_create_mask(sprite_area *, sprite_id *)`

sprite_remove_mask

Removes the mask definition from the specified sprite.

Syntax: `os_error *sprite_remove_mask(sprite_area *, sprite_id *)`

sprite_insert_row

Inserts a row into the specified sprite at the given row.

Syntax: `os_error *sprite_insert_row(sprite_area *, sprite_id *, int row)`

sprite_delete_row

Deletes the given row from the specified sprite.

Syntax: `os_error *sprite_delete_row(sprite_area *, sprite_id *, int row)`

sprite_insert_column

Inserts a column into the specified sprite at the given column.

Syntax: `os_error *sprite_insert_column(sprite_area *, sprite_id *, int column)`

sprite_delete_column

Deletes the given column from the specified sprite.

Syntax: `os_error *sprite_delete_column(sprite_area *, sprite_id *, int column)`

sprite_flip_x

Flips the specified sprite about the x axis.

Syntax: `os_error *sprite_flip_x(sprite_area *, sprite_id *)`

sprite_flip_y

Flips the specified sprite about the y axis.

Syntax: `os_error *sprite_flip_y(sprite_area *, sprite_id *)`

sprite_readsize

Reads the size information for the specified `sprite_id`.

Syntax: `os_error *sprite_readsize(sprite_area *, sprite_id *,
sprite_info *resultinfo)`

sprite_readpixel

Reads the colour of a given pixel in the specified `sprite_id`.

Syntax: `os_error *sprite_readpixel(sprite_area *, sprite_id *, int
x, int y, sprite_colour *resultcolour)`

sprite_writepixel

Writes the colour of a given pixel in the specified `sprite_id`.

Syntax: `os_error *sprite_writepixel(sprite_area *, sprite_id *, int
x, int y, sprite_colour *colour)`

sprite_readmask

Reads the state of a given pixel in the specified sprite mask.

Syntax: `os_error *sprite_readmask(sprite_area *, sprite_id *,
int x, int y, sprite_maskstate *resultmaskstate)`

sprite_writemask

Writes the state of a given pixel in the specified sprite mask.

Syntax: `os_error *sprite_writemask(sprite_area *, sprite_id *, int
x, int y, sprite_maskstate *maskstate)`

sprite_restorestate

Restores the old state after one of the sprite redirection calls.

Syntax: `os_error *sprite_restorestate(sprite_state state)`

sprite_outputtosprite

Redirects VDU output to a sprite, saving the old state.

Syntax: `os_error *sprite_outputtosprite(sprite_area *area,
sprite_id *id, int *save_area, sprite_state *state)`

sprite_outputtomask

Redirects output to a sprite's transparency mask, saving the old state.

Syntax: `os_error *sprite_outputtomask(sprite_area *area, sprite_id
*id, int *save_area, sprite_state *state)`

sprite_outputtoscreen

Redirects output back to screen, saving the old state.

Syntax: `os_error *sprite_outputtoscreen(int *save_area, sprite_state *state)`

sprite_sizeof_spritecontext

Gets the size of the save area needed to save the sprite context.

Syntax: `os_error *sprite_sizeof_spritecontext(sprite_area *area, sprite_id *id, int *size)`

sprite_sizeof_screencontext

Gets the size of the save area needed to save the screen context.

Syntax: `os_error *sprite_sizeof_screencontext(int *size)`

sprite_removewastage

Removes the lefthand wastage from a sprite.

Syntax: `os_error *sprite_removewastage(sprite_area *area, sprite_id *id)`

template

This file contains functions used for loading and manipulating templates (typically set up using the template editor, FormEd). The templates are assumed to be held in a file `Templates` in the application's directory. The dialogue box module of the RISC OS library uses these templates when creating dialogue boxes.

template_copy

Creates a copy of a template.

Syntax: `template *template_copy (template *from)`

Parameters: `template *from` – the original template

Returns: a pointer to a copy of `from`.

Other Information: Copying includes fixing up pointers into workspace for indirected icons/title, and the allocation of this space.

template_readfile

Reads the template file into a linked list of templates.

Syntax: `BOOL template_readfile (char *name)`

Parameters: `char *name` – name of template file

Returns: Non-zero if sprites are used in the template file.

Other Information: Note that a call is made to `resspr_area()`, in order to fix up a window's sprite pointers, so you must have already called `resspr_init`.

template_find

Finds a named template in the template list.

Syntax: `template *template_find(char *name)`

Parameters: `char *name` – the name of the template (as given in FormEd)

Returns: a pointer to the found template.

template_loaded

Sees if there is anything in the template list.

Syntax: `BOOL template_loaded(void)`

Parameters: `void`

Returns: Non-zero if there is something in the template list.

template_use_fancyfonts

Provides a font usage array for loading templates which use fonts other than 'system font'

Syntax: `void template_use_fancyfonts(void);`

Parameters: `void`

Returns: `void`

Other Information: This function should be called once BEFORE `template_init`. It allocates a font usage array, which it uses to 'lose' any fancy fonts used, when your program exits. It installs a C exit handler to do this. This function is useful if your dialogue boxes use fonts other than system font.

template_init

Initialises ready for the use of templates.

Syntax: `void template_init(void)`

Parameters: `void`

Returns: `void`.

Other Information: Should be called before any operations which use templates (such as dialogue box creation).

template_syshandle

Gets a pointer to the underlying window used to create a template.

Syntax: `wimp_wind *template_syshandle(char *name)`

Parameters: `char *templatename.`

Returns: Pointer to template's underlying window (0 if template not found).

Other Information: Any changes made to the `wimp_wind` structure will affect future windows generated using this template.

trace

These functions provide centralised control for trace/debug output.

tracef

Outputs tracing information.

Syntax: `void tracef(char*,...)`
`void tracef0(char*)`
`void tracef1(char*, int)`
`void tracef2(char*, int,int)`
`void tracef3(char*, int,int,int)`
`void tracef4(char*, int,int,int,int)`

Parameters: `char*` – printf-style format string
`...` – variable argument list.

Returns: `void.`

Other Information: called by `tracef0`, `tracef1` etc. Fixed-format ones will compile to nothing if `trace` is not set at compile time.

trace_is_on

`int trace_is_on(void)` returns True if tracing is turned on

trace_on

`void trace_on(void)` turns tracing on

trace_off

`void trace_off(void)` turns tracing off

txt

A `txt` is an array of characters, displayed in a window on the screen. It behaves in many ways similarly to a single buffer from Edit (see the *User Guide* for details of this application). It uses the system variable `Edit$Options` to set up colours, fonts and other features. You must call `flex_init` before calling `txt`.

txt: interface functions

txt_new

Creates a new `txt` object, containing no characters with a given title (to appear in its window).

Syntax: `txt txt_new(char *title)`

Parameters: `char *title` – the text title to appear in its window.

Returns: pointer to the newly created text.

Other Information: This function does not result in the text being displayed on the screen; it simply creates a new text object. 0 is returned if there is not enough space to create the object.

txt_show

Displays a given text object in a free-standing window of its own.

Syntax: `void txt_show(txt t)`

Parameters: `txt t` – the text to be displayed.

Returns: `void`.

Other Information: `t` should have been created using `txt_new`.

txt_hide

Hides a text which has been displayed.

Syntax: `void txt_hide(txt t)`

Parameters: `txt t` – the text to be hidden.

Returns: `void`.

txt_settitle

Changes the title of the window used to display a text object.

Syntax: `void txt_settitle(txt t, char *title)`

Parameters: `txt t` – the text object
`char *title` – new title of window.

Returns: `void`.

Other Information: Long titles may be truncated when displayed.

txt_dispose

Destroys a text and the window associated with it.

Syntax: `void txt_dispose(txt *t)`

Parameters: `txt *t` – pointer to the text.

Returns: `void`.

txt: general control operations

A text object's main data content is an array of characters. This resides in a buffer of known size. The characters of the array are not laid out precisely in the buffer; a gap is used in order to make insertion and deletion fast. When initially created, a text has `bufsize=0`.

txt_bufsize

Tells caller how many characters can be stored in the buffer before more memory needs to be requested from the operating system.

Syntax: `*int txt_bufsize(txt)`

Parameters: `txt t` – the text.

Returns: size of buffer.

txt_setbufsize

Allocates more space for the text buffer.

Syntax: `BOOL txt_setbufsize(txt, int)`

Parameters: `txt t` – the text
`int b` – new buffer size.

Returns: True if space could be allocated successfully.

Other Information: This call increases the buffer size, so that at least `b` characters can be stored before requiring more from the operating system.

The character array is displayed on the screen in a window. The characters travel horizontally from left to right. If a `\n` is encountered, this signifies the end of the current text line, and the start of a new one. All lines have the same height, although characters may be of differing widths. There is no limit on the number of characters allowed in a line. There is no restriction on the characters allowed in the array: any number from 0 to 255 is acceptable.

txt_charoptions

Informs the caller of the currently set charoptions.

Syntax: `txt_charoption txt_charoptions(txt)`

Parameters: `txt t` – text object.

Returns: Currently set charoptions.

Clearing the DISPLAY flag can be used during a long and complex sequence of edits, to reduce the overall amount of display activity. The UPDATED flag is set by the insertion or deletion of any characters in the array.

txt_setcharoptions

Sets the flags which are used to control the display of text in a screen window.

Syntax: `void txt_setcharoptions(txt, txt_charoption affect, txt_charoption values)`

Parameters: `txt t` – text object
`txt_charoption affect` – flags to affect
`txt_charoption values` – values to give to affected flags.

Returns: `void`.

Other Information: Only the flags named in `affect` are affected – they are set to the value `values`. This therefore has the meaning:

`(previousState & ~affect) | (affect & values)`

txt_lastref

Returns `last_ref` field (for `Message_DataSaved`).

Syntax: `int txt_lastref(txt);`

Parameters: `txt t` – text object.

Returns: Current value of `last_ref` (for `Message_DataSaved`).

Other Information: None

txt_setlastref

Sets value of `last_ref` (for `Message_DataSaved`).

Syntax: `void txt_setlastref(txt, int newvalue);`

Parameters: `txt t` – text object
`int newvalue` – new value

Returns: `void`.

Other Information: Sets the `last_ref` field in a `txt`, so that subsequently a `Message_DataSaved` can mark the data unmodified.

txt_setdisplayok

Sets the display flag in charoptions for a given text.

Syntax: `void txt_setdisplayok(txt)`

Parameters: `txt t` – text object

Returns: `void`.

Other Information: This asserts to the system that the display is up to date, preventing a redraw. It is useful only in very specialised circumstances.

txt: operations on the array of characters

`dot` is an index into the character array. If there are `n` characters in the array, with indices in `0...n-1`, then `dot` is in `0...n`. It is thought of as pointing just before the character with the same index, but it can also point just after the last one. When the text is displayed, the character after the `dot` is always visible. The caret is a visible indication of the position of the `dot` within the array. It can be made visible using `SetCharOptions` above.

txt_dot

Informs the caller of where the `dot` (current position) is in the array of characters.

Syntax: `txt_index txt_dot(txt t)`

Parameters: `txt t` – text object.

Returns: An index into the array of characters.

txt_size

Informs the caller as to the maximum value `dot` can take.

Syntax: `txt_index txt_size(txt t)`

Parameters: `txt t` – text object.

Returns: Maximum permissible value of `dot`.

txt_setdot

Sets the `dot` at a given index in the array of characters.

Syntax: `void txt_setdot(txt t, txt_index i)`

Parameters: `txt t` – text object.

`txt_index i` – index at which to set `dot`.

Returns: `void`.

Other Information: If `i` is outside the bounds of the array it is set to the beginning or end of the array, as appropriate.

txt_movedot

Moves the `dot` by a given distance in the array.

Syntax: `void txt_movedot(txt, int by)`

Parameters: `txt t` – text object
`int by` – distance to move by

Returns: `void`

Other Information: If the resulting `dot` is outside the bounds of the array it is set to the beginning or end of the array, as appropriate.

txt_insertchar

Inserts a character into the text just after the `dot`.

Syntax: `void txt_insertchar(txt t, char c)`

Parameters: `txt t` – text object
`char c` – the character to be inserted.

Returns: `void`.

Other Information: If the `DISPLAY` option flag is set, the window is redisplayed after insertion.

txt_insertstring

Inserts a given character string into a text.

Syntax: `void txt_insertstring(txt t, char *s)`

Parameters: `txt t` – text object
`char *s` – the character string.

Returns: `void`.

Other Information: If the `DISPLAY` option flag is set, the window is redisplayed after insertion.

txt_delete

Deletes `n` characters from the `dot` onwards.

Syntax: `void txt_delete(txt t, int n)`

Parameters: `txt t` – text object
`int n` – number of characters to delete.

Returns: `void`.

Other Information: If `dot+n` is beyond the end of the array, deletion is to the end of the array.

txt_replacechars

Deletes `ntodelete` characters from `dot`, and inserts `n` characters in their place, where the characters are pointed at by `a`.

Syntax: `void txt_replacechars(txt t, int ntodelete, char *a, int n)`
 Parameters: `txt t` – text object
`int ntodelete` – number of characters to delete
`char *a` – pointer to characters to insert
`int n` – number of characters to insert.
 Returns: `void`.

txt_charatdot

Informs the caller of the character held at `dot` in the array.

Syntax: `char txt_charatdot(txt t)`
 Parameters: `txt t` – text object.
 Returns: Character at `dot`.
 Other Information: Returns 0 if `dot` is at or beyond end of array.

txt_charat

Informs the caller of the character at a given index in the array.

Syntax: `char txt_charat(txt t, txt_index i)`
 Parameters: `txt t` – text object
`txt_index i` – the index into the array.
 Returns: Character at given index in array.
 Other Information: Returns 0 if index is at or beyond end of array.

txt_charsatdot

Copies at most `n` characters from `dot` in the array into a supplied buffer.

Syntax: `void txt_charsatdot(txt, char/*out*/ *buffer, int /*inout*/ *n)`
 Parameters: `txt t` – text object
`char *buffer` – the buffer
`int *n` – maximum characters to copy.
 Returns: `void`.
 Other Information: If you are close to the end of the array, `n` characters may not be available. In this case, characters up to the end of the array are copied, and `*n` is updated to report how many were copied.

txt_replaceatend

Deletes a specified number of characters from the end of the array and then inserts specified characters.

Syntax: `void txt_replaceatend(txt, int ndelete, char*, int)`
Parameters: `txt t` – text object
`int ndelete` – number of characters to delete
`char *s` – pointer to characters to insert
`int n` – number of characters to insert.
Returns: `void`.

txt: layout-dependent operations

These operations are provided specifically for the support of cursor-key-driven editing.

txt_movevertical

Moves the dot by a specified number of textual lines, with the caret staying in the same horizontal position on the screen.

Syntax: `void txt_movevertical(txt t, int by, int caretstill)`
Parameters: `txt t` – text object
`int by` – number of lines to move by
`int caretstill` – set to non-zero, if you want the text to move rather than the caret.
Returns: `void`.

txt_movehorizontal

Moves the caret (and dot) horizontally.

Syntax: `void txt_movehorizontal(txt, int by)`
Parameters: `txt t` – text object
`int by` – distance to move by.
Returns: `void`.
Other Information: This behaves like `txt_movedot()`, except that if `by` is positive and the end of the current text line is encountered, the caret will continue to move to the right on the screen.

txt_visiblelinecount

Gives the number of lines visible or partially visible on the display.

Syntax: `int txt_visiblelinecount(txt t)`

Parameters: `txt t` – text object.
Returns: Number of visible lines
Other Information: Takes into account current window size, font etc.

txt_visiblecolcount

Gives the number of columns currently visible.

Syntax: `int txt_visiblecolcount(txt t)`

Parameters: `txt t` – text object.

Returns: Visible column count.

Other Information: If a fixed pitch font is currently in use, this gives the number of display columns; otherwise, it makes a guess for average characters.

txt: operations on markers

Markers are indices into the array. Once set, a marker will point to the same character in the array regardless of insertions or deletions within the array. If the character pointed at by the marker is deleted, the marker will point to the next character. Markers never fall off the end of the array, but stay at the top or bottom of it, if that's where they end up.

txt_newmarker

Creates a new marker in the text.

Syntax: `void txt_newmarker(txt, txt_marker *mark)`

Parameters: `txt t` – text object
`txt_marker *mark` – pointer to your text marker.

Returns: `void`.

Other Information: The marker itself is kept by the client of this function, but the text object retains a pointer to it. The client's marker is updated by the text object whenever necessary. Its initial value is the same as `dot`. If the character at which a marker points is deleted, then the marker gets moved to the value of `dot` when the deletion occurred. If characters are inserted when the marker is at `dot`, the marker stays with `dot`.

txt_movemarker

Resets an existing marker.

Syntax: `void txt_movemarker(txt t, txt_marker *mark, txt_index to)`

Parameters: txt t – text object
 txt_marker *mark – the marker
 txt_index to – place to move the marker to.
Returns: void.
Other Information: The marker must already point into this text object.

txt_movedottomarker

Moves the dot to a given marker.

Syntax: void txt_movedottomarker(txt t, txt_marker *mark)
Parameters: txt t – text object
 txt_marker *mark – pointer to the marker.
Returns: void.

txt_indexofmarker

Gives the current index into the array of a given marker.

Syntax: txt_index txt_indexofmarker(txt t, txt_marker *mark)
Parameters: txt t – text object
 txt_marker *mark – pointer to the marker.
Returns: Index of marker.

txt_disposemarker

Delete a marker from a text object.

Syntax: void txt_disposemarker(txt, txt_marker*)
Parameters: txt t – text object
 txt_marker *mark – the marker to be deleted.
Returns: void.
Other Information: You should remember to dispose of a marker which logically ceases to exist, otherwise the text object will continue to update the location where it was.

txt: operations on a selection

The selection is a contiguous portion of the array which is displayed highlighted.

txt_selectset

Informs the caller whether there is a selection made in a text.

Syntax: BOOL txt_selectset(txt t)
Parameters: txt t – text object.

Returns: True if there is a selection in this text.

txt_selectstart

Gives the index into the array of the start of the current selection.

Syntax: `txt_index txt_selectstart(txt t)`

Parameters: `txt t` – text object.

Returns: Index of selection start.

txt_selectend

Gives the index into the array of the end of the current selection.

Syntax: `txt_index txt_selectend(txt t)`

Parameters: `txt t` – text object.

Returns: Index of selection end.

txt_setselect

Sets a selection in a given text, from start to end.

Syntax: `void txt_setselect(txt, txt_index start, txt_index end)`

Parameters: `txt t` – text object
`txt_index start` – array index of start of selection
`txt_index end` – array index of end of selection.

Returns: `void`.

Other Information: If `start >= end` then the selection will be unset.

txt: input from the user

Characters entered into the keyboard, and various mouse events, are buffered up by the text object for use by the client.

A call to the event handler registered with a text object will give an event code to the event handler, to say what sort of event has occurred. The following event codes are defined; any that are not understood should be ignored.

- Codes 0 – 255: key codes from the keyboard
- Codes 256 – 511: various function keys, etc; refer to `h.akbd` for the rules.
- Mouse events:

A mouse event occurs when the mouse is pointing in the text object and a button is pressed or released, or the mouse moves while any button is depressed. A mouse event will result in `Get` producing an `EventCode` with bit 31 set, bits 24..28 as a `mouseevent flags` value, and the rest of the word containing an index value.

The index shows where in the visible representation of the array the mouse event happened. If all three index bytes are 255, the event happened outside the window. The `mouseevent flags` show what button transitions occurred:

MSELECT	Select's new value
MEXTEND	Adjust's new value
MSELOLD	Select's old value
MEXTOLD	Adjust's old value
MEXACT	the event is in exactly the same place as the last one.

The byte gives the values of the select and extend buttons: 1 for depressed and 0 for not depressed. It gives their previous values, allowing transitions to be detected. It reports whether the position of the mouse is exactly the same as for the last event, so that multiple clicks may be detected. No assumptions should be made concerning the relationship of these bits to the last mouse event sent to the programmer, as polling delays etc. could cause any combinations to happen.

If `txt_EXTRACODE` is set, the identity of the event is not defined by this interface. This is used for any expansion. Clients of this interface which receive such events that they do not recognise, should ignore them without reporting an error.

The Menu button on the mouse is not transmitted through this interface, but caught elsewhere. Use `event_attach_menu` to attach a menu handler to the `txt_syshandle` of a `txt` object.

- Keyboard events:

```
txt_EXTRACODE + akbd_Fn + 1: - help request
txt_EXTRACODE + akbd_Fn + akbd_Sh + 2: insert drag line
txt_EXTRACODE + akbd_Fn + 127: - close icon
txt_EXTRACODE + akbd_Sh + akbd_Ctl + akbd_Upk: scroll up
one line
txt_EXTRACODE + akbd_Sh + akbd_Ctl + akbd_DownK: scroll
down one line
txt_EXTRACODE + akbd_Sh + akbd_UpK: scroll up one page
txt_EXTRACODE + akbd_Sh + akbd_DownK: scroll down one page
```

In the current implementation of `txt`, `txt_queue` never returns more than 1, so `wimpt_last_event()` can be accessed to get more information.

txt_get

Gives the next user event code to the caller.

Syntax: `txt_eventcode txt_get(txt t)`

Parameters: `txt t` - text object

Returns: The event code
 Other Information: The returned code can be ASCII, or various other (system-specific) values for function keys etc. This function can only be called within an event handler.

txt_queue

Informs the caller of how many event codes are currently buffered for a given text.

Syntax: `int txt_queue(txt t)`
 Parameters: `txt t` – text object
 Returns: Number of buffered event codes.
 Other information: This function can only be called within an event handler.

txt_unget

Puts an event code back on the front of the event queue for a given text.

Syntax: `void txt_unget(txt t, txt_eventcode code)`
 Parameters: `txt t` – text object
`txt_eventcode code` – the event code.
 Returns: `void`.
 Other information: This function can only be called within an event handler.

txt_eventhandler

Registers an eventhandler function for a given text, which will be called whenever there is a value ready which can be picked up by `txt_get()`.

Syntax: `void txt_eventhandler(txt, txt_event_proc, void *handle)`
 Parameters: `txt t` – text object
`txt_event_proc func` – event handler function
`void *handle` – caller-defined handle to be passed to `func`.
 Returns: `void`.
 Other Information: If `func==0`, no function is registered.

txt_readeventhandler

Informs the caller of the currently registered eventhandler function associated with a given text, and the handle which is passed to it.

Syntax: `void txt_readeventhandler(txt t, txt_event_proc *func, void **handle)`

Parameters: txt t – text object
 txt_event_proc *func – returned pointer to handler
 func
 void **handle – returned pointer to handle.
Returns: void.

txt: direct access to the array of characters

txt_arrayseg

Gives a direct pointer into the memory used to hold the characters in a text.

Syntax: void txt_arrayseg(txt t, txt_index at, char **a, int *n)

Parameters: txt t – text object
 txt_index at – index into the text
 char **a – *a will point at the character whose index in
 the text is at
 int *n – number of contiguous bytes after at.

Returns: void.

Other Information: It is permissible for the caller of this function to change
 the characters pointed at by *a, provided that a
 re-display is prompted (using `setcharoptions`).

txt: system hook

txt_syshandle

Obtains a `wimp_w` value for the window underlying a text.

Syntax: int txt_syshandle(txt t)

Parameters: txt t – text object.

Returns: System-dependent handle for the given text.

txt_init

Initialise the txt module of the library

Syntax: void txt_init(void);

Parameters: void

Returns: void

Other Information: None

txtdit

These functions provide text editing facilities.

txtdit_install

Installs an event handler for the `txt t`, thus making it an editable text.

Syntax: `txtdit_state *txtdit_install(txt t)`
 Parameters: `txt t` – the text object (created via `txt_new`)
 Returns: A pointer to the resulting `txtdit_state`.

txtdit_new

Creates a new text object and loads the given file into it. The text can then be edited.

Syntax: `txtdit_state *txtdit_new(char *filename)`
 Parameters: `char *filename` – the file to be loaded.
 Returns: a pointer to the `txtdit_state` for this text.
 Other Information: If the file cannot be found, then 0 is returned as a result, and no text is created. If `filename` is a null pointer, then an editor window with no given file name will be constructed. If the file is already being edited, then a pointer to the existing `txtdit_state` is returned.

txtdit_dispose

Destroys the given text being edited.

Syntax: `void txtdit_dispose(txtdit_state *s)`
 Parameters: `txtdit_state *s` – the text to be destroyed.
 Returns: `void`.
 Other Information: This will ask no questions of the user before destroying the text.

txtdit_mayquit

Check if we may safely quit editing.

Syntax: `BOOL txtdit_mayquit(void)`
 Parameters: `void`.
 Returns: True if we may safely quit, otherwise False.
 Other Information: If a text is being edited, then a dialogue box is displayed asking the user if he really wants to quit. This calls `dboxquery()`, and therefore requires the template `query` as described in `dboxquery.h`.

txtedit_prequit

Deals with a PREQUIT message from the Task Manager.

Syntax: `void txtedit_prequit(void)`

Parameters: `void`.

Returns: `void`.

Other Information: Calls `txtedit_mayquit()`, to see if we may quit, if text is being edited. If user replies that we may quit, then all texts are disposed of, and this function sends an acknowledgement to the Task Manager.

txtedit_menu

Sets up a menu structure for the text being edited, tailored to its current state.

Syntax: `menu txtedit_menu(txtedit_state *s)`

Parameters: `txtedit_state *s` – the text's current state.

Returns: a pointer to an appropriately formed menu structure.

Other Information: The menu created will have the same form as that displayed when Menu is clicked on an Edit window. (For Edit version 1.00). Entries in the menu are set according to the supplied `txtedit_state`.

txtedit_menuevent

Applies a given menu hit to a given text.

Syntax: `void txtedit_menuevent(txtedit_state *s, char *hit)`

Parameters: `txtedit_state *s` – the text to which hit should be applied
`char *hit` – a menu hit string.

Returns: `void`.

Other Information: This can be called from a menu event handler.

txtedit_doimport

Import data into the specified `txtedit` object, from a file of a given type.

Syntax: `BOOL txtedit_doimport(txtedit_state *s, int filetype, int estsize)`

Parameters: `txtedit_state *s` – the text object
`int filetype` – type of the file
`int estsize` – the file's estimated size.

Returns: True if the import is completed successfully.

txtdit_doinsertfile

Inserts a named file in a given text object.

Syntax: `void txtdit_doinsertfile(txtdit_state *s, char *filename, BOOL replaceifwasnull)`

Parameters: `txtdit_state *s` – the text object
`char *filename` – the given file
`BOOL replaceifwasnull` – if set to `True` then the text object will be considered to have come from `filename`, ie the window title is updated.

Returns: `void`.

txtdit_register_update_handler

Register a handler to be called when a text window is modified

Syntax: `txtdit_update_handler txtdit_register_update_handler(txtdit_update_handler h, void *handle);`

Parameters: `txtdit_handler h` – the handler function
`void *handle` – handle to be passed to the function

Returns: previous handler

Other Information: This routine will be called whenever a window's title bar is redrawn, and the text in the window has been modified. Note: this is not just called when the '*' first appears in a window's title bar, but every time the title bar of a modified text window is redrawn (eg when the filename changes or wordwrap is turned on/off etc).

The handler function will be passed:

- i) the filename for the window title
- ii) the address of this 'txtdit_state'
- iii) the handle registered with this function will be undone. This is only possible if the modification is not greater than ~5kb.

Calling with `h == 0` removes the handler.

txtdit_register_save_handler

Register a handler to be called when a text window is saved

Syntax: `txtdit_save_handler txtdit_register_save_handler(txtdit_save_handler h, void *handle);`

Parameters: `txtdit_handler h` – the handler function
`void *handle` – handle to be passed to the function

Returns: previous handler
Other Information: This routine will be called whenever a text window is saved to file (NOT via RAM transfer).
The handler function will be passed:
i) the filename for the window title
ii) the address of this 'txtedit_state'
iii) the handle registered with this function
Calling with h == 0 removes the handler.
Returning FALSE from your handler will abort the save operation.

txtedit_register_close_handler

Register a handler to be called when a modified text window is closed

Syntax: `txtedit_close_handler txtedit_register_close_handler(
txtedit_close_handler h, void *handle);`

Parameters: `txtedit_handler h` – the handler function
`void *handle` – handle to be passed to the function

Returns: previous handler

Other Information: This routine will be called whenever a text window is closed.
The handler function will be passed:
i) the filename for the window title
ii) the address of this 'txtedit_state'
iii) the handle registered with this function
Calling with h == 0 removes the handler.

txtedit_register_shutdown_handler

Register a handler to be called when `txtedit_prequit()` is called.

Syntax: `txtedit_shutdown_handler
txtedit_register_shutdown_handler(
txtedit_shutdown_handler h, void *handle);`

Parameters: `txtedit_handler h` – the handler function
`void *handle` – handle to be passed to the function

Returns: previous handler

Other Information: This routine will be called whenever `txtedit_prequit()` is called, and the user answers "yes" when asked if he really wants to quit edit, or no files have been modified. The

handler function will be passed the handle registered with this function. Calling with `h == 0` removes the handler.

txtdit_register_undofail_handler

Register a handler to be called when your `update_handler` returned `FALSE`, and the undo buffer overflowed.

Syntax: `txtdit_undofail_handler
txtdit_register_undofail_handler(
txtdit_undofail_handler h, void *handle);`

Parameters: `txtdit_handler h` – the handler function
`void *handle` – handle to be passed to the function

Returns: previous handler

Other Information: This will be called when the modification made to a edited file cannot be undone (only in conjunction with an update handler).
The handler function will be passed:
i) the filename for the window title
ii) the address of this 'txtdit_state'
iii) the handle registered with this function
Calling with `h == 0` removes the handler.

txtdit_register_open_handler

Register a handler to be called when a new `txtdit_state` is created.

Syntax: `txtdit_open_handler txtdit_register_open_handler(
txtdit_open_handler h, void *handle);`

Parameters: `txtdit_handler h` – the handler function
`void *handle` – handle to be passed to the function

Returns: previous handler

Other Information: The handler function will be passed:
i) the filename for the window title
ii) the address of this 'txtdit_state'
iii) the handle registered with this function
Calling with `h == 0` removes the handler.

txtdit_getstates

Get a pointer to the list of current `txtdit_states`

Syntax: `txtdit_state *txtdit_getstates(void);`
Parameters: `void.`

Returns: Pointer to the list of `txtedit_states`
 Other Information: The `txtedit` part of `RISC_OSLib` keeps a list of all `txtedit_states` created (via `txtedit_new`). This function allows access to this list.

txtedit_init

Initialise the `txtedit` module of the library

Syntax: `void txtedit_init(void);`
 Parameters: `void`
 Returns: `void`
 Other Information: `None`

txtopt

These functions set and read the name of the system variable used for text editing options.

txtopt_set_name

Set the name used as a system variable for setting text editing options

Syntax: `void txtopt_set_name(char *name);`
 Parameters: `char *name` – the name to be prepended to `$Options` to form the system variable name.
 Returns: `void`.
 Other Information: If this function is not called before using any of the `txt` and `txtedit` functions, the system variable name defaults to `Edit$Options`, eg
`txtopt_set_name("MyEdit")` sets the system variable name to `MyEdit$Options`.

txtopt_get_name

Get a pointer to the name currently prepended to `$Options` to form a system variable for use in setting text editing options.

Syntax: `char *txtopt_get_name(void);`
 Parameters: `void`
 Returns: pointer to name
 Other Information: If no name has been set, this will point to `"Edit"`, eg assuming option name is currently `MyEdit$Options` then `txtopt_get_name` will return a pointer to the string `"MyEdit"`.

txtscrap

These functions manage a single txt selection within an arbitrary number of txt objects.

txtscrap_setselect

Calls `txt_setselect(t, from, to)` and remembers `t`. If another txt object currently holds the selection then this is first cleared.

Syntax: `void txtscrap_setselect(txt t, txt_index from, txt_index to);`

Parameters: `txt t` – text object
`txt_index from` – array index of start of selection
`txt_index to` – array index of end of selection (ie first character not in the selection).

Returns: `void`

Other Information: If "from" >= "to" then the selection will be unset, and `t` will not be remembered as holding the current selection. A txt must not be destroyed while still holding the selection, please clear the selection first.

txtscrap_selectowner

Returns the current holder of the selection.

Syntax: `txt txtscrap_selectowner(void);`

Parameters: `void`

Returns: The txt that currently holds the selection, or 0 if none.

Other Information: None

txtwin

These functions give control of multiple windows on text objects. When the Text is updated, all the windows are updated in step. All the windows have the same title information.

txtwin_new

Creates an extra window on a given text object.

Syntax: `void txtwin_new(txt t)`

Parameters: `txt t` – the text to have a window added to it.

Returns: `void`

Other Information: The created window will be in the same style as for `txt_new()`, with the same title information. The window will be made visible.

txtwin_number

Informs the caller of the number of windows currently on a given text.

Syntax: `int txtwin_number(txt t)`
Parameters: `txt t` – the text.
Returns: The number of windows currently on `t`.

txtwin_dispose

Removes a window, previously on `t`.

Syntax: `void txtwin_dispose(txt t)`
Parameters: `txt t` – the text
Returns: `void`
Other Information: This call will have no effect if there is only one window on `t`.

txtwin_setcurrentwindow

Ensures that the last window to which the last event was delivered is the current window on a given text.

Syntax: `void txtwin_setcurrentwindow(txt t)`
Parameters: `txt t` – the text.
Returns: `void`.
Other Information: Call this when constructing menus, since the same menu structure is attached to each window on the same text object.

visdelay

These functions enable a visual indication of some delay.

visdelay_begin

Changes pointer to show user there will be some delay (currently the RISC OS hourglass).

Syntax: `void visdelay_begin(void)`
Parameters: `void`.
Returns: `void`.

Other Information: Under RISC OS, the hourglass will only appear if the delay is longer than 1/3 sec.

visdelay_percent

Indicates to the user that a delay is *p* percent complete.

Syntax: `void visdelay_percent(int p)`
Parameters: `int p` – percentage complete.
Returns: `void`.

visdelay_end

Removes the indication of delay.

Syntax: `void visdelay_end(void)`
Parameters: `void`.
Returns: `void`.

visdelay_init

Initialises ready for `visdelay` functions.

Syntax: `void visdelay_init(void)`
Parameters: `void`.
Returns: `void`.

werr

This function provides error reporting in Wimp programs, causing a (possibly fatal) error message to appear in a pop-up dialogue box.

Syntax: `void werr(int fatal, char* format, ...)`
Parameters: `int fatal` – non-zero indicates fatal error
`char *format` – printf-style format string
`...` – variable arg list of message to be printed.
Returns: `void`.

Other Information: The program exits if `fatal` is non-zero. The pointer is restricted to the displayed dialogue box to stop the user continuing until he has clicked on the OK button. The message should be divided into at most three lines, each of 40 characters or less.

wimp

This file provides a C interface to RISC OS Wimp SWIs, and the following useful type definitions.

```

wimp_flags
typedef enum{
wimp_WMOVEABLE = 0x00000002,      is moveable
wimp_REDRAW_OK = 0x00000010,     can be redrawn entirely by Wimp ie
                                   no user graphics
wimp_WPANE = 0x00000020,         window is stuck over tool window
wimp_WTRESPASS = 0x00000040,    window is allowed to go outside
                                   main area
wimp_WSCROLL_R1= 0x00000100,    scroll request returned when scroll
                                   button clicked – auto-repeat
wimp_SCROLL_R2 = 0x00000200,    as SCROLL_R1, debounced, no
                                   auto
wimp_REAL_COLOURS = 0x000000400, use real window colours.
wimp_BACK_WINDOW = 0x000000800,  this window is a background
                                   window.
wimp_HOT_KEYS = 0x000001000,    generate events for 'hot keys'
wimp_WOPEN = 0x00010000,        window is open
wimp_WTOP = 0x00020000,         window is on top (not covered)
wimp_WFULL = 0x00040000,        window is full size
wimp_WCLICK_TOGGLE = 0x00080000, open_window_request was due to
                                   click on Toggle size icon
wimp_WFOCUS = 0x00100000,       window has input focus
wimp_WBACK = 0x01000000,        window has Back icon
wimp_WQUIT = 0x02000000,        has a Close icon
wimp_WTITLE = 0x04000000,        has a title bar
wimp_WTOGGLE= 0x08000000,        has a Toggle size icon
wimp_WVSCR = 0x10000000,        has vertical scroll bar
wimp_WSIZE = 0x20000000,        has Adjust size icon

```

```
wimp_WHSCR = 0x40000000,      has horizontal scroll bar
wimp_WNEW = 0x80000000      use these new flags
}wimp_flags;
```

Note: Always set the WNEW flag.

wimp_wcolours

If the work area background is 255, it isn't painted. If the title foreground is 255, you get no borders, title etc. at all.

```
typedef enum{
wimp_WCTITLEFORE,
wimp_WCTITLEBACK,
wimp_WCWKAREAFORE,
wimp_WCWKAREABACK,
wimp_WCSCROLL OUTER,
wimp_WCSCROLL INNER,
wimp_WCTITLEHI,
wimp_WCRESERVED
}wimp_wcolours;
```

wimp_iconflags

If the icon contains anti-aliased text, the colour fields give the font handle

```
typedef enum{
wimp_ITEXT = 0x00000001,      icon contains text
wimp_ISPRITE = 0x00000002,    icon is a sprite
wimp_IBORDER = 0x00000004,    icon has a border
wimp_IHCENTRE = 0x00000008,   text is horizontally centred
wimp_IVCENTRE = 0x00000010,   text is vertically centred
wimp_IFILLED = 0x00000020,    icon has a filled background
wimp_IFONT = 0x00000040,      text is in an anti-aliased font
wimp_IREDRAW = 0x00000080,    redraw needs application's help
wimp_INDIRECT = 0x00000100,   icon data is 'indirected'
```

wimp_IRJUST = 0x00000200,	text right-justified in box
wimp_IESG_NOC = 0x00000400,	if selected by Adjust, don't cancel other icons in same ESG
wimp_IHALVESPRITE = 0x00000800,	plot sprites half-size
wimp_IBTYPE = 0x00001000,	4-bit field: button type
wimp_ISELECTED = 0x00200000,	icon selected by user (inverted)
wimp_INOSELECT = 0x00400000,	icon cannot be selected (shaded)
wimp_IDELETED = 0x00800000,	icon has been deleted
wimp_IFORECOL = 0x01000000,	4-bit field: foreground colour
wimp_IBACKCOL = 0x10000000	4-bit field: background colour
}wimp_iconflags;	

wimp_ibtype

Button types:

typedef enum{	
wimp_BIGNORE,	ignore all mouse ops
wimp_BNOTIFY,	
wimp_BCLICKAUTO,	
wimp_BCLICKDEBOUNCE,	
wimp_BSELREL,	
wimp_BSELDOUBLE,	
wimp_BDEBOUNCEDRAG,	
wimp_BRELEASEDRAG,	
wimp_BDOUBLEDRAG,	
wimp_BSELNOTIFY,	
wimp_BCLICKDRAGDOUBLE,	
wimp_BCLICKSEL,	useful for on/off and radio buttons
wimp_Bwritable = 15	
}wimp_ibtype;	

wimp_bbits**Button state bits**

```
typedef enum{
wimp_BRIGHT = 0x001,
wimp_BMID = 0x002,
wimp_BLEFT = 0x004,
wimp_BDRAGRIGHT = 0x010,
wimp_BDRAGLEFT = 0x040,
wimp_BCLICKRIGHT = 0x100,
wimp_BCLICKLEFT = 0x400
}wimp_bbits;
```

wimp_dragtype

<pre>typedef enum{</pre>	
<pre>wimp_MOVE_WIND = 1,</pre>	change position of window
<pre>wimp_SIZE_WIND = 2,</pre>	change size of window
<pre>wimp_DRAG_HBAR = 3,</pre>	drag horizontal scroll bar
<pre>wimp_DRAG_VBAR = 4,</pre>	drag vertical scroll bar
<pre>wimp_USER_FIXED = 5,</pre>	user drag box – fixed size
<pre>wimp_USER_RUBBER = 6,</pre>	user drag box – rubber box
<pre>wimp_USER_HIDDEN = 7</pre>	user drag box – invisible box
<pre>}wimp_dragtype;</pre>	

wimp_w

```
typedef int wimp_w;
Abstract window handle.
```

wimp_i

```
typedef int wimp_i;
Abstract icon handle.
```

wimp_t

```
typedef int wimp_t;
```

Abstract task handle.

wimp_icondata

The data field in an icon.

```
typedef union{
```

```
    char text[12];
```

up to 12 bytes of text

```
    char sprite_name[12];
```

up to 12 bytes of sprite name

```
    struct {
```

```
        char *name;
```

```
        void *spritearea;
```

0 → use the common sprite area

1 → use the Wimp sprite area

```
        BOOL nameisname;
```

if False, name is in fact a sprite pointer.

```
    } indirectsprite;
```

```
    struct {
```

if indirect

```
        char *buffer;
```

pointer to text buffer

```
        char *validstring;
```

pointer to validation string

```
        int bufflen;
```

length of text buffer

```
    } indirecttext;
```

```
    } wimp_icondata;
```

wimp_box

```
typedef struct{
```

```
    int x0, y0, x1, y1
```

```
    } wimp_box;
```

wimp_wind

If there are any icon definitions, they should follow this structure immediately in memory.

```
typedef struct{
wimp_box box;
int scx, scy;
wimp_w behind;

wimp_wflags flags;
char colours[8];

wimp_box ex;
wimp_iconflags titleflags;
wimp_iconflags workflags;
void *spritearea;

int minsize;

wimp_icondata title;
int nicons;
} wimp_wind;
```

screen coordinates of work area
scroll bar positions
handle to open window behind, or -1 if top
word of flag bits defined above
colours: index using wimp_wcolours.
maximum extent of work area
icon flags for title bar
just button type relevant
0 → use the common sprite area
1 → use the Wimp sprite area
two 16-bit OS-unit fields, (width/height) giving minimum size of window
0 → use title
title icon data
number of icons in window

wimp_winfo

Result of `get_info` call. Space for icons must follow.

```
typedef struct {
wimp_w w;
wimp_wind info;
} wimp_winfo;
```

wimp_icon

Icon description structure.

```
typedef struct {
    wimp_box box;
    wimp_iconflags flags;
    wimp_icondata data;
} wimp_icon;
```

bounding box – relative to window origin (work area top left)
word of flag bits defined above
union of bits & bobs as above

wimp_igate

Structure for creating icons.

```
typedef struct {
    wimp_w w;
    wimp_icon i;
} wimp_igate;
```

wimp_openstr

```
typedef struct {
    wimp_w w;
    wimp_box box;
    int x, y;
    wimp_w behind;
} wimp_openstr;
```

window handle
screen position of visible work area
'real' coordinates of visible work area
handle of window to go behind (-1 = top, -2 = bottom)

wimp_wstate

Result for window state enquiry.

```
typedef struct {
    wimp_openstr o;
    wimp_wflags flags;
} wimp_wstate;
```

wimp_etypes

Event types.

```
typedef enum {
wimp_ENULL,           null event
wimp_EREDRAW,        redraw event
wimp_EOPEN,
wimp_ECLOSE,
wimp_EPTRLEAVE,
wimp_EPTRENTER,
wimp_EBUT,           mouse button change
wimp_EUSERDRAG,
wimp_EKEY,
wimp_EMENU,
wimp_ESCROLL,
wimp_ELOSECARET,
wimp_EGAINCARET,
wimp_ESEND = 17,    send message, don't worry if it
                    doesn't arrive
wimp_ESENDWANTACK = 18, send message, return ack if not
                    acknowledged
wimp_EACK = 19      acknowledge receipt of message
} wimp_etype;
```

wimp_emask

Event type masks.

```
typedef enum {
wimp_EMNULL = 1 << wimp_ENULL,
wimp_EMREDRAW = 1 << wimp_EREDRAW,
wimp_EMOPEN = 1 << wimp_EOPEN,
wimp_EMCLOSE = 1 << wimp_ECLOSE,
```



```

wimp_EMPTRLEAVE = 1 << wimp_EPTRLEAVE,
wimp_EMPTREENTER = 1 << wimp_EPTREENTER,
wimp_EMBUT = 1 << wimp_EBUT,
wimp_EMUSERDRAG = 1 << wimp_EUSERDRAG,
wimp_EMKEY = 1 << wimp_EKEY,
wimp_EMENU = 1 << wimp_EMENU,
wimp_EMSCROLL = 1 << wimp_ESCROLL
wimp_EMLOSECARET = 1 << wimp_ELOSECARET
wimp_EMGAINCARET = 1 << wimp_EGAINCARET
wimp_ESEND = 1 << wimp_ESEND
wimp_EMSENDWANTACK = 1 << wimp_ESENDWANTACK
wimp_EMACK = 1 << wimp_EACK
} wimp_emask;

```

wimp_redrawstr

```

typedef struct {
wimp_w w;
wimp_box box;
int scx, scy;
wimp_box g;
} wimp_redrawstr;

```

work area coordinates
scroll bar positions
current graphics window

wimp_mousestr

```

typedef struct {
int x, y;
wimp_bbits bbits;
wimp_w w;
wimp_i i;
} wimp_mousestr;

```

mouse x and y
button state
window handle, or -1 if none
icon handle, or -1 if none

wimp_caretstr

```
typedef struct {
    wimp_w w;
    wimp_i i;
    int x, y;
    int height;

    int index;
} wimp_caretstr;
```

offset relative to window origin

-1 if calc within icon

bit 24 → VDU-5 type caret

bit 25 → caret invisible

bit 26 → bits 16...23 contain colour

bit 27 → colour is 'real' colour

position within icon

wimp_msgaction

Message action codes are allocated just like SWI codes.

```
typedef enum {
    wimp_MCLOSEDOWN = 0,

    wimp_MDATASAVE = 1,
    wimp_MDATASAVEOK = 2,
    wimp_MDATALOAD = 3,
    wimp_MDATALOADOK = 4,
    wimp_MDATAOPEN = 5,

    wimp_MRAMFETCH = 6,

    wimp_MRAMTRANSMIT = 7,

    wimp_MPREQUIT = 8,
    wimp_PALETTECHANGE = 9,
    wimp_FilerOpenDir = 0x0400,
```

reply if any dialogue with the user is required, and the closedown sequence will be aborted.

request to identify directory

reply to message type 1

request to load/insert dragged icon

reply that file has been loaded

warning that an object is to be opened

transfer data to buffer in my workspace

I have transferred some data to a buffer in your workspace

wimp_FilerCloseDir = 0x0401,	net filer notify broadcast
wimp_Notify = 0x40040	menu warning. Sent if wimp_MSUBLINKMSG set. Data sent is:
wimp_MMENUWARN = 0x400c0,	submenu field of relevant wimp_menuitem.
	screen x-coord
	screen y-coord
	list of menu selection indices (0..n-1 for each menu)
	terminating -1 word.
	Typical response is to call wimp_create_submenu.
wimp_MMODECHANGE = 0x400c1,	
wimp_MINITTASK = 0x400c2,	
wimp_MCLOSETASK = 0x400c3,	
wimp_MSLOTCHANGE = 0x400c4,	Slot size has altered
wimp_MSETSLOT = 0x400c5,	Task Manager requests application to change its slot size
wimp_MTASKNAMERQ = 0x400c6,	Request task name
wimp_MTASKNAMEIS = 0x400c7,	Reply to task name request
wimp_MHELPREQUEST = 0x502,	interactive help request
wimp_MHELPREPLY = 0x503,	interactive help message

Messages for dialogue with printer applications

wimp_MPrintFile = 0x80140,	Printer application's first response to a DATASAVE
wimp_MWillPrint = 0x80141,	Acknowledgement of PrintFile
wimp_MPrintTypeOdd = 0x80145,	Broadcast when strange files dropped on the printer
wimp_MPrintTypeKnown = 0x80146,	Acknowledgement to above

```
wimp_MPrinterChange = 0x80147    New printer application installed
} wimp_msgaction;
```

wimp_msghdr

Message block header. `size` is the size of the whole `msgstr`, see below.

```
typedef struct {
int size;                          20<=size<=256, multiple of 4
wimp_t task;                        task handle of sender (filled in by
                                   Wimp)
int my_ref;                          unique ref number (filled in by
                                   Wimp)
int your_ref;                        (0==>none) if non-zero,
                                   acknowledge
wimp_msgaction action;              message action code
} wimp_msghdr;
```

wimp_msgdatasave

```
typedef struct {
wimp_w w;                            window in which save occurs.
wimp_i i;                            icon there
int x; int y;                        position within that window of
                                   destination of save
int estsize;                          estimated size of data, in bytes
int type;                             file type of data to save
char leaf[12];                        proposed leaf-name of file,
                                   0-terminated
} wimp_msgdatasave;
```

wimp_msgdatasaveok

`w`, `i`, `x`, `y`, `type`, `estsize` copied unaltered from `DataSave` message.

```
typedef struct {
wimp_w w;                            window in which save occurs.
wimp_i i;                            icon there
```

<code>int x;int y;</code>	position within that window of destination of save
<code>int estsize;</code>	estimated size of data, in bytes
<code>int type;</code>	file type of data to save
<code>char name[212];</code>	the name of the file to save
<code>} wimp_msgdatasaveok;</code>	

wimp_msgdataload

For a data load reply, no arguments are required.

<code>typedef struct {</code>	
<code>wimp_w w;</code>	target window
<code>wimp_i i;</code>	target icon
<code>int x; int y;</code>	target coordinates in target window work area
<code>int size;</code>	must be 0
<code>int type;</code>	type of file
<code>char name[212];</code>	the filename follows.
<code>} wimp_msgdataload;</code>	

wimp_msgdataopen

`wimp_msgdataopen` derives its typedef from `wimp_msgdataload`, since the data provided when opening a file is exactly the same. The window, `x` and `y` refer to the bottom lefthand corner of the icon that represents the file being opened, or `w=-1` if there is no such icon.

wimp_msgramfetch

Transfer data in memory.

<code>typedef struct {</code>	
<code>char *addr;</code>	address of data to transfer
<code>int nbytes;</code>	number of bytes to transfer
<code>} wimp_msgramfetch;</code>	

wimp_msgramtransmit

'I have transferred some data to a buffer in your workspace'.

```
typedef struct {
    char *addr;                copy of value sent in RAMfetch
    int nbyteswritten;        number of bytes written
} wimp_msgramtransmit;
```

wimp_msghelprequest

```
typedef struct {
    wimp_mousestr m;          where the help is required
} wimp_msghelprequest;
```

wimp_msghelpreply

```
typedef struct {
    char text[200];           the helpful string
} wimp_msghelpreply;
```

wimp_msgprint

Structure used in all print messages.

```
typedef struct {
    int filler[5] ;
    int type ;                filetype
    char name[256-44] ;       filename
} wimp_msgprint;
```

wimp_msgstr

Message block.

```
typedef struct {
    wimp_msghdr hdr;
    union {
        char chars[236];
```

```

int words[59];                maximum data size
wimp_msgdatasave datasave;
wimp_msgdatasaveok datasaveok;
wimp_msgdataload dataload;
wimp_msgdataopen dataopen;
wimp_msgramfetch ramfetch;
wimp_msgramtransmit ramtransmit;
wimp_msghelprequest helprequest;
wimp_msghelpreply helpreply;
wimp_msgprint print;
} data;
} wimp_msgstr;

```

wimp_eventdata

```

typedef union {
wimp_openstr o;                for redraw, close, enter, leave events
struct {
    wimp_mousestr m;
    wimp_bbits b;} but;        for button change event
wimp_box dragbox;              for user drag box event
struct {wimp_caretstr c; int chcode;} key; for key events
int menu[10];                  for menu event: terminated by -1
struct {wimp_openstr o; int x, y;} scroll;for scroll request
                                x=-1 for left, +1 for right
                                y=-1 for down, +1 for up
                                scroll by +/-2 -> page scroll request
wimp_caretstr c;                for caret gain/lose
wimp_msgstr msg;                for messages
} wimp_eventdata;

```

wimp_eventstr

Wimp event description.

```
typedef struct {
wimp_etype e;                event type
wimp_eventdata data;
} wimp_eventstr;
```

wimp_menuhdr

```
typedef struct {
char title[12];              menu title (optional)
char tit_fcol, tit_bcol, work_fcol, work_bcol;  colours
int width, height;          size of following menu items
int gap;                     vertical gap between items
} wimp_menuhdr;
```

wimp_menuflags

Use `wimp_INOSELECT` to shade the item as unselectable, and the `button` type to mark it as writeable.

```
typedef enum {
wimp_MTICK = 1,
wimp_MSEPARATE = 2,
wimp_Mwriteable = 4,
wimp_MSUBLINKMSG = 8,      show a => flag, and inform program
                             when it is activated
wimp_MLAST = 0x80         signal last item in the menu
} wimp_menuflags;
```

wimp_menuptr

Only for the circular reference in `menuitem/str`.

```
typedef struct wimp_menustr *wimp_menuptr;
```


wimp_menuitem

Submenu can also be a `wimp_w`, in which case the window is opened as a dialogue box within the menu tree.

```
typedef struct {  
    wimp_menuflags flags;           menu entry flags  
    wimp_menuptr submenu;         wimp_menustr* pointer to sub  
                                   menu, or wimp_w dialogue box, or  
                                   -1 if no submenu  
    wimp_iconflags iconflags;     icon flags for the entry  
    wimp_icondata data;          icon data for the entry  
} wimp_menuitem;
```

wimp_menustr

```
typedef struct {  
    wimp_menuhdr hdr;             zero or more menu items follow in  
                                   memory  
} wimp_menustr;
```

wimp_dragstr

```
typedef struct {  
    wimp_w window;               initial position for drag box  
    wimp_dragtype type;          parent box for drag box  
    wimp_box box;                initial position for drag box  
    wimp_box parent;            parent box for drag box  
} wimp_dragstr;
```

wimp_which_block

```
typedef struct {  
    wimp_w window;               handle  
    int bit_mask;                bit set => consider this bit  
    int bit_set;                 desired bit setting  
} wimp_which_block;
```

wimp_pshapestr

```
typedef struct {
    int shape_num;

    char *shape_data;

    int width, height;

    int activex, activey;
} wimp_pshapestr;
```

pointer shape number (0 turn off pointer)

shape data, NULL pointer implies existing shape

Width and height in pixels Width = 4n, where n is an integer.

active point (pixels from top left)

wimp_font_array

```
typedef struct {
    char f[256];

} wimp_font_array;
```

initialise all to zero before using for first `load_template`, then just use repeatedly without altering

wimp_template

Template reading structure

```
typedef struct {
    int reserved;

    wimp_wind *buf;

    char *work_free;

    char *work_end;

    wimp_font_array *font;

    char *name;
```

ignore – implementation detail

pointer to space for putting template in

pointer to start of free Wimp workspace – you have to provide the Wimp system with workspace to store its redirected icons in end of workspace you are offering to the Wimp

points to font reference count array; 0 pointer implies fonts not allowed

name to match with (can be wildcarded)

```

int index;                                position in index to search from (0
                                           = start)

} wimp_template;

```

wimp_paletteword

The gcol char (least significant) is a gcol colour except in 8-bpp modes, when bits 0..2 are the tint and bits 3..7 are the gcol colour.

```

typedef union {
struct {char gcol; char red; char green; char blue;}
    bytes;
int word;
} wimp_paletteword;

```

wimp_palettestr

```

typedef struct {
wimp_paletteword c[16];                Wimp colours 0..15
wimp_paletteword screenborder, mouse1, mouse2, mouse3;
} wimp_palettestr;

```

Function prototypes**wimp_initialise**

```
os_error *wimp_initialise(int *v)
```

Closes and deletes all windows, returning Wimp version number.

wimp_taskinit

```
os_error *wimp_taskinit(char *name, int *version, wimp_t
*t)
```

name is the name of the program. Used instead of wimp_initialise. Returns your task handle. Version should be at least 200 on entry, and is set to the current wimp version number on return.

wimp_create_wind

```
os_error *wimp_create_wind(wimp_wind *, wimp_w *)
```

Defines (but does not display) window, returning window handle.

wimp_create_icon

```
os_error *wimp_create_icon(wimp_icode *, wimp_i *result)
```

Adds icon definition to that of window, returning icon handle.

wimp_delete_wind

```
os_error *wimp_delete_wind(wimp_w)
```

wimp_delete_icon

```
os_error *wimp_delete_icon(wimp_w, wimp_i)
```

wimp_open_wind

```
os_error *wimp_open_wind(wimp_openstr *)
```

Makes a window appear on the screen.

wimp_close_wind

```
os_error *wimp_close_wind(wimp_w)
```

Removes from the active list the window with its handle in the integer argument.

wimp_poll

```
os_error *wimp_poll(wimp_emask mask, wimp_eventstr  
*result)
```

Polls the next event from the Wimp.

wimp_save_fp_state_on_poll (void)

```
os_error *wimp_save_fp_state_on_poll(void)
```

Activates the saving of the floating point state on calls to `wimp_poll` and `wimp_pollidle`; this is needed if you do any floating point at all, as other programs may corrupt the FP status word, which is effectively a global in your program.

wimp_corrupt_fp_state_on_poll (void)

```
void *wimp_corrupt_fp_state_on_poll(void)
```

Disables the saving of the floating point state on calls to `wimp_poll` and `wimp_pollidle`; use only if you never use FP at all.

wimp_redraw_wind

```
os_error *wimp_redraw_wind(wimp_redrawstr*, BOOL*)
```

Draws a window outline and icons. Return False if there's nothing to draw.

wimp_update_wind

```
os_error *wimp_update_wind(wimp_redrawstr*, BOOL*)
```

Returns the visible portion of a window. Returns False if there's nothing to redraw.

wimp_get_rectangle

```
os_error *wimp_get_rectangle(wimp_redrawstr*, BOOL*)
```

Returns the next rectangle in the list, or False if done.

wimp_get_wind_state

```
os_error *wimp_get_wind_state(wimp_w, wimp_wstate  
*result)
```

Reads the current window state.

wimp_get_wind_info

```
os_error *wimp_get_wind_info(wimp_winfo *result)
```

On entry `result->w` gives the window in question. Space for any icons must follow `*result`.

wimp_set_icon_state

```
os_error *wimp_set_icon_state(wimp_w, wimp_i,  
wimp_iconflags value, wimp_iconflags mask)
```

Sets an icon's flags as `(old_state & ~mask) ^ value`.

wimp_get_icon_info

```
os_error *wimp_get_icon_info(wimp_w, wimp_i, wimp_icon  
*result)
```

Gets the current state of an icon.

wimp_get_point_info

```
os_error *wimp_get_point_info(wimp_mousestr *result)
```

Gives information regarding the state of the mouse.

wimp_drag_box

```
os_error *wimp_drag_box(wimp_dragstr *)
```

Starts the Wimp dragging a box.

wimp_force_redraw

```
os_error *wimp_force_redraw(wimp_redrawstr *r)
```

Marks an area of the screen as invalid. If `r->wimp_w == -1`, use screen coordinates. Only the first five fields of `r` are valid.

wimp_set_caret_pos

```
os_error *wimp_set_caret_pos(wimp_caretstr *)
```

Sets the position and size of the text caret.

wimp_get_caret_pos

```
os_error *wimp_get_caret_pos(wimp_caretstr *)
```

Gets the position and size of the text caret.

wimp_create_menu

```
os_error *wimp_create_menu(wimp_menstr *m, int x, int y)
```

'Pops up' a menu structure. Set `m==(wimp_menstr*)-1` to clear the menu tree.

wimp_decode_menu

```
os_error *wimp_decode_menu(wimp_menstr *, void *, void *)
```

wimp_which_icon

```
os_error *wimp_which_icon(wimp_which_block *, wimp_i  
*results)
```

The `results` appear in an array, terminated by a `(wimp_i) -1`.

wimp_set_extent

`os_error *wimp_set_extent(wimp_redrawstr *)`

Alters the extent of a window's work area – only the handle and the first set of four coordinates are looked at.

wimp_set_point_shape

`os_error *wimp_set_point_shape(wimp_pshapestr *)`

Sets the pointer shape on screen.

wimp_open_template

`os_error *wimp_open_template(char *name)`

Opens the named file to allow `load_template` to read a template from the file.

wimp_close_template

`os_error *wimp_close_template(void)`

Closes the currently open template file.

wimp_load_template

`os_error *wimp_load_template(wimp_template *)`

Loads a window template from an open file into buffer.

wimp_processkey

`os_error *wimp_processkey(int chcode)`

Hands back to the Wimp a key that you do not understand.

wimp_closedown

`os_error *wimp_closedown(void)`

wimp_taskclose

`os_error *wimp_taskclose(wimp_t)`

Calls `closedown` in the multi-tasking form.

wimp_starttask

`os_error *wimp_starttask(char *clicmd)`

Starts a new Wimp task, with the given CLI command.

wimp_getwindowoutline

```
os_error *wimp_getwindowoutline(wimp_redrawstr *r)
```

Sets $r \rightarrow w$ on entry. On exit, $r \rightarrow box$ will be the screen coordinates of the window, including border, title, scroll bars.

wimp_pollidle

```
os_error *wimp_pollidle(wimp_emask mask, wimp_eventstr  
*result, int earliest)
```

Like `wimp_poll`, but does not return before the earliest return time. This is a value produced by `OS_ReadMonotonicTime`.

wimp_ploticon

```
os_error *wimp_ploticon(wimp_icon*)
```

Called only within an update or redraw loop, and just does the plotting. This need not be a real icon attached to a window.

wimp_setmode

```
os_error *wimp_setmode(int mode)
```

Sets the screen mode. Palette colours are maintained, if possible.

wimp_readpalette

```
os_error *wimp_readpalette(wimp_palettestr*)
```

wimp_setpalette

```
os_error *wimp_setpalette(wimp_palettestr*)
```

The `bytes.gcol` values of each field of the `palettestr` are ignored; only the absolute colours are taken into account.

wimp_setcolour

```
os_error *wimp_setcolour(int colour)
```

bits 0...3 = Wimp colour (translate for current mode)

4...6 = gcol action

7 = foreground/background.

wimp_spriteop

```
os_error *wimp_spriteop(int reason_code, char *name)
```

Calls SWI Wimp_SpriteOp.

wimp_spriteop_full

```
os_error *wimp_spriteop_full(os_regset *)
```

Calls SWI Wimp_SpriteOp allowing full information to be passed.

wimp_baseofsprites

```
void *wimp_baseofsprites(void)
```

Returns a `sprite_area*`, which may be moved about by `mergespritefile`.

wimp_blockcopy

```
os_error *wimp_blockcopy(wimp_w, wimp_box *source, int x,  
int y)
```

Copies the source box (defined in window coordinates) to the given destination (in window coordinates). Invalidates any portions of the destination that cannot be updated using on-screen copy.

wimp_errflags

```
typedef enum {  
    wimp_EOK = 1,                put in OK box  
    wimp_ECANCEL = 2,           put in CANCEL box  
    wimp_EHICANCEL = 4,         highlight CANCEL rather than OK  
} wimp_errflags;
```

If OK and CANCEL are both 0 you get an OK.

wimp_reporterror

```
os_error *wimp_reporterror(os_error*, wimp_errflags, char  
*name)
```

Produces an error window. Uses sprite called `error` in the Wimp sprite pool. `name` should be the program name, appearing after `error` in at the head of the dialogue box.

wimp_sendmessage

```
os_error *wimp_sendmessage(wimp_etype code, wimp_msgstr*
msg, wimp_t dest)
```

dest can also be 0, in which case the message is sent to every task in turn, including the sender. msg can also be any other wimp_eventdata* value.

wimp_sendwmessage

```
os_error *wimp_sendwmessage(wimp_etype code, wimp_msgstr
*msg, wimp_w w, wimp_i i)
```

Sends a message to the owner of a specific window or icon. msg can also be any other wimp_eventdata* value.

wimp_create_submenu

```
os_error *wimp_create_submenu(wimp_menstr *sub, int x, int
y)
```

sub can also be a wimp_w, in which case it is opened by the Wimp as a dialogue box.

wimp_slotsize

```
os_error *wimp_slotsize (int *currentslot,
                          int *nextslot,
                          int *freepool)
```

currentslot/nextslot==0 -> just read setting.

wimp_transferblock

```
os_error *wimp_transferblock(
    wimp_t sourcetask,
    char *sourcebuf,
    wimp_t desttask,
    char *destbuf,
    int buflen)
```

Transfers memory between domains.

wimp_setfontcolours

```
os_error *wimp_setfontcolours(int foreground, int
background)
```

Sets font manager colours. The Wimp handles how many shades etc. to use.

wimp_readpixtrans

```
os_error *wimp_readpixtrans(sprite_area *area, sprite_id
*id, sprite_factors *factors, sprite_pixtrans *pixtrans)
```

Tells you how the Wimp will plot a sprite when asked to PutSpriteScaled.

wimp_command_tag

```
typedef enum {
wimp_command_TITLE = 0,
wimp_command_ACTIVE = 1,
wimp_command_CLOSE_PROMPT = 2,
wimp_command_CLOSE_NOPROMPT = 3
} wimp_command_tag;
```

wimp_commandwind

```
typedef struct {
wimp_command_tag tag;
char *title
} wimp_commandwind;
```

wimp_commandwindow

```
os_error *wimp_commandwindow(wimp_commandwind
commandwindow)
```

Opens a text window for normal VDU 4-type output. The tag types correspond to the four kinds of call to SWI wimp_CommandWindow described in the RISC OS *Programmer's Reference Manual*. `title` is only required if `tag == wimp_command_TITLE`. It is the application's responsibility to set the tag correctly.

wimpt

These functions provide low-level Wimp functionality.

wimpt_poll

Polls for an event from the Wimp (with extras to buffer one event).

Syntax: `os_error *wimpt_poll(wimp_emask mask, wimp_eventstr *result)`

Parameters: `wimp_emask mask` – ignore events in the mask
`wimp_eventstr *result` – the event returned from Wimp

Returns: possible error condition.

Other Information: If you want to poll at this low level (ie avoiding `event_process()`), use this function rather than `wimpt_poll`. Using `wimpt_poll` allows you to use the routines shown below.

wimpt_fake_event

Posts an event to be collected by `wimpt_poll`.

Syntax: `void wimpt_fake_event(wimpt_eventstr *)`

Parameters: `wimpt_eventstr` – the posted event

Returns: `void`

Other Information: use with care!

wimpt_last_event

Informs the caller of the last event returned by `wimpt_poll`.

Syntax: `wimpt_eventstr *wimpt_last_event(void)`

Parameters: `void`

Returns: pointer to last event returned by `wimpt_poll`.

wimpt_last_event_was_a_key

Informs the caller if the last event returned by `wimpt_poll` was a key stroke.

Syntax: `int wimpt_last_event_was_a_key(void)`

Parameters: `void`

Returns: non-zero if last event was a keystroke.

Other Information: retained for backwards compatibility. Use `wimpt_last_event` for preference, and test if `e` field of returned struct == `wimpt_EKEY`.

wimpt_noerr

Halts the program and reports an error in a dialogue box (if `e!=0`).

Syntax: `void wimpt_noerr(os_error *e)`

Parameters: `os_error *e` – error return from system call

Returns: `void`.

Other Information: Useful for 'wrapping up' system calls which are not expected to fail; if failure occurs, your program probably has a logical error. Call when an error would mean disaster; for example:

```
wimpt_noerr(some_system_call(.....));
```

The error message is:

```
ProgName has suffered a fatal internal error  
(errormessage) and must exit immediately.
```

wimpt_complain

Reports an error in a dialogue box (if e!=0).

Syntax: `os_error *wimpt_complain(os_error *e)`

Parameters: `os_error *e` – error return from system call

Returns: the error returned from the system call (ie e).

Other Information: Useful for 'wrapping up' system calls which may fail. Call when your program can still limp on regardless (taking some appropriate action).

wimpt: control of graphics environment

wimpt_checkmode

Registers the current screen mode with the wimpt module.

Syntax: `BOOL wimpt_checkmode(void)`

Parameters: `void`

Returns: True if screen mode has changed.

wimpt_mode

Reads the screen mode.

Syntax: `int wimpt_mode(void)`

Parameters: `void`

Returns: screen mode.

Other Information: faster than a normal OS call. Value is only valid if `wimpt_checkmode` is called at redraw events.

wimpt_dx/wimpt_dy

Informs the caller of OS x/y units per screen pixel.

Syntax: `int wimpt_dx(void)`

`int wimpt_dy(void)`

Parameters: `void`

Returns: OS x/y units per screen pixel.

Other Information: faster than a normal OS call. Value is only valid if `wimpt_checkmode` is called at redraw events.

wimpt_bpp

Informs the caller of bits per screen pixel.

Syntax: `int wimpt_bpp(void)`

Parameters: `void`

Returns: bits per screen pixel (in current mode)

Other Information: faster than a normal OS call. Value is only valid if `wimpt_checkmode` is called at redraw events.

wimpt_init

Set program up as a Wimp task.

Syntax: `int wimpt_init(char *programname)`

Parameters: `char *programname` – name of your program

Returns: the current wimp version number.

Other Information: Remembers screen mode, and sets up signal handlers so that task exits cleanly, even after fatal errors. Response to signals SIGABRT, SIGFPE, SIGILL, SIGSEGV and SIGTERM is to display error box with message:

```
progname has suffered an internal error (type = signal)
and must exit immediately
```

SIGINT (Escape) is ignored. *progname* will appear in the Task manager display and in error messages. Calls `wimpt_taskinit` and stores `task_id` returned. Also installs exit-handler to close down task when program calls `exit()` function.

wimpt_wimpversion

Tell wimpt what version of the wimp you understand.

Syntax: `void wimpt_wimpversion(int version);`

Parameters: `int` – the version number of the wimp that you understand.

Returns: `void`.

Other Information: Call this routine before calling `wimpt_init`, if you know about the features in a Wimp beyond version 2.00.

This argument will then be passed to `wimpt_init`, allowing the Wimp to understand what facilities you know about. Then call `wimpt_init`, allowing the wimp to return its current version number.

wimpt_programname

Informs the caller of the name passed to `wimpt_init`.

Syntax: `char *wimpt_programname(void)`

Parameters: `void`.

Returns: pointer to the program's name.

wimpt_reporterror

Reports an OS error in a dialogue box (including program name).

Syntax: `void wimpt_reporterror(os_error*, wimp_errflags)`

Parameters: `os_error*` – OS error block
`wimp_errflags` – flag whether to include OK and/or CANCEL (highlighted or not) button in dialogue box

Returns: `void`.

Other Information: similar to `wimp_reporterror()`, but includes the program name automatically (eg the one passed to `wimpt_init`).

wimpt_task

Informs the caller of its task handle.

Syntax: `wimp_t wimpt_task(void)`

Parameters: `void`

Returns: task handle.

wimpt_forceredraw

Causes the whole screen to be invalidated (running applications will be requested to redraw all windows).

Syntax: `void wimpt_forceredraw(void)`

Parameters: `void`.

Returns: `void`.

win

This file offers central management of RISC OS windows, constructing a very simple idea of 'window class' within RISC OS. RISC OS window class implementations register the existence of each window with this module.

This structure allows event-processing loops to be constructed that have no knowledge of what other modules are present in the program. For instance, the dialogue box module can contain an event-processing loop without reference to what other window types are present in the program.

Claiming Events

win_register_event_handler

Installs an event handler function for a given window.

Syntax: `void win_register_event_handler(wimp_w, win_event_handler, void *handle)`

Parameters: `wimp_w` – the window's handle
`win_event_handler` – the event handler function
`void *handle` – caller-defined handle.

Returns: `void`.

Other Information: This call has no effect on the window itself – it just informs the `win` module that the supplied function should be called when events are delivered to the window. To remove a handler, call with a null function pointer:

```
win_register_event_handler(w, (win_event_handler)0, 0)
```

To catch key events for an icon on the icon bar, register a handler for `win_ICONBAR`:

```
win_event_handler(win_ICONBAR, handler_func, handle)
```

To catch load events for an icon on the icon bar, register a handler for `win_ICONBARLOAD`:

```
win_event_handler(win_ICONBARLOAD, load_func, handle)
```

win_read_event_handler

Read current event handler for a given window, and the handle which it is passed.

Syntax: `BOOL win_read_eventhandler(wimp_w w, win_event_handler *p, void **handle);`

Parameters: `wimp_w w` – the window's handle
`win_event_handler *p` – the handler function
`void **handle` – the handle passed to the handler function

Returns: `TRUE` if given window is registered, `FALSE` otherwise

Other Information: This is useful for registering an alternative event handler which can vet events, before passing them on to the original handler.

win_claim_idle_events

Causes 'idle' events to be delivered to a given window.

Syntax: `void win_claim_idle_events(wimp_w)`

Parameters: `wimp_w` – the window's handle.

Returns: `void`.

Other Information: To cancel this, call with window handle `(wimp_w) - 1`.

Note that idle (or null) events will not be delivered to your window unless you also enable null events using `event_setmask(0)`.

win_add_unknown_event_processor

Adds a handler for unknown events onto the front of the queue of such handlers.

Syntax: `void win_add_unknown_event_processor
(win_unknown_event_processor, void *handle)`

Parameters: `win_unknown_event_processor` – handler function
`void *handle` – passed to handler on call.

Returns: `void`.

Other Information: The `win` module maintains a list of unknown event handlers. An unknown event results in the 'head of the list' function being called; if this function doesn't deal with the event it is passed on to the next in the list, and so on. Handler functions should return a Boolean result to show if they dealt with the event, or if it should be passed on. 'Known' events are as follows:

ENULL, EREDRAW, ECLOSE, EOPEN, EPTRELEAVE, EPTRENTER, EKEY, ESCROLL, EBUT and ESEND/ESENDWANTACK for the following msg types: MCLOSEDOWN, MDATASAVE, MDATALOAD, MHELPPREQUEST

All other events are considered 'unknown'. If none of the unknown event handlers deals with the event, then it is passed on to the unknown event claiming window (registered by `win_claim_unknown_events()`). If there is no such claimer, then the unknown event is ignored.

win_remove_unknown_event_processor

Removes the given unknown event handler with the given handle from the stack of handlers.

Syntax: `void win_remove_unknown_event_processor
(win_unknown_event_processor, void *handle)`

Parameters: `win_unknown_event_processor` – the handler to be removed
`void *handle` – its handle.

Returns: `void`.

Other Information: The handler to be removed can be anywhere in the stack (not necessarily at the top).

win_idle_event_claimer

Informs the caller of which window is claiming idle events.

Syntax: `wimp_w win_idle_event_claimer(void)`

Parameters: `void`

Returns: Handle of window claiming idle events.

Other Information: Returns `(wimp_w) - 1`, if no window is claiming idle events.

win_claim_unknown_events

Cause any unknown, or non-window-specific events to be delivered to a given window.

Syntax: `void win_claim_unknown_events(wimp_w)`

Parameters: `wimp_w` – handle of window to which unknown events should be delivered.

Returns: `void`.

Other Information: Calling with `(wimp_w) - 1` cancels this. See `win_add_unknown_event_processor()` for details of which events are 'known'.

win_unknown_event_claimer

Informs the caller of which window is claiming unknown events.

Syntax: `wimp_w win_unknown_event_claimer(void)`

Parameters: `void`

Returns: Handle of window claiming unknown events.

Other Information: Return of `(wimp_w) - 1` means no claimer registered.

win: menus

win_setmenuh

Attaches the given menu structure to the given window.

Syntax: `void win_setmenuh(wimp_w, void *handle)`

Parameters: `wimp_w` – handle of window
`void *handle` – pointer to menu structure.

Returns: `void`.

Other Information: Mainly used by higher level RISC_OSLib routines to attach menus to windows (eg `event_attachmenu()`).

win_getmenuh

Returns a pointer to the menu structure attached to the given window.

Syntax: `void *win_getmenuh(wimp_w)`

Parameters: `wimp_w` – handle of window

Returns: pointer to the attached menu (0 if no menu attached).

Other Information: As for `win_setmenuh()`, this is used mainly by higher level RISC OS routines (eg `event_attachmenu()`).

win: event processing

win_processevent

Delivers an event to its relevant window, if such a window has been registered with this module (via `win_register_event_handler()`).

Syntax: `BOOL win_processevent(wimp_eventstr*)`

Parameters: `wimp_eventstr*` – pointer to the event which has occurred

Returns: True if an event handler (registered with this module) has dealt with the event, False otherwise.

Other Information: the main client for this routine is `event_process()`, which uses it to deliver an event to its appropriate window. Keyboard events are delivered to the current owner of the caret.

win: termination

win_activeinc

Increment by one the win module's idea of the number of active windows owned by a program.

Syntax: `void win_activeinc(void)`

Parameters: `void`

Returns: `void`.

Other Information: `event_process()` calls `exit()` on behalf of the program when the number of active windows reaches zero. Programs which wish to remain running even when they have no active windows should ensure that `win_activeinc()` is called once before creating any windows, so that the number of active windows is always ≥ 1 . This is done for you if you use `baricon()` to install your program's icon on the icon bar.

win_activedec

Decrements by one the win module's idea of the number of active windows owned by a program.

Syntax: `void win_activedec(void)`

Parameters: `void`.

Returns: `void`.

Other Information: See the note in `win_activeinc()` regarding program termination.

win_activeo

Informs the caller of the number of active windows owned by your program.

Syntax: `int win_activeo(void)`

Parameters: `void`.

Returns: number of active windows owned by the program.

Other Information: This is given by (number of calls to `win_activeinc()`) minus (number of calls to `win_activedec()`). Note that modules in the RISC OS library itself may have made calls to `win_activeinc()` and `win_activedec()`.

win_give_away_caret

Gives the caret away to the open window at the top of the Wimp's window stack (if that window is owned by your program).

Syntax: `void win_give_away_caret(void)`

Parameters: `void`.

Returns: `void`.

Other Information: If the top window is interested it will take the caret. If not then nothing happens. This only works if polling is done using the `wimpt` module, which is the case if your main inner loop goes something like: `while (TRUE) event_process()`.

win_settitle

Changes the title displayed in a given window.

Syntax: `void win_settitle(wimp_w w, char *newtitle);`

Parameters: `wimp_w w` – given window's handle
`char *newtitle` – null-terminated string giving new title for window.

Returns: `void`.

Other information: The title icon of the given window must be indirected text. This will change the title used by all windows created from the given window's template if you have used the template module (since the Window Manager uses your address space to hold indirected text icons). To avoid this, the window can be created from a copy of the template, ie

```
template *t = template_copy(template_find("name"));  
wimp_create_wind(t->window, &w);
```

win_init

Initialise the centralised window event system

Syntax: `BOOL win_init(void);`

Parameters: `void`

Returns: `TRUE` if initialisation went OK.

Other Information: If you use `wimpt_init()`, to start your application, then this call is made for you.

xferrecv

This file covers the general purpose importing of data by dragging icons.

xferrecv_checkinsert

Sets up the acknowledge message for a MDATAOPEN or MDATALOAD and gets the filename to load from.

Syntax: `int xferrecv_checkinsert(char **filename)`
 Parameters: `char **filename` – returned pointer to filename.
 Returns: the file's type (eg 0x0fff for Edit).
 Other Information: This function checks to see if the last Wimp event was a request to import a file. If it was, the function returns file type and a pointer to file's name is put into `*filename`. Otherwise, it returns -1.

xferrecv_insertfileok

Deletes the scrap file (if used for transfer), and sends acknowledgement of MDATALOAD message.

Syntax: `void xferrecv_insertfileok(void)`
 Parameters: `void`
 Returns: `void`.

xferrecv_checkprint

Sets up an acknowledge message to a MPrintTypeOdd message and gets the filename to print.

Syntax: `int xferrecv_checkprint(char **filename)`
 Parameters: `char **filename` – returned pointer to filename.
 Returns: The file's type (eg 0x0fff for Edit).
 Other Information: The application can either print the file directly or convert it to Printer\$Temp for printing by the printer application.

xferrecv_printfileok

Sends an acknowledgement back to the printer application. If a file is sent to Printer\$Temp, this also fills in the file type in the message.

Syntax: `void xferrecv_printfileok(int type)`
 Parameters: `int type` – type of file sent to Printer\$Temp (eg 0x0fff for Edit).
 Returns: `void`.

xferrecv_checkimport

Sets up an acknowledgement message to a MDATASAVE message.

Syntax: `int xferrecv_checkimport(int *estsize)`
 Parameters: `int *estsize` – sender's estimate of file size
 Returns: File type.

xferrecv_buffer_processor

This is a typedef for the caller-supplied function to empty a full buffer during data transfer.

Syntax: `typedef BOOL (*xferrecv_buffer_processor)(char **buffer, int *size)`
 Parameters: `char **buffer` – new buffer to be used
`int *size` – updated size.
 Returns: False if unable to empty buffer or create new one.
 Other Information: This is the function, supplied by the application, which will be called when the buffer is full. It should empty the current buffer, or create more space and modify size accordingly, or return False. `*buffer` and `*size` are the current buffer and its size on function entry.

xferrecv_doimport

Loads data into a buffer, and calls the caller-supplied function to empty the buffer when full.

Syntax: `int xferrecv_doimport(char *buf, int size, xferrecv_buffer_processor)`
 Parameters: `char *buf` – the buffer
`int size` – buffer's size
`xferrecv_buffer_processor` – caller-supplied function to be called when the buffer is full.
 Returns: Number of bytes transferred on successful completion; –1 otherwise.

xferrecv_file_is_safe

Informs the caller if the file was received from a 'safe' source (see below for what this means).

Syntax: `BOOL xferrecv_file_is_safe(void)`
 Parameters: `void`
 Returns: True if file is safe.

Other Information: 'Safe' in this context means that the supplied filename will not change in the foreseeable future.

xfer_send

This file covers the general purpose export of data by dragging icons.

xfer_send: caller-supplied function types

xfer_send_saveproc

A function of this type should save to the given file and return True if successful. Handle is passed to the function by `xfer_send()`.

Syntax: `typedef BOOL (*xfer_send_saveproc)(char *filename, void *handle)`

Parameters: `char *filename` – file to be saved
`void *handle` – the handle you passed to `xfer_send()`.

Returns: True if the save was successful.

xfer_send_sendproc

A function of this type should call `xfer_send_sendbuf()` to send one buffer-full of data no bigger than `*maxbuf`.

Syntax: `typedef BOOL (*xfer_send_sendproc)(void *handle, int *maxbuf)`

Parameters: `void *handle` – handle which was passed to `xfer_send()`
`int *maxbuf` – size of receiver's buffer.

Returns: True if the data was successfully transmitted.

Other Information: Your `sendproc` will be called by functions in the `xfer_send` module to do an in-core data transfer, on receipt of MRAMFetch messages from the receiving application. If `xfer_send_sendbuf()` returns False, then return False **immediately**.

xfer_send_printproc

A function of this type should either print the file directly, or save it into the given filename, from where it will be printed by the printer application.

Syntax: `typedef int (*xfer_send_printproc)(char *filename, void *handle)`

Parameters: `char *filename` – file to save into, for printing
 `void *handle` – handle that was passed to
 `xfersend()`

Returns: Either the file type of the file it saved, or one of the reason
 codes #defined below.

Other Information: This is called if the file icon has been dragged onto a
 printer application.

Reason codes:

```
#define xfersend_printPrinted -1    file dealt with internally  
#define xfersend_printFailed -2    had an error along the way
```

The `saveproc` should report any errors it encounters itself. If saving to a file, it should convert the data into a type that can be printed by the printer application (ie text).

xfersend: library functions

xfersend

Allows the user to export application data, by icon drag.

Syntax: `BOOL xfersend(int filetype, char *name, int estsize,
 xfersend_saveproc, xfersend_sendproc, xfersend_printproc,
 wimp_eventstr *e, void *handle)`

Parameters: `int filetype` – type of file to save to
 `char *name` – suggested file name
 `int estsize` – estimated size of the file
 `xfersend_saveproc` – caller-supplied function for
 saving application data to a file
 `xfersend_sendproc` – caller-supplied function for
 in-core data transfer (if application is able to do this)
 `xfersend_printproc` – caller-supplied function for
 printing application data, if icon is dragged onto printer
 application
 `wimp_eventstr *e` – the event which started the
 export (usually mouse drag)
 `void *handle` – handle to be passed to handler
 functions.

Returns: True if data exported successfully.

Other Information: You should typically call this function in a window's event
 handler, when you get a mouse drag event. See the
 `saveas.c` code for an example of this. `xfersend` deals
 with the complexities of message-passing protocols to

achieve the data transfer. Refer to the above type definitions for an explanation of what the three caller-supplied functions should do.

If name is 0 then a default name of Selection is supplied.

If you pass 0 as the xfersend_sendproc, no in-core data transfer will be attempted.

If you pass 0 as the xfersend_printproc, the file format for printing is assumed to be the same as for saving. The estimated file size is not essential, but may improve performance.

xfersend_pipe

Allows the user to export application data, without an icon drag.

Syntax: `BOOL xfersend_pipe(int filetype, char *name, int estsize, xfersend_saveproc, xfersend_sendproc, xfersend_printproc, void *handle, wimp_t task);`

Parameters:

- `int filetype` – type of file to save to
- `char *name` – suggested file name
- `int estsize` – estimated size of the file
- `xfersend_saveproc` – caller-supplied function for saving application data to a file
- `xfersend_sendproc` – caller-supplied function for in-core data transfer (if application is able to do this)
- `xfersend_printproc` – caller-supplied function for printing application data, if "icon" is dragged onto printer application
- `void *handle` – handle to be passed to handler functions.
- `wimp_t task` – handle of task to pass data to.

Returns: TRUE if data exported successfully.

Other Information: This function works similarly to xfersend, except it is not normally used as the result of an icon drag. Typical use may be to export data to another application (using the same technique as xfersend), following a request for data from that application (maybe as a result of receiving an application-specific wimp message).

xfersend_sendbuf

Sends the given buffer to a receiver.

Syntax: `BOOL xfersend_sendbuf(char *buffer, int size)`

Parameters: `char *buffer` – the buffer to be sent
 `int size` – the number of characters placed in the
 buffer.
Returns: True if send was successful.
Other Information: This function should be called by the caller-supplied
 `xfersend_sendproc` (if such exists) to do in-core data
 transfer (see notes on `xfersend_sendproc` above).

xfersend_file_is_safe

 Informs the caller if the file's name can be reliably assumed not to change (during
 data transfer!).

Syntax: `BOOL xfersend_file_is_safe(void)`
Parameters: `void`.
Returns: True if file is 'safe'.
Other Information: See also the `xferrecv` module.
Returns: True if file recipient will not modify it; changing the
 window title of the file can be done conditionally on this
 result. This can be called within your
 `xfersend_saveproc`, `sendproc`, or `printproc`, or
 immediately after the main `xfersend`.

xfersend_set_fileissafe

 Allows the caller to set an indication of whether a file's name will remain
 unchanged during data transfer.

Syntax: `void xfersend_set_fileissafe(BOOL value)`
Parameters: `BOOL value` – True means the file is safe.
Returns: `void`.

xfersend_close_on_xfer

 Tells `xfersend` whether to close "parent" window after icon-drag export.

Syntax: `void xfersend_close_on_xfer(BOOL do_we_close, wimp_w w);`
Parameters: `BOOL do_we_close` – TRUE means close window after
 export.
 `wimp_w w` – handle of window to close (presumably
 "parent" window).
Returns: `void`
Other Information: The default is to not close the window after export. Once
 used, this function should be called before each call to
 `xfersend()`.

xfersend_clear_unknowns

Removes any unknown event processors registered by `xfersend` or `xfersend_pipe`.

Syntax: `void xfersend_clear_unknowns(void);`

Parameters: `void`

Returns: `void`

Other Information: `xfersend` and `xfersend_pipe` use unknown event processors to deal with inter-application data transfer. These may be left around after completion of the transfer (especially if the transfer failed). This function should be called when it is known that the transfer has ended.

xfersend_read_last_ref

Returns the `my_ref` value of the last `wimp_MDATASAVE` or `wimp_MDATALOAD` message sent by `xfersend` or `xfersend_pipe`.

Syntax: `int xfersend_read_last_ref(void);`

Parameters: `void`.

Returns: integer message reference

Other Information: After saving a file to another application (ie where the resulting file is not 'safe', the `my_ref` value of the final `wimp_MDATALOAD` should be stored with the document data, so that if a `wimp_MDATASAVED` is received, the document can be marked unmodified. If the document is modified after being saved, the `last_ref` value should be reset to 0, so that a subsequent `wimp_MDATASAVED` message will not cause the document to be marked unmodified. NB: If RAM transfer is used, the `my_ref` of the `datasave` message should be stored instead.

13 Assembly language interface

Interworking assembly language and C – writing programs with both assembly language and C parts – requires use of both the Acorn Desktop Assembler and Acorn Desktop C products for anything other than trying the examples supplied with Acorn Desktop C. Further explanation of examples is provided in the chapter entitled *Interworking assembler with C* in *Acorn Assembler Release 2* supplied with Acorn Desktop Assembler.

Interworking assembly language and C can be very useful for construction of top quality RISC OS applications. Using this technique you can take advantage of many of the strong points of both languages. Writing most of the bulk of your application in C allows you to take advantage of the portability of C, the maintainability of a high level language and the power of the C libraries and language. Writing critical portions of code in assembler allows you to take advantage of all the speed of the Archimedes and all the features of the machine (eg use the complete floating-point instruction set).

The key to interworking C and assembler is writing assembly language procedures that obey the ARM Procedure Call Standard (APCS). This is a contract between two procedures, one calling the other. The called procedure needs to know which ARM and floating-point registers it can freely change without restoring them before returning, and the caller needs to know which registers it can rely on not being corrupted over a procedure call.

Additionally, both procedures need to know which registers contain input arguments and return arguments, and the arrangement of the stack has to follow a pattern that debuggers and so on can understand. For the specification of the APCS, see *Appendix F - ARM procedure call standard* in the *Acorn Desktop Development Environment* user guide.

This chapter explains how C uses the APCS, in terms of the appearance of assembly language optionally output by CC and the way the stack set up by the C run-time library works.

Register names

The following names are used in referring to ARM registers:

a1	R0	Argument 1, also integer result, temporary
a2	R1	Argument 2, temporary
a3	R2	Argument 3, temporary
a4	R3	Argument 4, temporary
v1	R4	Register variable
v2	R5	Register variable
v3	R6	Register variable
v4	R7	Register variable
v5	R8	Register variable
v6	R9	Register variable
s1	R10	Stack limit
fp	R11	Frame pointer
ip	R12	Temporary work register
sp	R13	Lower end of current stack frame
lr	R14	Link address on calls, or workspace
pc	R15	Program counter and processor status
f0	F0	Floating point result
f1	F1	Floating-point work register
f2	F2	Floating-point work register
f3	F3	Floating-point work register
f4	F4	Floating-point register variable (must be preserved)
f5	F5	Floating-point register variable (must be preserved)
f6	F6	Floating-point register variable (must be preserved)
f7	F7	Floating-point register variable (must be preserved)

In this section, 'at [r]' means at the location pointed to by the value in register r; 'at [r, #n]' refers to the location pointed to by r+n. This accords with ObjAsm's syntax.

Register usage

The following points should be noted about the contents of registers across function calls.

- Calling a function (potentially) corrupts the argument registers a1 to a4, ip, lr, and f0–f3. The calling function should save the contents of any of these registers it may need.
- Register lr is used at the time of a function call to pass the return link to the called function; it is not necessarily preserved during or by the function call.

- The stack pointer `sp` is not altered across the function call itself, though it may be adjusted in the course of pushing arguments inside a function. The limit register `s1` may change at any time, but should always represent a valid limit to the downward growth of `sp`. User code will not normally alter this register.
- Registers `v1` to `v6`, and the frame pointer `fp`, are expected to be preserved across function calls. The called procedure is responsible for saving and restoring the contents of any of these registers which it may need to use.

Control arrival

At a procedure call, the convention is that the registers are used as follows:

- `a1` to `a4` contain the first four arguments. If there are fewer than four arguments, just as many of `a1` to `a4` as are needed are used.
- If there are more than four arguments, `sp` points to the fifth argument; any further arguments will be located in succeeding words above `[sp]`.
- `fp` points to a backtrace structure.
- `sp` and `s1` define a temporary workspace of at least 256 bytes available to the procedure.
- `s1` contains a stack chunk handle, which is used by stack handling code to extend the stack in a non-contiguous manner.
- `lr` contains the value which should be restored into `pc` on exit from the called procedure.
- `pc` contains the entry address of the called procedure.

Passing arguments

All integral and pointer arguments are passed as 32-bit words. Floating point 'float' arguments are 32-bit values, 'double'-argument 64-bit values. These follow the memory representation of the IEEE single and double precision formats.

Arguments are passed **as if** by the following sequence of operations:

- Push each argument onto the stack, last argument first.
- Pop the first four words (or as many as were pushed, if fewer) of the arguments into registers `a1` to `a4`.
- Call the function, for example by the `branch with link` instruction:

```
BL functionname
```

In many cases it is possible to use a simplified sequence with the same effect (eg load three argument words into `a1-a3`).

If more than four words of arguments are passed, the calling procedure should adjust the stack pointer after the call, incrementing it by four for each argument word which was pushed and not popped.

Return link

On return from a procedure, the registers are set up as follows:

- `fp`, `sp`, `s1`, `v1` to `v6` and `f4` to `f7` have the same values that they contained at the procedure call.
- Any result other than a floating point or a multi-word structure value is placed in register `a1`.
- A floating point result should be placed in register `f0`.

Structure values returned as function results are discussed below.

Structure results

A C function which returns a multi-word structure result is treated in a slightly different manner from other functions by the compiler. A pointer to the location which should receive the result is added to the argument list as the first argument, so that a declaration such as the following:

```
s_type afunction(int a, int b, int c)
{
    s_type d;
    /* ... */
    return d;
}
```

is in effect converted to this form:

```
void afunction(s_type *p, int a, int b, int c)
{
    s_type d;
    /* ... */
    *p = d;
    return;
}
```

Any assembler-coded functions returning structure results, or calling such functions, must conform to this convention in order to interface successfully with object code from the C compiler.

Storage of variables

The code produced by the C compiler uses argument values from registers where possible; otherwise they are addressed relative to `fp`, as illustrated in *Examples* below.

Local variables, by contrast, are always addressed with positive offsets relative to `sp`. In code which alters `sp`, this means that the offset for the same variable will differ from place to place. The reason for this approach is that it permits the stack overflow procedure to recover by changing `sp` and `s1` to point to a new stack segment as necessary.

Function workspace

The values of `sp` and `s1` passed to a called function define an area of readable, writable memory available to the called function as workspace. All words below `[sp]` and at or above `[s1, #-512]` are guaranteed to be available for reading and writing, and the minimum allowed value of `sp` is `s1-256`. Thus the minimum workspace available is 256 bytes.

The C run-time system, in particular the stack extension code, requires up to 256 bytes of additional workspace to be left free. Accordingly, all called functions which require no more than 256 bytes of workspace should test that `sp` does not point to a location below `s1`, in other words that at least 512 bytes remain. If the value in `sp` is less than that in `s1`, the function should call the stack extension function `x$stack_overflow`. Functions which need more than 256 bytes of workspace should amend the test accordingly, and call `x$stack_overflow1`, as described below. The following examples illustrate a method of performing this test.

Note that these are the C-specific aliases for the kernel functions `_kernel_stkovf_split_0frame` and `_kernel_stkovf_split_frame` respectively, described in the section entitled *How to use the C library kernel* on page 367.

Examples

The following fragments of assembler code illustrate the main points to consider in interfacing with the C compiler. If you want to examine the code produced by the compiler in more detail for particular cases, you can request an assembler listing by enabling the **Assembler** option on the CC SetUp menu.

This is a function `gggg` which expects two integer arguments and uses only one register variable, `v1`. It calls another function `ffff`.

```

        AREA    |C$$code|, CODE, READONLY
        IMPORT  |ffff|
        IMPORT  |x$stack_overflow|
        EXPORT  |gggg|
gggx    DCB    "gggg", 0          ;name of function, 0 terminated
        ALIGN  ;padded to word boundary
gggy    DCD    &ff000000 + gggy - ggxx
        ;dist. to start of name
;Function entry: save necessary regs. and args. on stack
gggg    MOV    ip, sp
        STMPD  sp!, {a1, a2, v1, fp, ip, lr, pc}
        SUB    fp, ip, #4        ;points to saved pc
;Test workspace size
        CMPS  sp, sl
        BLLT  |x$stack_overflow|
;Main activity of function
; ....
        ADD    v1, v1, 1        ;use a register variable
        BL    |ffff|           ;call another function
        CMP    v1, 99          ;rely on reg. var. after call
; ....
;Return: place result in a1, and restore saved registers
        MOV    a1, result
        LDMEA  fp, {v1, fp, sp, pc}^

```

If a function will need more than 256 bytes of workspace, it should replace the two-instruction workspace test shown above with the following:

```

        SUB    ip, sp, #n
        CMP    ip, sl
        BLLT  |x$stack_overflow1|

```

where *n* is the number of bytes needed. Note that `x$stack_overflow1` must be called if more than 256 bytes of frame are needed. `ip` must contain `sp_needed`, as shown in the example above.

A function which expects a variable number of arguments should store its arguments in the following manner, so that the whole list of arguments is addressable as a contiguous array of values:

```

        MOV    ip, sp          ;copy value of sp
        STMPD  sp!, {a1, a2, a3, a4};save 4 words of args.
        STMPD  sp!, {v1, v2, fp, ip, lr, pc}
        ;save v1-v6 needed
        SUB    fp, ip, #20;fp points to saved pc
        CMPS  sp, sl          ;test workspace
        BLLT  |x$stack_overflow|

```

Some complete program examples are described in the chapter entitled *Interworking assembler with C* in *Acorn Assembler Release 2* supplied with Acorn Desktop Assembler.

14 How to write relocatable modules in C

Relocatable modules are the basic building blocks of RISC OS and the means by which RISC OS can be extended by a user. The archetypal use for RISC OS extensions is the provision of device drivers for devices attached to Archimedes hardware.

Relocatable modules also provide mechanisms which can be exploited to:

- extend RISC OS's repertoire of built-in commands (* commands) (analogous to plugging additional ROMs into a BBC microcomputer of pre-Archimedes vintages)
- provide services to applications (for example, as does the shared C library module)
- implement 'terminate and stay resident' (TSR) applications.

The idea of TSR applications will be most familiar to PC users, whereas extending the * command set (via 'software ROM modules') will seem most familiar to those with a background in the BBC computer. A complete discussion of these topics is beyond the scope of this chapter.

For modules which provide services, the principal mechanism for accessing those services from user code is the SoftWare Interrupt (SWI). For example, the shared C library implements a handler for a single SWI which, when called from the library stubs linked with the application, returns the address of the C library module which in turn allows the library stubs to be initialised to point to the correct addresses within the library module. Thereafter, library services are accessed directly by procedure call, rather than by SWI call. All this illustrates is the rich variety of mechanism available to be exploited.

Getting started

To write a module in C you will need:

- the CC and CMHG tools supplied with Acorn Desktop C
- the C Shared Library module and shared C library stubs supplied with Acorn Desktop C

- a thorough understanding of RISC OS modules (read the chapter of the RISC OS *Programmer's Reference manual* entitled *Modules*).

Constraints on modules written in C

A module written in C **must** use the shared C library module via the library stubs. Use of the stand-alone C library (ANSILib) is **not** a supported option.

All components of a module written in C **must** be compiled with the compiler SetUp menu option **Module code** enabled. This allows the module's static data to be separated from its code and multiply instantiated.

Modules written in C should **not** be compiled with stack limit checking disabled. The stack limit check is small and fast, and can save your machine from crashing.

Overview of modules written in C

A module written in C includes the following:

- a Module Header (described in the *Modules* chapter of the RISC OS *Programmer's Reference manual*), constructed using CMHG;
- a set of entry and exit 'veneers', interfacing the module header to the C run-time environment (also constructed using CMHG);
- the stubs of the shared C library;
- code written by you to implement the module's functionality – for example: *command handlers, SWI handlers and service call handlers.

These parts must be linked together using the Link tool with the SetUp box **Module** option enabled.

The next section describes:

- how to write a CMHG input file to make a module header and any necessary entry veneers
- the interface definitions to which each component of your module must conform
- how to write a CMHG input file to generate entry veneers for IRQ handlers written in C.

Functional components of modules written in C

The following components may be present in a module written in C (all are optional except for the title string and the help string which are obligatory):

- Runnable application code (called start code in the module header description). This will be present if you tell CMHG that the module is runnable and include a `main()` function amongst your module code.
- Initialisation code. 'System' initialisation code is always present, as the shared library must be initialised. Your initialisation function will be called after the system has been initialised if you declare its name to CMHG.
- Finalisation code. The C library has to be closed down properly on module termination. Your own finalisation code will be called on `exit()` if you register it with the C library by using the `atexit()` library function.
- Service call handler. This will be present if you declare the name of a handler function to CMHG. In addition, you can give a list of service call numbers which you wish to deal with and CMHG will generate fast code to ignore other calls without calling your handler.
- A title string in the format described in the *RISC OS Programmer's Reference manual*. CMHG will insist that you give it a valid title string.
- A help string in the format described in the *RISC OS Programmer's Reference manual*. Again, CMHG will insist that you give a valid help string.
- Help and command keyword table. This section is optional and will be present only if you describe it to CMHG and declare the names of the command handlers to CMHG. Obviously, their implementations must be included in the linked module.
- SWI chunk base number. Present only if declared to CMHG.
- SWI handler code. Present if you declare the name of a handler function to CMHG.
- SWI decoding table. Present only if described to CMHG.
- SWI decoding code. Present only if you declare the name of your decoding function to CMHG.
- IRQ handlers. Though not associated with the module header, CMHG will generate entry veneers for IRQ handlers. You can register these veneers with RISC OS using `SWI_OS_Claim`, etc; you have to provide implementations of the handlers themselves. The names of the handler functions and of the entry veneers have to be given to CMHG.

Each component that you wish to use must be described in your input to CMHG. Use of most components also requires that you write some C code which must conform to the interface descriptions given in the sections below.

The C module header generator

The C Module Header Generator (CMHG) is a special-purpose assembler of module headers. It accepts as input a text file describing which module facilities you wish to use and generates as output a linkable object module (in Acorn Object Format). For details of how to run the CMHG tool, see the chapter entitled CMHG earlier in this manual.

The format of input to CMHG

Input to CMHG is in free format and consists of a sequence of 'logical lines'. Each logical line starts with a keyword which is followed by some number of parameters and (sometimes) keywords. The precise form of each kind of logical input line is described in the following sections.

A logical line can be continued on the next line of input immediately after a comma (that is, if the next non-white-space character after a comma is a newline then the line is considered to be continued).

Lists of parameters can be separated by commas or spaces, but use of comma is required if the line is to be continued.

A comment begins with a `;` and continues to the end of the current line. A comment is valid anywhere that trailing white space is valid (and, in particular, after a comma).

A keyword consists of a sequence of alphabetic characters and minus signs. Often, a keyword is the same as the description of the corresponding field of the module header (as described in the RISC OS *Programmer's Reference manual*) but with spaces replaced by minus signs. For example: `initialisation-code;`
`title-string;` `service-call-handler.`

Keywords are always written entirely in lower case and are always immediately followed by a `:`. Character case is significant in all contexts: in keywords, in identifiers, and in strings.

Numbers used as parameters are unsigned. Three formats are recognised:

- unsigned decimal
- `0xhhh...` (up to 8 hex digits)
- `&hhh...` (up to 8 hex digits).

In the following sections, the parts headed *CMHG description* tell you what you have to describe to CMHG in order to use the facility described in that section; the parts headed *C interface* introduce a description of the interface to which the handler function you write must conform. You may omit any trailing arguments that you don't need from your handler implementations.

Runnable application code

CMHG description:

```
module-is-runnable:                ; No parameters.
```

C interface:

```
int main(int argc, char *argv[]);
/*
 * Entered in user-mode with argc and argv
 * set up as for any other application. Malloc
 * obtains storage from application workspace.
 */
```

To be useful (ie re-runnable) as a 'terminate and stay resident' application, a runnable application must implement at least one * command handler (see below) for its command line, which, when invoked, enters the module (calls SWI OS_Module with the Enter reason code).

Initialisation code

CMHG description:

```
initialisation-code: user_init    ; The name of your initialisation function.
                                ; Any valid C function name will do.
```

C interface:

```
_kernel_oserror *user_init(char *cmd_fail, int podule_base, void *pw);
/*
 * Return NULL if your initialisation succeeds; otherwise return a pointer to an
 * error block. cmd_tail points to the string of arguments with which the
 * module is invoked (may be "").
 * podule_base is 0 unless the code has been invoked from a podule.
 * pw is the 'r12' value established by module initialisation. You may assume
 * nothing about its value (in fact it points to some RMA space claimed and
 * used by the module veneers). All you may do is pass it back for your module
 * veneers via an intermediary such as SWI OS_Call Every (use _kernel_swi() to
 * issue the SWI call).
 */
```

Note that you can choose any valid C function name as the name of your initialisation code (CMHG insists on no more than 31 characters).

Finalisation code

User finalisations are handled by using `atexit()` to register finalisation functions. A call to library finalisation code is inserted automatically by CMHG; the C library finalisation code will call these registered functions immediately before closing down the library (on module finalisation).

Service call handler

CMHG description:

```
service-call-handler: sc_handler <number> <number> ...
```

C interface:

```
void sc_handler(int service_number, _kernel_swi_regs *r, void *pw);
/*
 * Return values should be poked directly into r->r[n];
 * the right value/register to use depends on the service number
 * (see the relevant RISC OS Programmer's Reference Manual section for details).
 * pw is the private word (the 'r12' value.
 */
```

Service calls provide a generic mechanism. Some need to be handled quickly; others are not time critical. Because of this, you may give a list of service numbers in which you are interested and CMHG will generate code to ignore the rest quickly. The fast recognition code looks like:

```
CMPS    r1, #FirstInterestingServiceNumber
CMPNES  r1, #SecondInterestingServiceNumber
...
CMPNES  r1, #NthInterestingServiceNumber
MOVNES  pc, lr
; drop into service call entry veneer.
```

If you give no list of interesting service numbers then all service calls will be passed to your handler.

In order to construct a relocatable module which implements a RISC OS application (a TSR application) you must claim and deal with the Service_Memory service call. See the relevant section in the Programmer's Reference Manual for details of this service call.

The following is a suitable handler written in C for this service call:

```
#define Service_Memory 0x11
extern void FrontEnd_services(int service_number,
_kernel_swi_regs *r, void
*pw)
{
    IGNORE(pw);
    /* keep application workspace (r2 holds CAO pointer) */
    if (service_number == Service_Memory && r->r[2] ==
(int)Image__RO_Base)
    {
        r->r[1] = 0; /* refuse to relinquish app. workspace */
    }
}
```

The above handler needs to compare the contents of `r[2]` with the address of the base of your module containing it. This is not a value directly available in C, so the following assembly language fragment can be used to gain access to the symbol `Image$$RO$$Base`, which is defined by `Link` when your module is linked together:

```

        IMPORT    |Image$$RO$$Base|
        EXPORT   Image__RO_Base

        AREA     Code_Description, DATA, REL
Image__RO_Base      DCD      |Image$$RO$$Base|

        END

```

As an assembler is not supplied with Acorn Desktop C, this assembly language fragment is supplied as a ready to link object file `User.RMBase.o.Base`, which has been already assembled using `ObjAsm`.

Title string

CMHG description:

```
title-string: <title>
```

`<title>` must consist entirely of printable, non-space ASCII characters.

Any underscores in the title are replaced by spaces. CMHG will fault any title longer than 31 characters and warn if the length of the title string is more than 16.

Help string

CMHG description:

```
help-string: <help> d.dd <comment> ; help string and version number
```

The help string is restricted to 15 or fewer alphanumeric, ASCII characters and underscores. Longer strings are truncated (with a warning) to 15 characters then padded with a single space. Shorter titles are padded with one or two TAB characters so they will appear exactly 16 characters long.

The version number must consist of a digit, a dot, then 2 consecutive digits. Conventionally, the first digit denotes major releases; the second digit minor releases; and the third digit bug-fix or technical changes. If the version number is omitted, 0.00 is used.

CMHG automatically inserts the current date into the version string, as required by RISC OS convention.

A 'comment' of up to 34 characters can also be included after the version number. It will appear in the tail of the module's help string, after the date. A typical use is for annotating the help string in the following style:

```
SomeModule      0.91 (27 JUN 1989) Experimental version
```

CMHG refuses to generate a help string longer than 79 characters and warns if it has to truncate your input.

Help and command keyword table

CMHG description:

```
command-keyword-table: cmd_handler command-description+
```

(Here *command-description*+ denotes one or more command descriptions).

A command-description has the format:

```
<star-command-name> " ("
    min-args:      <unsigned-int> ; default 0
    max-args:      <unsigned-int> ; default 0
    gstrans-map:   <unsigned-int> ; default 0
    fs-command:    ;>flag bits in
    status:        ;>the flag byte
    configure:     ;>of the cmd table
    help:          ;>info word.
    invalid-syntax: <text>
    help-text:     <text>
    " ) "
```

Each sub-argument is optional. A comma after any item allows continuation on the next line.

A <text> item follows the conventions of ANSI C string constants: it is a sequence of implicitly concatenated string segments enclosed in " and ".

Segments may be separated by white space or newlines (no continuation comma is needed following a string segment).

Within a string segment \ introduces an escape character. All the single character ASCII escapes are implemented, but hexadecimal and octal escape codes are not implemented. A \ immediately preceding a newline allows the string segment to be continued on the following line (but does **not** include a newline in the string; if a newline is required, it must be explicitly included as \n).

`min-args` and `max-args` record the minimum and maximum number of arguments the command may accept; `gstrans-map` records, in the least significant 8 bits, which of the first 8 arguments should be subject to expansion by `OS_GSTrans` before calling the command handler.

The keywords `fs-command`, `status`, `configure` and `help` set bits in the command's information word which mark the command as being of one of those classes.

`invalid-syntax` and `help-text` messages are (should be) self-explanatory.

Example CMHG description:

```
command-keyword-table: cmd_handler
    tm0(min-args: 0, max-args: 255,
        help-text: "Syntax\ttm1 <filenames>\n"),
    tm1(min-args:1, max-args:1,
        help-text: "Syntax\ttm2" " <integer>"
        "\n")
```

This describes two * commands, *tm0 and *tm1, which are to be handled by the C function `cmd_handler`. The handler function will be called with 0 as its third argument if it is being called to handle the first command (tm0, above), 1 as its third argument if it is being called to handle the second command (tm1, above), etc. The programmer must keep the CMHG description in step with the implementation of `cmd_handler`.

C interface:

```
_kernel_oseror *cmd_handler(char *arg_string, int argc, int cmd_no, void *pw);
/*
 * If cmd_no identifies a *HELP entry, then cmd_handler must return
 * arg_string or NULL (if arg_string is returned, the NUL-terminated
 * buffer will be printed).
 * Return NULL if if the command has been successfully handled;
 * otherwise return a pointer to an error block describing the failure
 * (in this case, the veneer code will set the 'V' bit).
 * *STATUS and *CONFIGURE handlers will need to cast 'arg_string' to
 * (possibly unsigned) long and ignore argc. See the RISC OS Programmer's
 * Reference Manual for details.
 * pw is the private word pointer ('r12') value passed into the entry veneer
 */
```

SWI chunk base number

CMHG description:

```
swi-chunk-base-number: <number>
```

You should use this entry if your module provides any SWI handlers. It denotes the base of a range of 64 values which may be passed to your SWI handler. SWI chunks are allocated by Acorn: read the documentation carefully to discover which chunks you may use safely. In some cases you may need to write to Acorn to get a chunk allocated uniquely to your product (though this should not be undertaken lightly and should only be done when all alternatives have been exhausted). See the chapter entitled *An introduction to SWIs* in the *RISC OS Programmer's Reference manual* for more details.

SWI handler code

CMHG description:

```
swi-handler-code: swi_handler ; any valid C function name will do
```

C interface:

```
_kernel_oserror *swi_handler(int swi_no, _kernel_swi_regs *r, void *pw);
/*
 * Return: NULL if the SWI is handled successfully; otherwise return
 * a pointer to an error block which describes the error.
 * The veneer code sets the 'V' bit if the returned value is non-NULL.
 * The handler may update any of its input registers (r0-r9).
 * ps is the private word pointer ('r12') value passed into the
 * swi_handler entry veneer.
 */
```

If your module is to handle SWIs then it must include both `swi-handler-code` and `swi-chunk-base`.

Example CMHG description:

```
swi-chunk-base-number: 0x88000
swi-handler-code:      widget_swi
```

SWI decoding table

CMHG description:

```
swi-decoding-table: <swi-base-name> <swi-name>*
```

This table, if present, is used by `OS_SWINumberTo/FromString`.

Example CMHG description:

```
swi-chunk-base-number: 0x88000
swi-handler-code:      widget_swi
swi-decoding-table:    Widget,
                       Init Read Write Close
```

This would be appropriate for the following name/number pairs:

Widget_Init	0x88000
Widget_Read	0x88001
Widget_Write	0x88002
Widget_Close	0x88003

SWI decoding code

CMHG description:

```
swi-decoding-code: swi_decoder ; any valid C
                        function name will do
```

C interface:

```
void swi_decode(int r[4], void *pw);
/*
 * On entry, r[0] < 0 means a request to convert from text to a number.
 * In this case r[1] points to the string to convert (terminated by a
 * control character, NOT necessarily by NUL).
 * Set r[0] to the offset (0..63) of the SWI within the SWI chunk if
 * you recognise its name; set r[0] < 0 if you don't recognise the name.
 *
 * On entry, r[0] >= 0 means a request to convert from a SWI number to
 * a SWI string:
 *   r[0] is the offset (0..63) of th SWI within the SWI chunk.
 *   r[1] is a pointer to a buffer;
 *   r[2] is the offset within the buffer at which to place the text;
 *   r[3] points to the byte beyond the end of the buffer.
 * You should write th SWI name into the buffer at th position given
 * by r[2] then update r[2] by the length of the text written (excluding
 * any terminating NUL, if you add one).
 *
 * pw is the private word pointer ('r12') passed into the swi_decode
 * entry veneer.
 */
```

If you omit a SWI decoding table then your SWI decoding code will be called instead. Of course, you don't have to provide either.

IRQ handlers

CMHG description:

```
irq-handlers: entry_name/handler_name ...
```

Any number of entry_name/handler_name pairs may be given. If you omit the / and the handler name, CMHG constructs a handler name by appending _handler to the entry name.

C interface:

```
extern int entry_name(_kernel_swi_regs *r, void *pw);
/*
 * This is name of the IRQ handler entry veneer compiled by CMHG.
 * Use this name as an argument to, for example, SWI OS_Claim, in
 * order to attach your handler to IrqV.
 */
int handler_name(_kernel_swi_regs *r, void *pw);
/*
 * This is the handler function you must write to handle the IRQ for
 * which entry_name is the veneer function.
 *
 * Return 0 if you handled the interrupt.
 * Return non-0 if you did NOT handle the interrupt (because,
 * for example, it wasn't for your handler, but for some other
 * handler further down the stack of handlers).
 *
 * 'r' points to a vector of words containing the values of r0-r9 on
 * entry to the veneer. Pure IRQ handlers do not require these, though
 * event handlers and filing system entry points do. If r is updated,
 * the updated values will be loaded into r0-r9 on return from the
 * handler.
 *
 * pw is the private word pointer ('r12') value with which
 * the IRQ entry veneer is called.
 */
```

Handlers must be installed from some part of the module which runs in SVC mode (eg initialisation code, a SWI handler, etc). The name to use at installation time is the `entry_name` (**not** the name of the handler function). This is because C functions cannot be entered directly from IRQ mode, but have to be entered and exited via a veneer which switches to SVC mode. Running in SVC mode gives your handler maximum flexibility.

IRQ handlers can also be used as event handlers and filing system entry points. A full discussion of these topics is beyond the scope of this Guide; refer to the *RISC OS Programmer's Reference manual* for details and for information on how to install and remove handlers.

Turning interrupts on and off

The following (`<kernel.h>`) library functions support the control of the interrupt enable state:

```
int _irqs_disabled(void);
/*
 * Returns non-0 if IRQs are currently disabled.
 */
void _irqs_off(void);
/*
 * Disable IRQs.
 */
void _irqs_on(void);
/*
 * Enable IRQs.
 */
```

These functions suffice to allow saving, restoring and setting of the IRQ state. Ground rules for using these functions are beyond the scope of this document. However, general advice is to leave the IRQ state alone in SWI handlers which terminate quickly, but to enable it in long-running SWI handlers.

What a SWI handler does to the IRQ state is part of its interface contract with its clients: you, the implementor, control that interface contract.

Overlays are a very old technique for squeezing quart-sized programs into pint-sized memories: a kind of poor man's paging.

In common with paged programs, an overlaid program is stored on some backing store medium such as a floppy disc or a hard disc and its components (called overlay segments) are loaded into memory only as required. In theory, this reduces the amount of memory required to run a program at the expense of increasing the time taken to load it and repeatedly re-load parts of it. It is a classic space-time trade-off. In practice, except in rather special circumstances, the saving in memory accruing from the use of overlays is rather modest and less than you might expect. Indeed, as discussed below, overlays have rather restricted applicability under RISC OS. Nonetheless, their use can occasionally be a 'life saver'.

Paging vs overlays

In a paged system, a program and its workspace is broken up into fixed size chunks called *pages*. A combination of special hardware and operating system support ensures that pages are loaded only when needed and that un-needed pages are soon discarded. In principle, the author of a paged program need not be aware that it will be paged (but this is often not true in practice if the author wishes the program to run at maximum speed). Both code and data are paged, automatically. In general, for single programs which re-use their workspace whenever possible, one sees a ratio of program size plus workspace size to occupied memory size in the region 1.5 to 3. One can always increase the ratio arbitrarily by integrating several sequentially used programs into a single image and by never re-using workspace. But, fundamentally, paging rarely squeezes more than a quart-sized program into a pint-sized memory. Of course, there are other benefits of paging, but these are beyond the scope of this section.

RISC OS is not a paged system, but Acorn's sister product (the Unix-based RISC iX operating system) is.

In contrast, an overlaid program is broken up into variable sized chunks (called overlay segments) by the user, who also determines which of these chunks may share the same area of memory. As the overlay system permits two code fragments which share the same area of memory to call one another and return successfully to the caller, this is merely a matter of performance. However, if data is included in an overlaid segment the situation becomes more complicated and the user has

more work to do. For example, it must be ensured that all code which uses the data resides in the same segment as the data. Furthermore, it must be acceptable that the data is re-initialised every time the segment is re-loaded. Thus, in general, it is possible to overlay two work areas each of which is private to two distinct sets of functions which are not simultaneously resident in memory. Overall, it would be unusual to overlay more than a quart-sized program into a pint-sized memory, much as with paging (you may achieve a factor as high as four for code, but non-overlaid data will usually dilute the overall factor substantially; it all depends on the details of your application).

A more detailed description of the low-level aspects of overlays is given in the section entitled *Generating overlaid programs* in the chapter entitled *Link* in the *Acorn Desktop Development Environment* user guide. If you are especially interested in using overlays you may prefer to read that section next. Otherwise, if you are more interested in when to use overlays, please read on.

When to use overlays

Overlays work best when a program has several semi-independent parts. A good model for purposes of understanding is to think of a special-purpose command interpreter (the root segment) which can invoke separate commands (overlay segments) in response to user input. Consider, for example, a word processor which consists of a text editor and a collection of printer drivers. It is clear that each of the printer drivers can be overlaid (you are unlikely to have more than one printer); it may even be plausible to overlay each with the editor itself (you may not be able to edit while printing – depending on how fast the printer goes and on how much CPU time is required to drive it). Furthermore, if the time taken to load an overlay segment can be tacked on to an interaction with the user, it is probable that the program will feel little slower than if it were memory-resident. In summary: overlays work best if your program has many independent sub-functions.

On the other hand, if your program has many semi-independent parts, it may be better to structure it as several independent programs, each called from a control program. By using the shared C library, each program can be relatively small, and the Squeeze utility can be used to reduce the space taken by it on backing store by nearly a factor of 2. (See the chapter entitled *Squeeze* in the *Acorn Desktop Development Environment* user guide for details). In contrast, overlay segments cannot be squeezed (though the root program can be). Consider, for example, the following command line programs from this release of C:

Program	Squeezed Size	Unsqueezed Size
CMHG	9Kb	16Kb
Link	22Kb	41Kb
Squeeze	8Kb	14Kb

So, if you can structure your application as independent, squeezed programs it may take up less precious floppy disc space and load faster, especially from a floppy disc, than if you structure it using overlays.

If adopted, this strategy will force the independent programs to communicate via files. Provided the data to be communicated has a simple structure this causes no problems for the application; provided it is not too voluminous, use of the RAM filing system (RamFS) is suggested as this is fast and requires no special application code in order to use it.

So, overlays are most appropriate for applications which manipulate very large amounts of highly structured data – Computer Aided Design applications are archetypal here – whereas multiple independent programs are most appropriate for applications which manipulate relatively small amounts of simply structured data and are otherwise dominated by large amounts of code.

Naturally, if you are porting an existing application to RISC OS, your view will be coloured by whether or not it is already structured to use overlays. If it is, it will probably be best to stick to using overlays, rather than attempting to split the application up into semi-independent sub-applications.

On the other hand, if you are writing an application from scratch, you probably want to ponder this question in more depth. For example, to what other systems will the application be targetted? Using multiple semi-independent applications may work very nicely under Unix or OS/2 where the output of one process can be piped into another, but less well under MS-DOS where use of overlays is much more the norm.

This chapter provides basic information on memory management by RISC OS applications. It is intended to provide some specialist knowledge to help you write efficient programs for RISC OS, and to provide some practical hints and tips.

All the information in this chapter relating to programs written in C refers to the Acorn Desktop C product.

Guidelines

Follow the guidelines in this section to make the best use of available memory. The guidelines are explained in more detail on the following pages.

- **Use recovery procedures** – Your program should keep the machine operational. Don't allow your program to lock up when memory runs out; your program should indicate that it has run out of memory (with an error or warning message) and only stop subsequent actions that use more memory. Ideally, ensure that actions which free up memory have enough reserved memory to run in.
- **Return unwanted memory** – You should return any memory you have no further use for. Claiming memory then not returning it can tie up memory unnecessarily until the machine is re-booted. RISC OS has no garbage collection, so once you have asked for memory RISC OS assumes that you want it until you explicitly return it, even if your program terminates execution. Language libraries often provide you with protection from this, as long as memory is claimed from them.
- **Don't waste memory** – You should avoid wasting memory. It is a finite resource, often wasted in two ways:
 - by permanently claiming memory for infrequent operations
 - by fragmenting it, so that although there is enough unused memory, it is either in the wrong place, or it is not in large enough blocks to use.

Recovery from lack of memory

An important consideration when designing programs for RISC OS is the recovery process, not just from user errors, but also from lack of system resources.

An example of a technique that can be designed into an application is to make an algorithm more disc-based and less RAM-based on detection of lack of memory. This could allow you to continue using an application on a small machine (especially one with a hard disc) at the expense of some speed.

When implementing your code, expect the unexpected and program defensively. Be sure that when the system resources you need (memory, windows, files etc) are not available, your program can cope. Make sure that, when a document managed by your program expands and memory runs out, the document is still valid and can be saved. Don't just check that your main document expansion routines work; check that **all** routines which require memory (or in fact any system resource) fail gracefully when there is no more.

Centralising access to system resources can help: write your program as if every operating system interface is likely to return an error.

Avoiding permanent loss of memory

Permanent loss of memory is mainly a problem for applications or modules written entirely in assembly language. When interworking assembler routines with C or another high level language you should use memory handed to you by the high level language library (eg use `malloc` to get a memory area from C and pass a pointer to it as an argument to your assembler routine). The language library automatically returns such areas to RISC OS on program exit. Additional types of program requiring care to avoid memory loss are those expected to run for a long time (eg a printer spooler) and those making use of RMA directly through SWI calls.

When using the RMA for storage directly through SWI calls, especially for items in linked lists, consider using the first word as a check word containing four characters of text to identify it as belonging to your program. When a block of RMA is deallocated, the heap manager puts it back into a list of free blocks, and in so doing overwrites the first word of the block.

This technique therefore serves two purposes:

- 1 after your program has been run and exited, your check word can be searched for, showing up any blocks you have failed to deallocate
- 2 it avoids problems when accidentally referencing deallocated memory.

A typical problem of referencing deallocated blocks results from using the first word as a pointer to your program's next block, then accidentally referencing a wild pointer when it is overwritten.

You can use the following BASIC routine to search for any lost blocks:

```

100 REM > LostMemory checks for un-released blocks
110 RMA%=&01800000: RMAEnd% = RMA% + (RMA%!12)
120 FOR PossibleBlock% = RMA%+20 TO RMAEnd%-12 STEP 16
130     REM Now loop looking for "Prog"
140     IF PossibleBlock%!0 = &676F7250 THEN
150         PRINT "Block found at &";~PossibleBlock%
160     ENDIF
170 NEXT PossibleBlock%
180 END

```

Avoiding memory wastage

The key factor in writing programs that use memory efficiently and don't waste it is understanding the following:

- how SWI XOS_Module and SWI XOS_Heap work if you are constructing a relocatable module or are using the RMA from an application
- how C flex and malloc work when writing a C program (parts of which may be written in assembler).

This understanding will lead you to writing programs that will work in harmony with the storage allocator. See the following section for a description of C memory allocation.

The C storage manager

Normal C applications (ie those not running as modules) claim memory blocks in two main ways:

- from malloc
- from flex.

The malloc heap storage manager is the standard interface from which to claim small areas of memory. It is tuned to give good performance to the widest variety of programs.

In the following sections, the word *heap* refers to the section of memory currently under the control of the storage manager (usually referred to as malloc, or the malloc heap).

The flex facility is available as part of RISC_OSLib, and can be useful for claiming large areas of data space. It manages a shifting set of areas, so its operation can be slow, and address-dependent data cannot be stored in it. However, it has the following advantages:

- it doesn't waste memory by fragmenting free space
- it returns deallocated memory to RISC OS for use by other applications.

Allocation of malloc blocks

All block sizes allocated are in bytes and are rounded up to a multiple of four bytes. All blocks returned to the user are word-aligned. All blocks have an overhead of eight bytes (two words). One word is used to hold the block's length and status, the other contains a guard constant which is used to detect heap corruptions. The guard word may not be present in future releases of the ANSI C library. When the stack needs to be extended, blocks are allocated from the `malloc` heap.

When an allocation request is received by the storage manager, it is categorised into one of three sizes of blocks

- small 0 → 64
- medium 65 → 512
- large 513 → 16777216.

The storage manager keeps track of the free sections of the heap in two ways. The medium and large sized blocks are chained together into a linked list (overflow list) and small blocks of the same size are chained together into linked lists (bins). The overflow list is ordered by ascending block address, while the bins have the most recently freed block at the start of the list.

When a small block is requested, the bin which contains the blocks of the required size is checked, and, if the bin is not empty, the first block in the list is returned to the user. If there was not a block of the exact size available, the bin containing blocks of the next size up is checked, and so on until a block is found. If a block is not found in the bins, the last block (highest address) on the overflow list is taken. If the block is large enough to be split into two blocks, and the remainder is a usable size (> 12 including the overhead) then the block is split, the top section returned to the user and the remainder, depending on its size, is either put in the relevant bin at the front of the list or left in the overflow list.

When a medium block is requested, the search ignores the bins and starts with the overflow list. This is searched in reverse order for a block of usable size, in the same way as for small blocks.

When a large block is requested, the overflow list is searched in increasing address order, and the first block in the list which is large enough is taken. If the block is large enough to be split into two blocks, and the size of the remainder is larger than a small block (> 64) then the block is split, the top section is returned to the overflow list, and bottom section given to the user.

Should there not be a block of the right size available, the C storage manager has two options:

- 1 Take all the free blocks on the heap and join adjacent free blocks together (coalescing) in the hope that a block of the right size will be created which can then be used
- 2 Ask the operating system for more heap, put the block returned in the overflow list, and try again.

The heap will only be coalesced if there is at least enough free memory in it to make it worthwhile (ie four times the size of the requested block, and at least one sixth of the total heap size) or if the request for more heap was denied. Coalescing causes the following:

- the bins and overflow list are emptied;
- the heap is scanned;
- adjacent free blocks are merged;
- the free blocks are scattered into the bins and overflow list in increasing address order.

Deallocation of malloc blocks

When a block is freed, if it will fit in a bin then it is put at the start of the relevant bin list, otherwise it is just marked as being free and effectively taken out of the heap until the next coalesce phase, when it will be put in the overflow list. This is done because the overflow list is in ascending block address order, and it would have to be scanned to be able to insert the freed block at the correct position. Fragmentation is also reduced if the block is not reusable until after the next coalesce phase. It is worth noting that deallocating a block and then reallocating a block of the same size can not be relied upon to deliver the original block.

Reallocation of malloc blocks

You should be cautious when using `realloc`. Reallocating a block to a larger size will usually require another block of memory to be used and the data to be copied into it. This means that you cannot use the whole of the heap as both blocks need to be present at the same time.

If consecutive calls keep increasing the block size until all memory is used up, then only about a third of the heap is likely to be available in one block. A typical course of events is:

- 1 The first block is present (block A).
- 2 It is extended to a larger sized block (block B). Block A must still be present (see above).

- 3 It is again extended to a larger sized block (block C). Block B must still be present (see above). However, block A also still exists because it is too small to use, and cannot be coalesced with another block because block B is in the way.

Wimp slots and the C flex system

A typical C application running under the Wimp has a single contiguous application area (wimp slot) into which are placed the following:

- program image
- static data
- stack
- malloc data.

The initial wimp slot size is set by the size of the Next slot (in the Task display window) when the application is started, or by *WimpSlot commands in the !Run file associated with the C application. If the malloc heap is full, and the flex system has not been initialised and the operating system has free memory, the wimp slot grows, raising its highest address. Once enlarged by malloc, the wimp slot never reduces again until program termination.

The stack is allocated on the heap, in 4K (or as big as needed) chunks: the ARM procedure call standard means that disjoint extension of the stack is possible. The only other use that the ANSI library makes of the malloc heap is in allocating file buffers, but even this usage can be prevented by making the appropriate calls to the ANSI library buffer handling facilities (setvbuf). The operation of the malloc heap is described above and is designed to provide good performance under heavy use. Its design is such that small blocks can be allocated and freed rapidly.

Any malloc heap tends to fragment over time. This is particularly serious in the following circumstances:

- no virtual memory
- multitasking – if memory is not in use, it should be handed to other applications
- if a program runs out of memory it must not crash, but must recover and continue.

These are just the conditions under which a desktop application operates!

Because of this, the flex facilities are available as part of RISC_OSLib (the RISC OS-specific C library provided with Acorn Desktop C). These provide a shifting heap, intended for the allocation of large blocks of memory which might otherwise destroy the structure of a malloc-style heap.


Flex works by increasing the size of the application area, using space above that reserved for use by `malloc`. Once the flex system is initialised the `malloc` heap cannot grow, unless you enable this (see later). The benefits of using flex can be seen in Draw, Paint and Edit, which are all written in C using early versions of RISC_OSLib. Their application areas expand when new files are added, contract when files are discarded, and do not suffer from needless incremental application area growth over time.

The implementation of flex is quite simple. There is no free list as memory is shifted whenever a block is destroyed or changed in size. New blocks are always allocated at the top. When blocks are deallocated or resized, those above are moved. This means that deallocating or changing the size of a block can take quite a long time (proportional to the sum of the sizes of the blocks above it in memory). Flex is also not recommended for allocation of small blocks. Its other limitation is that as flex blocks can be shifted, you should not use them for address-dependent data (eg pointers or indirected icon data).


In addition to the facilities described above, RISC_OSLib also provides an obsolete `malloc`-like allocator of non-shifting blocks called heap.

Two facilities are provided, because no one storage manager can solve all problems in the absence of Virtual Memory. A program which works adequately with `malloc` should feel no compulsion to use anything else. The use of flex, however, particularly in desktop applications such as editors (which are likely to be resident on the desktop for a long period of time) can go a long way towards improving their memory usage.

The model of a C application's memory layout is as follows:



If the application uses flex store as supported by RISC_OSLib, the model is:



To expand the `malloc` heap when a flex store area is being used the flex area has to be moved. To achieve this, `malloc` calls a flex function to move the flex blocks. The flex function called is registered with the C library, and may be a dummy function which does not move flex. If a dummy function is registered or flex cannot be successfully moved, then `malloc` itself returns a 0 to indicate failure.

The Acorn Desktop C version of RISC_OSLib registers a dummy flex-moving function during `flex_init()`, inhibiting `malloc` heap expansion after `flex_init()` has been called. This is registered with a call to the function `_kernel_register_slotextend()`.

A functional flex-moving function performs the relocation, sets a pointer to the newly available space, and returns the size of the memory thus obtained (which may be less than that requested by `malloc`).

Allowing `malloc` heap expansion to move flex makes the use of pointers into flex blocks potentially hazardous when the pointers are set before, but used after, the following:

- calls to `flex_alloc`, `flex_free`, `flex_extend`
- calls to `malloc` and `_kernel_alloc`
- calls to any functions which may cause stack extension (since stack extension uses the `malloc`-heap for this purpose)

Consider the following code fragment:

```
#define FLEX_SIZE 1024 /* for example */
#define OFFSET    42  /* for example */

static void nonleaf_function(char *p)
{
    /* declaration of local vars, and calls to other functions here */
    /* use of p happens here ... */
}

static void access_flex_store(void)
{
    char *message;

    flex_alloc((flex_ptr)&message, FLEX_SIZE);
    nonleaf_function(message+OFFSET);
}
```

Notice that when the value of the char pointer `message+OFFSET` is passed by value to the function `nonleaf_function()`, use of `p` in this function may no longer be valid, since stack extension may have happened during the function call, which may have caused the allocated flex store to move.

Working in this Environment

- 1 If you have an existing binary, linked with a version of stubs pre-dating the 3.1b intermediate release, such as that included with ANSI C Release 3, then you do not get an extending wimpslot, and hence no new problems arise (the shared C library 'knows' which stubs the application was linked with). You must make your initial wimpslot large enough to accommodate your stack/heap needs. This is important for old applications which rely on `malloc` returning 0 when the application's initial wimpslot is exhausted.
- 2 If you link with the Acorn Desktop C version of stubs, but do not use the flex functions in RISC_OSLib, you get a wimpslot extendable by `malloc`, and have no new problems. When more heap is required your wimpslot may be increased by the C library (but will not shrink when `free()` is called).
- 3 If you link with the Acorn Desktop C version of stubs, and use the flex functions in RISC_OSLib, then your malloc-heap will (by default) not be allowed to grow. You must make your initial wimpslot large enough to accommodate your stack/heap needs.

Note: `flex_init()` makes the call:

```
_kernel_register_slotextend(flex_dont_budge);
```

This means that when the C library attempts to acquire more wimpslot, the extension will fail. This gives you the guarantee that flex store will only be relocated due to `flex_alloc`, `flex_extend`, and `flex_free`. Your wimpslot will grow or shrink to satisfy flex requests, but your malloc-heap will have a bound fixed by the size of your initial wimpslot.

- 4 If you link with the Acorn Desktop C version of stubs, use the flex functions in RISC_OSLib, and require malloc to extend the application's wimpslot, you must be prepared to exist in a world where flex store may move as described in the section above.

After calling `flex_init()`, you can make the call:

```
_kernel_register_slotextend(flex_budge);
```

This registers a function which will relocate flex store whenever the C library needs to grow its malloc-heap.

If you choose to do this, then the following guidelines will be of use to you:

- Always pass `flex_ptr`'s (`void **`'s) to your own functions, with an integral offset.
Avoid passing direct flex block pointers.
- Direct calls to `malloc` may cause the flex store to move in the same way that calls to `flex_alloc`, `flex_extend` and `flex_free` do.

- You can safely make SWI calls which require pointer arguments where these arguments point into flex blocks, by using `_kernel_swi()`, since `_kernel_swi` **cannot** cause stack extension. This state **must** be guaranteed by the C library, since `flex_budge()` uses `_kernel_swi()` and may be called during stack extension.
- Using the Acorn Desktop C version of RISC_OSLib, you can also call any SWI 'vener' functions, with the knowledge that the stack will not be extended. These functions have been compiled with stack checking turned off. The functions (which are all in RISC_OSLib) are:

```
bbc.h
colourtran.h
drawmod.h
font.h
os.h
print.h
sprite.h
visdelay.h
wimp.h
```

- You can turn stack checking off in your own code using pragmas, thus:

```
#pragma no_stack_checks
/* functions defined after here are compiled without stack checks */
#pragma stack_checks
/* functions defined after here are compiled with stack checks */
```

Or for a whole source file by compiling using the flag `-zps1`

Note that functions which are compiled with stack checking off have only 512 bytes of stack available to them, and any 'non-stack-check' functions which they call.

- You can toggle whether the malloc-heap is permitted to extend, using calls to `_kernel_register_slotextend()` with arguments `flex_budge` or `flex_dont_budge`. This can be used to surround critical regions of code, where you may wish to temporarily stop flex blocks moving due to malloc-heap extension.

You can set the root stack segment size using:

```
int __root_stack_size = 16*1024; /* to get a 16kb stack size */
```

Using heap_alloc and heap_free

Since when malloc heap expansion is inhibited (as it is by default with the Acorn Desktop C version of flex) the bottom flex block is static, it is valid to retain pointers into it, and useful to manage a malloc style heap of fixed blocks within it. The `heap_alloc()` and `heap_free()` functions provide facilities to perform this.

Using the heap functions to do memory allocation is similar to `malloc()` in that a pointer to the block allocated is returned to the caller: the routine to do this is called `heap_alloc()`. Memory may be released with `heap_free()`. Before you use heap, you must call `heap_init()`; if `heap_init()` is called with a non-zero parameter, then the heap will be shrunk when it is possible to do so after a call to `heap_free()`. The call to `heap_init()` must be made after flex has been initialised with `flex_init()`. Since the heap functions support a heap in the first flex block allocated, `heap_init` must be called before any calls to flex allocation functions, and you must **not** allow the C heap to extend thus causing all flex blocks to be relocated (ie you must not have registered `flex_budge` with `_kernel_slot_extend()`).

Using memory from C relocatable modules

All memory allocated by `malloc` comes from the RMA when your program is executing in non-user mode. So remember to free it up when you've finished with it. If your module allocates any RMA blocks by calling SWI XOS_Module directly, the C run-time system does not clear them out when your module finalises, so make sure you do!

There are two sets of `atexit()` routines, the ones which you registered during initialisation ie before your module was entered via the `main()` entry point (because the module was RMRun for instance), and the ones you registered after. The ones registered before will be executed when your module is finalised – this is how to clear up after yourself; the ones after will be called when your module exits from being run, ie when `main()` terminates.

When you are writing a C module, use `exit()`, not `SWI XOS_Exit`.

Never access your application's workspace from an interrupt routine. During interrupts, the state of the application area is effectively random. Since your interrupt routine could execute at any time, it could happen while some other application is switched in. If this did happen, and the interrupt routine updated application space, then some other application could be affected. To get around this problem, allocate some RMA space for your interrupt routine to use when it needs to; this memory will be visible when your application is running. Remember to free up the RMA space when you've finished with it.

When executing as C module SVC mode code (during initialisation, finalisation, service or interrupt entry) your stack will be small. Also, your stack, unlike when in USR mode (ie running as an application) will not extend dynamically. It is therefore very important to be extremely economical with stack space; avoiding large auto arrays, using `malloc` where larger spaces are required, then freeing at the routine end.

Static variables (and arrays etc.) in a C module are extant for the lifetime of the module, ie the entire time it is loaded. If they are only needed when it is running as an application, then they should be claimed using `malloc` instead.

This chapter describes the following machine-specific features of the Acorn C compiler:

- the C library kernel
- calling other programs from C
- the shared C library
- `#pragma` directives
- storage management
- handling host errors.

How to use the C library kernel

C library structure

The C library is organised into layers, like the skins of an onion. At the centre is the language-independent library kernel. This is implemented in assembly language and provides basic support services, described below, to language run-time systems and, directly, to client applications.

One level out from the library kernel is a thin, C-specific layer, also implemented in assembly language. This provides compiler support functions such as structure copy, interfaces to stack-limit checking and stack extension, `setjmp` and `longjmp` support, etc. Everything above this level is written in C.

Finally, there is the C library proper. This is implemented in C and, with the exception of one module which interfaces to the library kernel and the C-specific veneer, is highly portable.

The library kernel

The library kernel is designed to allow run-time libraries for different languages to co-reside harmoniously, so that inter-language calling can be smooth. At the present time, the Fortran-77 library uses the run-time kernel, but the Pascal library does not. Currently, code compiled by the F77 compiler does not adhere to the ARM Procedure-Call Standard, so inter-working with C is not possible in this release.

The library kernel provides the following facilities:

- a generic, status-returning, procedural interface to SWIs
- a procedural interface to the following commonly used SWIs:
 - OS_Byte
 - OS_Rdch
 - OS_Wrch
 - OS_BGet
 - OS_BPut
 - OS_GBPB
 - OS_Word
 - OS_Find
 - OS_File
 - OS_Args
 - OS_CLI /* use is not advised – use `_kernel_system()` */
- a procedural interface to the following arithmetic functions:
 - unsigned integer division
 - unsigned integer remainder
 - unsigned divide by 10 (much faster than general division)
 - signed integer division
 - signed integer remainder
 - signed divide by 10 (much faster than general division).
- a procedural interface to the following miscellaneous functions:
 - finding the identity of the host system (RISC OS, Arthur, etc)
 - determining whether the floating point instruction set is available
 - getting the command string with which the program was invoked
 - returning the identity of the last OS error
 - reading an environmental variable
 - setting an environmental variable
 - invoking a sub-application
 - claiming memory to be managed by a heap manager
 - unwinding the stack
 - finding the name of a function containing a given address
 - finding the source language associated with code at a given address.
- support for manipulating the IRQ state from a relocatable module:
 - getting the processor mode
 - determining if IRQs are enabled
 - enabling IRQs
 - disabling IRQs.

- support for allocating and freeing memory in the RMA area:
 - allocating a block of memory in the RMA
 - extending a block of memory in the RMA
 - freeing a block of memory in the RMA.
- support for stack-limit checking and stack extension:
 - finding the current stack chunk
 - four kinds of stack extension –
 - small-frame and large-frame extension,
 - number of actual arguments known (eg Pascal), or unknown (eg C) by the callee.
- trap handling, error handling, event handling and escape handling.

Most of these functions are described in the C library header file `<kernel.h>`. This header also declares the data structures you will need to use in order to call these functions or to interpret their results. See *Appendix C: kernel.h* for a detailed description.

Interfacing a language run-time system to the Acorn library kernel

In order to use the kernel, a language run-time system must provide an area named `RTSK$$DATA`, with attributes `READONLY`. The contents of this area must be a `_kernel_languagedescription` as follows:

```
typedef enum { NotHandled, Handled } _kernel_HandledOrNot

typedef struct {
    int regs [16];
} _kernel_registerset;

typedef struct {
    int regs [10];
} _kernel_eventregisters;

typedef void (*PROC) (void);
typedef _kernel_HandledOrNot
    (*_kernel_trapproc) (int code, _kernel_registerset *regs);
typedef _kernel_HandledOrNot
    (*_kernel_eventproc) (int code, _kernel_registerset *regs);

typedef struct {
    int size;
    int codestart, codeend;
```

```
char *name;
PROC (*InitProc)(void); /* that is, InitProc returns a PROC */
PROC FinaliseProc;
_kernel_trapproc TrapProc;
_kernel_trapproc UncaughtTrapProc;
_kernel_eventproc EventProc;
_kernel_eventproc UnhandledEventProc;
void (*FastEventProc) (_kernel_eventregisters *);
int (*UnwindProc) (_kernel_unwindblock *inout, char **language);
char * (*NameProc) (int pc);
} _kernel_languagedescription;
```

Any of the procedure values may be zero, indicating that an appropriate default action is to be taken. Procedures whose addresses lie outside [codestart...codeend] also cause the default action to be taken.

codestart, codeend

These values describe the range of program counter (PC) values which may be taken while executing code compiled from the language. The linker ensures that this is describable with just a single base and limit pair if all code is compiled into areas with the same unique name and same attributes (conventionally, *Language\$\$code*, CODE, READONLY. The values required are then accessible through the symbols *Language\$\$code\$\$Base* and *Language\$\$code\$\$Limit*).

InitProc

The kernel contains the entrypoint for images containing it. After initialising itself, the kernel calls (in a random order) the *InitProc* for each language RTS present in the image. They may perform any required (language-library-specific) initialisation: their return value is a procedure to be called in order to run the main program in the image. If there is no main program in its language, an RTS should return 0. (An *InitProc* may not itself enter the main program, otherwise other language RTSs might not be initialised. In some cases, the returned procedure may be the main program itself, but mostly it will be a piece of language RTS which sets up arguments first.)

It is an error for all *InitProcs* in a module to return 0. What this means depends on the host operating system; if RISC OS, SWI OS_GenerateError is called (having first taken care to restore all OS handlers). If the default error handlers are in place, the difference is marginal.

FinaliseProc

On return from the entry call, or on call of the kernel's *Exit* procedure, the *FinaliseProc* of each language RTS is called (again in a random order). The kernel then removes its OS handlers and exits setting any return code which has been specified by call of *_kernel_setreturncode*.

TrapProc, UncaughtTrapProc

If an image is not being run under a debugger, the kernel installs OS trap and error handlers. On occurrence of a trap, or of a fatal error, all registers are saved in an area of store belonging to the kernel. These are the registers at the time of the instruction causing the trap, except that the PC is wound back to address that instruction rather than pointing a variable amount past it.

The PC at the time of the trap together with the call stack are used to find the TrapHandler procedure of an appropriate language. If one is found, it is invoked in user mode. It may return a value (Handled or NotHandled), or may not return at all. If it returns Handled, execution is resumed using the dumped register set (which should have been modified, otherwise resumption is likely just to repeat the trap). If it returns NotHandled, then that handler is marked as failed, and a search for an appropriate handler continues from the current stack frame.

If the search for a trap handler fails, then the same procedure is gone through to find a 'uncaught trap' handler.

If this too fails, it is an error. It is also an error if a further trap occurs while handling a trap. The procedure `_kernel_exittraphandler` is provided for use in the case the handler takes care of resumption itself (eg via `longjmp`).

(A language handler is appropriate for a PC value if `LanguageCodeBase <= PC` and `PC < LanguageCodeLimit`, and it is not marked as failed. Marking as 'failed' is local to a particular kernel trap handler invocation. The search for an appropriate handler examines the current PC, then RI4, then the link field of successive stack frames. If the stack is found to be corrupt at any time, the search fails).

EventProc, UnhandledEventProc

The kernel always installs a handler for OS events and for Escape flag change. On occurrence of one, all registers are saved and an appropriate EventProc, or failing that an appropriate UnhandledEventProc is found and called. Escape pseudo-events are processed exactly like Traps. However, for 'real' events, the search for a handler terminates as soon as a handler is found, rather than when a willing handler is found (this is done to limit the time taken to respond to an event). If no handler is willing to claim the event, it is handed to the event handler which was in force when the program started. (The call happens in `CallBack`, and if it is the result of an Escape, the Escape has already been acknowledged.)

In the case of escape events, all side effects (such as termination of a keyboard read) have already happened by the time a language escape handler is called.

FastEventProc

The treatment of events by EventProc isn't too good if what the user level handler wants to do is to buffer events (eg conceivably for the key up/down event), because there may be many to one event handler call. The FastEventProc allows a call at the time of the event, but this is constrained to obey the rules for writing interrupt code (called in IRQ mode; must be quick; may not call SWIs or enable interrupts; must not check for stack overflow). The rules for which handler gets called in this case are rather different from those of (uncaught) trap and (unhandled) event handlers, partly because the user PC is not available, and partly because it is not necessarily quick enough. So the FastEventProc of each language in the image is called in turn (in some random order).

UnwindProc

UnwindProc unwinds one stack frame (see description of `_kernel_unwindproc` for details). If no procedure is provided, the default unwind procedure assumes that the ARM Procedure Call Standard has been used; languages should provide a procedure if some internal calls do not follow the standard.

NameProc

NameProc returns a pointer to the string naming the procedure in whose body the argument PC lies, if a name can be found; otherwise, 0.

How the run-time stack is managed and extended

The run-time stack consists of a doubly-linked list of stack chunks. The initial stack chunk is created when the run-time kernel is initialised. Currently, the size of the initial chunk is 4Kb. Subsequent requests to extend the stack are rounded up to at least this size, so the granularity of chunking of the stack is fairly coarse. However, clients may not rely on this.

Each chunk implements a portion of a descending stack. Stack frames are singly linked via their frame pointer fields within (and between) chunks. See *Appendix F - ARM procedure call standard* in the *Acorn Desktop Development Environment* user guide for more details.

In general, stack chunks are allocated by the storage manager of the master language (the language in which the root procedure – that containing the language entry point – is written). Whatever procedures were last registered with `_kernel_register_allocs()` will be used (each chunk 'remembers' the identity of the procedure to be called to free it). Thus, in a C program, stack chunks are allocated and freed using `malloc()` and `free()`.

In effect, the stack is allocated on the heap, which grows monotonically in increasing address order.

The use of stack chunks allows multiple threading and supports languages which have co-routine constructs (such as Modula-2). These constructs can be added to C fairly easily (provided you can manufacture a stack chunk and modify the `fp`, `sp` and `s1` fields of a `jmp_buf`, you can use `setjmp` and `longjmp` to do this).

Stack chunk format

A stack chunk is described by a `_kernel_stack_chunk` data structure located at its low-address end. It has the following format:

```
typedef struct stack_chunk {
    unsigned long sc_mark;          /* == 0xf60690ff */
    struct stack_chunk *sc_next, *sc_prev;
    unsigned long sc_size;
    int (*sc_deallocate)();
} _kernel_stack_chunk;
```

`sc_mark` is a magic number; `sc_next` and `sc_prev` are forward and backward pointers respectively, in the doubly linked list of chunks; `sc_size` is the size of the chunk in bytes and includes the size of the stack chunk data structure; `sc_deallocate` is a pointer to the procedure to call to free this stack chunk – often `free()` from the C library. Note that the chunk lists are terminated by NULL pointers – the lists are not circular.

The seven words above the stack chunk structure are reserved to Acorn. The stack-limit register points 512 bytes above this (ie 560 bytes above the base of the stack chunk).

Stack extension

Support for stack extension is provided in two forms:

- `fp`, arguments and `sp` get moved to the new chunk (Pascal/Modula-2-style)
- `fp` is left pointing at arguments in the old chunk, and `sp` is moved to the new chunk (C-style).

Each form has two variants depending on whether more than 4 arguments are passed (Pascal/Modula-2-style) or on whether the required new frame is bigger than 256 bytes or not (C-style). See *Appendix F - ARM procedure call standard* in the *Acorn Desktop Development Environment* user guide for more details.

`_kernel_stkovf_copyargs`

Pascal/Modula-2-style stack extension, with some arguments on the stack (ie stack overflow in a procedure with more than four arguments). On entry, `ip` must contain the number of argument words on the stack.

_kernel_stkovf_copy0args

Pascal/Modula-2-style stack extension, without arguments on the stack (ie stack overflow in a procedure with four arguments or fewer).

_kernel_stkovf_split_frame

C-style stack extension, where the procedure detecting the overflow needs more than 256 bytes of stack frame. On entry, `ip` must contain the value of `sp` – the required frame size (ie the desired new `sp` which would be below the current stack limit).

_kernel_stkovf_split_0frame

C-style stack extension, where the procedure detecting the overflow needs 256 or fewer bytes of stack frame.

Stack chunks are deallocated on returning from procedures which caused stack extension, but with one chunk of latency. That is, one extra stack chunk is kept in hand beyond the current one, to reduce the expense of repeated call and return when the stack is near the end of a chunk; others are freed on return from the procedure which caused the extension.

Calling other programs from C

The C library procedure `system()` provides the means whereby a program can pass a command to the host system's command line interpreter. The semantics of this are undefined by the draft ANSI standard.

RISC OS distinguishes two kinds of commands, which we term *built-in commands* and *applications*. These have different effects. The former always return to their callers, and usually make no use of application workspace; the latter return to the previously set-up 'exit handler', and may use the currently-available application workspace. Because of these differences, `system()` exhibits three kinds of behaviour. This is explained below.

Applications in RISC OS are loaded at a fixed address specified by the application image. Normally, this is the base of application workspace, 0x8000. While executing, applications are free to use store between the base and end of application workspace. The end is the value returned by `SWI_OS_GetEnv`. They terminate with a call of `SWI_OS_Exit`, which transfers control to the current exit handler.

When a C program makes the call `system("command")` several things are done:

- The calling program and its data are copied to the top end of application workspace and all its handlers are removed.

- The current end of application workspace is set to just below the copied program and an exit handler is installed in case "command" is another application.
- "command" is invoked using SWI OS_Cli.

When "command" returns, either directly (if it is a built-in command) or via the exit handler (if it is an application), the caller is copied back to its original location, its handlers are re-installed and it continues, oblivious of the interruption.

The value returned by `system()` indicates

- whether the command or application was successfully invoked
- if the command is an application which obeys certain conventions, whether or not it ran successfully.

The value returned by `system()` (with a non-NULL command string) is as follows:

- < 0 – couldn't invoke the command or application (eg command not found);
- >=0 – invoked OK and set `Sys$ReturnCode` to the returned value.

By convention, applications set the environmental variable `Sys$ReturnCode` to 0 to indicate success and to something non-0 to indicate some degree of failure. Applications written in C do this for you, using the value passed as an argument to the `exit()` function or returned from the `main()` function.

If it is necessary to replace the current application by another, use:

```
system("CHAIN:command");
```

If the first characters of the string passed to `system()` are "CHAIN:" or "chain:", the caller is not copied to the top end of application workspace, no exit handler is installed, and there can be no return (return from a built-in command is caught by the C library and turned into a SWI OS_Exit).

Typically, `CHAIN:` is used to give more memory to the called application when no return from it is required. The C compiler invokes the linker this way if a link step is required. On the other hand, the Acorn Make Utility (AMU) calls each command to be executed. Such commands include the C compiler (as both use the shared C library, the additional use of memory is minimised). Of course, a called application can call other applications using `system()`. A callee can even `CHAIN:` to another application and still, eventually, return to the caller. For example, AMU might execute:

```
system("cc hello.c");
```

to call the C compiler. In turn, `cc` executes:

```
system("CHAIN:link -o hello o.hello $.CLib.o.Stubs");
```

to transfer control to the linker, giving link all the memory cc had.

However, when Link terminates (calls `exit()`, returns from `main()` or aborts) it returns to AMU, which continues (providing `Sys$ReturnCode` is good).

The shared C library

Release 4 of C makes extensive use of the shared C library module, first introduced with Release 2 of C and subsequently used by the RISC OS applications Edit, Paint, Draw and Configure.

The shared C library is a RISC OS relocatable module (called SharedCLibrary) which contains the whole of the ANSI C library. Once installed in your computer it can be used by every program written in C. Consequently, it save both RAM space and disc space.

In fact, this is as much as you really need to know about the shared C library and probably as far as you should delve at first reading. So, if you are eager to try your first practical work with this release of C, skip the rest of this section. However, if you are curious and would like to know more about what it really costs to use it, its benefits, and a little of how it works, then read on.

Costs involved in using the shared C library

The SharedCLibrary modules occupies about 60Kb. Each program that uses it must be linked with the **library stubs**, a small object module containing space for a copy of the shared C library's data and an entry vector via which functions in the shared library can be called. The stubs occupy just 5Kb. Thus a single program linked with the shared C library consumes about 65Kb of RAM for C library. However, two programs in memory at the same time use only 70Kb for library and three programs, only 75Kb.

In contrast, a program linked with Release 3 of ANSILib will include a minimum of 40Kb. So, as soon as you have two or more C programs in memory at the same time, it is cheaper to use the SharedCLibrary. Usually, you will have SrcEdit resident (which uses the shared C library anyway) and then you may want to run CC under Make. In this situation, use of the shared C library saves 45Kb of RAM.

Efficient use of RAM is not the only consideration. The C compiler includes 48Kb of ANSILib and when squeezed occupies 172Kb on disc. However, when linked with Stubs and squeezed it occupies only 140Kb. There are similar savings from Link, AMU, and Squeeze, as well as for the programs you compile (the 'hello world' program is reduced in size from over 40Kb to just 5.5Kb).

Without using the shared C library it would not be possible to use C Release 4 on a system with only a single floppy disc drive (imagine the loss of 150Kb of work space, together with a minimum image size of 40Kb). And, of course, smaller programs load much faster from a floppy disc.

If you have a larger Acorn system, use of the shared C library still brings benefits:

- Small programs load noticeably faster, even from a hard disc.
- No hard disc is ever big enough; saving 25-40Kb per program is not to be sneezed at if you have 40 or 50 programs (1-2Mb saved).
- Much more can be packed into the RAMFS – perhaps all the tools you ever use, giving almost instantaneous loading of them.

Execution time costs

It costs only 4 cycles (0.5 μ s) per function call and a very small penalty on access to the library's static data by the library (the user program's access to the same data is unpenalised). In general, the difference in performance between using the shared C library and linking a program stand-alone with ANSILib is less than 1%. For the important Dhrystone-2.1 benchmark the performance difference cannot be measured (you can try this experiment for yourself using the sources provided in User).

How it works

The shared C library module implements a single SWI which is called by code in the library stubs when your program linked with the stubs starts running. That SWI call tells the stubs where the library is in the machine. This allows the vector of library entry points contained in the stubs to be patched up in order to point at the relevant entry points in the library module.

The stubs also contain your private copy of the library's static data. When code in the library executes on your behalf, it does so using your stack and relocates its accesses to its static data by a value stored in your stack-chunk structure by the stubs initialisation code and addressed via the stack-limit register (this is why you must preserve the stack-limit register everywhere if you use the shared C library and call your own assembly language sub-routines). The compiler's register allocation strategy ensures that the real dynamic cost of the relocation is almost always low: for example, by doing it once outside a loop that uses it many times.

If you go on to write your own relocatable modules in C, you'll use the **Module code** SetUp option of the compiler which causes similar code to be generated.

#pragma directives

Pragmas recognised by the compiler come in two forms:

```
#pragma <letter><optional-digit>
```

and

```
#pragma [no]<feature-name>
```

A short-form pragma given without a digit resets that pragma to its default state; otherwise to the state specified.

For example,

```
#pragma -s1
#pragma nocheck_stack

#pragma -p2
#pragma profile_statements
```

The current list of recognised pragmas is:

pragma name	short form	'no' form
warn_implicit_fn_decls	a1	a0
warn_implicit_casts	b1	b0
check_memory_accesses	c1	c0
warn_deprecated	d1	d0
continue_after_hash_error	e1	e0
optimise_crossjump	j1	j0
optimise_multiple_loads	m1	m0
profile	p1	p0
profile_statements	p2	p0
check_stack	s0	s1
check_printf_formats	v1	v0
check_scanf_formats	v2	v0
check_formats	v3	v0
side_effects	y0	y1
optimise_cse	z1	z0

The set of pragmas recognised by the compiler, together with their default settings, varies from release to release of the compiler. In general, the only pragmas you should need to use are `check_stack` and `nocheck_stack`. These enable and disable, respectively, the generation of code to check the stack limit on function entry and exit. In reality there is little advantage to turning stack checks off: they cost at most two instructions and two machine cycles (about 0.25µs) per function call. The one occasion when `nocheck_stack` would be used is in writing a signal

handler for the SIGSTAK event. When this occurs, stack overflow has already been detected, so checking for it again in the handler would result in a fatal circular recursion.

Storage management (malloc, calloc, free)

The aim of the storage manager is to manage the heap in as 'efficient' a manner as possible. However, 'efficient' does not mean the same to all programs and since most programs differ in their storage requirements, certain compromises have to be made. The main issues to be considered are **speed** and **heap fragmentation**.

You should also try to keep the peak amount of heap used to a minimum so that, for example, a C program may invoke another C program leaving it the maximum amount of memory. This implementation has been tuned to hold the overhead due to fragmentation to less than 50%, with a fast turnover of small blocks.

The heap can be used in many different ways. For example it may be used to hold data with a long life (persistent data structures) or as temporary work space; it may be used to hold many small blocks of data or a few large ones or even a combination of all of these allocated in a disorderly manner. The storage manager attempts to address all of these problems but like any storage manager, it cannot succeed with all storage allocation/deallocation patterns. If your program is unexpectedly running out of storage, see the chapter entitled *Using memory efficiently*, earlier in this guide. This gives you information on the storage manager's strategy for managing the heap, and may help you to remedy the problem.

Note the following:

- The word *heap* refers to the section of memory currently under the control of the storage manager.
- All block sizes are in bytes and are rounded up to a multiple of four bytes.
- All blocks returned to the user are word-aligned.
- All blocks have an overhead of eight bytes (two words). One word is used to hold the block's length and status, the other contains a guard constant which is used to detect heap corruptions. The guard word may not be present in future releases of the ANSI C library.

Handling host errors

Calls to RISC OS can be made via one of the functions in the C header file `kernel.h`, (such as `_kernel_osfind(64, ".....")`). If the call causes an operating system error, the function will return the value `_kernel_ERROR`. To find out what the error was, a call to `_kernel_last_oserror` should be made. This will return a pointer to a `_kernel_oserror` block containing the error

number and any associated error string. If there has been no error since `_kernel_last_oserror` was last called, the function returns the NULL pointer. Some functions in the ANSI C library call `_kernel` functions, so if an ANSI C library function (such as `fopen(". . . .", "r")`) fails, try calling `_kernel_last_oserror` to find out what the error was.

For more details about operating system calls, refer to the `kernel.h` header (reproduced as *Appendix C: kernel.h* in this Guide), and for more information about RISC OS error handling, refer to the chapter entitled *Generating and handling errors* in the *RISC OS Programmer's Reference manual*.

Appendices

18 Appendix A: New features of Desktop C

Acorn Desktop C is the fourth release of an Acorn C compiler product for RISC OS, and replaces the ANSI C Release 3 product. For the first time you can develop C programs within a desktop development environment provided by 19 programming tools working together in the RISC OS desktop.

Some of these tools are completely new, such as the windowed DDT debugger and desktop Make, while some are improved versions of old tools such as the FormEd template file editor and the SrcEdit desktop program editor. The C compiler is one of many tools which previously was operated from the command line, and has now been provided with a standardised desktop interface. Many tools such as the text searching programs Diff and Find which were previously part of the Software Developer's Toolbox are now included in Acorn Desktop C.

In addition to the new support for developers while working on their host systems, additions and improvements have been made to the support of finished applications running on their target machines. The FrontEnd relocatable module supplied with Acorn Desktop C can provide a complete RISC OS desktop interface for a non-interactive command line program knowing nothing of the RISC OS desktop. This has been used to add interfaces to tools such as the C compiler, and you can use it for your programs. The performance of some tools such as the C compiler has been enhanced, and some additions made to library support.

The user guides supplied with Acorn Desktop C now extend to two volumes, incorporating coverage of the new tools and interfaces, an increased number of worked examples and details of how the Desktop Development Environment fits together so that you can extend it with your own tools if you wish.

Acorn Desktop C tools

The programming tools included with Acorn Desktop C are all either new or enhanced since release 3. They are:

- DDT – the new Desktop Debugging tool, capable of debugging RISC OS desktop applications.
- SrcEdit – the new programmer's editor operating in the desktop (a development of Edit).

- Make – the new desktop make and automatic makefile generator.
- FormEd – an improved version of the old FormEd template file editor.
- CC – the C compiler. This has the addition of a desktop front end, plus further optimisation yielding a few percent more compact code.
- CMHG – the C Module Header Generator, now with a desktop front end.
- ToANSI – the C dialect conversion program, now with a desktop front end.
- ToPCC – the C dialect conversion program, now with a desktop front end.
- AMU – the make utility with a new desktop front end.
- Common – a utility to count the most common words in a text file. This was previously part of the Software Developer's Toolbox, and has an added desktop front end.
- Diff – a utility to compare two text files. This was previously part of the Software Developer's Toolbox, and has an added desktop front end.
- DecAOF – a utility to decode Acorn Object Format files, with a desktop front end.
- DecCF – a utility to decode Chunk Files, with a desktop front end.
- Find – a utility to find text in file names and text file contents. This was previously part of the Software Developer's Toolbox, and now has a desktop front end.
- Link – the linker, now with a desktop front end.
- LibFile – the tool for constructing and managing linkable libraries of object files. This performs the job of two utilities provided as part of the Software Developer's Toolbox, and has an added desktop front end.
- ObjSize – a tool extracting size information from an Acorn Object Format file, with an added desktop front end.
- Squeeze – a tool for compacting executable binary files, now also capable of compacting relocatable modules, with an added desktop front end.
- WC – a utility to count the number of words in a text file. This was previously part of the Software Developer's Toolbox, and has an added desktop front end.

The only tool supplied with release 3 which is no longer included is ASD, a command line debugger superseded by DDT, which has greater versatility and a windowed display.

New technical features

The C compiler has been enhanced in the following ways since release 3:

- Conformance with the newly ratified ANSI standard (see the Introduction of this manual for more details).
- Slightly improved space efficiency of generated code – code files generated are on average a few percent smaller than those of the release 3 compiler.

The RISC_OSLib linkable library of OS specific routines assisting the production of applications has been extended, with the addition of over 30 new functions.

The procedure call standard has not changed since release 3.

Since only a very low percentage of Archimedes users still run the Arthur operating system which was superseded by RISC OS, the Arthurlib library and its header (which were documented as obsolescent in release 3) are no longer included.

User guides

The user guide for Acorn Desktop C has been split into two volumes – Desktop Development Environment and ANSI C Release 4. The former volume covers all tools and other information not specific to programming in C, whereas the latter covers the C processing tools, language issues and programming information.

The volume of material in the user guides has increased significantly as a result of describing the new tools, their user interfaces and information to extend the development environment with your own tools. Each tool now has a chapter devoted to it.

19 Appendix B: Errors and warnings

This appendix gives a brief description of the intended purposes of error and warning messages from the CC tool, along with some hints for interpreting them. It then lists most of them in alphabetical order. Since the messages are designed as far as possible to be self-explanatory, some of the more simple ones are not listed here.

Interpreting CC errors and warnings

The compiler can produce error and warning messages of several degrees of severity. They are as follows:

- **Warnings** indicating curious, but legal, program constructs, or constructs that are indicative of potential error;
- **Non-serious errors** that still allow code to be produced;
- **Serious errors** that may cause loss of code;
- **Fatal errors** that may stop the compiler from compiling;
- **System errors** that signal faults in the system itself.

Warnings from CC are intended to provide a helpful level of checking, in addition to the level required by the ANSI standard. On some other systems, such as UNIX, separate facilities (like lint) are provided to perform this checking. Warnings flag program constructs that may indicate potential errors, or those not recommended because they may function differently on other machines, and hence hinder the portability of code.

Some warnings point out the use of facilities provided in this ANSI C implementation which are above the minimum required by ANSI – for example, use of external identifiers that are identical in the first six characters, which may not be differentiated by other systems which conform to the ANSI standard.

Programs ported from other machines may cause large numbers of warning messages from CC, which you can disable with the **Suppress** option (see the chapter entitled CC for more information). When you have finished producing high-quality software, you can also enable additional checks with the CC **Feature** option.

Errors and serious errors collectively respond to ANSI 'diagnostics'; whether an error is serious or not reflects the compiler's view, not yours, or that of the ANSI committee.

After issuing a warning, non-serious, or serious error, CC continues compiling, sometimes producing more such messages in the process. Compilation of C by CC can be thought of as a pipeline process, starting with preprocessing, syntax analysis, then semantic analysis (when the structure of a portion of code is analysed). When syntax errors in C are encountered by CC, the compiler can often guess what the error was, correct it, and continue. When semantic errors are found, portions of your code are often ignored before continuing, and serious error messages are reported.

Unfortunately, the compact and powerful nature of C leads to a high proportion of semantic errors. Ignoring portions of your code is likely to make subsequent portions incorrect, so one serious error can often start a cascade of error messages. Often, therefore, it is sensible to ignore a set of error messages following a serious error message.

If the compiler produces any message more serious than a warning, it will set a bad return code, usually terminating any 'make' of which it is a part in the process. Any serious error will cause the output object file to be deleted; fatal and system errors cause immediate termination of compilation, with loss of the object file and bad return code set.

Future releases of the compiler may distinguish further forms of error, or produce slightly different forms of wording.

In pcc mode, constructs that are erroneous in ANSI mode are reported, even though legal in pcc mode.

Warnings

Warning messages indicate legal but curious C programs, or possibly unintended constructs (unless warnings are suppressed). On detection of such a condition, the compiler issues a warning message, then continues compilation.

Warning messages

'&' unnecessary for function or array xx

This is a reminder that if `xx` is defined as `char xx[10]` then `xx` already has a pointer type. There is a similar reminder for function names too. Example:

```
static char mesg[] = "hello\n";
int main ()
{
    char *p = &mesg; /* mesg is already compatible with char * */
    ...
}
```

actual type 'xx' mismatches format '%x'

A type error in a `printf` or `scanf` format string. Example:

```
{
    int i;
    printf("%s\n", i); /* %s need char* not int */
    ...
}
```

ANSI 'xx' trigraph for 'x' found - was this intended?

This helps to avoid inadvertent use of ANSI trigraphs. Example:

```
printf("Type ??/!!: "); /* "??/" is trigraph for "\" */
```

argument and old-style parameter mismatch : xx

A function with a non-ANSI declaration has been called using a parameter of a wrong data type. Example:


```
int fn1(a , b)
int a;
int b;
{
    return a * b;
}
...
int main()
{
    int l; float m;
    fn1(l , m);          /* m should be 'int' */
    ...
}
```

character sequence /* inside comment

You cannot nest comments in C. Example:

```
/* comment out func() for now...
/* func() returns a random number */
int func(void)
{
    ...
    return i;
}
*/
```

Dangling 'else' indicates possible error

This hints that you may have mis-matched your ifs and elses. Remember an else always refers to the most recent un-matched if. Use braces to avoid ambiguity. Example:

```
if (a)
    if (b)
        return 1;
    else if (c)
        return 2;
    else /* this belongs to the if (a). Or does it?*/
        return 3;
```

Deprecated declaration of xx() - give arg types

A feature of the ANSI standard is that argument types should be given in function declarations (prototypes). 'No arguments' is indicated by void. Example:

```
extern int func(); /* should have 'void' in the parentheses */
```

extern clash xx , xx clash (ANSI 6 char monospace)

Using compiler **Feature** option e, it was found that two external names were not distinct in the first six characters. Some linkers provide only six significant characters in their symbol table. Example:

```
extern double function1 (int i);
extern char * function2 (long l);
```

extern 'main' needs to be 'int' function

This is a reminder that main() is expected to return an integer. Example:

```
void main()
{
    ...
}
```

extern xx not declared in header

Compiling with **Feature** h, an external object was discovered which was not declared in any included header file.

floating point constant overflow

This is typically caused by a division by zero in a floating point constant expression evaluated at compile time. Example:

```
#define lim 1
#define eps 0.01
static float a = eps/(lim-1); /* lim-1 yields 0 */
```

floating to integral conversion failed

A cast (possibly implicit) of a floating point constant to an integer failed at compile time. Example:

```
static int i = (int) 1.0e20; /* INT_MAX is about 2e10 */
```

formal parameter 'xx' not declared - 'int' assumed

The declaration of a function parameter is missing. Example:

```
int func(a)
/*a should be declared here or within the parentheses*/
{
    ...
}
```

Format requires nn parameters, but mm given

Mismatch between a printf or scanf format string and its other arguments.

Example:

```
printf("%d, %d\n",1); /* should be two ints */
```

function xx declared but not used

When compiling with **Feature v**, the function xx was declared but not used within the source file.

Illegal format conversion '%x'

Indicates an illegal conversion implied by a printf or scanf format string.

Example:

```
printf("%w\n",10); /* no such thing as %w */
```

implicit narrowing cast : xx

An arithmetic operation or bit manipulation is attempted involving assignment from one data type to another, where the size of the latter is naturally smaller than that of the assigned value. Example:

```
double d = 1.0; long l = 2L; int i = 3;
i = d * i;
i = l | 3;
i = l & ~1;
```

implicit return in non-void function

A non-void function may exit without using a return statement, but won't return a meaningful result. Example:

```
int func(int a)
{
    int b=a*10;
    .../* no return <expr> statement */
}
```

incomplete format string

A mistake in a `printf` or `scanf` format string. Example:

```
printf("Score was %d%",score); /* 2nd % should be %% */
```

'int xx()' assumed - 'void' intended?

If the definition of a function omits its return type – it defaults to `int`. You should be explicit about the type, using `void` if the function doesn't return a result.

Example:

```
main()
{
    ...
}
```

inventing 'extern int xx();'

The declaration of a function is missing. Example:

```
printf("Type your name: ");
/* forgot to #include <stdio.h> */
```

lower precision in wider context: xx

An arithmetic operation or bit manipulation is attempted involving assignment from `int`, `short` or `char` to `long`. Example:

```
long l = 1L; int i = 2; short j = 3;
l = i & j;
l = i | 5;
l = i * j;
```

One circumstance in which this causes problems is when code like

```
long f(int x){return 1<<x;}
```

(which fails if `int` has 16 bits) is moved to machines such as the IBM PC.

No side effect in void context: 'op'

An expression which does not yield any side effect was evaluated; it will have no effect at run-time. Example:

```
a+b;
```

no type checking of enum in this compiler

Compiling with **Feature x**, an enum declaration was found, and this message refers to the ANSI stipulation that enum values be integers, less strictly typed than in some earlier dialects of C.

Non-ANSI #include <xx>

A header file has been #included which is not defined in the ANSI standard. <> should be replaced by "".

non-portable - not 1 char in 'xx'

Assigning character constants containing more than one character to an int will produce non-portable results. Example:

```
static int exitCode = 'ABEX';
```

non-value return in a non-void function

The expression was omitted from a return statement in a function which was defined with a non-void return type. Example:

```
int func(int a)
{
    int b=a*10;
    ...
    return; /* no <expr> */
}
```

odd unsigned comparison with 0 : xx

An attempt has been made to determine whether an unsigned variable is negative. Example:

```
unsigned u , v;
if (u < 0) u = u * v;
if (u >= 0) u = u / v;
```

Old-style function: xx

Compiling with **Feature o**, it was noted that the code contains a non-ANSI function declaration. Example:

```
void fn2(a , b)
int a;
int b;
{ b = a; }
```

omitting trailing '\0' for char[nn]

The character array being equated to a string is one character too short for the whole string, so the trailing zero is being omitted. Example:

```
static char mesg[14] = "(C)1988 Acorn\n"; /* needs 15 */
```

repeated definition of #define macro xx

When compiling with **Feature h**, a macro has been repeatedly #defined to take the same value.

shift by nn illegal in ANSI C

This is given for negative constant shifts or shifts greater than 31. On the ARM, the bottom byte of the number given is used, ie it is treated as (unsigned char) nn. NB: negative shifts are not treated as positive shifts in the other direction. Example:

```
printf("%d\n", 1<<-2);
```

'short' slower than 'int' on this machine (see manual)

For speed you are advised to use ints rather than shorts where possible. This is because of the overhead of performing implicit casts from short to int in expression evaluation. However, shorts are half the size of ints, so arrays of shorts can be useful. Example:

```
{
    short i,j; /* quicker to use ints */
    ...
}
```

spurious {} around scalar initialiser

Braces are only required around structure and array initialises. Example:

```
static int i = {INIT_I}; /* don't need braces */
```

static xx declared but not used

A static variable was declared in a file but never used in it. It is therefore redundant.

Unrecognised #pragma (no '-' or unknown word)

#pragma directives are of the form

```
#pragma -xd  
or  
#pragma long_spelling
```

where *x* is a letter and *d* is an optional digit. These messages warn against unknown letters and missing minus signs.

use of 'op' in condition context

Warns of such possible errors as = and not == in an if or looping statement. Example:

```
if (a=b) {  
    ...  
}
```

variable xx declared but not used

This refers to an automatic variable which was declared at the start of a block but never used within that block. It is therefore redundant. Example:

```
int func(int p)  
{  
    int a; /* this is never used */  
    return p*100;  
}
```

xx may be used before being set

Compiling with **Feature a**, an automatic variable is found to have been used before any value has been assigned to it.

xx treated as xxul in 32-bit implementation

This message warns of two's complement arithmetic's dependence on assigning negative constants to unsigned ints, and it explains that ints and long ints are both 32 bits.

Non-serious errors

These are errors which will allow 'working' code to be produced – they will not produce loss of code. On detection of such an error the compiler issues an error message, if enabled, then continues compilation.

',' (not ';') separates formal parameters

Incorrect punctuation between function parameters. Example:

```
extern int func(int a;int b);
```

ANSI C does not support 'long float'

This used to be a synonym for double, but is not allowed in ANSI C.

ancient form of initialisation, use '='

An obsolete syntax for initialisation was used, or incorrectly nested brackets have been found. Example:

```
int i{1};          /* use int i=1; */
```

array [0] found

The minimum subscript count allowed is 1. (Remember that the subscripts go from 0..n-1.) Example:

```
static int a[0];
```

array of xx illegal - assuming pointer

Illegal objects have been declared to occupy an array. Examples:

```
int fn2[5]();          /* array of functions */
void v[10];           /* array of voids */
```

assignment to 'const' object 'xx'

You can't assign to objects declared as const. Example:


```

{
    const int ic = 42; /* initialisation ok */
    ic = 69; /* can't change it now */
    ...

```

comparison 'op' of pointer and int:

literal 0 (for == and !=) is the only legal case

You cannot use the comparison operators between an integer and a pointer type. As the message implies, you can only check for a pointer being (not) equal to NULL (int 0). Example:

```

{
    int i,j,*ip;
    j = i>ip; /* can't compare an int and an int * */
    ...

```

declaration with no effect

The compiler detected what appeared to be a declaration statement, but which resulted in no store being allocated. This may imply that a data type name was omitted.

differing pointer types: 'xx'

An illegal implicit type cast was detected in a comparison operation between two pointers of different types. Example:

```

{
    int    *ip;
    char   *cp;
    printf("%d\n", ip==cp); /* can't compare these */
    ...

```

differing redefinition of #define macro xx

#define gives a definition contradicting that already assigned to the named macro.

ellipsis (...) cannot be only parameter

Although C allows variable length argument lists, the '...' parameter cannot stand alone in this function declaration. Example:

```
void fn1(...) { }
```

expected 'xx' or 'x' - inserted 'x' before 'yy'

Often caused by omitting a terminating symbol in a statement when the compiler is able to insert this symbol for you, and then to recover. Example:

```
int f(int j)
{
    return j;
}
int main()
{
    int i=f(10;    /* ')' omitted here */
    return i;
}
```

formal name missing in function definition

This error occurs when a comma in a function definition led the compiler to suspect a further formal parameter was going to follow, but none did. Example:

```
int a(int b,) /* missing parameter */
{
    ...
}
```

function prototype formal 'xx' needs type or class - 'int' assumed

A formal parameter in a function prototype was not given a type or class. It needs at least one of these (register being the only allowed class). Example:

```
void func(a); /* I mean int a or perhaps register a */
```

function returning xx illegal - assuming pointer

A function apparently intends to return an illegal object. Example:

```
int fn3() [] /* hoping to return an array */
{
    int list[3] = {1,2,3};
    return list;
}
```

function *xx* may not be initialised - assuming function pointer

A function is not a variable, so cannot be initialised. As an attempt to initialise *xx* has been made, *xx* is treated as of type function *. Example:

```
extern int func(void);
static int fn() = func; /* the compiler will use
                        static int (*fn)() = func; instead */
```

<int> *op* <pointer> treated as <int> *op* (int)<pointer>

Warns of an illegal implicit cast within an expression. Typically *op* is an operator which has no business being used on pointers anyway, such as `|` or dyadic `*`. Example:

```
{
    int i, *ip;
    i = i | ip; /* bitwise-or on a pointer?! */
    ...
}
```

junk at end of #*xx* line - ignored

The *xx* is either `else` or `endif`. These directives should not have anything following them on the line. Example:

```
/* text after the #else should be a comment */
#else if it isn't defined
...
```

L'...' needs exactly 1 wide character

The `wchar_t` declaration of a wide character names an identifier comprising other than one character. Example:

```
wchar_t wc = L'abc';
```

linkage disagreement for '*xx*' - treated as '*xx*'

There was a linkage type disagreement for declarations, eg a function was declared as `extern` then defined later in the file as `static`. Example:

```
int func(int a); /* compiler assumes extern here */
...
static func(int a) /* but told static here */
{
    ...
}
```

more than 4 chars in character constant

A character constant of more than four characters cannot be assigned to a 32 bit int. Example:

```
{
    int i = '12345'; /* more than four chars */
    ...
}
```

no chars in character constant ''

At least one character should appear in a character constant. The empty constant is taken as zero. Example:

```
{
    int i = ''; /* less than one char == '\0' */
    ...
}
```

objects that have been cast are not l-values

The programmer tried to use a cast expression as an l-value. Example:

```
char *p;
*((int *)p)=10;          /* (int *)p is NOT an l-value */
```

omitted <type> before formal declarator - 'int' assumed

This is given in a formal parameter declaration where a type modifier is given but no base type. Example:

```
int func(*a); /* a is a pointer, but to what? */
```

'op': cast between function pointer and non-function object

Casts between function and object pointers can be very dangerous! One possibly valid (but still very suspect) use is in casting an array of int into which machine code has been loaded into a function pointer. Example:

```
static int mcArray[100];
/*pointer to function returning void*/
typedef void (*pfv)(void);
...
((pfv)mcArray)();/* convert to fn type and apply */
```

'op': implicit cast of non-0 int to pointer

Zero, equal to a NULL pointer, is the only int which can be legally implicitly cast to a pointer type. Example:

```
{
    int i, *ip;
    ip = i; /* only the constant int 0 can be
            implicitly cast to a pointer type */
    ...
```

'op': implicit cast of pointer to non-equal pointer

An illegal implicit cast has been detected between two different pointer types. The type casting must be made explicit to escape this error. Example:

```
{
    int *ip;
    char *cp;
    ip = cp; /* differing pointer types */
    ...
```

'op': implicit cast of pointer to 'int'

An illegal implicit cast has been detected between an integer and a pointer. Such casts must be made explicitly. Example:

```
{
    int i, *ip;
    i = ip; /* pointer must be cast explicitly */
    ...
```

overlarge escape '\\xxxx' treated as '\\xxx'

A hexadecimal escape sequence is too large. Example:

```
int novalue()
{
    if (seize) return '\xffff';          /* \xffff' too large */
    else return '\xff';
}
```

overlarge escape '\\x' treated as '\\x'

An octal escape sequence is too large. Example:

```
int novalue()
{
    if (huit) return '\777';            /* \777 too large */
    else return '\77';
}
```

<pointer> op <int> treated as (int)<pointer> op <int>

The only legal operators allowed in this context are + and -.

prototype and old-style parameters mixed

Use has been made of both the ANSI style function/definition (including a type name for formal parameters in a function's heading) and pcc style parameters lists.

Example:

```
void fn4(a, int b)
int a;
{
    a = b;
}
```

'register' attribute for 'xx' ignored when address taken

Addresses of register variables cannot be calculated, so an address being taken of a variable with a register storage class causes that attribute to be dropped.

Example:

```
{
    register int i, *ip;
    ip = &i; /* & forces i to lose its register attribute */
    ...
}
```

return <expr> illegal for void function

A function declared as void must not return with an expression. Example:

```
void a(void)
{
    ...
    return 0;
}
```

size of 'void' required - treated as 1

This indicates an attempt to do pointer arithmetic on a void *, probably indicating an error. Example:

```
{
    void    *vp;
    vp++; /* how many bytes to increment by ? */
    ...
}
```

size of a [] array required - treated as [1]

If an array is declared as having an empty first subscript size, the compiler cannot calculate the array's size. It therefore assumes the first subscript limit to be 1 if necessary. This is unlikely to be helpful.

```
extern int array[][10];
static int s = sizeof(array); /*can't determine this*/
```

sizeof <bit field> illegal - sizeof(int) assumed

Bitfields do not necessarily occupy an integral number of bytes but they are always parts of an int, so an attempt to take the size of a bitfield will return sizeof(int). Example:

```
struct s {
    int exp : 8;
    int mant : 23;
    int s : 1;
};
int main(void)
{
    struct s st;
    int i = sizeof(st.exp); /* can't obtain this in
                             bytes */
    ...
}
```

Small (single precision) floating value converted to 0.0
Small floating point value converted to 0.0

A floating point constant was so small that it had to be converted to 0.0. Example:

```
static float f = 1.0001e-38 - 1.0e-38; /* 1e-42 too
                                         small for
                                         float */
```

Spurious #elif ignored
Spurious #else ignored
Spurious #endif ignored

One of these three directives was encountered outside any #if or #ifdef scope. Example:

```
#if defined sym
...
#endif
#else /* this one is spurious */
...

```

static function xx not defined - treated as extern

A prototype declares the function to be static, but the function itself is absent from this compilation unit.

string initialiser longer than char [nn]

An attempt was made to initialise a character array with a string longer than the array. Example:

```
static char str[10] = "1234567891234";
```

struct component xx may not be function - assuming function pointer

A variable such as a structure component cannot be declared to have type function, only function *. Example:

```
struct s {
    int fn(); /* compiler will use int (*fn)(); */
    char c;
};
```


type or class needed (except in function definition) - int assumed

You can't declare a function or variable with neither a return type nor a storage class. One of these must be present. Examples:

```
func(void); /* need, eg, int or static */
x;
```

Undeclared name, inventing 'extern int xx'

The name xx was undeclared, so the default type `extern int` was used. This may produce later spurious errors, but compilation continues. Example:

```
int main(void) {
    int i = j; /*j has not been previously declared*/
    ...
}
```

unprintable character xx found - ignored

An unrecognised character was found embedded in your source – this could be file corruption, so back up your sources! Note that 'unprintable character' means any non-whitespace, non-printable character.

variable xx may not be function - assuming function pointer

A variable cannot be declared to have type `function`, only `function *`. Example:

```
int main(void)
{
    auto void fn(void); /* treated as void (*fn)(void);*/
    ...
}
```

xx may not have whitespace in it

Tokens such as the compound assignment operators (`+=` etc) may not have embedded whitespace characters in them. Example:

```
{
    int i;
    ...
    i += 4; /* space not allowed between + and = */
    ...
}
```

Serious errors

These are errors which will cause loss of generated code. On detection of such an error, the compiler will attempt to continue and produce further diagnostic messages, which are sometimes useful, but will delete the partly produced object file.

'...' must have exactly 3 dots

This is caused by a mistake in a function prototype where a variable number of arguments is specified. Example:

```
extern int printf(const char *format,...); /*one . too many*/
```

'{' of function body expected - found 'xx'

This is produced when the first character after the formal parameter declarations of a function is not the { of the function body. Example:

```
int func(a)
int a;
    if (a) ... /* omitted the { */
```

'{' or <identifier> expected after 'xx', but found 'yy'

xx is typically struct or union, which must be followed either by the tag identifier or the open brace of the field list. Example:

```
struct *fred; /* Missed out the tag id */
```

'xx' variables may not be initialised

A variable is of an inappropriate class for initialisation. Example:

```
int main()
{
    extern int n=1;
    return 1;
}
```

'op': cast to non-equal 'xx' illegal

'op': illegal cast of 'xx' to pointer

'op': illegal cast to 'xx'

These errors report various illegal casting operations. Examples:

```

struct s {
    int a,b;
};
struct t {
    float ab;
};
int main(void)
{
    int i;
    struct s s1;
    struct t s2;
    /* '=': illegal cast to 'int' */
    i = s1;
    /* '=': illegal cast to non-equal 'struct' */
    s1 = s2;
    /* <cast>: illegal cast of 'struct' to pointer */
    i = *(int *) s1;
    /* <cast>: illegal cast to 'int' */
    i = (int) s2;
    ...
}

```

'op': illegal use in pointer initialiser

(Static) pointer initialisers must evaluate to a pointer or a pointer constant plus or minus an integer constant. This error is often accompanied by others. Example:

```

extern int count;
static int *ip = &count*2;

```

{ } must have 1 element to initialise scalar

When a scalar (integer or floating type) is initialised, the expression does not have to be enclosed in braces, but if they are present, only one expression may be put between them. Example:

```

static int i = {1,2}; /* which one to use? */

```

Array size nn illegal - 1 assumed

Arrays have a maximum dimension of 0xffffffff. Example:

```

static char dict[0x1000000]; /* Too big */

```

attempt to apply a non-function

The function call operator () was used after an expression which did not yield a pointer to function type. Example:

```
{
    int i;
    i();
    ...
}
```

Bit fields do not have addresses

Bitfields do not necessarily lie on addressable byte boundaries, so the & operator cannot be used with them. Example:

```
struct s {
    int h1,h2 : 13;
};
int main(void)
{
    struct s s1;
    short *sp = &s1.h2; /* can't take & of bit field */
    ...
}
```

Bit size *nn* illegal - 1 assumed

Bitfields have a maximum permitted width of 32 bits as they must fit in a single integer. Example:

```
struct s {
    int f1 : 40; /* This one is too big */
    int f2 : 8;
};
```

'break' not in loop or switch - ignored

A break statement was found which was not inside a for, while or do loop or switch. This might be caused by an extra }, closing the statement prematurely. Example:

```
int main(int argc)
{
    if (argc == 1)
        break;
    ...
}
```

'case' not in switch - ignored

A case label was found which was not inside a switch statement. This might be caused by an extra }, closing the switch statement prematurely. Example:

```
void fn(void)
{
    case '*': return;
    ...
}
```

<command> expected but found a 'op'

This error occurs when a (binary) operator is found where a statement or side-effect expression would be expected. Example:

```
if (a) /10; /* mis-placed ) perhaps? */
...
```

'continue' not in loop - ignored

A continue statement was found which was not inside a for, while or do loop. This might be caused by an extra }, closing the loop statement prematurely. Example:

```
while (cc) {
    if (dd) /* intended a { here */
        error();
    } /*this closes the while */
    if (ee)
        continue;
}
```

'default' not in switch - ignored

A default label was found which was not inside a switch statement. This might be caused by an extra }, closing the switch statement prematurely. Example:

```

switch (n) {
  case 0:
    return fn(n);
  case 1: if (cc)
    return -1;
  else
    break;
} /* spurious } closes the switch */
default:
  error();
}

```

duplicated case constant: *nn*

The case label whose value is *nn* was found more than once in a `switch` statement. Note that *nn* is printed as a decimal integer regardless of the form the expression took in the source. Example:

```

switch (n) {
  case ' ':
  ...
  case ' ':
  ...
}

```

duplicate 'default' case ignored

Two cases in a single `switch` statement were labelled `default`. Example:

```

switch (n) {
  default:
  ...
  default:
  ...
}

```

duplicate definition of 'struct' tag 'xx'

There are duplicate definitions of the type `struct xx {...}`; Example:

```

struct s { int i,j;};
struct s {float a,b;};

```

duplicate definition of 'union' tag 'xx'

There are duplicate definitions of the type union `xx {...}` ;. Example:

```
union u {int i; char c[4];};
union u {double d; char c[8];};
```

duplicate type specification of formal parameter 'xx'

A formal function parameter had its type declared twice, once in the argument list and once after it. Example:

```
void fn(int i)
int i;          /* this one is redundant */
{
    ...
}
```

EOF in comment

EOF in string

EOF in string escape

These all refer to unexpected occurrences of the end of the source file.

Expected <identifier> after 'xx' but found 'xx'
expected 'xx' - inserted before 'yy'

This typically occurs when a terminating semi-colon has been omitted before a `}`. (Common amongst Pascal programmers) Another case is the omission of a closing bracket of a parenthesised expression. Examples:

```
int fn(int a, int b, int c)
{
    int d = a*(b+c;          /* missing ) */
    return d                /* missing ; */
}
```

Expecting <declarator> or <type>, but found 'xx'

`xx` is typically a punctuation character found where a variable or function declaration or definition would be expected (at the top level). Example:

```
static int i = MAX;+1;    /* spurious ; ends expression */
```

<expression> expected but found 'op'

Similar to above. An operator was found where an operand might reasonably be expected. Example:

```
func(>>10); /* missing left hand side of >> */
```

grossly over-long floating point number

Only a certain number of decimal digits are needed to specify a floating point number to the accuracy that it can be stored to. This number of digits was exceeded by an unreasonable amount.

grossly over-long number

A constant has an excessive number of leading zeros, not affecting its value.

hex digit needed after 0x or 0X

Hexadecimal constants must have at least one digit from the set 0..9, a..f, A..F following the 0x. Example:

```
int i = 0xg; /* illegal hex char */
```

<identifier> expected but found 'xx' in 'enum' definition

An unexpected token was found in the list of identifiers within the braces of an enum definition. Example:

```
enum colour {red, green, blue,;}; /* spurious ; */
```

identifier (xx) found in <abstract declarator> - ignored

The `sizeof()` function and cast expressions require abstract declarators, ie types without an identifier name. This error is given when an identifier is found in such a situation. Examples:

```
i = (int j) ip; /* trying to cast to integer */
l = sizeof(char str[10]); /* probably just mean
                           sizeof(str) */
```

illegal bit field type 'xx' - 'int' assumed

Int (signed or unsigned) is the only valid bitfield type in ANSI-conforming implementations. Example:


```
struct s { char a : 4; char b : 4;};
```

illegal in case expression (ignored): xx

illegal in constant expression: xx

illegal in floating type initialiser: xx

All of these errors occur when a constant is needed at compile time but a variable expression was found.

illegal in l-value: 'enum' constant 'xx'

An incorrect attempt was made to assign to an enum constant. This could be caused by misspelling an enum or variable identifier. Example:

```
enum col {red, green, blue};
int fn()
{
    int read;
    red = 10;
    ...
}
```

illegal in the context of an l-value: 'xx'

illegal in lvalue: function or array 'xx'

An incorrect attempt was made to assign to xx, where the object in question is not assignable (an l-value). You can't, for example, assign to an array name or a function name. Examples:

```
{
    int a,b,c;
    a ? b : c = 10;        /* ?: can't yield l-values. */
    if (a)                /* use this instead */
        b = 10;
    else
        c = 10;
    ...
}
```

or, in the same context,

```
*(a ? &b: &c) = 10;
```

illegal in static integral type initialiser: xx

A constant was needed at compile time but a suitable expression wasn't found.

illegal types for operands : 'op'

An operation was attempted using operands which are unsuitable for the operator in question. Examples:

```
{
    struct {int a,b;} s;
    int i;
    i = *s;      /* can't indirect through a struct */
    s = s+s;    /* can't add structs */
    ...
}
```

incomplete type at tentative declaration of 'xx'

An incomplete non-static tentative definition has not been completed by the end of the compilation unit. Example:

```
int incomplete[];
...
/* should be completed with a declaration like: */
/* int incomplete[SOMESIZE];                      */
```

```
junk after #if <expression>
junk after #include "xx"
junk after #include <xx>
```

None of these directives should have any other non-whitespace characters following the expression/filename. Example:

```
#include <stdio.h> this isn't allowed
```

label 'xx' has not been set

An attempt has been made to use a label that has not been declared in the current scope, after having been referenced in a `goto` statement. Example:

```
int main(void)
{
    goto end;
}
```

misplaced '{' at top level - ignoring block

{ } blocks can only occur within function definitions. Example:

```
/* need a function name here */
{
    int i;
    ...
}
```

misplaced 'else' ignored

An else with no matching if was found. Example:

```
if (cc)          /* should have used { } */
    i = 1;
    j = 2;
else
    k = 3;
...
```

misplaced preprocessor character 'xx'

Usually a typing error; one of the characters used by the preprocessor was detected out of context. Example:

```
char #str[] = "string";          /* should be char *str[] */
```

missing #endif at EOF

A #if or #ifdef was still active at end of the source file. These directives must always be matched with a #endif.

missing '"' in pre-processor command line

A line such as #include "name has the second " missing.

missing ')' after xx(... on line nn

The closing bracket (or comma separating the arguments) of a macro call was omitted. Example:

```
#define rdch(p) {ch=*p++;}
...
{
    rdch(p          /* missing ) */
    ...
}
```

missing ', ' or ')' after #define xx(...

One of the above characters was omitted after an identifier in the macro parameter list. Example:

```
#define rdch(p {ch = *p++;}
```

missing '<' or ''' after #include

A #include filename should be within either double quotes or angled brackets.

missing hex digit(s) after \x

The string escape \x is intended to be used to insert characters in a string using their hexadecimal values, but was incorrectly used here. It should be followed by between one and three hexadecimal digits. Example:

```
printf("\xxx/"); /* probably meant "\\xxx/" */
```

missing identifier after #define**missing identifier after #ifdef****missing identifier after #undef**

Each of these directives should be followed by a valid C identifier. Example:

```
#define @ at
```

missing parameter name in #define xx(...

No identifier was found after a , in a macro parameter list. Example:

```
#define rdch(p,) {ch=*p++;}
```

no ')' after #if defined (...

The defined operator expects an identifier, optionally enclosed within brackets. Example:

```
#if defined(debug
```

no identifier after #if defined

See above.

non static address 'xx' in pointer initialiser

An attempt was made to take the address of an automatic variable in an expression used to initialise a `static` pointer. Such addresses are not known at compile-time. Example:

```
{
    int i;
    static int *ip = &i; /*&i not known to compiler*/
    ...
}
```

non-formal 'xx' in parameter-type-specifier

A parameter name used to declare the parameter types did not actually occur in the parameter list of the function. Example:

```
void fn(a)
int a,b;
{
    ...
}
```

number nn too large for 32-bit implementation

An integer constant was found which was too large to fit in a 32 bit int. Example:

```
static int mask = 0x800000000; /*0x80000000 intended?*/
```

objects or arrays of type void are illegal

`void` is not a valid data type.

overlarge floating point value found

overlarge (single precision) floating point value found

A floating point constant has been found which is so large that it will not fit in a floating point variable. Examples:

```
float f = 1e40; /* largest is approx 1e38 for float */
double d = 1e310; /* and 1e308 for double */
```

quote (" or ') inserted before newline

Strings and character constants are not allowed to contain unescaped newline characters. Use `\<nl>` to allow strings to span lines. Example:

```
printf("Total =
```

re-using 'struct' tag 'xx' as 'union' tag

There are conflicting definitions of the type `struct xx {...}` ; and `union xx {...}` ;. Structure and union tags currently share the same name-space in C.

Example:

```
struct s {int a,b;};
...
union s (int a; double d;};
```

re-using 'union' tag 'xx' as 'struct' tag

As above.

size of struct 'xx' needed but not yet defined

An operation requires knowledge of the size of the struct, but this was not defined. This error is likely to accompany others. Example:

```
{
    struct s;          /* forward declaration */
    struct s *sp;     /* pointer to s */
    sp++;             /* need size for inc operation */
    ...
```

size of union 'xx' needed but not yet defined

See above.

storage class 'xx' incompatible with 'xx' - ignored

An attempt was made to declare a variable with conflicting storage classes.

Example:

```
{
    static auto int i; /* contradiction in terms */
    ...
```

storage class 'xx' not permitted in context xx - ignored

An attempt was made to declare a variable whose storage class conflicted with its position in the program. Examples:

```

register int i;          /* can't have top-level regs */
void fn(a)
static int a;          /* or static parameters */
{
    ...

```

struct 'xx' must be defined for (static) variable declaration

Before you can declare a static structure variable, that structure type must have been defined. This is so the compiler knows how much storage to reserve for it. Examples:

```

static struct s s1;          /* s not defined */
struct t;
static struct t t1;          /* t not defined */

```

struct/union 'xx' not yet defined - cannot be selected from

The structure or union type used as the left operand of a . or → operator has not yet been defined so the field names are not known. Example:

```

{
    struct s s1;          /* forward reference          */
    s1.a = 12;          /* don't know field names yet */
    ...

```

**too few arguments to macro xx(... on line nn
too many arguments to macro xx(... on line nn**

The number of arguments used in the invocation of a macro must match exactly the number used when it was defined. Example:

```

#define rdch(ch,p) while((ch = *p++)==' ');
...
    rdch(ptr); /* need ptr and ch */
...

```

too many initialisers in {} for aggregate

The list of constants in a static array or structure initialiser exceeded the number of elements/fields for the type involved. Example:

```

static int powers[8] = {0,1,2,4,8,16,32,64,128};

```

type 'xx' inconsistent with 'xx'
type disagreement for 'xx'

Conflicting types were encountered in function declaration (prototype) and its definition. Example:

```
void fn(int);
...
int fn(int a)
{
    ...
}
```

A pernicious error of this type is caused by mixing ANSI and old-style function declarations. Example:

```
int f(char x);
int f(x)char x;
{
    ...
}
```

typedef name 'xx' used in expression context

A typedef name was used as a variable name. Example:

```
typedef char flag;
...
{
    int i = flag;
}
```

undefined struct/union 'xx' cannot be member

A struct/union not already defined cannot be a member of another struct/union. In particular this means that a struct/union cannot be a member of itself: use pointers for this. Example:

```
struct s1 {
    struct s2 type; /* s2 not defined yet */
    int count;
};
```

unknown preprocessor directive : #xx

The identifier following a # did not correspond to any of the recognised pre-processor directives. Example:

```
#asm /* not an ANSI directive */
```


uninitialised static [] arrays illegal

Static [] arrays must be initialised to allow the compiler to determine their size.
Example:

```
static char str[];           /* needs {} initialiser */
```

union 'xx' must be defined for (static) variable declaration

Before you can declare a static union variable, that union type must have been defined. Example:

```
static union u u1; /* compiler can't ascertain size */
```

'while' expected after 'do' - found 'xx'

The syntax of the do statement is do statement while (expression).
Example:

```
do      /* should put these statements in {} */
    l = inputLine();
    err = processLine(l); /* finds err, not while */
while (!err);
```

Fatal errors

These are causes for the compiler to give up compilation. Error messages are issued and the compiler stops.

couldn't create object file 'file'

The compiler was unable to open or write to the specified output code file, perhaps because it was locked or the o directory does not exist.

macro args too long

Grossly over-long macro arguments, possibly as a result of some other error.

macro expansion buffer overflow

Grossly over-complicated macros were used, possibly as a result of some other error.

**no store left
out of store (in cc_alloc)**

The compiler has run out of memory – either shorten your source programs or free some RAM using the *UNPLUG command. To do this, first check which modules are present in your machine by typing *MODULES. If there is a module that you do not currently need, you can release its space by typing

```
*UNPLUG modulename
*RMTidy
```

It can later be restored using the *RMREINIT command. For further details, refer to the chapter entitled *Modules* in the *Programmer's Reference Manual*, (second edition).

If running under the desktop, you can use the Task Manager to increase your wimpslot size.

too many errors

More than 100 serious errors were detected.

too many file names

An attempt was made to compile too many files at once. 25 is the maximum that will be accepted.

System errors

There are some additional error messages that can be generated by the compiler if it detects errors in the compiler itself. It is very unusual to encounter this type of error. If you do, note the circumstances under which the error was caused and contact your Acorn supplier.

These error messages all look like this:

```
*****
* The compiler has detected an internal inconsistency. This can occur *
* because it has run out of a vital resource such as memory or disk *
* space or because there is a fault in it. If you cannot easily alter *
* your program to avoid causing this rare failure, please contact your *
* Acorn dealer. The dealer may be able to help you immediately and will *
* be able to report a suspected compiler fault to Acorn Computers. *
*****
```



```
#pragma force_top_level
#pragma include_only_once

/*
 * Interface to host OS.
 * Copyright (C) Acorn Computers Ltd., 1988
 */

#ifndef __kernel_h
#define __kernel_h

#ifndef __size_t
# define __size_t 1
typedef unsigned int size_t; /* from <stddef.h> */
#endif

typedef struct {
    int r[10]; /* only r0 - r9 matter for swi's */
} _kernel_swi_regs;

typedef struct {
    int load, exec; /* load, exec addresses */
    int start, end; /* start address/length, end address/attributes */
} _kernel_osfile_block;

typedef struct {
    void * dataptr; /* memory address of data */
    int nbytes, fileptr;
    int buf_len; /* these fields for Arthur gbbp extensions */
    char * wild_fld; /* points to wildcarded filename to match */
} _kernel_osgbbp_block;

typedef struct {
    int errnum; /* error number */
    char errmsg[252]; /* error message (zero terminated) */
} _kernel_oserror;

typedef struct stack_chunk {
    unsigned long sc_mark; /* == 0xf60690ff */
    struct stack_chunk *sc_next, *sc_prev;
    unsigned long sc_size;
    int (*sc_deallocate)();
} _kernel_stack_chunk;

extern _kernel_stack_chunk *_kernel_current_stack_chunk(void);

extern void _kernel_setreturncode(unsigned code);
```

```

extern void _kernel_exit(int);

extern void _kernel_raise_error(_kernel_oserror *);
/* return the specified error to the parent */

extern void _kernel_exittraphandler(void);

#define _kernel_HOST_UNDEFINED -1
#define _kernel_BBC_MOS1_0 0
#define _kernel_BBC_MOS1_2 1
#define _kernel_BBC_ACW 2
#define _kernel_BBC_MASTER 3
#define _kernel_BBC_MASTER_ET 4
#define _kernel_BBC_MASTER_COMPACT 5
#define _kernel_ARTHUR 6
#define _kernel_SPRINGBOARD 7
#define _kernel_A_UNIX 8

extern int _kernel_hostos(void);
/*
 * Returns the identity of the host OS
 */

extern int _kernel_fpavailable(void);
/*
 * Returns 0 if floating point is not available (no emulator nor hardware)
 */

#define _kernel_NONX 0x80000000
extern _kernel_oserror *_kernel_swi(int no, _kernel_swi_regs *in,
    _kernel_swi_regs *out);
/*
 * Generic SWI interface. Returns NULL if there was no error.
 * The SWI called normally has the X bit set. To call a non-X bit set SWI,
 * _kernel_NONX must be orred into no (in which case, if an error occurs,
 * _kernel_oserror does not return).
 */

extern _kernel_oserror *_kernel_swi_c(int no, _kernel_swi_regs *in,
    _kernel_swi_regs *out, int *carry);
/*
 * As _kernel_swi, but for use with SWIs which return status in the C flag.
 * The int to which carry points is set to reflect the state of the C flag on
 * exit from the SWI.
 */

extern char *_kernel_command_string(void);
/*
 * Returns the address of (maybe a copy of) the string used to run the program
 */

/*
 * The int value returned by the following functions may have value:
 * >= 0 if the call succeeds (significance then depends on the function)
 */

```

```

    * -1 if the call fails but causes no os error (eg escape for rdch). Not
    * all functions are capable of generating this result. This return
    * value corresponds to the C flag being set by the SWI.
    * -2 if the call causes an os error (in which case, _kernel_last_oserror
    * must be used to find which error it was)
    */

#define _kernel_ERROR (-2)

extern int _kernel_osbyte(int op, int x, int y);
/*
 * Performs an OSByte operation.
 * If there is no error, the result contains
 * the return value of R1 (X) in its bottom byte
 * the return value of R2 (Y) in its second byte
 * 1 in the third byte if carry is set on return, otherwise 0
 * 0 in its top byte
 * (Not all of these values will be significant, depending on the
 * particular OSByte operation).
 */

extern int _kernel_osrdch(void);
/*
 * Returns a character read from the currently selected OS input stream
 */

extern int _kernel_oswrch(int ch);
/*
 * Writes a byte to all currently selected OS output streams
 * The return value just indicates success or failure.
 */

extern int _kernel_osbget(unsigned handle);
/*
 * Returns the next byte from the file identified by 'handle'.
 * (-1 => EOF).
 */

extern int _kernel_osbput(int ch, unsigned handle);
/*
 * Writes a byte to the file identified by 'handle'.
 * The return value just indicates success or failure.
 */

extern int _kernel_osgbpb(int op, unsigned handle, _kernel_osgbpb_block
*inout);/*
 * Reads or writes a number of bytes from a filing system.
 * The return value just indicates success or failure.
 * Note that for some operations, the return value of C is significant,
 * and for others it isn't. In all cases, therefore, a return value of -1
 * is possible, but for some operations it should be ignored.
 */

extern int _kernel_osword(int op, int *data);
/*

```

```

* Performs an OSWord operation.
* The size and format of the block *data depends on the particular OSWord
* being used; it may be updated.
*/

extern int _kernel_osfind(int op, char *name);
/*
* Opens or closes a file.
* Open returns a file handle (0 => open failed without error)
* Close the return value just indicates success or failure
*/

extern int _kernel_osfile(int op, const char *name, _kernel_osfile_block
*inout);
/* Performs an OSFile operation, with values of r2 - r5 taken from the osfile
* block. The block is updated with the return values of these registers,
* and the result is the return value of r0 (or an error indication)
*/

extern int _kernel_osargs(int op, unsigned handle, int arg);
/*
* Performs an OSArgs operation.
* The result is an error indication, or
* the current filing system number (if op = handle = 0)
* the value returned in R2 by the OSArgs operation otherwise
*/

extern int _kernel_oscli(const char *s);
/*
* Hands the argument string to the OS command line interpreter to execute
* as a command. This should not be used to invoke other applications:
* _kernel_system exists for that. Even using it to run a replacement
* application is somewhat dubious (abort handlers are left as those of the
* current application).
* The return value just indicates error or no error
*/

extern _kernel_oserror *_kernel_last_oserror(void);
/*
* Returns a pointer to an error block describing the last os error since
* _kernel_last_oserror was last called (or since there program started
* if there has been no such call). If there has been no os error, returns
* a null pointer. Note that occurrence of a further error may overwrite the
* contents of the block.
* If _kernel_swi caused the last os error, the error already returned by that
* call gets returned by this too.
*/

extern _kernel_oserror *_kernel_getenv(const char *name, char *buffer,
unsigned size);
/*
* Reads the value of a system variable, placing the value string in the
* buffer (of size size).
* (This just gives access to OS_ReadVarVal).
*/

```

```

extern _kernel_oserror *_kernel_setenv(const char *name, const char *value);
/*
 * Updates the value of a system variable to be string valued, with the
 * given value (value = NULL deletes the variable)
 */

extern int _kernel_system(const char *string, int chain);
/*
 * Hands the argument string to the OS command line interpreter to execute.
 * If chain is 0, the calling application is copied to the top of memory
 first,
 * then handlers are installed so that if the command string causes an
 * application to be invoked, control returns to _kernel_system, which then
 * copies the calling application back into its proper place - the command
 * is executed as a sub-program. Of course, since the sub-program executes
 * in the same address space, there is no protection against errant writes
 * by it to the code or data of the caller. And if there is insufficient
 * space to load the sub-program, the manner of the subsequent death is
 * unlikely to be pretty.
 * If chain is 1, all handlers are removed before calling the CLI, and if it
 * returns (the command is built-in) _kernel_system Exits.
 * The return value just indicates error or no error
 */

extern unsigned _kernel_alloc(unsigned minwords, void **block);
/*
 * Tries to allocate a block of sensible size >= minwords. Failing that,
 * it allocates the largest possible block (may be size zero).
 * Sensible size means at least 2K words.
 * *block is returned a pointer to the start of the allocated block
 * (NULL if 'a block of size zero' has been allocated).
 */

typedef void freeproc(void *);
typedef void * allocproc(unsigned);

extern void _kernel_register_allocs(allocproc *malloc, freeproc *free);
/*
 * Registers procedures to be used by the kernel when it requires to
 * free or allocate storage. The allocproc may be called during stack
 * extension, so may not check for stack overflow (nor may any procedure
 * called from it), and must guarantee to require no more than 41 words
 * of stack.
 */

typedef int _kernel_ExtendProc(int /*n*/, void** /*p*/);
extern _kernel_ExtendProc *_kernel_register_slotextend(_kernel_ExtendProc
*proc);
/* When the initial heap is full, the kernel is normally capable of extending
 it (if the OS will allow). However, if the heap limit is not the same as
 the OS limit, it is assumed that someone else has acquired the space between,
 and the procedure registered here is called to request n bytes from it.
 Its return value is expected to be >= n or = 0 (failure); if success,

```



```

    *p is set to point to the space returned.
    */

extern int _kernel_escape_seen(void);
/*
 * Returns 1 if there has been an escape since the previous call of
 * _kernel_escape_seen (or since program start, if there has been no
 * previous call). Escapes are never ignored with this mechanism,
 * whereas they may be with the language EventProc mechanism since there
 * may be no stack to call it on.
 */

typedef union {
    struct {int s:1, u:16, x: 15; unsigned mhi, mlo; } i;
    int w[3]; } _extended_fp_number;

typedef struct {
    int r4, r5, r6, r7, r8, r9;
    int fp, sp, pc, sl;
    _extended_fp_number f4, f5, f6, f7; } _kernel_unwindblock;

extern int _kernel_unwind(_kernel_unwindblock *inout, char **language);
/*
 * Unwinds the call stack one level.
 * Returns >0 if it succeeds
 * 0 if it fails because it has reached the stack end
 * <0 if it fails for any other reason (eg stack corrupt)
 * Input values for fp, sl and pc must be correct.
 * r4-r9 and f4-f7 are updated if the frame addressed by the input value
 * of fp contains saved values for the corresponding registers.
 * fp, sp, sl and pc are always updated
 * *language is returned a pointer to a string naming the language
 * corresponding to the returned value of pc.
 */

extern char *_kernel_procname(int pc);
/*
 * Returns a string naming the procedure containing the address pc.
 * (or 0 if no name for it can be found).
 */

extern char *_kernel_language(int pc);
/*
 * Returns a string naming the language in whose code the address pc lies.
 * (or 0 if it is in no known language).
 */

/* divide and remainder functions.
 * The signed functions round towards zero.
 *
 * The div functions actually also return the remainder in a2, and use of
 * this by a code-generator will be more efficient than a call to the rem
 * function.
 *
 * Language RTSs are encouraged to use these functions rather than providing

```

```

    * their own, since considerable effort has been expended to make these fast.
    */

extern unsigned _kernel_udiv(unsigned divisor, unsigned dividend);
extern unsigned _kernel_urem(unsigned divisor, unsigned dividend);
extern unsigned _kernel_udiv10(unsigned dividend);

extern int _kernel_sdiv(int divisor, int dividend);
extern int _kernel_srem(int divisor, int dividend);
extern int _kernel_sdiv10(int dividend);

/*
 * Description of a 'Language description block'
 */

typedef enum { NotHandled, Handled } _kernel_HandledOrNot;

typedef struct {
    int regs [16];
} _kernel_registerset;

typedef struct {
    int regs [10];
} _kernel_eventregisters;

typedef void (*PROC) (void);
typedef _kernel_HandledOrNot (*_kernel_trapproc) (int code,
    _kernel_registerset *regs);
typedef _kernel_HandledOrNot (*_kernel_eventproc) (int code,
    _kernel_registerset *regs);

typedef struct {
    int size;
    int codestart, codeend;
    char *name;
    PROC (*InitProc)(void);
    PROC FinaliseProc;
    _kernel_trapproc TrapProc;
    _kernel_trapproc UncaughtTrapProc;
    _kernel_eventproc EventProc;
    _kernel_eventproc UnhandledEventProc;
    void (*FastEventProc) (_kernel_eventregisters *);
    int (*UnwindProc) (_kernel_unwindblock *inout, char **language);
    char * (*NameProc) (int pc);
} _kernel_languagedescription;

typedef int _kernel_ccproc(int, int, int);

extern int _kernel_call_client(int a1, int a2, int a3, _kernel_ccproc callee);
/* This is for shared library use only, and is not exported to shared library
 * clients. It is provided to allow library functions which call arbitrary
 * client code (C library signal, exit, _main) to do so correctly if the
 * client uses the old calling standard.
 */

```

```
extern int _kernel_client_is_module(void);
/* For shared library use only, not exported to clients. Returns a
 * non-zero value if the library's client is a module executing in user
 * mode (ie module run code).
 */

extern int _kernel_processor_mode(void);

extern void _kernel_irqs_on(void);

extern void _kernel_irqs_off(void);

extern int _kernel_irqs_disabled(void);
/* returns 0 if interrupts are enabled; some non-zero value if disabled. */

extern void *_kernel_RMAalloc(size_t size);

extern void *_kernel_RMAextend(void *p, size_t size);

extern void _kernel_RMAfree(void *p);

#endif
/* end of kernel.h */
```

21 Appendix D: The floating point emulator

The floating point emulator is a relocatable module which provides support for floating point instructions. It must be resident in memory to run programs which perform operations on real numbers.

Its primary function is to emulate floating point instructions in software on machines which do not have the optional hardware floating point co-processor attached.

However, even with the co-processor attached, the floating point emulator is still required:

- to interface with the co-processor;
- to perform range reduction on trigonometric function arguments;
- for a few floating point instructions that the co-processor does not directly support.

There are two current variants of the floating point emulator:

FPE280	software-only floating point support – v 2.80 and earlier
FPEmulator	hardware-assisted and software-only support – v 3.10 and later.

Both have the same module name, namely `FPEmulator`. You can find out the version number of the module currently resident in your computer by typing the following at the `*` prompt:

```
*help modules
```

FPE280

On initialisation, this module disables the floating point co-processor if it finds one present. It occupies 25Kb. This is the version supplied with Acorn Desktop C as the file `!System.Modules.FPEmulator`.

FPEmulator

This behaves just like FPE280 if no co-processor is attached (ie it emulates all floating point instructions in software), but it makes use of the co-processor if it is present. It occupies 37Kb.

Using the compiler

Without the floating point maths co-processor

If your machine does not have the floating point co-processor attached, the floating point emulator is required to run any C program which performs operations on real numbers.

The floating point emulator supplied with Acorn Desktop C is FPE280, and is the file `FPE280` in the `$.!System` directory.

Before loading the emulator, it is a good idea to issue a command that will check that no more recent version of the module is already present, by typing

```
*RMEnsure FPEmulator 2.80
```

Then load the emulator:

```
*RMLoad $.Modules.fpe280
```

When setting up your working environment, it is recommended that you place the above module (as `!System.modules.FPEmulator`) in your `!System` directory and arrange for it to be loaded automatically on power up.

Observe the change of file name to `FPE280`, since existing applications will incorporate the earlier name in their start-up sequence.

With the floating point maths co-processor

In order to make use of the speed increase given by the floating point co-processor, you will need to use the FPEmulator module.

This is supplied with the co-processor, and you will find it convenient to copy the module into your `!System` directory and arrange for it to be loaded automatically on power up.

Floating point requirements of applications

Applications should cater for both floating point environments: with and without the co-processor. In general, programs do not need to know whether a co-processor is fitted; the only effective difference is in the speed of execution. However, the combined hardware and software variant, FPEmulator, is larger than

the software-only variant, FPE280, since it includes the code for interfacing with the co-processor. Therefore, to cater for both environments, an application must be able to accommodate the extra 12Kb RAM taken up by FPEmulator.

Software products do not have to supply either version of the floating point emulator. FPE280 is supplied with new machines and FPEmulator is supplied with the co-processor itself. It is then up to you to have the appropriate emulator in your !System directory; this should be covered in the manual accompanying the application.

Subject Index

Symbols

- #include directives 29
- #include files
 - including directives 39
 - searching for in CC 33
- *Wimplot 360
- :mem
 - use in reinstating in-memory filing system in search path 32

A

- absolute machine addresses 96
- active count 152
- akbd 175
- alarm 176 - 178
 - in desktop applications 173
- alarm.h 173
- ANSI standard
 - vs K&R C 94 - 97
- ANSI standard 6 6
- application
 - access of workspace 365
- application resources 151
- applications, desktop 149 - 174
 - error reporting 158 - 159
 - general form of 152
 - initialising 152 - 153
 - loading 164 - 166
 - saving 166 - 167
 - see also* alarm
 - see also* memory management
 - standards for 149
 - terminate and stay resident 337
 - tracing 173 - 174

- arguments
 - passing to assembler 333
- arithmetic conversions in ANSI standard 96
- arithmetic functions 368
- arithmetic operations 77 - 78
- ARM Procedure Call Standard 367, 372
- arrays 84, 97
 - lifetime 366
- assembly language interface 331 - 336
- assembly list files 29
- assembly listing of code from CC 43
- assert.h 105
- atexit() 365

B

- baricon 178
- BASIC
 - routine to search for lost memory blocks 357
- bbc 180 - 185
- bins (linked lists) 358
- bitfields 85
- buffering of input/output 116
- byte
 - limits 107
 - ordering *see* portability, byte ordering

C

- C compiler 2
- C libraries 24 - 25
- C library kernel 367 - 374
 - interfacing to 369 - 372
- C Module Header Generator (CMHG) *see* CMHG

- C storage manager 357
- C\$Libroot, use 32
- casts
 - reporting to CC 40
- CC 9, 21 - 58
 - adding #include lines to listing file 39
 - adding to the command line 48
 - allowing writeable string literals 40
 - altering the command line 37
 - application menu 49
 - assembly list files 29
 - auto running and saving 49
 - Browser window 10
 - C portability 27
 - checking data flow anomalies 39
 - checking uniqueness of external names 39
 - code generation flag options 52
 - command line 36, 51 - 53
 - command line extension 48
 - compilation list files 29
 - compiler checks 39
 - controlling execution 51
 - counting function executions (profiling) 36
 - default output location 26
 - entry vector to program 25
 - error browser 35
 - errors and warnings 51
 - file naming and placing conventions 25
 - filename conventions 27
 - filename validity 29
 - getting started 21 - 24
 - keyword options 52
 - Limited pcc option
 - see* Feature menu option
 - link step options 49
 - linking code to relocatable modules 36
 - linking object files 34
 - listing unexpanded source text 40
 - loading 9
 - location of CLib 25
 - managing with Make 21
 - menu options
 - Assembler 29, 43
 - Command line 36
 - Debug options 42
 - Default path 28, 32, 40
 - Feature 37, 38
 - Keep comments 34, 38
 - Libraries 43
 - Listing 29, 37
 - Module code 36
 - Other 37, 48
 - Predefine 44
 - Profile 36
 - Suppress 46
 - Undefine 45
 - UNIX pcc 38
 - Work directory 26, 47
 - naming include files 28
 - object files 28
 - output 50
 - Output window and Summary box -
 - toggling 50
 - preprocessor flag options 52
 - preprocessor options 48
 - producing executable image files 21
 - program files 29
 - removing function names from code area 39
 - reporting unused declarations 40
 - rooted filenames 27
 - saving options 49
 - search path 31
 - searching for include files 29, 30
 - searching for nested #includes 39
 - searching for nested includes 31
 - setting your own system search path 32
 - SetUp dialogue box 33
 - SetUp menu 35
 - SetUp options
 - Compile only 28, 34
 - Debug 34
 - Debug option, limiting 42
 - Include 28, 30, 32, 33
 - Preprocess only 34, 38
 - Source 33
 - Throwback 10, 35

- specifying libraries 43
- specifying search path (default) 40
- specifying source files 33
- starting (unmanaged by Make) 21
- stubs 25
- suppressing messages 46
- switching between ANSI C and 'portable C compiler' C 38
- switching on reporting 39
- turning on warnings 39
- turning on warnings about short integers and enums 40
- using pcc-style include files in ANSI mode *see* Feature menu option
- using with Link 21
- warning feature flag options 53
- work directory 26
- worked examples 53 - 58
- chaining memory blocks 358
- char, limits 107
- characters 83
 - testing and mapping 105
- check words
 - using RMA 356
- CLib module 25
- CMHG 9, 340 - 348
 - application menu 60
 - command descriptions 344 - 345
 - command line 61
 - finalisation code 341
 - help string 343
 - initialisation code 341
 - input to 340
 - IRQ handler 347 - 349
 - loading 9
 - menu options 60
 - runnable code 341
 - service call handler 342
 - SetUp options 60
 - SWI chunks 345 - 346
 - SWI decoding code 347
 - SWI decoding table 346
 - SWI handler code 346

- title string 343
- codeend 370
- codestart 370
- colour translation 169, 186 - 191
- colourmenu 185 - 186
- colours, desktop 170, 287
- colourtran 169, 186 - 191
- compilation list files 29
- compiling and linking applications without Link 43
- const 95
- control statements 97
- coordinates, work area vs screen 169
- coords 191 - 194
- ctype.h 100, 105 - 106

D

- data
 - elements 73 - 76
 - export 253 - 254, 325 - 328
 - import 323 - 325
- data types
 - see also* portability, data types
 - structured 76
- dbox 195 - 205
- DDE
 - C tools 9
- DDT
 - debugging C programs 13
 - example session 13 - 18
 - using CC Debug option to enable source-level debugging 34
- DecCF
 - use with libraries 19
- declarators 85
- default search paths
 - SWIs, kernel, pragmas and varargs 30
- desktop applications
 - examples 56
- desktop C
 - new features 383

desktop compiler *see* CC
device drivers 337
diagnostic messages 81
diagnostics 105
dialogue box 170 - 172
 creation 158, 195
 deletion 195
 events 199 - 205
 fields 197 - 199
 functions 195 - 205
 writable fields in 171
dot *see* pointer into text
Draw files 167 - 168, 205 - 210
 coordinates in 168
 data types 205
 object level interface 212 - 215
drawdiag 167, 205 - 210
drawferror 210 - 211
drawfobj 167, 212 - 215
drawftypes 216
drawmod 216 - 219

E

Edit 173
Editor throwback
 SrcEdit
 throwback 10
EDOM 106
empty comments in ANSI standard 95
empty declarations in ANSI standard 97
entry vector in C program 25
enumerations 85
ERANGE 106
errno.h 100, 106
error
 domain 106
 operating system 368
 range 106
error handling 371, 379
ESIGNUM 106
evaluation

 expression 78
 function argument 93
event 220 - 222
event handling 149 - 151, 155 - 158, 170 - 171,
 199 - 205, 371
event types 293
 masks 293
EventProc 371
executable image files 21

F

FastEventProc 372
FILE 113
file
 buffering 88
 creation 114
 deletion 114
 opening 115 - 116
 position indicators 124 - 126
 renaming 114
file buffers
 allocation 360
fileicon 222
filename
 validity 88
filename components, length of 102
filename conventions 27
filename extension 102
filename generation 114
files
 zero-length 88
FinaliseProc 370
flex 160 - 162, 223 - 225, 357
 advantages 357
 description 360
 limitations 361
 shifting heaps 361
float 96
float.h 107
floating point
 co-processor 434

- emulator 433 - 435
- instruction set 368
- types 84
- floating types 78
- font 225 - 232
- font management 225 - 232
- fontlist 232
- fonts
 - counting into tree structure 232
- fontselect 233 - 235
 - interface for choosing fonts 233
- FormEd 151
- fpos_t 113
- fragmentation 359
 - of malloc heap 360
- FrontEnd 18
 - producing new RISC OS applications 18
- function
 - call, bypassing 109 - 110
 - declarations 96
 - definitions 96
 - prototypes 96
 - workspace 335

G

- graphical data *see* Draw files
- graphics output functions
 - BBC-style 181 - 184
- guard constant, in memory blocks 358

H

- headers
 - searching for with CC 33
- headers for RISC_OSLib 30
- heap 235
 - coalescing 359
- heap allocation 235, 379
- help 235 - 238
- hourglass 166, 167, 285

- HUGE_VAL 100

I

- I/O
 - functions 113 - 127
 - redirection 82
- icon 151
 - button state bits 289
 - button types 288
 - creating 292
 - data fields 290
 - flags 287
 - placing on icon bar 154, 178
 - placing on left of icon bar 178
- icon bar 151
- identifiers 73, 82, 86
 - limits 74 - 76
- implementation limits 78
- include files
 - naming in C program source 28
 - searching for 29
- initialisation 400
- InitProc 370
- input 121, 122, 123
 - functions 119 - 120
- int, limits 107
- integers 84
- interactive device 82
- IRQ state
 - manipulating 368
 - see also* cmhg (IRQ handler)

K

- kernel 369
 - default search path in CC 30
- kernel.h 102
- keyboard polling 175

L

- language libraries
 - recovering memory 355
- LibFile
 - use with linkable libraries 19
- libraries
 - ANSI vs BSD UNIX 99 - 102
- library functions 86 - 90
- lifetimes
 - static variables and arrays 366
- limits.h 107
- Link
 - using with CC 21
- linkable libraries 19
- listing file
 - CC 37
- locale.h 107 - 108
- long 95
- long double 95, 96
- long float 95
- long int, limits 107

M

- macros 44
- magnifier 172, 239
- magnify 238 - 239
- Make
 - managing CC 21
 - use 9
 - using C tools 10
- Makefiles 9, 21
 - command lines and desktop interface 10
- malloc 357
 - deallocation of blocks 359
 - use when designing programs 356
- malloc heap 357
- math.h 100, 108 - 109
- mathematical functions 87, 108 - 109, 133 - 134
- memory
 - alignment 358

- allocation in C 357
- allocation of block sizes 358
- allocation of file buffers 360
- allocation with flex and malloc 357
- avoiding permanent loss 356
- avoiding references to deallocated blocks 356
- avoiding wastage 357
- BASIC routine to search for lost blocks 357
- coalescing blocks 359
- efficient use 355
- fragmentation 359
- malloc allocation 358
- splitting blocks 358
- memory allocation functions 130
- memory management 160 - 162, 223 - 225, 355, 368
- memory models in MS-DOS 103
- menu 239 - 242
 - amending 240 - 242
 - creating 154, 240
 - responding to user choice 157 - 158
- menu syntax 239
- message action codes 295 - 300
- messages 242
- modules, relocatable *see* relocatable modules
- msgs 242
- multibyte character
 - functions 134 - 135
 - limits 107
- multibyte string functions 135 - 136

N

- NameProc 372
- nested include files
 - searching 31

O

- object files

- linking 34
 - naming conventions 28
- object files, producing with CC 21
- offsetof 113
- operating system interface 102 - 103, 132
- operating system library 24
- os 244 - 245
- output 121, 123, 124
 - functions 117 - 119, 120
- output from CC 50
- overlays 351 - 353
 - alternatives to 353

P

- paging 351
- pathname
 - separator in 102
- pcc mode 98 - 99
 - preprocessor in 99
 - type checking in 99
- pipng 103
- pointer 84, 93, 96, 160 - 161, 245 - 246
 - into text 172
 - subtraction 77
 - types 77
- pointer, desktop 245 - 246
 - see also* hourglass
- portability 91 - 103, 353
 - byte ordering 92
 - data types 91
 - hexadecimal constants 92
 - operating system calls 94
 - pointers 93
 - store alignment 93
- portability of C programs 27
- pragmas
 - default search path in CC 30
- predefining 44
- preprocessor
 - controlling with CC 34
 - options in CC 48

- output 34
 - predefining macros 44
 - undefining macros 45
- preprocessor directives 86, 97
- print 246 - 252
- procedures
 - storing in linkable libraries 19
- profiling 36
- program design
 - for efficient use of memory 355
- program files 29
- program termination functions 131
- programs, calling from C 374 - 376
- ptrdiff_t 113

Q

- qualifiers 85

R

- RAM filing system 353
- random number generating functions 129
- register 84
 - names 332
 - usage 332 - 333
- relocatable modules 337 - 349
 - components of 338 - 339
 - constraints on 338
 - linking code compiled by CC 36
- res 252 - 253
- resspr 253
- RISCOS
 - command types 374
 - library 149
 - library, initialising 152
- RMA
 - clearing blocks 365
 - deallocation 356
 - using for storage though SWI calls 356

S

- saveas 253 - 254
- saving multiple object files in C 34
- screen units *see* Draw files, coordinates in 168
- search
 - for allocated memory blocks 356
 - functions 131, 132
- searching for #include files 33
- setjmp.h 109 - 110
- setvbuf 360
- shared C library 24, 376 - 377
 - modules 337
 - when to use 376 - 377
- short int, limits 107
- signal handling 110 - 111
- signal.h 100, 110 - 111
- signed 95
- signed char, limits 107
- size_t 113
- Software Interrupt *see* SWI
- sort functions 133
- source files
 - naming conventions, *see* filename conventions
 - paths and naming conventions 27
 - specifying to CC 33
- source-level debugging
 - enabling 34
- spooling output 173
- sprite 254 - 261
- sprite facilities 254 - 261
- stack
 - allocation 360
 - extension 358, 362
- stack extension 362, 373 - 374
- stack, run-time 372 - 374
- stack-limit checking 369
- static variables
 - lifetime 366
- stdarg.h 111 - 112
- stddef.h 113
- stdio.h 101, 111, 113 - 127
- stdlib.h 101, 127 - 136
- storage management 379
- storage manager
 - description 358
- stream
 - closing 115
 - flushing 115
- string functions
 - appending 137
 - comparison 137 - 138
 - conversion 127 - 129
 - copying 136 - 137
 - error message mapping 141
 - length 139, 140, 141
 - locating 139 - 140
 - time 144 - 145
 - tokenising 140
 - transformation 138 - 139
- string literal 97
- string.h 101, 136 - 141
- strings in ANSI standard 95
- struct 95
 - alignment of members 93
- structure result 334
- structures 85
 - see also* struct
- stubs 25
- SVC mode 366
- SWI 102, 337, 368
 - XOS_Heap 357
 - XOS_Module 357
- SWIs
 - default search path in CC 30
- switch statement 97

T

- template 261 - 263
 - editor *see* FormEd
- text
 - displaying and editing 172 - 173
 - highlighting 172

text functions 264 - 279
 text output functions
 BBC-style 180
 text streams 88
 Throwback
 example session 10 - 12
 throwback
 controlling in CC 35
 time zones 88
 time.h 141 - 145
 ToANSI 63, 98
 application menu 65
 command line 66
 loading 9
 menu options 65
 SetUp options 64
 use in DDE 9
 token-pasting 97
 ToPCC 67, 98
 application menu 69
 command line 70
 loading 9
 menu options 69
 SetUp options 68
 use in DDE 9
 trace 263
 trace.h 173
 translation
 ordering of phases 97
 trap handling 371
 TrapProc 371
 tree structures
 fonts 232
 txt 264 - 276
 txt.h 172
 txtedit 277 - 279
 txtedit.h 172
 txtscrap 283
 txtwin 283 - 284
 txtwin.h 172

U

UncaughtTrapProc 371
 UnhandledEventProc 371
 union 95
 unions 85
 uniqueness of names
 checking 39
 unsigned 95
 unsigned char, limits 107
 unsigned int, limits 107
 unsigned long int 96
 limits 107
 unsigned short int, limits 107
 UnwindProc 372
 user files
 naming conventions 30
 USR mode 366

V

values, limits on 107
 varargs
 default search path in CC 30
 variable
 environmental 368
 storage of 335
 variadic functions in ANSI standard 95
 visdelay 284 - 285
 void 95
 void * 95
 volatile 94

W

wchar_t 113
 werr 285
 wimp 286 - 312
 Wimp polling 149
 wimp slot
 contents 360
 wimpt 312 - 316

- win 316 - 322
- window
 - creating 153 - 154
 - maintaining 155 - 156
 - opening 155
- window coordinates 191 - 194
- window drag types 289
- window flags 286 - 287
- window manager 316 - 322
 - function prototypes 304 - 312
 - idle events 163
 - unknown events 163 - 164
- work directory 26
 - specifying in CC 47

X

- xfer_send 167, 254
- xferrecv 164, 323 - 325
- xfersend 166 - 167, 325 - 328

Function Index

The main entry for each function is printed in bold type.

Symbols

`_kernel_exittraphandler` 371
`_kernel_register_allocs` 372
`_kernel_setreturncode` 370
`_kernel_swi` 102

A

`abort` 89, **130**
`abs` **133**
`acos` **108**
`akbd_pollctl` **175**
`akbd_pollkey` **175**
`akbd_pollsh` **175**
`alarm_anypending` **177**
`alarm_callnext` **178**
`alarm_init` **176**
`alarm_next` **177**
`alarm_remove` **177**
`alarm_removeall` **177**
`alarm_set` **176**
`alarm_timedifference` **176**
`alarm_timenow` **176**
`asctime` **143**
`asin` **108**
`atan` **108**
`atan2` **108**
`atexit` **131**, 341
`atof` **127**
`atoi` **127**
`atol` **127**

B

`baricon` 152, **178**
`baricon_left` **178**
`baricon_newsprite` **179**
`baricon_textandsprite` **179**
`baricon_textandsprite_left` **180**
`bbc_adval` **184**
`bbc_circle` **182**
`bbc_circlefill` **182**
`bbc_clg` **183**
`bbc_cls` **181**
`bbc_colour` **181**
`bbc_cursor` **184**
`bbc_draw` **182**
`bbc_drawby` **182**
`bbc_fill` **183**
`bbc_gcol` **183**
`bbc_get` **184**
`bbc_getbeat` **184**
`bbc_getbeats` **184**
`bbc_gettempo` **184**
`bbc_gwindow` **183**
`bbc_inkey` **185**
`bbc_mode` **182**
`bbc_modevar` **184**
`bbc_move` **182**
`bbc_moveby` **182**
`bbc_origin` **183**
`bbc_palette` **183**
`bbc_plot` **181**
`bbc_point` **183**
`bbc_pos` **181**
`bbc_rectangle` **182**
`bbc_rectanglefill` **182**
`bbc_setbeats` **185**
`bbc_settempo` **185**

bbc_sound **185**
 bbc_soundoff **185**
 bbc_soundon **185**
 bbc_stereo **185**
 bbc_stringprint **181**
 bbc_tab **181**
 bbc_tint **183**
 bbc_vdu **180**
 bbc_vduq **181**
 bbc_vduvar **183**
 bbc_vduvars **184**
 bbc_vduw **180**
 bbc_voices **185**
 bbc_vpos **181**
 bsearch **132**

C

calloc 89, **130**
 ceil **109**
 clearerr **126**
 clock 90, **142**
 colourmenu_make 170, **186**
 colourtran_colournumbertoGCOL **190**
 colourtran_GCOL_tocolournumber **190**
 colourtran_invalidate_cache **191**
 colourtran_return_colourformode **188**
 colourtran_return_colournumber **187**
 colourtran_return_GCOLformode **188**
 colourtran_return_Oppcolourformode **189**
 colourtran_return_Oppcolournumber **189**
 colourtran_return_OppGCOL **188**
 colourtran_return_OppGCOLformode **189**
 colourtran_returnfontcolours **190**
 colourtran_returnGCOL **187**
 colourtran_select_GCOLtable **186**
 colourtran_select_table **186**
 colourtran_setfontcolours **191**
 colourtran_setGCOL **187**
 colourtran_setOppGCOL **189**
 coords_box_toscreen **192**
 coords_box_toworkarea **192**

coords_boxesoverlap **194**
 coords_intersects **194**
 coords_offsetbox **194**
 coords_point_toscreen **193**
 coords_point_toworkarea **193**
 coords_withinbox **193**
 coords_x_toscreen/coords_y_toscreen **192**
 coords_x_toworkarea/coords_y_toworkarea **192**
 cos **108**
 cosh **108**
 ctime **143**

D

dbox_dispose 158, **195**
 dbox_eventhandler 170, **200**
 dbox_fadefield 170, **199**
 dbox_field/dbox_fieldtype **197**
 dbox_fillin 158, 170, **201**
 dbox_get **199**
 dbox_getfield 170, **198**
 dbox_getnumeric 170, **198**
 dbox_hide **196**
 dbox_init 172, **203**
 dbox_new 158, **195**
 dbox_persist 170, **202**
 dbox_popup **202**
 dbox_raweventhandler **200**
 dbox_read **200**
 dbox_setfield 158, 170, **197**
 dbox_setnumeric 170, **198**
 dbox_show 158, 170, **196**
 dbox_showstatic 170, **196**
 dbox_syshandle **202**
 dbox_unfadefield 170, **199**
 dboxfile 172, **203**
 dboxquery 172, **204**
 dboxtcol **204**
 difftime **142**
 div **133**
 draw_append_diag **206**
 draw_convertBox **209**

draw_create_diag **212**
 draw_createObject **213**
 draw_deleteObjects **214**
 draw_doObjects **212**
 draw_extractObject **215**
 draw_querybox **208**
 draw_rebind_diag **209**
 draw_registerMemoryFunctions **207**
 draw_render_diag **207**
 draw_set_unknown_object_handler **210**
 draw_setFontTable **213**
 draw_shift_diag 168, **208**
 draw_translateText **215**
 draw_verify_diag **206**
 draw_verifyObject **213**
 drawex_load_ram 168
 drawex_ram_loader 168
 drawdiag_init **210**
 drawmod_ask_flattenpath **218**
 drawmod_ask_strokepath **217**
 drawmod_buf_transformpath **218**
 drawmod_do_flattenpath **217**
 drawmod_do_strokepath **217**
 drawmod_fill **216**
 drawmod_insitu_transformpath **218**
 drawmod_processpath **219**
 drawmod_stroke **216**

E

event_anywindows **220**
 event_attachmenu 154, 155, 157, **220**
 event_attachmenumaker 154, **221**
 event_clear_current_menu **221**
 event_getmask **222**
 event_is_menu_being_recreated **221**
 event_process 155, **220**
 event_setmask 156, **222**
 exit 89, **131**, 375
 exp **109**

F

fabs **109**
 fclose **115**
 feof **126**
 ferror **126**
 fflush **115**
 fgetc **121**
 fgetpos 89, **124**
 fgets **121**
 flex_alloc **223**
 flex_budge **224**
 flex_dont_budge **225**
 flex_extend **224**
 flex_free **223**
 flex_init **225**
 flex_midextend **224**
 flex_size **223**
 float.h 101
 floor **109**
 fmod 87, **109**
 font_cacheaddress **225**
 font_caret **228**
 font_charbbox **230**
 font_converttoos **228**
 font_converttopoints **228**
 font_current **229**
 font_find **226**
 font_findcaret **229**
 font_findcaretj **232**
 font_future **229**
 font_list **230**
 font_lose **226**
 font_paint **227**
 font_readdef **226**
 font_readinfo **227**
 font_readscalefactor **230**
 font_readthresholds **231**
 font_setcolour **231**
 font_setfont **229**
 font_setpalette **231**
 font_setscalefactor **230**
 font_setthresholds **231**

font_stringbbox **232**
font_strwidth **227**
fontlist_free_font_tree **233**
fontlist_list_all_fonts **233**
fontselect_attach_menu **234**
fontselect_closedown **233**
fontselect_closewindows **234**
fontselect_init **233**
fontselect_selector **234**
fopen **115, 116**
fprintf **89, 117, 118**
fputc **121**
fputs **122**
fread **124**
free **130, 160**
freopen **116**
frexp **109**
fscanf **89, 119, 120**
fseek **125**
fsetpos **125**
ftell **89, 125**
fwrite **124**

G

getc **122**
getchar **122**
getenv **89, 131**
gets **122**
gmtime **143**

H

heap_alloc **235**
heap_free **235**
heap_init **235**
help_dboxrawevents **237**
help_genmessage **236**
help_process **236**
help_register_handler **236**
help_reply **238**

help_simplehandler **237**

I

isalnum **87, 105**
isalpha **87, 105**
iscntrl **87, 105**
isdigit **105**
isgraph **105**
islower **105**
islowert **87**
isprint **87, 105**
ispunct **87**
isspace **106**
isupper **87**
isxdigit **106**

L

labs **133**
lconv **108**
ldexp **109**
ldiv **133**
limits.h **101**
locale.h **102**
localtime **144**
log **109**
log10 **109**
longjmp **109, 110, 373**

M

magnify_select **172, 238, 239**
main **81, 97, 102, 375**
malloc **89, 130, 160**
mblen **134**
mbstowcs **135**
mbtowc **134**
memchr **139**
memcmp **137**
memcpy **136**

memmove **136**
 memset **141**
 menu_dispose **240**
 menu_extend **240**
 menu_make_sprite **242**
 menu_make_writeable **241**
 menu_new **154, 240**
 menu_setflags **241**
 menu_submenu **241**
 menu_syshandle **242**
 mktime **142**
 modf **109**
 msgs_init **242**
 msgs_lookup **243**
 msgs_readfile **243**

O

os_args **245**
 os_byte **244**
 os_cli **245**
 os_file **166, 245**
 os_find **245**
 os_gbbp **245**
 os_read_var_val **245**
 os_swi **244**
 os_swix **244**
 os_word **245**

P

perror **89, 126**
 pointer_reset_shape **246**
 pointer_set_shape **245**
 pow **109**
 print_abortjob **250**
 print_box **247**
 print_canceljob **250**
 print_checkfeatures **248**
 print_currentjob **249**
 print_drawpage **251**

print_endjob **250**
 print_features **246**
 print_getrectangle **251**
 print_giverectangle **251**
 print_identity **246**
 print_info **248**
 print_infostr **247**
 print_pagesize **249**
 print_pagesizestr **247**
 print_positionstr **248**
 print_reset **250**
 print_screendump **252**
 print_selectillustration **250**
 print_selectjob **249**
 print_setinfo **248**
 print_setpagesize **249**
 print_transmatstr **248**
 printf **118, 174**
 putc **123**
 putchar **123**
 puts **123**

Q

qsort **132**

R

rand **129**
 realloc **89, 130**
 remove **88, 114**
 rename **88, 114**
 res_findname **252**
 res_init **153, 252**
 res_openfile **252**
 resspr_area **155, 253**
 resspr_init **253**
 rewind **126**

S

- saveas_read_leafname_during_send **254**
- scanf **120**
- setbuf **116**
- setjmp **373**
- setlocale **87**
- setvbuf **117**
- signal **87**
- sin **108**
- sinh **109**
- sprintf **119**
- sprite_area_initialise **255**
- sprite_area_load **255**
- sprite_area_merge **255**
- sprite_area_readinfo **255**
- sprite_area_reinit **255**
- sprite_area_save **256**
- sprite_copy **257**
- sprite_create **257**
- sprite_create_mask **259**
- sprite_create_rp **257**
- sprite_delete **257**
- sprite_delete_column **259**
- sprite_delete_row **259**
- sprite_flip_x **259**
- sprite_flip_y **259**
- sprite_get **256**
- sprite_get_given **256**
- sprite_get_given_rp **256**
- sprite_get_rp **256**
- sprite_getname **256**
- sprite_insert_column **259**
- sprite_insert_row **259**
- sprite_outputtomask **260**
- sprite_outputtoscreen **261**
- sprite_outputtosprite **260**
- sprite_put **258**
- sprite_put_char_scaled **258**
- sprite_put_given **258**
- sprite_put_greyscaled **258**
- sprite_put_mask **258**
- sprite_put_mask_given **258**
- sprite_put_mask_scaled **258**
- sprite_put_scaled **258**
- sprite_readmask **260**
- sprite_readpixel **260**
- sprite_readsize **260**
- sprite_remove_mask **259**
- sprite_removewastage **261**
- sprite_rename **257**
- sprite_restorestate **260**
- sprite_screenload **255**
- sprite_screensave **255**
- sprite_select **257**
- sprite_select_rp **257**
- sprite_sizeof_screencontext **261**
- sprite_sizeof_spritecontext **261**
- sprite_writemask **260**
- sprite_writepixel **260**
- sqrt **109**
- rand **129**
- sscanf **120**
- strcat **137**
- strchr **139**
- strcmp **138**
- strcoll **138**
- strcpy **136**
- strcspn **139**
- strerror **90, 141**
- strftime **144, 145**
- strlen **141**
- strncat **137**
- strncmp **138**
- strncpy **137**
- strpbrk **139**
- strrchr **139**
- strspn **140**
- strstr **140**
- strtod **127**
- strtok **140**
- strtol **128**
- strtoul **128, 129**
- strxfrm **138**
- switch statement **85**
- system **89, 132**

T

- tan **108**
- tanh **109**
- template_copy **261**
- template_find **262**
- template_init **172, 253, 262**
- template_loaded **262**
- template_readfile **261**
- template_syshandle **263**
- template_use_fancyfonts **262**
- time **143**
- tmpfile **114**
- tmpnam **114**
- tolower **106**
- toupper **106**
- trace_is_on **263**
- trace_off **174, 263**
- trace_on **174, 263**
- tracef **263**
- txt_arrayseg **276**
- txt_bufsize **265**
- txt_charat **269**
- txt_charatdot **269**
- txt_charoptions **266**
- txt_charsatdot **269**
- txt_delete **268**
- txt_dispose **265**
- txt_disposemarker **272**
- txt_dot **267**
- txt_eventhandler **275**
- txt_get **274**
- txt_hide **264**
- txt_indexofmarker **272**
- txt_init **276**
- txt_insertchar **268**
- txt_insertstring **268**
- txt_lastref **266**
- txt_movedot **268**
- txt_movedottomarker **272**
- txt_movehorizontal **270**
- txt_movemarker **271**
- txt_movevertical **270**
- txt_new **264**
- txt_newmarker **271**
- txt_queue **275**
- txt_readeventhandler **275**
- txt_replaceatend **270**
- txt_replacechars **269**
- txt_selectend **273**
- txt_selectset **272**
- txt_selectstart **273**
- txt_setbufsize **265**
- txt_setcharoptions **266**
- txt_setdisplayok **267**
- txt_setdot **267**
- txt_setlastref **266**
- txt_setselect **273**
- txt_show **264**
- txt_size **267**
- txt_syshandle **276**
- txt_unget **275**
- txt_visiblecolcount **271**
- txt_visiblelinecount **270**
- txtedit_dispose **277**
- txtedit_doimport **278**
- txtedit_doinsertfile **279**
- txtedit_getstates **281**
- txtedit_init **282**
- txtedit_install **277**
- txtedit_mayquit **277**
- txtedit_menu **278**
- txtedit_menuevent **278**
- txtedit_new **277**
- txtedit_prequit **278**
- txtedit_register_close_handler **280**
- txtedit_register_open_handler **281**
- txtedit_register_save_handler **279**
- txtedit_register_save_handler **279**
- txtedit_register_shutdown_handler **280**
- txtedit_register_undofail_handler **281**
- txtedit_register_update_handler **279**
- txtedit_register_update_handler **279**
- txtopt **282**
- txtopt_get_name **282**
- txtopt_set_name **282**

txtscrap_selectowner **283**
txtscrap_setselect **283**
txtwin_dispose **284**
txtwin_new **283**
txtwin_number **284**
txtwin_setcurrentwindow **284**

U

ungetc **123**

V

va_arg **112**
va_list **112**
vfprintf **111, 120**
visdelay_begin **284**
visdelay_end **285**
visdelay_init **285**
visdelay_percent **285**
vprintf **120**
vsprintf **121**

W

wcstombs **135**
wctomb **135**
wimp_baseofsprites **310**
wimp_bbits **289**
wimp_blockcopy **310**
wimp_box **290**
wimp_caretstr **295**
wimp_close_template **308**
wimp_close_wind 156, **305**
wimp_closedown **308**
wimp_command_tag **312**
wimp_commandwind **312**
wimp_commandwindow **312**
wimp_corrupt_fp_state_on_poll **306**
wimp_create_icon **305**
wimp_create_menu **307**

wimp_create_submenu **311**
wimp_create_wind 153, **305**
wimp_decode_menu **307**
wimp_delete_icon **305**
wimp_delete_wind **305**
wimp_drag_box **307**
wimp_dragstr **302**
wimp_dragtype **289**
wimp_emask **293**
wimp_errflags **310**
wimp_etypes **293**
wimp_eventdata **300**
wimp_eventstr **301**
wimp_flags **286**
wimp_font_array **303**
wimp_force_redraw **307**
wimp_get_caret_pos **307**
wimp_get_icon_info **306**
wimp_get_point_info **307**
wimp_get_rectangle **306**
wimp_get_wind_info **306**
wimp_get_wind_state 155, **306**
wimp_getwindowout **309**
wimp_i **289**
wimp_ibtype **288**
wimp_icon **292**
wimp_icondata **290**
wimp_iconflags **287**
wimp_igate **292**
wimp_initialise **304**
wimp_load_template **308**
wimp_menuflags **301**
wimp_menuhdr **242, 301**
wimp_menuitem **242, 302**
wimp_menustr **242, 301, 302**
wimp_mousestr **294**
wimp_msgaction **295, 297**
wimp_msgdataload **298**
wimp_msgdataopen **298**
wimp_msgdatasave **297**
wimp_msgdatasaveok **297**
wimp_msghdr **297**
wimp_msghelpreply **299**

- wimp_msghelprequest **299**
- wimp_msgprint **299**
- wimp_msggramfetch **298**
- wimp_msggramtransmit **299**
- wimp_msgstr **299**
- wimp_open_template **308**
- wimp_open_wind 155, **305**
- wimp_openstr 169, **292**
- wimp_palettestr **304**
- wimp_paletteword **304**
- wimp_ploticon **309**
- wimp_poll **305**
- wimp_pollidle **309**
- wimp_processkey **308**
- wimp_pshapestr **303**
- wimp_readpalette **309**
- wimp_readpixtrans **312**
- wimp_redraw_wind **306**
- wimp_redrawstr 169, **294**
- wimp_reporterror **310**
- wimp_save_fp_state_on_poll **305**
- wimp_sendmessage 103, 164, **311**
- wimp_sendwmessage **311**
- wimp_set_caret_pos **307**
- wimp_set_extent **308**
- wimp_set_icon_state **306**
- wimp_set_point_shape **308**
- wimp_setcolour **309**
- wimp_setfontcolours **311**
- wimp_setmode **309**
- wimp_setpalette **309**
- wimp_slotsize **311**
- wimp_spriteop **310**
- wimp_spriteop_full **310**
- wimp_starttask **308**
- wimp_t **290**
- wimp_taskclose **308**
- wimp_taskinit **304**
- wimp_template **303**
- wimp_transferblock 103, **311**
- wimp_update_wind **306**
- wimp_w **289**
- wimp_wcolours **287**
- wimp_which_block **302**
- wimp_which_icon **307**
- wimp_wind **291**
- wimp_winfo **291**
- wimp_wstate **292**
- wimpt_bpp **315**
- wimpt_checkmode **314**
- wimpt_complain 159, **314**
- wimpt_dx/wimpt_dy **314**
- wimpt_fake_event **313**
- wimpt_forceredraw **316**
- wimpt_init 153, **315**
- wimpt_last_event **313**
- wimpt_last_event_was_a_key **313**
- wimpt_mode **314**
- wimpt_noerr 159, **313**
- wimpt_poll **312**
- wimpt_programname **316**
- wimpt_reporterror **316**
- wimpt_task **316**
- wimpt_wimpversion **315**
- win_activatedec **321**
- win_activeinc 152, **321**
- win_activeno **321**
- win_add_unknown_event_processor 163, **318**
- win_claim_idle_events 163, **318**
- win_claim_unknown_events 164, **319**
- win_getmenuh **320**
- win_give_away_caret **322**
- win_idle_event_claimer **319**
- win_init **322**
- win_processevent **320**
- win_read_event_handler **317**
- win_register_event_handler **317**
- win_remove_unknown_event_processor 164, **319**
- win_setmenuh **320**
- win_settitle **322**
- win_unknown_event_claimer **319**
- win_unknown_event_processor 163

X

xferrecv_buffer_processor 165, **324**
xferrecv_checkimport **324**
xferrecv_checkinsert 166, **323**
xferrecv_checkprint **323**
xferrecv_do_import 165
xferrecv_doimport **324**
xferrecv_file_is_safe 166, **324**
xferrecv_insertfileok **323**
xferrecv_printfileok **323**
xferrecv_send **326**
xferrecv_clear_unknowns **329**
xferrecv_close_on_xfer **328**
xferrecv_file_is_safe **328**
xferrecv_pipe **327**
xferrecv_printproc 166, **254, 325**
xferrecv_read_last_ref **329**
xferrecv_saveproc 166, **325**
xferrecv_sendbuf **327**
xferrecv_sendproc 166, **254, 325**
xferrecv_set_fileissafe **328**

Reader's Comment Form

Acorn ANSI C Release 4

We would greatly appreciate your comments about this Manual, which will be taken into account for the next issue:

Did you find the information you wanted?

Do you like the way the information is presented?

General comments:

If there is not enough room for your comments, please continue overleaf

How would you classify your experience with computers?

Used computers before

Experienced User

Programmer

Experienced Programmer

Cut out (or photocopy) and post to:
Dept RC, Technical Publications
Acorn Computers Limited
645 Newmarket Road
Cambridge CB5 8PB
England

Your name and address:

This information will only be used to get in touch with you in case we wish to explore your comments further

