

## SYSTEM V APPLICATION BINARY INTERFACE Motorola 68000 Processor Family Supplement



**UNIX Software Operation** 

Copyright 1990 AT&T All Rights Reserved Printed in USA

Published by Prentice-Hall, Inc. A Division of Simon & Schuster Englewood Cliffs, New Jersey 07632

No part of this publication may be reproduced or transmitted in any form or by any means—graphic, electronic, electrical, mechanical, or chemical, including photocopying, recording in any medium, taping, by any computer or information storage and retrieval systems, etc., without prior permissions in writing from AT&T.

#### IMPORTANT NOTE TO USERS

While every effort has been made to ensure the accuracy of all information in this document, AT&T assumes no liability to any party for any loss or damage caused by errors or omissions or by statements of any kind in this document, its updates, supplements, or special editions, whether such errors are omissions or statements resulting from negligence, accident, or any other cause. AT&T further assumes no liability arising out of the application or use of any product or system described herein; nor any liability for incidental or consequential damages arising from the use of this document. AT&T disclaims all warranties regarding the information contained herein, whether expressed, implied or statutory, including implied warranties of merchantability or fitness for a particular purpose. AT&T makes no representation that the interconnection of products in the manner described herein will not infringe on existing or future patent rights, nor do the descriptions contained herein imply the granting or license to make, use or sell equipment constructed in accordance with this description.

AT&T reserves the right to make changes without further notice to any products herein to improve reliability, function, or design.

#### TRADEMARKS

UNIX is a registered trademark of AT&T.

The publisher offers discounts on this book when ordered in bulk quantities. For more information, write:

Special Sales Prentice-Hall, Inc. College Technical and Reference Division Englewood Cliffs, New Jersey 07632 or

call 201-592-2498 For single copies, call 201-767-5937

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-877663-6



4,070.350



## Contents

1	INTRODUCTION  Motorola 68000 Family and the System V ABI  How to Use the Motorola 68000 Family ABI Supplement	1-1 1-2
2	SOFTWARE INSTALLATION Software Distribution Formats	2-1
3	LOW-LEVEL SYSTEM INFORMATION  Machine Interface Function Calling Sequence Operating System Interface Coding Examples	3-1 3-10 3-19 3-31
4	OBJECT FILES ELF Header Sections Symbol Table Relocation	4-1 4-2 4-3 4-4
5	PROGRAM LOADING AND DYNAMIC LINKING Program Loading Dynamic Linking	5-1 5-5

6	LIBRARIES	
O	System Library	
	C Library	6-1
	System Data Interfaces	6-3
	System Bata interfaces	6-4

# Figures and Tables

	Scalar Types	3-2
	Structure Smaller Than a Long Word	3-3
	No Padding	3-4
Figure 3-4:	Internal Padding	3-4
	Internal and Tail Padding	3-5
	union Allocation	3-5
Figure 3-7:	Bit-Field Ranges	3-6
Figure 3-8:	Bit Numbering	3-7
Figure 3-9:	Left-to-Right Allocation	3-7
Figure 3-10:	Boundary Alignment	3-7
Figure 3-11:	Storage Unit Sharing	3-8
Figure 3-12:	union Allocation	3-8
Figure 3-13:	Unnamed Bit-Fields	3-8
Figure 3-14:	Processor Registers	3-11
Figure 3-15:	Standard Stack Frame	3-12
Figure 3-16:	Function Prologue	3-13
Figure 3-17:	Integral and Pointer Arguments	3-15
Figure 3-18:	Floating-Point Arguments	3-15
Figure 3-19:	Structure and Union Arguments	3-16
Figure 3-20:	Function Epilogue	3-16
Figure 3-21:	Function Epilogue	3-17
Figure 3-22:	Virtual Address Configuration	3-20
Figure 3-23:	Exceptions and Signals	3-23
Figure 3-24:	Declaration for main	3-24
Figure 3-25:	Condition Code Register Fields	3-25
Figure 3-26:	Initial Process Stack	3-26
Figure 3-27:	Auxiliary Vector	3-27
Figure 3-28:	Auxiliary Vector Types, a type	3-27
		3-30
Figure 3-30:	A CARLOS IN TANK IN CONTROL OF THE CARLOS AND THE CONTROL OF THE C	3-33
Figure 3-31:	Absolute Load and Store	3-34
Figure 3-32:	Position-Independent Load and Store	3-35
Figure 3-33:	Absolute Direct Function Call	3-35
Figure 3-34:	Position-Independent Direct Function Call	3-36
Figure 3-35:	Absolute Indirect Function Call	3-36
Figure 3-36:	. como: marponaria manostr anotion can	3-37
Figure 3-37:	Branch Instruction, Both Models	3-37

Table of Contents iii

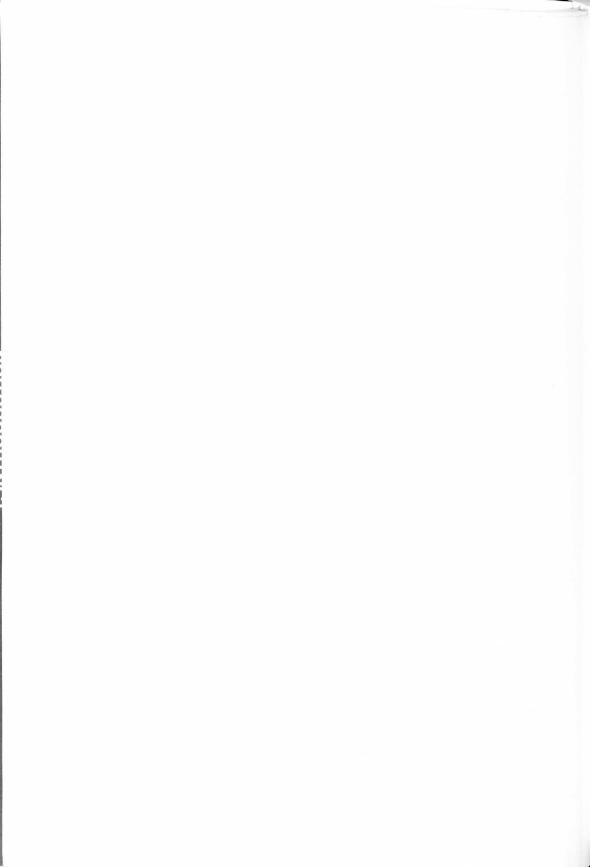
Figure 3-38: Absolute switch Code	0.00
Figure 3-39: Position-Independent switch Code	3-38
Figure 3-40: C Stack Frame	3-39
Figure 4-1: Motorola 68000 Family Identification, e_ident	3-40
Figure 4-2: Special Sections	4-1 4-2
Figure 4-3: Relocatable Fields	4-2 4-4
Figure 4-4: Relocation Types	4-4 4-7
Figure 5-1: Executable File	5-1
Figure 5-2: Program Header Segments	5-1 5-2
Figure 5-3: Process Image Segments	5-3
Figure 5-4: Example Shared Object Segment Addresses	5-4
Figure 5-5: Initial Procedure Linkage Table	5-7
Figure 6-1: libsys Support Routines	6-1
Figure 6-2: libsys, Global External Data Symbols	6-2
Figure 6-3: <assert.h></assert.h>	6-4
Figure 6-4: <ctype.h></ctype.h>	6-5
Figure 6-5: <dirent.h></dirent.h>	6-6
Figure 6-6: <errno.h>, Part 1 of 4</errno.h>	6-7
Figure 6-7: <errno.h>, Part 2 of 4</errno.h>	6-8
Figure 6-8: <errno.h>, Part 3 of 4</errno.h>	6-9
Figure 6-9: <errno.h>, Part 4 of 4</errno.h>	6-10
Figure 6-10: <fcntl.h>, Part 1 of 2</fcntl.h>	6-11
Figure 6-11: <fcntl.h>, Part 2 of 2</fcntl.h>	6-12
Figure 6-12: <float.h></float.h>	6-12
Figure 6-13: <fmtmsg.h></fmtmsg.h>	6-13
Figure 6-14: <ftw.h></ftw.h>	6-14
Figure 6-15: <grp.h></grp.h>	6-14
Figure 6-16: <sys ipc.h=""></sys>	6-15
Figure 6-17: <langinfo.h>, Part 1 of 2</langinfo.h>	6-16
Figure 6-18: <langinfo.h>, Part 2 of 2</langinfo.h>	6-17
Figure 6-19: <limits.h></limits.h>	6-18
Figure 6-20: <locale.h></locale.h>	6-19
Figure 6-21: <math.h></math.h>	6-20
Figure 6-22: <sys mman.h=""></sys>	6-20
Figure 6-23: <mon.h></mon.h>	6-21
Figure 6-24: <sys mount.h=""></sys>	6-21
Figure 6-25: <sys msg.h=""></sys>	6-22
Figure 6-26: <netconfig.h>, Part 1 of 2</netconfig.h>	6-23
Figure 6-27: <netconfig.h>, Part 2 of 2</netconfig.h>	6-24
Figure 6-28: <netdir.h></netdir.h>	6-25
Figure 6-29: <nl types.h=""></nl>	6-26

Figure 6-30:	<sys param.h=""></sys>	6-26
Figure 6-31:	<poll.h></poll.h>	6-27
Figure 6-32:	<sys procset.h=""></sys>	6-28
Figure 6-33:	<pwd.h></pwd.h>	6-29
Figure 6-34:	<sys regset.h=""></sys>	6-30
Figure 6-35:	<sys resource.h=""></sys>	6-31
Figure 6-36:	<pre><pre><pre><pre><pre>, Part 1 of 12</pre></pre></pre></pre></pre>	6-32
Figure 6-37:	<pre><rpc.h>, Part 2 of 12</rpc.h></pre>	6-33
Figure 6-38:	<rpc.h>, Part 3 of 12</rpc.h>	6-34
Figure 6-39:	<rpc.h>, Part 4 of 12</rpc.h>	6-35
Figure 6-40:	<rpc.h>, Part 5 of 12</rpc.h>	6-36
•	<rpc.h>, Part 6 of 12</rpc.h>	6-37
5	<rpc.h>, Part 7 of 12</rpc.h>	6-38
-	<rpc.h>, Part 8 of 12</rpc.h>	6-39
_	<rpc.h>, Part 9 of 12</rpc.h>	6-40
	<rpc.h>, Part 10 of 12</rpc.h>	6-41
•	<rpc.h>, Part 11 of 12</rpc.h>	6-42
_	<rpc.h>, Part 12 of 12</rpc.h>	6-43
9	<search.h></search.h>	6-44
	<sys sem.h=""></sys>	6-45
-	<set jmp.h=""></set>	6-46
_	<sys shm.h=""></sys>	6-47
	<sigaction.h></sigaction.h>	6-48
	<sys siginfo.h="">, Part 1 of 3</sys>	6-49
	<sys siginfo.h="">, Part 2 of 3</sys>	6-50
0	<sys siginfo.h="">, Part 3 of 3</sys>	6-51 6-52
-	<signal.h>, Part 1 of 2</signal.h>	6-53
_	<pre><signal.h>, Part 2 of 2</signal.h></pre>	6-54
	<pre><sys stat.h="">, Part 1 of 2</sys></pre>	6-55
	<sys stat.h="">, Part 2 of 2</sys>	6-56
_	<pre><sys statvfs.h=""></sys></pre>	6-56
	<stddef.h></stddef.h>	6-57
	<stdio.h></stdio.h>	6-58
0	<pre><stdlib.h> </stdlib.h></pre>	6-59
	<pre><stropts.h>, Part 1 of 4</stropts.h></pre>	6-60
	<pre>&lt; stropts.h&gt;, Part 2 of 4</pre>	6-61
-	<pre>&lt; <stropts.h>, Part 4 of 4</stropts.h></pre>	6-62
-	: <stropts.h>, Part 4 of 4 : <termios.h>, Part 1 of 6</termios.h></stropts.h>	6-63
	: <termios.n>, Part 1016 : <termios.h>, Part 2 of 6</termios.h></termios.n>	6-64
	: <termios.h>, Part 3 of 6</termios.h>	6-65
rigule 0-70.	, NUMBER OF THE PARTY OF THE PROPERTY OF THE P	0 00

Figure 6-71:	<termios.h>, Part 4 of 6</termios.h>	
	<termios.h>, Part 5 of 6</termios.h>	6-66
	<termios.h>, Part 6 of 6</termios.h>	6-67
	<pre><sys time.h="">, Part 1 of 2</sys></pre>	6-68
	<pre><sys time.h="">, Part 2 of 2</sys></pre>	6-69
	<pre><sys times.h=""></sys></pre>	6-70
	<pre><sys tiuser.h="">, Service Types</sys></pre>	6-70
	<pre><sys tiuser.h="">, Transport Interface States</sys></pre>	6-71
	<pre><sys tiuser.h="">, Transport interface States <sys tiuser.h="">, User-level Events</sys></sys></pre>	6-71
		6-72
Figure 6-80.	<pre><sys tiuser.h="">, Error Return Values</sys></pre>	6-73
Figure 6-01.	<sys tiuser.h="">, Transport Interface Data Structures, 1 of 2</sys>	6-74
Figure 6-82:	<pre><sys tiuser.h="">, Transport Interface Data Structures, 2 of 2</sys></pre>	6-75
	<sys tiuser.h="">, Structure Types</sys>	6-75
	<sys tiuser.h="">, Fields of Structures</sys>	6-76
	<sys tiuser.h="">, Events Bitmasks</sys>	6-76
	<sys tiuser.h="">, Flags</sys>	6-77
	<sys types.h=""></sys>	6-77
Figure 6-88:	<ucontext.h></ucontext.h>	6-78
Figure 6-89:		6-78
	<ul><li><ulimit.h></ulimit.h></li></ul>	6-79
	<pre><unistd.h>, Part 1 of 3</unistd.h></pre>	6-79
Figure 6-92:	<unistd.h>, Part 2 of 3</unistd.h>	6-80
Figure 6-93:	<unistd.h>, Part 3 of 3</unistd.h>	6-81
Figure 6-94:	<utime.h></utime.h>	6-81
Figure 6-95:	<utsname.h></utsname.h>	6-82
Figure 6-96:	<wait.h></wait.h>	6-83

# 1 INTRODUCTION

Motorola 68000 Family and the System V	
ABI	1-1
 How to Use the Motorola 68000 Family Al	BI
	1-2
Supplement	
Evolution of the ABI Specification	1-2



#### Motorola 68000 Family and the System V ABI

The **System V Application Binary Interface**, or **ABI**, defines a system interface for compiled application programs. Its purpose is to establish a standard binary interface for application programs on systems that implement UNIX System V Release 4.0 or some other operating system that complies with the **System V Interface Definition**, **Issue 3**.

This document is a supplement to the generic **System V ABI**, and it contains information specific to System V implementations built on the Motorola 68000 processor architecture family. The generic term "Motorola 68000 family" is used in this specification to mean the Motorola MC68020, MC68030, and MC68040 processor architectures; it does not refer to the MC68000, MC68008, or MC68010.

Together, these two specifications, the generic **System V ABI** and the **Motorola 68000 Family System V ABI Supplement**, constitute a complete *System V Application Binary Interface* specification for systems that implement the architecture of the Motorola 68000 family.

INTRODUCTION 1-1

# How to Use the Motorola 68000 Family ABI Supplement

This document is a supplement to the generic **System V ABI** and contains information referenced in the generic specification that may differ when System V is implemented on different processors. Therefore, the generic **ABI** is the prime reference document, and this supplement is provided to fill gaps in that specification.

As with the **System V ABI**, this specification references other publicly-available reference documents, including the

- MC68020 32-Bit Microprocessor User's Manual, MC68020UM/AD
- MC68030 Enhanced 32-Bit Microprocessor User's Manual, MC68030UM/AD
- MC68040 32-Bit Microprocessor User's Manual, MC68040UM/AD
- M68000 Programmer's Reference Manual, M68000PM/AD

available from Motorola.

All the information referenced by this supplement should be considered part of this specification, and just as binding as the requirements and data explicitly included here.

### **Evolution of the ABI Specification**

The **System V Application Binary Interface** will evolve over time to address new technology and market requirements, and will be reissued at intervals of approximately three years. Each new edition of the specification is likely to contain extensions and additions that will increase the potential capabilities of applications that are written to conform to the **ABI**.

As with the **System V Interface Definition**, the **ABI** will implement **Level 1** and **Level 2** support for its constituent parts. **Level 1** support indicates that a portion of the specification will continue to be supported indefinitely, while **Level 2** support means that a portion of the specification may be withdrawn or altered after the next edition of the **ABI** is made available. That is, a portion of the specification will remain in effect at least until the following edition of the specification is published.

How to Use the Motorola 68000 Family ABI Supp	plement
---	---------

These Level 1 and Level 2 classifications and qualifications apply to this Supplement, as well as to the generic specification. All components of the ABI and of this supplement have Level 1 support unless they are explicitly labeled as Level 2.

INTRODUCTION 1-3



# 2. SOFTWARE INSTALLATION

# 2 SOFTWARE INSTALLATION

Software Distribution Formats	2-1
Physical Distribution Media	2-1



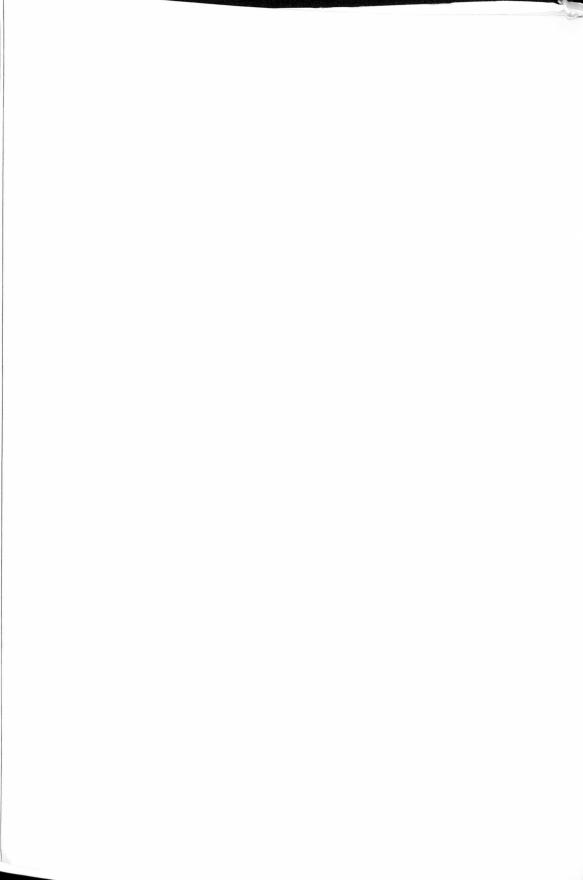
#### **Software Distribution Formats**

#### **Physical Distribution Media**

Approved media for physical distribution of ABI-conforming software are listed below. Inclusion of a particular medium on this list does not require an ABI-conforming system to accept that medium. For example, a conforming system may install all software through its network connection and accept none of the listed media.

- 5.25-inch floppy disk: 96 TPI (80 tracks/side) doubled-sided, 15 sectors/track, 512 bytes/sector, total format capacity of 1.2 megabytes per disk.
- 3.5-inch floppy disk: 135 TPI (80 tracks/side) double sided, 18 sectors/track, 512 bytes/sector, total format capacity of 1.44 megabytes per disk.
- 1/2-inch reel-to-reel tape: conforms to ANSI-standard reel-to-reel tape standard which consists of 9 tracks, 1600 BPI, no label.
- 60 MB quarter-inch cartridge tape in QIC-24 format.

The QIC-24 cartridge tape data format is described in *Proposed Standard for Data Interchange on the Streaming 1/4 Inch Magnetic Tape Cartridge Using Group Code Recording at 10000 FRPI*, Revision D, April 22, 1983. This document is available from the Quarter-Inch Committee (QIC) through Freeman Associates, 311 East Carillo St., Santa Barbara, CA 93101.



# 3 LOW-LEVEL SYSTEM INFORMATION

Machine Interface	3-1
Processor Architecture	3-1
Data Representation	3-1
■ Fundamental Types	3-1
Aggregates and Unions	3-3
Function Calling Sequence	3-10
Registers and the Stack Frame	3-10
Integral and Pointer Arguments	3-14
Floating-Point Arguments	3-15
Structure and Union Arguments	3-16
Functions Returning Scalars or No Value	3-16
Functions Returning Structures or Unions	3-17
Operating System Interface	3-19
Virtual Address Space	3-19
■ Page Size	3-19
Virtual Address Assignments	3-19
Managing the Process Stack	3-21
<ul><li>Coding Guidelines</li></ul>	3-21
Processor Execution Modes	3-22
Exception Interface	3-23
Process Initialization	3-24
Registers	3-24
■ Process Stack	3-25

Coding Examples	
Code Model Overview	3-31
	3-32
Position-Independent Function Prologue	3-33
Data Objects	3-34
Function Calls	3-35
Branching	3-37
C Stack Frame	
Variable Argument List	3-39
Allocating Stack Space Dynamically	3-40
- maraning attack opacie by halfileally	2 /1

#### **Machine Interface**

#### **Processor Architecture**

The MC68020 32-Bit Microprocessor User's Manual, the MC68030 Enhanced 32-Bit Microprocessor User's Manual, and the MC68040 32-Bit Microprocessor User's Manual define the 68000 family processor architecture. The M68000 Programmer's Reference Manual defines the programming model. An MC68851 Paged Memory Management Unit (PMMU) may be present with the MC68020. An MC68881 or MC68882 Floating Point Coprocessor (FPCP) is assumed to be present with the MC68020 or MC68030. Programs intended to execute directly on the processor use the instruction set, instruction encodings, and instruction semantics of this architecture. A program may use only the instructions defined for the MC68040. Refer to the M68000 Programmer's Reference Manual for information regarding proper instruction usage for the MC68020, MC68030, and MC68040 processors.

To be ABI-conforming, the processor must implement the architecture's instructions, perform the specified operations, and produce the specified results. The ABI neither places performance constraints on systems nor specifies what instructions must be implemented in hardware. A software emulation of the architecture could conform to the ABI.

Some processors might support the 68000 and MC68881/2 FPCP architectures as subsets, providing additional instructions or capabilities. Programs that use those capabilities explicitly do not conform to the 68000 ABI. Executing those programs on machines without the additional capabilities gives undefined behavior.

#### **Data Representation**

#### **Fundamental Types**

Figure 3-1 shows the correspondence between ANSI C's scalar types and the processor's.

Figure 3-1: Scalar Types

	Alignment					
Type	C	sizeof	(bytes)	68000		
	signed char char	1	1	signed byte		
	unsigned char	1	1	unsigned byte		
Integral	short signed short	2	2	signed word (2 bytes)		
	unsigned short	2	2	unsigned word		
	int signed int long signed long enum	4	4	signed long word (4 bytes)		
	unsigned int unsigned long	4	4	unsigned long word		
Pointer	any-type * any-type (*) ()	4	4	unsigned long word		
Floating- point	float	4	4	single-precision		
	double	8	8	double-precision		
	long double	16	8	extended-precision		

A null pointer (for all types) has the value zero.

NOTE

If a double or long double appears on the stack, not as a part of an aggregate or union, then it is aligned to 4 bytes.

#### **Aggregates and Unions**

An array assumes the alignment of its elements' type. The size of any object, including arrays, structures, and unions, always is a multiple of the object's alignment. Structure and union objects may, therefore, require padding to meet size and alignment constraints.

- The alignment of a structure or a union is the maximum of the alignment of its elements.
- Each member is assigned to the lowest available offset with the appropriate alignment. This may require *internal padding*, depending on the previous member.
- A structure's size is increased, if necessary, to make it a multiple of the structure's alignment. This may require *tail padding*, depending on the last member.



Aggregates or unions residing on the stack only require 4-byte alignment.

In the following examples, members' byte offsets appear in the upper left corners.

Figure 3-2: Structure Smaller Than a Long Word

```
struct {
    char c;
};

Byte aligned, sizeof is 1

C

C
```

#### Figure 3-3: No Padding

```
char c;
char d;
short s;
long n;
};
```

Long word aligned, sizeof is 8

0 C	1	d 2	s
4		n	

Figure 3-4: Internal Padding

```
struct {
    char c;
    short s;
};
```

Word aligned, sizeof is 4

) C	<sup>1</sup> pad	<sup>2</sup> S

Figure 3-5: Internal and Tail Padding

```
struct {
    char c;
    double d;
    short s;
};
```

8-byte aligned (4-byte on stack), sizeof is 24

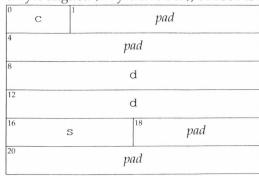
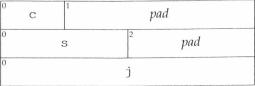


Figure 3-6: union Allocation

```
union {
    char c;
    short s;
    int j;
};
```

Long word aligned, size of is 4



#### **Bit-Fields**

C struct and union definitions may have *bit-fields*, defining integral objects with a specified number of bits.

Figure 3-7: Bit-Field Ranges

Bit-field Type	Width $w$	Range
signed char		$-2^{w-1}$ to $2^{w-1}-1$
char	1 to 8	0 to $2^{w}-1$
unsigned char		0 to $2^{w}-1$
signed short		$-2^{w-1}$ to $2^{w-1}-1$
short	1 to 16	0 to $2^{w}-1$
unsigned short		0 to $2^{w}-1$
signed int		$-2^{w-1}$ to $2^{w-1}-1$
int	1 to 32	0 to $2^{w}-1$
enum		0 to $2^{w}-1$
unsigned int		0 to $2^{w}-1$
signed long	1 to 32	$-2^{w-1}$ to $2^{w-1}-1$
long		0 to $2^{w}-1$
unsigned long		0 to $2^{w}-1$

"Plain" bit-fields always have non-negative values. Although they may have type char, short, int, or long (which can have negative values), these bit-fields are extracted into a long word with zero fill. Bit-fields obey the same size and alignment rules as other structure and union members, with the following additions.

- Bit-fields are allocated from left to right (most to least significant).
- A bit-field must entirely reside in a storage unit appropriate for its declared type. Thus a bit-field never crosses its unit boundary.
- Bit-fields may share a storage unit with other struct/union members, including members that are not bit-fields. Of course, struct members occupy different parts of the storage unit.
- Unnamed bit-fields' types do not affect the alignment of a structure or union, although individual bit-fields member offsets obey the alignment constraints.

The following examples show struct and union members' byte offsets in the upper left corners; bit numbers appear in the lower corners.

Figure 3-8: Bit Numbering



Figure 3-9: Left-to-Right Allocation

```
struct {
    int j:5;
    int k:6;
    int m:7;
};
Long word aligned, sizeof is 4

order

long worder

long word aligned, sizeof is 4

order

long word aligned, si
```

Figure 3-10: Boundary Alignment

```
Long word aligned, sizeof is 12
struct {
    short
            s:9;
                                                 pad
                             S
                                                          C
    int
            j:9;
                        31
    char
            C;
                                      pad
                                                          pad
                             t
                                                 u
                        31
    short t:9;
    short u:9;
                                                pad
                             d
    char
            d;
};
```

#### Figure 3-11: Storage Unit Sharing

```
struct {
    char c;
    short s:8;
};
```

```
Word aligned, sizeof is 2 \begin{bmatrix} 0 & & & 1 \\ 0 & & & & 1 \\ 15 & & & & 7 \end{bmatrix}
```

#### Figure 3-12: union Allocation

```
union {
    char c;
    short s:8;
};
```

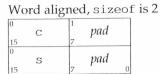
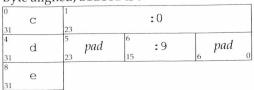


Figure 3-13: Unnamed Bit-Fields

```
char c;
int :0;
char d;
short :9;
char e;
char :0;
};
```

Byte aligned, sizeof is 9



As the examples show, int bit-fields (including signed and unsigned) pack more densely than smaller base types. One can use char and short bit-fields to force particular alignments, but int generally works better.

## **Function Calling Sequence**

This section discusses the standard function calling sequence, including stack frame layout, register usage, parameter passing, etc. The operating system interface and C programs use this calling sequence. See "Coding Examples" later in this chapter for more information on C.



All of the coding examples in this section are simply illustrations of a sample implementation.

# Registers and the Stack Frame

The 68000 provides 8 data and 8 address registers, which are global to a running program. Additionally, the MC68040, or the MC68881/2 Floating Point Coprocessor, provides 8 global floating-point registers. Brief register descriptions appear in Figure 3-14; more complete information appears later.

Figure 3-14: Processor Registers

Туре	N	ames	Usage
	%d0 %d1 %a0 %a1		Scratch registers
68000	%d2  %d7 %a2  %a5		Local register variables
	%a6	%fp	Frame pointer (if implemented)
	%a7	%sp	Stack pointer
		%pc	Program counter
		%ccr	Condition code register
	%fp0 %fp1		Scratch registers
MC68040,	%fp2  %fp7		Local register variables
MC68881/2		%fpcr	Floating Point Control Register
		%fpsr	Floating Point Status Register
		%fpiar	Floating Point Instruction Address Register

In addition to the registers, each function has a frame on the run-time stack. This stack grows downward from high addresses. Figure 3-15 shows the stack frame organization right after function prologue processing has allocated the frame and saved the registers (see below).

Figure 3-15: Standard Stack Frame

Base	Offset	Contents	Purpose	_
	+4+4*n +4	argument long word $n$ argument long word $0$	incoming arguments	High Addresses
		return address		_
%fp	(optional)	previous %fp (optional)		_
		unspecified 	local storage and register	
SP	(SP after prologue)	variable size	save area	
		stack top, unused		Low addresses

Several key points about the stack frame deserve mention.

- The stack frame is long word aligned.
- Functions may save registers %d2 through %d7, %a2 through %a6, and %fp2 through %fp7 *as necessary*, without saving unused registers.
- All incoming arguments reside on the stack. "Coding Examples" below explains how variable argument lists may be implemented.
- A frame pointer may be implemented.
- Other areas depend on the compiler and the code being compiled. The standard calling sequence does not define a maximum stack frame size, nor does it restrict how a language system uses the "unspecified" areas of the standard stack frame.

Registers %d2 through %d7, %a2 through %a6, and %fp2 through %fp7, which are visible to both a calling and a called function, "belong" to the calling function. In other words, a called function must save these registers' values before it changes them, restoring their values before it returns. Remaining registers "belong" to the called function. If a calling function wants to preserve such a register value

across a function call, it must save the value in its local stack frame.

Across function boundaries, the standard function prologue saves registers %d2 through %d7, %a2 through %a6, and %fp2 through %fp7 (as needed) and allocates stack space, including the required areas of Figure 3-15 and any private space it needs. The example below illustrates this, saving registers %fp, %d7, %a5, and %fp2 and allocating 80 bytes for local storage.

Figure 3-16: Function Prologue

fcn:
link.l %fp,&-80
movm.l %d7/%a5,-(%sp)
fmovm.x %fp2,-(%sp)

The movm instruction manipulates registers as part of the normal function prologue and epilogue. As explained later, the function epilogue executes a movm instruction to unwind the stack and restore the saved registers to their original condition.

NOTE

Strictly speaking, a function does not need the  ${\tt movm}$  instructions if it preserves the registers as described above. Although some functions can be optimized to eliminate the  ${\tt movm}$  instructions, the general case uses the standard prologue and epilogue.

Some registers have assigned roles.

%fp or %a6 The frame pointer, if used by an individual function, holds

the address of the local storage within a stack frame, refer-

enced as negative offsets from %fp.

%sp or %a7 The *stack pointer* holds the limit of the current stack frame,

which is the address of the stack's topmost valid long word. The long word to which \$sp points is "in" the valid

stack.

#### **Function Calling Sequence**

%d0	Integral return values appear in %d0.
%a0	Pointer return values appear in %a0. When calling a function that returns a structure or union, the caller allocates space for the return value and sets %a0 to its address. A function that returns a structure or union value places the same address in %a0 before it returns.
%fp0	Floating-point return values appear in this register.

Except as specified here, d0, d1, a0, a1, fp0, and fp1 are scratch registers. Functions do not need to preserve their values for the caller.

Local registers %d2 through %d7, %a2 through %a5, and %fp2 through %fp7 have no specified role in the standard calling sequence.

Signals can interrupt processes [see signal(BA\_OS)]. Functions called during signal handling have no unusual restrictions on their use of registers. Moreover, if a signal handling function returns, the process resumes its original execution path with registers restored to their original values. Thus programs and compilers may freely use all registers without the danger of signal handlers changing their values.

## **Integral and Pointer Arguments**

As mentioned, a function receives all integral and pointer argument long words on the stack. Functions pass all integer-valued arguments as long words, expanding signed or unsigned bytes and words as needed.



The following examples use the frame pointer. Functions not using  $\$ {\rm fp}$  would find arguments at different locations.

Figure 3-17: Integral and Pointer Arguments

Call	Argument	Callee	
	1	8(%fp)	
g(1, 2, 3,	2	12 (%fp)	
(void *)0);	3	16(%fp)	
	(void *)0	20 (%fp)	

## **Floating-Point Arguments**

The stack also holds floating-point arguments: single-precision values use one long word, double-precision use two, and extended-precision use four. The example below uses only double-precision arguments.

Figure 3-18: Floating-Point Arguments

Call	Argument	Callee
h(1.414, 1, 2.998e10);	long word 0, 1.414	8 (%fp)
	long word 1, 1.414	12 (%fp)
	1	16(%fp)
2.996010);	long word 0, 2.998e10	20 (%fp)
	long word 1, 2.998e10	24(%fp)

## **Structure and Union Arguments**

As described in the data representation section, structures and unions always have the alignment of their most strictly aligned member. When passed as arguments, sizes are rounded up to the next long word size. Structure and union objects appear directly on the stack, occupying as many long words as necessary.

Figure 3-19: Structure and Union Arguments

Call	Argument	Callee
i(1, s);	long word 0, s long word 1, s	8 (%fp) 12 (%fp) 16 (%fp)

## **Functions Returning Scalars or No Value**

A function that returns an integral value places its result in %d0. A function that returns a pointer value places its result in %a0. A function that returns a floating-point value places its result in %fp0.

Functions that return no value put no specified value in any return register.

Just as the function prologue may save registers, the epilogue must restore those same registers before returning to the caller. To complete the example from Figure 3-16, the following function epilogue restores %d7, %a5, and %fp2 and then returns.

Figure 3-20: Function Epilogue

fmovm.x	(%sp)+,%fp2
movm.1	(%sp)+,%d7/%a5
unlk	%fp
rts	

## **Functions Returning Structures or Unions**

As mentioned above, when a function returns a structure or union, it expects the caller to provide space for the return value and to place its address in register %a0. Having the caller supply the return object's space allows re-entrancy.



Structures and unions in this context have fixed sizes. The ABI does not specify how to handle variable sized objects.

A function returning a structure or union also sets %a0 to the value it finds in %a0. Thus when the caller receives control again, the address of the returned object resides in register %a0. Both the calling and the called functions must cooperate to pass the return value successfully. Failure of either side to meet its obligations leads to undefined program behavior.

The following example assumes the return object has been copied, and its address is in register %a5.

Figure 3-21: Function Epilogue

mov.l	%a5,%a0
fmovm.x	(%sp)+,%fp2
movm.1	(%sp)+,%d7/%a5
unlk	%fp
rts	

## **Operating System Interface**

## **Virtual Address Space**

Processes execute in a 32-bit virtual address space. Memory management hardware translates virtual addresses to physical addresses, hiding physical addressing and letting a process run anywhere in the system's real memory. Processes typically begin with three logical segments, commonly called text, data, and stack. As Chapter 5 describes, dynamic linking creates more segments during execution, and a process can create additional segments for itself with system services.

#### Page Size

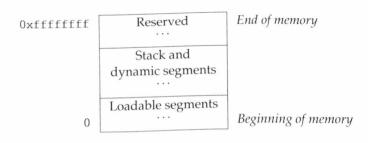
Memory is organized by pages, which are the system's smallest units of memory allocation. Page size can vary from one system to another, depending on the processor, memory management unit and system configuration. Allowable page sizes are 2K, 4K, or 8K. Processes may call <code>sysconf(BA\_OS)</code> to determine the system's current page size.

#### **Virtual Address Assignments**

Conceptually, processes have the full 32-bit address space available. In practice, however, several factors limit the size of a process.

- The system reserves a configuration-dependent amount of virtual space.
- A tunable configuration parameter limits process size.
- A process whose size exceeds the system's available, combined physical memory and secondary storage cannot run. Although some physical memory must be present to run any process, the system can execute processes that are bigger than physical memory, paging them to and from secondary storage. Nonetheless, both physical memory and secondary storage are shared resources. System load, which can vary from one program execution to the next, affects the available amounts.

Figure 3-22: Virtual Address Configuration





Programs that dereference null pointers are erroneous. Although such programs may appear to work on the 68000, they might fail or behave differently on other systems. To enhance portability, programmers are strongly cautioned not to rely on dereferencing null pointers.

#### Loadable segments

Processes' loadable segments may begin at 0. The exact addresses depend on the executable file format (see Chapters 4 and 5).

#### Stack and dynamic segments

A process's stack and dynamic segments reside below the reserved area. Processes can control the amount of virtual memory allotted for stack space, as described below.

Reserved

A reserved area resides at the top of virtual space.



Although application programs may begin at virtual address 0, they conventionally begin above  $0 \times 10000$  (64K), leaving the initial 64K with an invalid address mapping. Processes that reference this invalid memory (for example, by dereferencing a null pointer) generate an access exception trap, as described in the "Exception Interface" section of this chapter.

As the figure shows, the system reserves the high end of virtual space, with a process's stack and dynamic segments below that. Although the exact boundary between the reserved area and a process depends on the system's configuration, the reserved area shall not consume more than 512 MB from the virtual address

space. Thus the user virtual address range has a minimum upper bound of 0xdfffffff. Individual systems may reserve less space, increasing processes' virtual memory range. More information follows in the section "Managing the Process Stack."

Although applications may control their memory assignments, the typical arrangement follows the diagram above. Loadable segments reside at low addresses; dynamic segments occupy the higher range. When applications let the system choose addresses for dynamic segments (including shared object segments), it chooses high addresses. This leaves the "middle" of the address spectrum available for dynamic memory allocation with facilities such as malloc(BA OS).

#### Managing the Process Stack

Section "Process Initialization" in this chapter describes the initial stack contents. Stack addresses can change from one system to the next—even from one process execution to the next on a single system. Processes, therefore, should *not* depend on finding their stack at a particular virtual address. The stack segment has read and write permissions.

A tunable configuration parameter controls the system maximum stack size. A process also can use setrlimit(BA\_OS), to set its own maximum stack size, up to the system limit. Changes in the stack virtual address and size affect the virtual addresses for dynamic segments. Consequently, processes should *not* depend on finding their dynamic segments at particular virtual addresses. Facilities exist to let the system choose dynamic segment virtual addresses.

#### Coding Guidelines

Operating system facilities, such as mmap(KE\_OS), allow a process to establish address mappings in two ways. First, the program can let the system choose an address. Second, the program can force the system to use an address the program supplies. This second alternative can cause application portability problems, because the requested address might not always be available. Differences in virtual address space can be particularly troublesome between different architectures, but the same problems can arise within a single architecture.

Processes' address spaces typically have three segment areas that can change size from one execution to the next: the stack [through setrlimit(BA\_OS)], the data segment [through malloc(BA\_OS)], and the dynamic segment area [through mmap(KE\_OS)]. Consequently, an address that is available in one process

execution might not be available in the next. A program that used mmap(KE\_OS) to request a mapping at a specific address thus could appear to work in some environments and fail in others. For this reason, programs that wish to establish a mapping in their address space should let the system choose the address.

Despite these warnings about requesting specific addresses, the facility can be used properly. For example, a multiprocess application might map several files into the address space of each process and build relative pointers among the files' data. This could be done by having each process ask for a certain amount of storage at an address chosen by the system. After each process receives its own, private address from the system, it would map the desired files into memory, at specific addresses within the original area. This collection of mappings could be at different addresses in each process but their *relative* positions would be fixed. Without the ability to ask for specific addresses, the application could not build shared data structures, because the relative positions for files in each process would be unpredictable.

### **Processor Execution Modes**

Two execution modes exist in the 68000 architecture: user and supervisor. Processes run in user mode (the least privileged). The operating system kernel runs in supervisor mode. A program executes the trap instruction to change execution modes.



The ABI does not define the implementation of individual system calls. Instead, programs shall use the system libraries that Chapter 6 describes. Programs with embedded system call trap instructions do not conform to the ABI.

## **Exception Interface**

As the 68000 user manuals describe, the processor also changes mode to handle *exceptions*. Exceptions can be explicitly generated by a process as a result of instruction execution. The operating system defines the following correspondence between hardware exceptions and the signals specified by signal (BA\_OS).

Figure 3-23: Exceptions and Signals

<b>Exception Name</b>	Signal
trap #1 (breakpoint trap)	SIGTRAP
external memory fault	see below
address error	SIGBUS
all floating-point exceptions	SIGFPE
illegal instruction	SIGILL
integer zero-divide	SIGFPE
privileged opcode	SIGILL
trace	SIGTRAP
chk, chk2 instruction	SIGFPE
cptrapcc, trapcc, trapv	SIGFPE
line 1010 emulator	SIGSYS
line 1111 emulator	SIGSYS
trap #2-15	SIGSYS

An external memory fault exception can generate various signals, depending on why the exception occurred.

#### SIGBUS or SIGEMT

The process accessed a memory location in a way disallowed by the current mapping's permissions. As an example, the process tried to store a value into a location without write permission.

SIGSEGV

The process referenced a memory address for which no valid mapping existed.

#### **Process Initialization**

This section describes the machine state that exec(BA\_OS) creates for "infant" processes, including argument passing, register usage, stack frame layout, etc. Programming language systems use this initial program state to establish a standard environment for their application programs. As an example, a C program begins executing at a function named main, conventionally declared in the following way.

Figure 3-24: Declaration for main

```
extern int main(int argc, char *argv[], char *envp[]);
```

Briefly, argc is a non-negative argument count; argv is an array of argument strings, with argv[argc]==0; and envp is an array of environment strings, also terminated by a null pointer.

Although this section does not describe C program initialization, it gives the information necessary to implement the call to main or to the entry point for a program in any other language.

#### Registers

Registers %d0 through %d7, %a0 through %a6, and all MC68040 or MC68881/2 FPCP floating-point data registers have unspecified values at process entry. The floating-point control registers are set to provide IEEE default behavior. Consequently, a program that requires registers to have specific values must set them explicitly during process initialization. It should *not* rely on the operating system to set all registers to 0. See "Process Stack" below for information about the initial values of the stack registers.

As the architecture defines, the status register controls and monitors the processor. Application programs cannot access the entire status register directly; they run in the processor's *user mode*, and the instructions to access the entire status register are privileged. Nonetheless, a program can access the condition code register, which initially has the following value.

Figure 3-25: Condition Code Register Fields

Field Value		Note	
XNZVC	unspecified	Condition codes unspecified	

#### **Process Stack**

When a process receives control, its stack holds the arguments and environment from exec(BA\_OS).

Figure 3-26: Initial Process Stack

	Unspecified	1
	Information block, including	
	argument strings	
	environment strings	
	auxiliary information	
	•••	
	size varies	
	Unspecified	
	Null auxiliary vector entry	
	Auxiliary vector	
	2-long word entries	
	0 long word	
	Environment pointers	
	one long word each	
	0 long word	
	Argument pointers	
	•••	
4(%sp)	Argument count long words	
)(%sp)	Argument count	
	Undefined	

High Addresses

Low Addresses

Every process has a stack, but the system defines *no* fixed stack address. Furthermore, a program's stack address can change from one system to another—even from one process invocation to another. Thus the process initialization code must use the address in %sp.

Argument strings, environment strings, and auxiliary information appear in no specific order within the information block; the system makes no guarantees about their arrangement. The system may leave an unspecified amount of memory between the null auxiliary vector entry and the start of the information block.

Whereas the argument and environment vectors transmit information from one application program to another, the auxiliary vector conveys information from the operating system to the program. This vector is an array of the following structures, interpreted according to the a type member.

Figure 3-27: Auxiliary Vector

```
typedef struct
{
    int a_type;
    union {
        long a_val;
        void *a_ptr;
        void (*a_fcn)();
    } a_un;
} auxv_t;
```

Figure 3-28: Auxiliary Vector Types, a type

Name	Value	a_un
AT NULL	0	ignored
AT IGNORE	1	ignored
AT EXECFD	2	a val
AT_PHDR	3	a ptr
AT_PHENT	4	a_val
AT_PHNUM	5	a_val
AT_PAGESZ	6	a_val

Figure 3-28: Auxiliary Vector Types, a type (continued)

Name	Value	a_un	
AT BASE	7	a_ptr	
AT FLAGS	8	a_val	
AT ENTRY	9	a_ptr	

The auxiliary vector has no fixed length; instead its last entry's a type member has this value.

AT\_IGNORE This type indicates the entry has no meaning. The corresponding value of a un is undefined.

As Chapter 5 describes, exec(BA\_OS) may pass control to an interpreter program. When this happens, the system places either an entry of type AT\_EXECFD or one of type AT\_PHDR in the

auxiliary vector. The entry for type AT\_EXECFD uses the a\_val member to contain a file descriptor open to read the application

program's object file.

AT\_PHDR Under some conditions, the system creates the memory image

of the application program before passing control to the interpreter program. When this happens, the a\_ptr member of the AT\_PHDR entry tells the interpreter where to find the program header table in the memory image. If the AT\_PHDR entry is present, entries of types AT\_PHENT, AT\_PHNUM, and AT\_ENTRY must also be present. See Chapter 5 in both the System V ABI and the processor supplement for more information about the program

header table.

AT PHENT The a val member of this entry holds the size, in bytes, of one

entry in the program header table to which the AT\_PHDR entry

points.

AT\_PHNUM The a\_val member of this entry holds the number of entries in

the program header table to which the AT PHDR entry points.

AT_PAGESZ	If present, this entry's a_val member gives the system page
	size, in bytes. The same information also is available through
	sysconf(BA_OS).

The a\_ptr member of this entry holds the base address at which the interpreter program was loaded into memory. See "Program Header" in the System V ABI for more information about

the base address.

AT\_FLAGS If present, the a\_val member of this entry holds one-bit flags.

Bits with undefined semantics are set to zero. No flags are

defined for the 68000.

AT ENTRY The a ptr member of this entry holds the entry point of the

application program to which the interpreter program should

transfer control.

Other auxiliary vector types are reserved.

In the following example, the stack resides at an address below 0xf0000000, growing toward lower addresses. The process receives three arguments.

- **С**р
- src
- dst

It also inherits two environment strings (this example is not intended to show a fully configured execution environment).

- HOME=/home/dir
- PATH=/home/dir/bin:/usr/bin:

Its auxiliary vector holds one non-null entry, a file descriptor for the executable file.

■ {AT EXECFD, 13}

Figure 3-29: Example Process Stack

0xeffffffc	\0	С	р	\0	High addresses
Oxelllile	\0	S	r	С	
	\0	d	S	t	
oxeffffff0	/	d	i	r	
01102	h	0	m	е	
	М	Е	=	/	
	:	\0	Н	О	
oxefffffe0	/	b	i	n	
	/	u	S	r	
	b	i	n	:	
	d	i	r	/	
0xefffffd0	0	m	е	/	
	Н	=	/	h	
0xefffffc8	pad	Р	A	Т	
			0		
0xefffffc0	0				
		1	13		
			2		Auxiliary Vector
			0		
0xefffffb0	0xefffffc9				
		0xef1	Efffe	6	Environment vector
			0		
		0xef	fffff	5	
0xefffffa0	0xefffffa0 0xefff		fffff	9	
	0xeffffffd		d	Argument Vector	
%sp, 0xefffff98			3		Argument Count
-4(%sp), 0xefffff94		Und	efined		Low addresses

## **Coding Examples**

This section discusses example code sequences for fundamental operations such as calling functions, accessing static objects, and transferring control from one part of a program to another. Previous sections discuss how a program may use the machine or the operating system, and they specify what a program may and may not assume about the execution environment. The information here illustrates how operations *may* be done, not how they *must* be done.

As before, examples use the ANSI C language. Other programming languages may use the same conventions displayed below, but failure to do so does *not* prevent a program from conforming to the ABI. Two main object code models are available.

- Absolute code. Instructions can hold absolute addresses under this model. To execute properly, the program must be loaded at a specific virtual address, making the program's absolute addresses coincide with the process's virtual addresses.
- Position-independent code. Instructions under this model hold relative addresses, not absolute addresses. Consequently, the code is not tied to a specific load address, allowing it to execute properly at various positions in virtual memory.

The following sections describe the differences between these models. Code sequences for the models (when different) appear together, allowing easier comparison.

NOTE

Examples below show code fragments with various simplifications. They are intended to explain addressing modes, not to show optimal code sequences nor to reproduce compiler output.



When other sections of this document show assembly language code sequences, they typically show only the absolute versions. Information in this section explains how position-independent code would alter the examples.

#### **Code Model Overview**

When the system creates a process image, the executable file portion of the process has fixed addresses, and the system chooses shared object virtual addresses to avoid conflicts with other segments in the process. To maximize text sharing, shared object libraries conventionally use position-independent code, in which instructions contain no absolute addresses. Shared object text segments can be loaded at various virtual addresses without having to change the segment images. Thus multiple processes can share a single shared object text segment, even though the segment resides at a different virtual address in each process.

Position-independent code relies on two techniques.

- Control transfer instructions hold addresses relative to the program counter (PC). A PC-relative branch or function call computes its destination address in terms of the current program counter, *not* relative to any absolute address.
- When the program requires an absolute address, it computes the desired value. Instead of embedding absolute addresses in the instructions, the compiler generates code to calculate an absolute address during execution.

Because the processor architecture provides PC-relative call and branch instructions, compilers can satisfy the first condition easily.

A global offset table and a procedure linkage table provide information for address calculation. Position-independent object files (executable and shared object files) have these tables in their data segment. When the system creates the memory image for an object file, the table entries are relocated to reflect the absolute virtual addresses as assigned for an individual process. Because data segments are private for each process, the table entries can change—unlike text segments, which multiple processes share.

Assembly language examples below show the explicit notation needed for position-independent code.

name@GOT This expression denotes the displacement in the global offset table of the entry for the symbol *name*.

name@PLT This expression denotes the displacement in the procedure linkage table of the entry for the symbol *name*.

#### name@GOTPC

This expression denotes a PC-relative reference to the global offset table entry for the symbol *name*.

#### name@PLTPC

This expression denotes a PC-relative reference to the procedure linkage table entry for the symbol *name*.

## Position-Independent Function Prologue

This section describes the function prologue for position-independent code. A function's prologue first allocates the local stack space. Position-independent functions also set register %a5 to the global offset table's address, accessed with the symbol \_GLOBAL\_OFFSET\_TABLE\_. Because %a5 is private for each function and preserved across function calls, a function calculates its value once at the entry.



As a reminder, this entire section contains examples. Using %a5 is a convention, not a requirement; moreover, this convention is private to a function. Not only could other registers serve the same purpose, but different functions in a program could use different registers.



When an instruction uses \_GLOBAL\_OFFSET\_TABLE\_@GOTPC, it sees the offset between the current instruction and the global offset table as the symbol value.

Figure 3-30: Position-Independent Function Prologue

```
name:
link.1 %fp,&-80
movm.1 %a5,-(%sp)
lea (%pc,_GLOBAL_OFFSET_TABLE_@GOTPC),%a5
```

## **Data Objects**

This discussion excludes stack-resident objects, because programs always compute their virtual addresses relative to the stack and frame pointers. Instead, this section describes objects with static storage duration. Symbolic references in absolute code put the symbols' values—or absolute virtual addresses—into instructions.

Figure 3-31: Absolute Load and Store

C	Assembly		
extern int src; extern int dst;	.global	src,dst,ptr	
<pre>extern int *ptr; ptr = &amp;dst</pre>	mov.1	&dst,ptr	
*ptr = src;	mov.1	src,([ptr])	

Position-independent code cannot contain absolute addresses. Referencing global symbols must be done with a base register and global offset table index (as in these examples). Alternatively, instructions that reference symbols hold the PC-relative offsets into the global offset table. Combining the offset with the PC gives the absolute address of the table entry holding the desired address.

Figure 3-32: Position-Independent Load and Store

C Assembly

extern int src;
extern int dst;
extern int \*ptr;
ptr = &dst;

\*ptr = src;

mov.l ([%a5,ptr@GOT]), %a4
mov.l ([%a5,src@GOT]), (%a4)

#### **Function Calls**

Function calls in absolute code put the symbols' values—or absolute virtual addresses—into instructions. Programs use the <code>jsr</code> or <code>bsr</code> instructions to make function calls. For absolute code, the destination operand is an absolute address. Even when the code for a function resides in a shared object, the caller uses the same assembly language instruction sequence, although in that case control passes from the original call, through an indirection sequence, to the desired destination. See ''Procedure Linkage Table'' in Chapter 5 for more information on the indirection sequence.

Figure 3-33: Absolute Direct Function Call

For position-independent code, the bsr instruction's destination operand is a PC-relative value. Typically, the destination will be an entry in the procedure linkage table mentioned above.

Figure 3-34: Position-Independent Direct Function Call

C	Assembly		
<pre>extern void function();</pre>	.global	function	
<pre>function();</pre>	bsr	function@PLTPC	

Indirect function calls also use the jsr instruction.

Figure 3-35: Absolute Indirect Function Call

С	Assembly		
extern void (*ptr)();	.global	ptr,name	
<pre>extern void name(); ptr = name;</pre>	mov.1	&name,ptr	
(*ptr)();	jsr	([ptr])	

For position-independent code, the global offset table supplies absolute addresses for all required symbols, whether the symbols name objects or functions.

Figure 3-36: Position-Independent Indirect Function Call

extern void (\*ptr)();
extern void name();
ptr = name;
(\*ptr)();

	Assembly
.global	ptr,name
mov.1	(%a5,name@GOT),([%a5,ptr@GOT])
mov.l jsr	([%a5,ptr@GOT]),%a0 (%a0)

## **Branching**

Programs use branch instructions to control their execution flow. As defined by the architecture, branch instructions can hold a PC-relative value with a range that covers the entire address space.

Figure 3-37: Branch Instruction, Both Models



C switch statements provide multiway selection. When the case labels of a switch statement satisfy grouping constraints, the compiler implements the selection with an address table. The following examples use several simplifying conventions to hide irrelevant details:

- The selection expression resides in register %d0;
- case label constants begin at zero;
- case labels, default, and the address table use assembly names .Lcasei, .Ldef, and .Ltab, respectively.

Address table entries for absolute code contain virtual addresses; the selection code extracts an entry's value and jumps to that address. Position-independent table entries hold offsets; the selection code computes a destination's absolute address.

Figure 3-38: Absolute switch Code

C	Assembly			
switch (j) {     case 0:		cmp.1 bhi asl.1 jmp	%d0,&3 .Ldef &2,%d0 ([%pc,%d0.1,.Ltab])	
case 2:  case 3:  default:	.Ltab:	long .long .long	.Lcase0 .Ldef .Lcase2 .Lcase3	
}				

Figure 3-39: Position-Independent switch Code

C	Assembly			
switch (j)		cmp.1	%d0,&3	
{		bhi	.Ldef	
case 0:		mov.1	(%pc,%d0.w*4,.Ltab),%d0	
		jmp	(%pc,%d0,.Ltab)	
case 2:				
case 3:				
default:				
	.Ltab:	.long	.Lcase0Ltab	
}		.long	.LdefLtab	
		.long	.Lcase2Ltab	
		.long	.Lcase3Ltab	

## C Stack Frame

Figure 3-40 shows the C stack frame organization. It conforms to the standard stack frame with designated roles for unspecified areas in the standard frame.

#### Figure 3-40: C Stack Frame

Base	Offset	Contents	Purpose	
Duse	+4+4*n	argument long word n	incoming	High Addresses
	+4	 argument long word 0	arguments	
SP'	(SP before call)	return address		
		unspecified 	local storage and register save area	
SP	(SP after call)	variable size		_
		outgoing arguments		_
		stack top, unused		Low addresses

# Variable Argument List

Previous sections describe the rules for passing arguments. Unfortunately, some otherwise portable C programs depend on the argument passing scheme, implicitly assuming that 1) all arguments reside on the stack, and 2) arguments appear in increasing order on the stack. Programs that make these assumptions never have been portable, but they have worked on many machines, including the 68000. Nonetheless, portable C programs should use the header files <std>stdarg.h> or <varargs.h> to deal with variable argument lists (on 68000 and other machines as well).

## **Allocating Stack Space Dynamically**

Unlike some other languages, C does not need dynamic stack allocation *within* a stack frame. Frames are allocated dynamically on the program stack, depending on program execution. The architecture supports dynamic allocation for those languages that require it, and the standard calling sequence and stack frame support it as well. Thus languages that need dynamic stack frame sizes can call C functions, and vice versa.

Figure 3-40 shows the layout of the C stack frame. The double line divides the area allocated by the compiler from the dynamically allocated memory. Dynamic space is allocated below the line as a downward growing heap whose size changes as required. Typical C functions have no space in the heap. All areas above the double line in the current frame have a known size to the compiler. Dynamic stack allocation thus takes the following steps.

- Stack frames are long word aligned; dynamic allocation should preserve this property. Thus the program rounds (up) the desired byte count to a multiple of 4.
- 2. The program decreases the stack pointer by the rounded byte count, increasing its frame size. At this point, the "new" space resides just below the register save area at the bottom of the stack.

Even in the presence of signals, dynamic allocation is "safe." If a signal interrupts allocation, one of three things can happen.

- The signal handler can return. The process then resumes the dynamic allocation from the point of interruption.
- The signal handler can execute a non-local goto, or longjmp [see set jmp(BA\_LIB)]. This resets the process to a new context in a previous stack frame, automatically discarding the dynamic allocation.
- The process can terminate.

Regardless of when the signal arrives during dynamic allocation, the result is a consistent (though possibly dead) process.

Coding	Examp	es

Existing stack objects reside at fixed offsets from the frame pointer; stack heap allocation doesn't move them. No special code is needed to free dynamically allocated stack memory. The function epilogue resets the stack pointer and removes the entire stack frame, including the heap, from the stack. Naturally, a program should not reference heap objects after they have gone out of scope.

# 4 OBJECT FILES

ELF Header	4-1
Machine Information	4-1
Sections	4-2
Special Sections	4-2
Symbol Table	4-3
Symbol Values	4-3
Relocation	4-4
Relocation Types	4-4

i

**Table of Contents** 



#### **ELF Header**

#### **Machine Information**

For file identification in e\_ident, the Motorola 68000 Family requires the following values.

Figure 4-1: Motorola 68000 Family Identification, e ident

Position		Value
е	ident[EI CLASS]	ELFCLASS32
e	ident[EI_DATA]	ELFDATA2MSB

The ELF header's e\_flags member holds bit flags associated with the file. Motorola 68000 Family defines no flags; so this member contains zero. Processor identification resides in the ELF header's e\_machine member and must have the value 4, defined as the name EM 68K.

OBJECT FILES 4-1

#### **Sections**

#### **Special Sections**

Various sections hold program and control information. Sections in the list below are used by the system and have the indicated types and attributes.

Figure 4-2: Special Sections

Name	Туре	Attributes	
.got	SHT PROGBITS	SHF_ALLOC + SHF_WRITE	
.plt	SHT_PROGBITS	SHF_ALLOC + SHF_EXECINSTR	

- .got This section holds the global offset table. See "Coding Examples" in Chapter 3 and "Global Offset Table" in Chapter 5 for more information.
- .plt This section holds the procedure linkage table. See "Procedure Linkage Table" in Chapter 5 for more information.

## **Symbol Table**

#### **Symbol Values**

If an executable file contains a reference to a function defined in one of its associated shared objects, the symbol table section for that file will contain an entry for that symbol. The st\_shndx member of that symbol table entry contains SHN\_UNDEF. This signals to the dynamic linker that the symbol definition for that function is not contained in the executable file itself. If that symbol has been allocated a procedure linkage table entry in the executable file, and the st\_value member for that symbol table entry is non-zero, the value will contain the virtual address of the first instruction of that procedure linkage table entry. Otherwise, the st\_value member contains zero. This procedure linkage table entry address is used by the dynamic linker in resolving references to the address of the function. See "Function Addresses" in Chapter 5 for details.

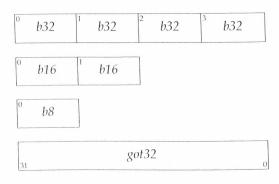
OBJECT FILES 4-3

## Relocation

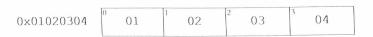
# **Relocation Types**

Relocation entries describe how to alter the following instruction and data fields (bit numbers appear in the lower box corners; byte numbers appear in the upper box corners).

Figure 4-3: Relocatable Fields



This specifies a 32-bit field occupying 4 bytes with arbitrary alignment. These values use the byte order illustrated below.



b16 This specifies a 16-bit field occupying 2 bytes with arbitrary alignment.



*b8* This specifies a 8-bit field occupying 1 byte with arbitrary alignment.



got32 This specifies a 32-bit field occupying 4 bytes with long word alignment. These bytes represent values in the same byte order as *b*32.

Calculations below assume the actions are transforming a relocatable file into either an executable or a shared object file. Conceptually, the link editor merges one or more relocatable files to form the output. It first decides how to combine and locate the input files, then updates the symbol values, and finally performs the relocation. Relocations applied to executable or shared object files are similar and accomplish the same result. Descriptions below use the following notation.

- A This means the addend used to compute the value of the relocatable field.
- B This means the base address at which a shared object has been loaded into memory during execution. Generally, a shared object file is built with a 0 base virtual address, but the execution address will be different.
- This means the address of the global offset table entry that will contain the address of the relocation entry's symbol during execution.

  See "Coding Examples" in Chapter 3 and "Global Offset Table" in Chapter 5 for more information.

OBJECT FILES 4-5

#### Relocation

- G' This means the address of entry zero in the global offset table.
- This means the place (section offset or address) of the procedure linkage table entry for a symbol. A procedure linkage table entry redirects a function call to the proper destination. The link editor builds the initial procedure linkage table, and the dynamic linker modifies the entries during execution. See "Procedure Linkage Table" in Chapter 5 for more information.
- L' This means the address of entry zero in the procedure linkage table.
- P This means the place (section offset or address) of the storage unit being relocated (computed using  $r_offset$ ).
- This means the value of the symbol whose index resides in the relocation entry.

A relocation entry's r\_offset value designates the offset or virtual address of the first byte of the affected storage unit. The relocation type specifies which bits to change and how to calculate their values. Because the Motorola 68000 Family uses only E1f32\_Rela relocation entries, the relocation table entry holds the addend. In all cases, the addend and the computed result use the same byte order.

Figure 4-4: Relocation Types

Name	Value	Field	Calculation
R 68K NONE	0	none	попе
R 68K 32	1	b32	S + A
R 68K 16	2	b16	S + A
R 68K 8	3	<i>b8</i>	S + A
R 68K PC32	4	<i>b</i> 32	S + A - P
R 68K PC16	5	b16	S + A - P
R 68K PC8	6	<i>b8</i>	S + A - P
R 68K GOT32	7	<i>b</i> 32	G + A - P
R 68K GOT16	8	b16	G + A - P
R 68K GOT8	9	<i>b8</i>	G + A - P
R 68K GOT320	10	<i>b</i> 32	G - G'
R 68K GOT160	11	b16	G - G'
R 68K GOT8O	12	<i>b8</i>	G - G'
R 68K PLT32	13	<i>b</i> 32	L + A - P
R 68K PLT16	14	b16	L + A - P
R 68K PLT8	15	<i>b</i> 8	L + A - P
R 68K PLT320	16	<i>b</i> 32	L - L'
R 68K PLT160	17	b16	L - L'
R 68K PLT80	18	b8	L - L'
R 68K COPY	19	none	none
R 68K GLOB DAT	20	got32	S
R 68K JMP SLOT	21	got32	S
R_68K_RELATIVE	22	b32	B + A

Some relocation types have semantics beyond simple calculation.

R 68K GOT32

This relocation type resembles R\_68K\_PC32, except it refers to the address of the symbol's global offset table entry and additionally instructs the link editor to build a global offset table.

OBJECT FILES 4-7

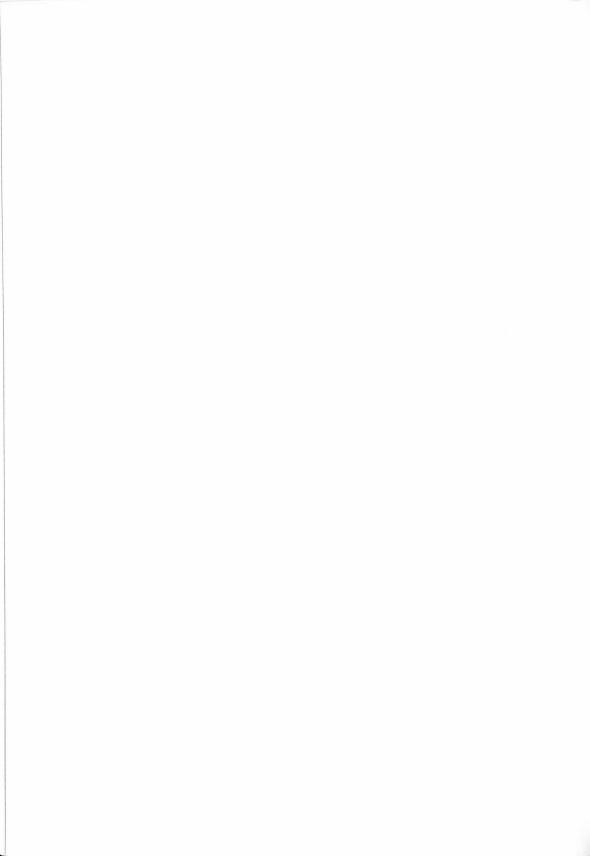
R_68K_GOT16	This relocation type resembles R_68K_PC16, except it refers to the address of the symbol's global offset table entry and additionally instructs the link editor to build a global offset table.
R_68K_GOT8	This relocation type resembles R_68K_PC8, except it refers to the address of the symbol's global offset table entry and additionally instructs the link editor to build a global offset table.
R_68K_GOT32O	This relocation type resembles R_68K_GOT32, except it refers to the address of the symbol's global offset table entry relative to the address of entry zero in the GOT.
R_68K_GOT160	This relocation type resembles R_68K_GOT16, except it refers to the address of the symbol's global offset table entry relative to the address of entry zero in the GOT.
R_68K_GOT8O	This relocation type resembles R_68K_GOT8, except it refers to the address of the symbol's global offset table entry relative to the address of entry zero in the GOT.
R_68K_PLT32	This relocation type resembles R_68K_PC32, except it refers to the address of the symbol's procedure linkage table entry and additionally instructs the link editor to build a procedure linkage table.
R_68K_PLT16	This relocation type resembles R_68K_PC16, except it refers to the address of the symbol's procedure linkage table entry and additionally instructs the link editor to build a procedure linkage table.
R_68K_PLT8	This relocation type resembles R_68K_PC8, except it refers to the address of the symbol's procedure linkage table entry and additionally instructs the link editor to build a procedure linkage table.
R_68K_PLT320	This relocation type resembles R_68K_PLT32, except it refers to the address of the symbol's procedure linkage table entry relative to the address of entry zero in the PLT.
R_68K_PLT160	This relocation type resembles R_68K_PLT16, except it refers to the address of the symbol's procedure linkage table entry relative to the address of entry zero in the PLT.

This relocation type resembles R 68K PLT8, except it refers to R 68K PLT8O the address of the symbol's procedure linkage table entry relative to the address of entry zero in the PLT. R 68K COPY This relocation type assists dynamic linking. Its offset member refers to a location in a writable segment. The symbol table index specifies a symbol that should exist both in the current object file and in a shared object. During execution, the dynamic linker copies data associated with the shared object's symbol to the location specified by the offset. This relocation type resembles R 68K 32, except it is used to R 68K GLOB DAT set a global offset table entry to the specified symbol's value. The relocation type allows one to determine the correspondence between symbols and global offset table entries. The relocated field should be aligned on a long word boundary. This relocation type does *not* use the addend. This relocation type assists dynamic linking. Its offset R 68K JMP SLOT member gives the location of a global offset table entry. This relocation type does not use the addend. This relocation type assists dynamic linking. The addend R 68K RELATIVE

member contains a value representing a relative address within a shared object. The offset member gives a location within the shared object for the final virtual address. The dynamic linker computes the corresponding virtual address by adding the virtual address at which the shared object was loaded to the relative address. Relocation entries for this

type must specify 0 for the symbol table index.

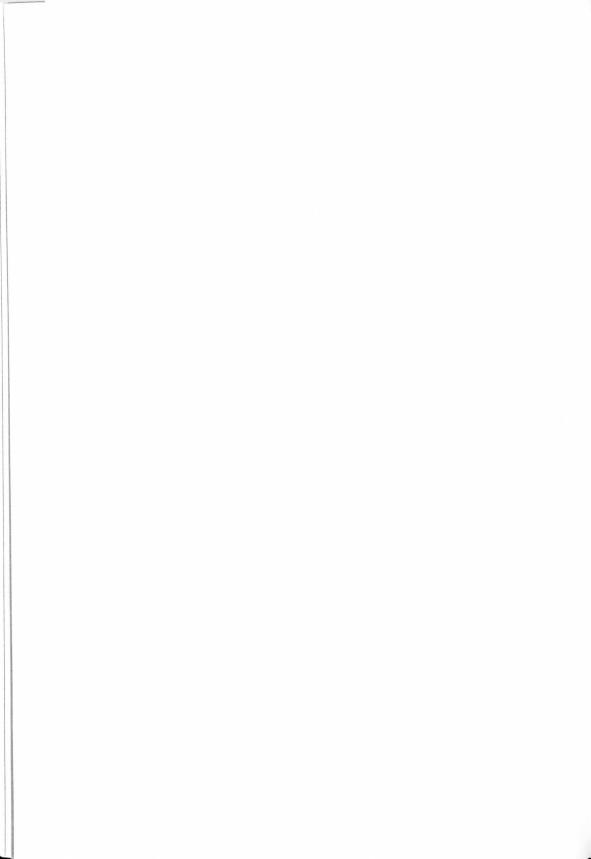
OBJECT FILES 4-9



# 5 PROGRAM LOADING AND DYNAMIC LINKING

Program Loading	5-1
Dynamic Linking	5-5
Dynamic Section	5-5
Global Offset Table	5-5
Function Addresses	5-6
Procedure Linkage Table	5-7

Table of Contents



## **Program Loading**

As the system creates or augments a process image, it logically copies a file's segment to a virtual memory segment. When—and if—the system physically reads the file depends on the program's execution behavior, system load, etc. A process does not require a physical page unless it references the logical page during execution, and processes commonly leave many pages unreferenced. Therefore delaying physical reads frequently obviates them, improving system performance. To obtain this efficiency in practice, executable and shared object files must have segment images whose file offsets and virtual addresses are congruent, modulo the page size.

Virtual addresses and file offsets for Motorola 68000 Family segments are congruent modulo 8 K ( $0\times2000$ ) or larger powers of 2. Because 8 KB is the maximum page size, the files will be suitable for paging regardless of physical page size. Figure 5-1 is an example of an executable file.

Figure 5-1: Executable File

File	Virtual Address
ELF header	
Program header table	
Other information	
Text segment	0x80000100
0x2bo00 bytes	0x8002beff
0x2be00 bytes	UXOUUZDEII
Data segment 	0x8004bf00
0x4e00 bytes	0x80050cff
Other information	
	ELF header Program header table Other information Text segment 0x2be00 bytes Data segment 0x4e00 bytes

Figure 5-2: Program Header Segments

Member	Text	Data
p_type p_offset p_vaddr p_paddr p_filesz p_memsz p_flags p_align	PT_LOAD	PT_LOAD

Although the example's file offsets and virtual addresses are congruent modulo 8 K for both text and data, up to four file pages hold impure text or data (depending on page size and file system block size).

- The first text page contains the ELF header, the program header table, and other information.
- The last text page holds a copy of the beginning of data.
- The first data page has a copy of the end of text.
- The last data page may contain file information not relevant to the running process.

Logically, the system enforces the memory permissions as if each segment were complete and separate; segments' addresses are adjusted to ensure each logical page in the address space has a single set of permissions. In the example above, the region of the file holding the end of text and the beginning of data will be mapped twice: at one virtual address for text and at a different virtual address for data.

The end of the data segment requires special handling for uninitialized data, which the system defines to begin with zero values. Thus if a file's last data page includes information not in the logical memory page, the extraneous data must be set to zero, not the unknown contents of the executable file. "Impurities" in

the other three pages are not logically part of the process image; whether the system expunges them is unspecified. The memory image for this program follows, assuming  $4~{\rm KB}~(0{\times}1000)$  pages.

Figure 5-3: Process Image Segments

Virtual Address	Contents	Segment
0x80000000	Header padding 0×100 bytes	
0x80000100	Text segment	
	•••	Text
	0x2be00 bytes	
0x8002bf00	Data padding 0×100 bytes	
0x8004b000	Text padding 0xf00 bytes	
0x8004bf00	Data segment	
	0×4e00 bytes	Data
0x80032d00	Uninitialized data 0x1024 zero bytes	
0x80033d24	Page padding 0x2dc zero bytes	

One aspect of segment loading differs between executable files and shared objects. Executable file segments typically contain absolute code [see "Coding Examples" in Chapter 3]. To let the process execute correctly, the segments must reside at the virtual addresses used to build the executable file. Thus the system uses the p\_vaddr values unchanged as virtual addresses.

On the other hand, shared object segments typically contain position-independent code. This lets a segment's virtual address change from one process to another, without invalidating execution behavior. Though the system chooses virtual addresses for individual processes, it maintains the segments' *relative positions*. Because position-independent code uses relative addressing between segments, the difference between virtual addresses in memory must match the difference between virtual addresses in the file. The following table shows possible shared object virtual address assignments for several processes, illustrating constant relative positioning.

Figure 5-4: Example Shared Object Segment Addresses

Source	Text	Data	Base Address
File	0x200	0x2a400	0x0
Process 1	0xc0080200	0xc00aa400	0xc0080000
Process 2	0xc0082200	0xc00ac400	0xc0082000
Process 3	0xd00c0200	0xd00ea400	0xd00c0000
Process 4	0xd00c6200	0xd00f0400	0xd00c6000

# **Dynamic Linking**

#### **Dynamic Section**

Dynamic section entries give information to the dynamic linker. Some of this information is processor-specific, including the interpretation of some entries in the dynamic structure.

DT PLTGOT

On the 68000, this entry's d\_ptr member gives the address of the first entry in the global offset table. As mentioned below, the first three global offset table entries are reserved, and two are used to hold procedure linkage table information.

#### Global Offset Table

Position-independent code cannot, in general, contain absolute virtual addresses. Global offset tables hold absolute addresses in private data, thus making the addresses available without compromising the position-independence and sharability of a program's text. A program references its global offset table using position-independent addressing and extracts absolute values, thus redirecting position-independent references to absolute locations.

Initially, the global offset table holds information as required by its relocation entries [see "Relocation" in Chapter 4]. When the dynamic linker creates memory segments for a loadable object file, it processes the relocation entries, some of which will be type R\_68K\_GLOB\_DAT referring to the global offset table. The dynamic linker determines the associated symbol values, calculates their absolute addresses, and sets the appropriate memory table entries to the proper values. Although the absolute addresses are unknown when the link editor builds an object file, the dynamic linker knows the addresses of all memory segments and can thus calculate the absolute addresses of the symbols contained therein.

If a program requires direct access to the absolute address of a symbol, that symbol will have a global offset table entry. Because the executable file and shared objects have separate global offset tables, a symbol may appear in several tables. The dynamic linker processes all the global offset table relocations before giving control to any code in the process image, thus ensuring the absolute addresses

are available during execution.

The table's entry zero is reserved to hold the address of the dynamic structure, referenced with the symbol \_DYNAMIC. This allows a program, such as the dynamic linker, to find its own dynamic structure without having yet processed its relocation entries. This is especially important for the dynamic linker, because it must initialize itself without relying on other programs to relocate its memory image.

The dynamic linker may choose different memory segment addresses for the same shared object in different programs; it may even choose different library addresses for different executions of the same program. Nonetheless, memory segments do not change addresses once the process image is established. As long as a process exists, its memory segments reside at fixed virtual addresses.

A global offset table's format and interpretation are processor-specific. For the Motorola 68000 Family an offset into the table is an *unsigned* value, allowing only non-negative "subscripts" into the array of addresses.

#### **Function Addresses**

References to the address of a function from an executable file and the shared objects associated with it might not resolve to the same value. References from within shared objects will normally be resolved by the dynamic linker to the virtual address of the function itself. References from within the executable file to a function defined in a shared object normally will be resolved by the link editor to the address of the procedure linkage table entry for that function within the executable file.

To allow comparisons of function addresses to work as expected, if an executable file references a function defined in a shared object, the link editor will place the address of the procedure linkage table entry for that function in its associated symbol table entry. [See "Symbol Values" in Chapter 4]. The dynamic linker treats such symbol table entries specially. If the dynamic linker is searching for a symbol, and encounters a symbol table entry for that symbol in the executable file, it normally follows the rules below.

- 1. If the st\_shndx member of the symbol table entry is not SHN\_UNDEF, the dynamic linker has found a definition for the symbol and uses its st\_value member as the symbol's address.
- 2. If the st\_shndx member is SHN\_UNDEF and the symbol is of type STT\_FUNC and the st\_value member is not zero, the dynamic linker recognizes this entry as special and uses the st\_value member as the symbol's address.
- 3. Otherwise, the dynamic linker considers the symbol to be undefined within the executable file and continues processing.

Some relocations are associated with procedure linkage table entries. These entries are used for direct function calls rather than for references to function addresses. These relocations are not treated in the special way described above because the dynamic linker must not redirect procedure linkage table entries to point to themselves.

#### **Procedure Linkage Table**

Much as the global offset table redirects position-independent address calculations to absolute locations, the procedure linkage table redirects position-independent function calls to absolute locations. The link editor cannot resolve execution transfers (such as function calls) from one executable or shared object to another. Consequently, the link editor arranges to have the program transfer control to entries in the procedure linkage table. On the 68000, procedure linkage tables reside in shared text, but they use addresses in the private global offset table. The dynamic linker determines the destinations' absolute addresses and modifies the global offset table's memory image accordingly. The dynamic linker thus can redirect the entries without compromising the position-independence and sharability of the program's text. Executable files and shared object files have separate procedure linkage tables.

Figure 5-5: Initial Procedure Linkage Table

```
got_plus_4, - (%sp)
         mov.1
.PLTO:
                   ([got_plus_8])
          qmr
          nop
          nop
                   ([name1@GOTPC,%pc])
.PLT1:
          qmp
                   &offset, - (%sp)
          mov.1
                    .PLTO
          bra
                    ([name2@GOTPC,%pc])
.PLT2:
          qmr
                   &offset, - (%sp)
          mov.l
                    .PLTO
          bra
```

Following the steps below, the dynamic linker and the program "cooperate" to resolve symbolic references through the procedure linkage table and the global offset table.

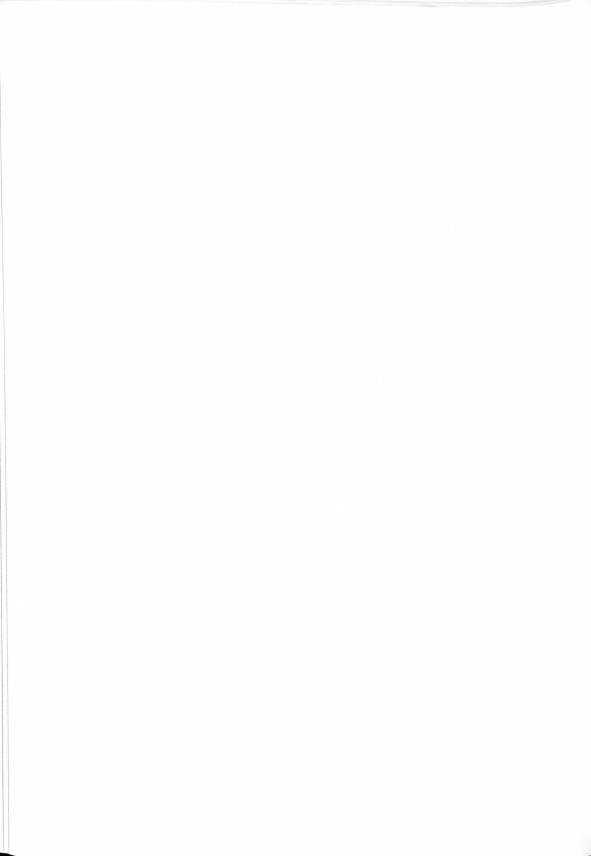
- When first creating the memory image of the program, the dynamic linker sets the second and the third entries in the global offset table to special values. Steps below explain more about these values.
- 2. For illustration, assume the program calls name1, which transfers control to the label .PLT1.
- 3. The first instruction jumps to the address in the global offset table entry for name1. Initially, the global offset table holds the address of the following instruction, not the real address of name1.
- 4. Consequently, the program pushes a relocation offset (offset) on the stack. The relocation offset is a 32-bit, non-negative byte offset into the relocation table. The designated relocation entry will have type R\_68K\_JMP\_SLOT, and its offset will specify the global offset table entry used in the previous jmp instruction. The relocation entry also contains a symbol table index, thus telling the dynamic linker what symbol is being referenced, name1 in this case.

- 5. After pushing the relocation offset, the program then jumps to .PLTO, the first entry in the procedure linkage table. The mov.l instruction places the value of the second global offset table entry (got\_plus\_4) on the stack, thus giving the dynamic linker one long word of identifying information. The program then jumps to the address in the third global offset table entry (got\_plus\_8), which transfers control to the dynamic linker.
- 6. When the dynamic linker receives control, it unwinds the stack, looks at the designated relocation entry, finds the symbol's value, stores the "real" address for name1 in its global offset table entry, and transfers control to the desired destination.
- 7. Subsequent executions of the procedure linkage table entry will transfer directly to name1, without calling the dynamic linker a second time. That is, the jmp instruction at .PLT1 will transfer to name1, instead of "falling through" to the next instruction.

The LD\_BIND\_NOW environment variable can change dynamic linking behavior. If its value is non-null, the dynamic linker evaluates procedure linkage table entries before transferring control to the program. That is, the dynamic linker processes relocation entries of type R\_68K\_JMP\_SLOT during process initialization. Otherwise, the dynamic linker evaluates procedure linkage table entries lazily, delaying symbol resolution and relocation until the first execution of a table entry.

NOTE

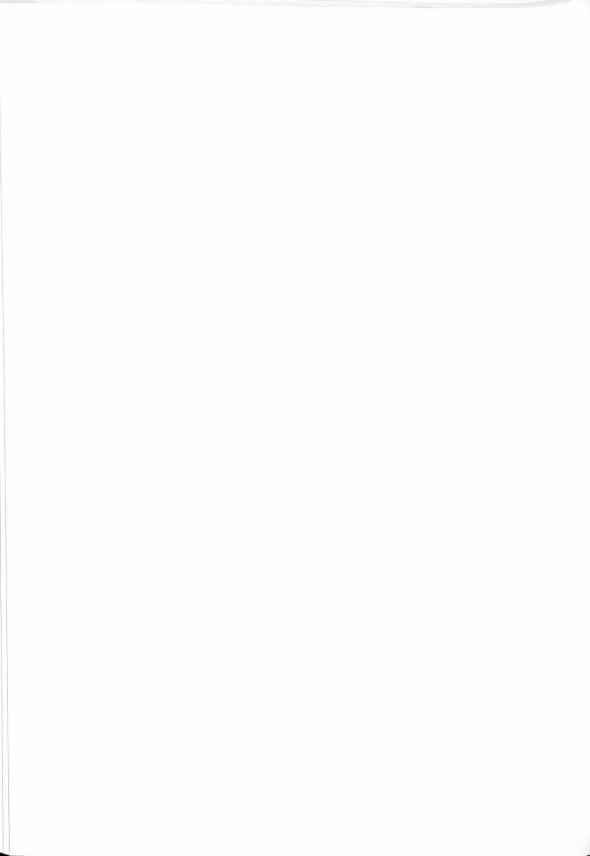
Lazy binding generally improves overall application performance, because unused symbols do not incur the dynamic linking overhead. Nevertheless, two situations make lazy binding undesirable for some applications. First, the initial reference to a shared object function takes longer than subsequent calls, because the dynamic linker intercepts the call to resolve the symbol. Some applications cannot tolerate this unpredictability. Second, if an error occurs and the dynamic linker cannot resolve the symbol, the dynamic linker will terminate the program. Under lazy binding, this might occur at arbitrary times. Once again, some applications cannot tolerate this unpredictability. By turning off lazy binding, the dynamic linker forces the failure to occur during process initialization, before the application receives control



# 6 LIBRARIES

System Library	6-1
Additional Entry Points	6-1
Support Routines	6-1
Global Data Symbols	6-2
Application Constraints	6-2
C Library	6-3
Additional Support Routines	6-3
System Data Interfaces	6-4
Data Definitions	6-4

Table of Contents



## **System Library**

#### **Additional Entry Points**

There are no additional entry points required by the Motorola 68000 Family Processor Supplement.

#### **Support Routines**

Besides operating system services, **libsys** contains the following processor-specific support routine. The routine is also accessible named with a leading underscore.

Figure 6-1: libsys Support Routines

sbrk

char \*sbrk(int incr);

This function adds incr bytes to the break value and changes the allocated space accordingly. Incr can be negative, in which case the amount of allocated space is decreased. The break value is the address of the first location beyond the end of the data segment. The amount of allocated space increases as the break value increases. Newly allocated space is set to zero. If, however, the same memory space is reallocated to the same process, its contents are undefined. Upon successful completion, sbrk returns the old break value. Otherwise, it returns –1 and sets errno to indicate the error.

LIBRARIES 6-1

#### **Global Data Symbols**

The **libsys** library requires that some global external data objects be defined for the routines to work properly. In addition to the corresponding data symbols listed in the **System V ABI**, the following symbols must be provided in the system library on all ABI-conforming systems implemented with the Motorola 68000 Family processor architecture. Declarations for the data objects listed below can be found in the Data Definitions section of this chapter or immediately following the table.

Figure 6-2: libsys, Global External Data Symbols

\_flt\_rounds \_ huge val

#### **Application Constraints**

As described above, **libsys** provides symbols for applications. In a few cases, however, an executable is obliged to provide symbols for the library. In addition to the application-provided symbols listed in this section of the **System V ABI**, conforming applications on the Motorola 68000 Family processor architecture are also required to provide the following symbols.

extern end;

This symbol refers neither to a routine nor to a location with interesting contents. Instead, its address must correspond to the beginning of a program's dynamic allocation area, called the heap. Typically, the heap begins immediately after the data segment of the program's executable file. This value is normally provided by the static linker.

extern const int lib version;

This variable's value specifies the compilation and execution mode for the program. If the value is zero, the program wants to preserve the semantics of older (pre-ANSI) C, where conflicts exist with ANSI. Otherwise, the value is non-zero, and the program wants ANSI C semantics. This value is normally provided by the compiler.

# **C** Library

## **Additional Support Routines**

There are no additional support routines required by the Motorola 68000 Family Processor Supplement.

LIBRARIES 6-3

# **System Data Interfaces**

#### **Data Definitions**

This section contains standard header files that describe system data. These files are referred to by their names in angle brackets: <name.h> and <sys/name.h>. Included in these headers are macro definitions and data definitions.

The data objects described in this section are part of the interface between an **ABI**-conforming application and the underlying **ABI**-conforming system where it will run. While an **ABI**-conforming system must provide these interfaces, it is not required to contain the actual header files referenced here.

ANSI C serves as the ABI reference programming language, and data definitions are specificed in ANSI C format. The C language is used here as a convenient notation. Using a C language description of these data objects does *not* preclude their use by other programming languages.

Figure 6-3: <assert.h>

```
extern void __assert(const char *, const char *, int);
#define assert(EX) \
   (void)((EX)||(_assert(#EX, __FILE__, __LINE__), 0))
```

Figure 6-4: <ctype.h>

```
#define _U
          01
#define L 02
#define N 04
#define S 010
#define P 020
#define C 040
#define B 0100
#define X
           0200
extern unsigned char __ctype[];
#define isalpha(c) ((__ctype+1)[c]&(_U|_L))
\#define isupper(c) ((__ctype+1)[c]&_U)
#define islower(c) ((__ctype+1)[c]&_L)
#define isdigit(c) ((__ctype+1)[c]&_N)
#define isxdigit(c) ((__ctype+1)[c]&_X)
#define isprint(c) ((__ctype+1)[c]&(_P|_U|_L|_N|_B))
#define isgraph(c) (( ctype+1)[c]&(P|U|L|N))
#define iscntrl(c) ((__ctype+1)[c]&_C)
#define isascii(c) (!(c)&~0177))
#define _toupper(c) ((__ctype+258)[c])
#define _tolower(c) ((__ctype+258)[c])
#define toascii(c) ((c)&0177)
```

LIBRARIES 6-5

#### Figure 6-5: <dirent.h>

```
struct dirent {
    ino_t d_ino;
    off t d off;
    unsigned short d_reclen;
    char d_name[1];
};
```

Figure 6-6: <errno.h>, Part 1 of 4

```
extern int errno;
                         . 1
#define EPERM
#define ENOENT
#define ESRCH
#define EINTR
#define EIO
#define ENXIO 6
#define E2BIG 7
#define ENOEXEC 8
#define ERADE 9
#define EBADF
#define ECHILD 10
#define EAGAIN 11
#define ENOMEM
#define ENOMEM 12
#define EACCES 13
#define EFAULT 14
#define ENOTBLK 15
#define EBUSY 16
#define EEXIST 17
#define EXDEV 18
#define ENODEV 19
#define ENOTDIR 20
#define EISDIR 21
#define EISDIR 21
#define EINVAL 22
                              22
#define EINVAL
                              23
#define ENFILE
#define EMFILE
                              24
#define ENOTTY
                              25
#define ETXTBSY
#define EFBIG
                              27
                              28
#define ENOSPC
#define ESPIPE
                               29
#define EROFS
                               30
                               31
#define EMLINK
#define EPIPE
                               32
```

LIBRARIES 6-7

Figure 6-7: <errno.h>, Part 2 of 4

```
#define EDOM 33
#define ERANGE 34
#define ENOMSG 35
#define ELDRM 36
#define ECHRNG 37
#define ELZNSYNC 38
#define ELJSHLT 39
#define ELJSHLT 40
#define ELNRNG 41
#define ENONCH 42
#define ENOCSI 43
#define ENOCSI 43
#define EDEADLK 45
#define ENOLCK 46
#define ENONCT 61
#define ENONCT 62
#define ENONCT 63
#define ENONCT 64
#define ENONCT 64
#define ENONCT 65
#define ENONCT 66
#define ENONCT 66
#define ENONCT 66
#define ENOLINK 67
#define ERADV 68
#define ERADV 68
#define ERADV 68
#define ESONMT 69
#define EROCMM 70
#define EPROTO 71
```

Figure 6-8: <errno.h>, Part 3 of 4

```
#define EMULTIHOP
#define EBADMSG
                      77
#define ENAMETOOLONG 78
#define EOVERFLOW 79
                      80
#define ENOTUNIQ
#define EBADFD 81
#define EREMCHG 82
#define ENOSYS 89
#define ELOOP 90
#define ERESTART 91
#define ESTRPIPE 92
#define EBADFD
                      81
                      158
#define ENOTEMPTY
#define ESTALE
                      162
/* The following errno values are optional. */
#define EWOULDBLOCK EDEADLK
#define EBADE
#define EBADR
                      51
                      52
#define EXFULL
                      53
#define ENOANO
#define EBADRQC
                       54
#define EBADSLT
                       55
#define EDEADLOCK
                       56
#define EBFONT
                       57
                       76
#define EDOTDOT
                      83
#define ELIBACC
#define ELIBBAD
                     84
#define ELIBSCN
                      85
```

Figure 6-9: <errno.h>, Part 4 of 4

```
#define ELIBMAX
#define ELIBEXEC 87
#define EINPROGRESS 128
#define EALREADY
#define ENOTSOCK 130
#define EDESTADDRREQ 131
#define EMSGSIZE 132
#define EPROTOTYPE 133
#define ENOPROTOOPT 134
                           135
#define EPROTONOSUPPORT
#define ESOCKTNOSUPPORT
                           136
#define EOPNOTSUPP 137
#define EPFNOSUPPORT 138
#define EAFNOSUPPORT
#define EADDRINUSE
#define EADDRNOTAVAIL 141
 #define ENETDOWN 142
 #define ENETUNREACH 143
 #define ENETRESET
 #define ECONNABORTED 145
 #define ECONNRESET 146
 #define ENOBUFS
                     147
                     148
 #define EISCONN
 #define EISCONN 148
#define ENOTCONN 149
#define ESHUTDOWN 150
 #define ETOOMANYREFS 151
 #define ETIMEDOUT 152
 #define ECONNREFUSED 153
 #define EHOSTDOWN 156
 #define EHOSTUNREACH 157
 #define EPROCLIM 159
 #define EUSERS 160
#define EDQUOT 161
 #define EPOWERFAIL 163
```

Figure 6-10: <fcntl.h>, Part 1 of 2

```
#define O_ROWLY 1
#define O_ROWR
                          0
#define O_NDELAY 04
#define O_APPEND 010
#define O_SYNC 020
#define O NONBLOCK 0100
#define O_CREAT 00400
#define O_TRUNC 01000
#define O_EXCL 02000
#define O_NOCTTY 04000
#define F DUPFD
#define F GETFD
                         1
#define F_SETFD
                          2
#define F_GETFL
                         3
#define F_SETFL
                         4
#define F_GETLK
                        14
#define F SETLK
                         6
#define F SETLKW
                         7
#define F FREESP
                        11
#define FD CLOEXEC
                         1
#define O ACCMODE
                         03
```

## Figure 6-11: <fcntl.h>, Part 2 of 2

```
typedef struct flock {
    short 1_type;
    short 1_whence;
    off_t 1_start;
    off_t 1_len;
    long 1_sysid;
    pid_t 1_pid;
    long pad[4];
} flock_t;

#define F_RDLCK 01
#define F_WRLCK 02
#define F_UNLCK 03
```

# Figure 6-12: <float.h>

```
extern int __flt_rounds;
#define FLT_ROUNDS __flt_rounds
```

Figure 6-13: <fmtmsg.h>

```
#define MM NULL
                             0L
#define MM_HARD 0x00000001L
#define MM_SOFT 0x00000002L
#define MM_FIRM 0x00000004L
#define MM RECOVER 0x00000100L
#define MM_NRECOV 0x00000200L
#define MM_APPL 0x00000008L
#define MM_UTIL 0x00000010L
#define MM_OPSYS 0x00000020L
#define MM_PRINT 0x00000040L
#define MM CONSOLE 0x00000080L
#define MM_NOSEV
#define MM HALT
#define MM ERROR
#define MM_WARNING
                             3
#define MM INFO
#define MM_NULLLBL ((char *) 0)
#define MM_NULLSEV MM_NOSEV
#define MM_NULLTXT 0L
#define MM_NULLTXT ((char *) 0)
#define MM NULLACT ((char *) 0)
#define MM_NULLTAG ((char *) 0)
#define MM NOTOK
                            -1
#define MM OK
                           0 \times 00
#define MM NOMSG
                            0 \times 01
#define MM NOCON
                            0 \times 04
```

#### Figure 6-14: <ftw.h>

### Figure 6-15: <grp.h>

```
struct group {
    char *gr_name;
    char *gr_passwd;
    gid_t gr_gid;
    char **gr_mem;
};
```

### Figure 6-16: <sys/ipc.h>

```
struct ipc_perm {
                  uid;
gid;
cuid;
cgid;
         uid t
          gid_t
          uid t
          gid_t
         gid_t cgid;
mode_t mode;
          unsigned long seq;
          key t
                           key;
          long
                            pad[4];
};
#define IPC_CREAT 0001000
#define IPC_EXCL 0002000
#define IPC_NOWAIT 0004000
#define IPC_ALLOC 0100000
#define IPC PRIVATE
                            (key t)0
#define IPC RMID
                             10
#define IPC SET
                             11
#define IPC_STAT
                             12
```

Figure 6-17: <langinfo.h>, Part 1 of 2

```
#define DAY 1
#define DAY 2 2
#define DAY 3 3
#define DAY 4 4
#define DAY_6 6 #define DAY_7 7
#define ABDAY 1 8
#define ABDAY_2 9
#define ABDAY_3 10
#define ABDAY_4 11
#define ABDAY 5 12
#define ABDAY 6 13
#define ABDAY 7 14
#define MON 1 15
#define MON 2 16
#define MON 3 17
#define MON 4 18
#define MON 5 19
#define MON 6 20
#define MON 7 21
#define MON 8 22
 #define MON_9 23
 #define MON_10 24
 #define MON_11 25
 #define MON_12 26
```

Figure 6-18: <langinfo.h>, Part 2 of 2

```
#define ABMON 1
                  27
#define ABMON 2
                  28
                  29
#define ABMON 3
                   30
#define ABMON 4
                   31
#define ABMON 5
                  32
33
34
35
#define ABMON 6
#define ABMON 7
#define ABMON 8
#define ABMON 9
                  36
#define ABMON 10
                37
#define ABMON 11
                  38
#define ABMON 12
#define RADIXCHAR 39
#define THOUSEP
#define YESSTR
                  41
#define NOSTR
                  42
#define CRNCYSTR
                  43
                 44
#define D_T_FMT
                  45
#define D FMT
                  46
#define T FMT
#define AM_STR
                   47
#define PM STR
                   48
```

Figure 6-19: imits.h>

```
#define MB LEN MAX 5
#undef ARG MAX
#undef CHILD MAX
#undef MAX CANON
#undef NGROUPS MAX
#undef LINK MAX
#undef NAME MAX
#undef OPEN_MAX
#undef PASS MAX
#undef PATH MAX
#undef PIPE BUF
#undef MAX INPUT
/* the #undef-fed values vary and should be
       retrieved using sysconf() or pathconf() */
#define _POSIX_ARG_MAX
                           4096
#define POSIX_CHILD_MAX
#define POSIX LINK MAX
#define POSIX MAX CANON
                           255
#define POSIX MAX INPUT
                           255
#define POSIX_NAME_MAX
                           14
#define POSIX_NGROUPS_MAX 0
#define POSIX OPEN MAX 16
#define POSIX PATH MAX
                           255
#define POSIX PIPE BUF
                           512
#define NL ARGMAX
#define NL LANGMAX 14
#define NL MSGMAX 32767
#define NL NMAX 1
#define NL SETMAX 255
#define NL_TEXTMAX 255
                     20
 #define NZERO
#define TMP_MAX 17576
#define FCHAR_MAX 1048576
```

Figure 6-20: <locale.h>

```
struct lconv {
       char
             *decimal point;
       char *thousands sep;
       char *grouping;
       char
             *int curr symbol;
       char
             *currency_symbol;
       char
             *mon_decimal_point;
       char
             *mon_thousands_sep;
       char
             *mon_grouping;
       char
             *positive sign;
       char
             *negative sign;
       char
             int frac digits;
       char
             frac digits;
       char p cs precedes;
       char p_sep_by_space;
       char n_cs_precedes;
       char n_sep_by_space;
       char p_sign_posn;
       char n sign posn;
} lconv;
#define LC CTYPE
#define LC NUMERIC
#define LC TIME
#define LC COLLATE 1
#define LC MONETARY 4
#define LC MESSAGES 5
#define LC ALL
                    6
#define NULL
```

### Figure 6-21: <math.h>

```
typedef union _h_val {
    unsigned long i[2];
    double d;
} _h_val;

extern const _h_val __huge_val;
#define HUGE_VAL __huge_val.d
```

# Figure 6-22: <sys/mman.h>

```
#define PROT_READ 0x1
#define PROT_WRITE 0x2
#define PROT_EXEC 0x4
#define PROT_NONE 0x0

#define MAP_SHARED 1
#define MAP_PRIVATE 2
#define MAP_FIXED 0x10

#define MS_SYNC 0x0
#define MS_SYNC 0x1
#define MS_ASYNC 0x1
#define MS_INVALIDATE 0x2
```

## Figure 6-23: <mon.h>

```
struct hdr {
    char *lpc;
    char *hpc;
    int nfns;
};

struct cnt {
    char *fnpc;
    long mcnt;
};

typedef unsigned short WORD;
```

## Figure 6-24: <sys/mount.h>

```
#define MS_RDONLY 0x01
#define MS_DATA 0x04
#define MS_NOSUID 0x10
#define MS_REMOUNT 0x20
```

Figure 6-25: <sys/msg.h>

Figure 6-26: <netconfig.h>, Part 1 of 2

```
struct netconfig {
                 *nc netid;
      char
      unsigned long nc semantics;
      unsigned long nc flag;
      unsigned long nc nlookups;
      char **nc lookups;
      unsigned long nc_unused[8];
};
#define NC TPI CLTS
                       1
#define NC_TPI_COTS
                      2
#define NC_TPI_COTS_ORD 3
#define NC TPI RAW
                      4
#define NC NOFLAG
                      00
#define NC VISIBLE
                      01
#define NC BROADCAST
```

Figure 6-27: <netconfig.h>, Part 2 of 2

```
#define NC_NOPROTOFMLY "-"
#define NC_LOOPBACK "loopback"
#define NC_INET "inet"
#define NC_IMPLINK "implink"
#define NC_PUP "pup"
#define NC_CHAOS "chaos"
#define NC_NS "ns"
#define NC_ECMA "ecma"
#define NC_ECMA "ecma"
#define NC_CITT "ccitt"
#define NC_DECNET "decnet"
#define NC_DECNET "decnet"
#define NC_LAT "lat"
#define NC_HYLINK "hylink"
#define NC_NIT "nit"
#define NC_NIT "nit"
#define NC_IEEE802 "ieee802"
#define NC_OSI "osi"
#define NC_OSIP "gosip"
#define NC_GOSIP "gosip"
#define NC_NOPROTO "-"
#define NC_NOPROTO "-"
#define NC_ICMP "icmp"
```

Figure 6-28: <netdir.h>

```
struct nd addrlist {
        int n cnt;
          struct netbuf *n addrs;
struct nd hostservlist {
        int h_cnt;
         struct nd_hostserv *h_hostservs;
};
struct nd hostserv {
        char *h host;
         char *h_serv;
#define ND_BADARG
                             -2
-1
0
                                  -2
#define ND_NOMEM
                                  -1
#define ND OK
                            1
2
3
#define ND NOHOST
#define ND_NOSERV
#define ND_NOSYM
#define ND OPEN
                                 4
                           5
6
#define ND_ACCESS
#define ND UKNWN
#define ND_NOCTRL 7
#define ND_FAILCTRL 8
#define ND_SYSTEM 9
#define ND_HOSTSERV 0
#define ND_HOSTSERVLIST 1
#define ND_ADDR 2
                                 2
#define ND ADDR 2
#define ND ADDRLIST 3
#define ND_SET_BROADCAST 1
#define ND_SET_RESERVEDPORT 2
#define ND CHECK RESERVEDPORT 3
#define ND MERGEADDR 4
#define HOST_SELF "\\1"
#define HOST_ANY "\\2"
#define HOST_BROADCAST "\\3"
```

## Figure 6-29: <nl types.h>

```
#define NL_SETD 1

typedef short nl_item;

typedef void *nl_catd;
```

## Figure 6-30: <sys/param.h>

```
#define HZ sysconf(3)

#define NGROUPS_UMIN 0

#define MAXPATHLEN 1024

#define MAXSYMLINKS 20

#define MAXNAMELEN 256

#define NADDR 13

#define NBBY 8

#define NBPSCTR 512
```

## **Figure 6-31:** <poll.h>

```
struct pollfd {
        int fd;
        short events;
        short revents;
};
#define POLLIN
                      0x0001
#define POLLPRI
#define POLLOUT
                      0x0002
                      0x0004
#define POLLRDNORM 0x0040
#define POLLWRNORM POLLOUT
#define POLLRDBAND 0x0080
#define POLLWRBAND 0x0100
#define POLLNORM
                      POLLRDNORM
#define POLLERR
#define POLLHUP
                        0x0008
                        0x0010
#define POLLNVAL
                        0x0020
```

Figure 6-32: <sys/procset.h>

```
#define P_INITPID 1
#define P_INITUID 0
#define P_INITPGID 0
#define P_MYID (-1)
typedef long id t;
typedef enum idtype {
      P PID,
       P PPID,
       P PGID,
       P SID,
       P CID,
       P UID,
       P GID,
       P ALL
} idtype t;
typedef enum idop {
       POP DIFF,
        POP AND,
        POP OR,
        POP XOR
} idop t;
typedef struct procset {
        } procset_t;
```

# Figure 6-33: <pwd.h>

```
struct passwd {
    char *pw_name;
    char *pw_passwd;
    uid_t pw_uid;
    gid_t pw_gid;
    char *pw_age;
    char *pw_comment;
    char *pw_gecos;
    char *pw_dir;
    char *pw_shell;
};
```

Figure 6-34: <sys/regset.h>

```
typedef int greg_t;
#define NGREG 18
typedef greg_t gregset_t[NGREG];
#define R D0 0
#define R D1 1
#define R_D2 2
#define R_D3 3
#define R_D4 4
#define R_D5 5
#define R_D6 6
#define R_D7 7
#define R A0 8
#define R_A1 9
#define R_A2 10
#define R_A3 11
#define R A4 12
#define R A5 13
#define R A6 14
#define R A7 15
#define R SP 15
#define R PC 16
#define R_PS 17
typedef struct fpregset {
        int f_pcr;
int f_psr;
        int f_fpiaddr;
int f_fpregs[8][3];
 } fpregset t;
```

Figure 6-35: <sys/resource.h>

```
#define RLIMIT_CPU 0
#define RLIMIT_FSIZE 1
#define RLIMIT_DATA 2
#define RLIMIT_STACK 3
#define RLIMIT_CORE 4
#define RLIMIT_NOFILE 5
#define RLIMIT_WMEM 6
#define RLIMIT_AS RLIMIT_VMEM

struct rlimit {
    rlim_t rlim_cur;
    rlim_t rlim_max;
};

typedef unsigned long rlim_t;
```

# 

```
#define MAX_AUTH_BYTES 400
#define MAXNETNAMELEN 255
#define HEXKEYBYTES 48
enum auth_stat {
       AUTH OK=0,
        AUTH BADCRED=1,
        AUTH REJECTEDCRED=2,
        AUTH BADVERF=3,
        AUTH REJECTEDVERF=4,
        AUTH_TOOWEAK=5,
        AUTH INVALIDRESP=6,
        AUTH_FAILED=7
};
union des_block {
        struct {
          unsigned long high;
          unsigned long low;
        } key;
        char c[8];
 };
```

## Figure 6-37: c.h>, Part 2 of 12

```
struct opaque_auth {
int oa flavor;
 char *oa base;
 unsigned int oa length;
typedef struct {
       struct opaque_auth ah_cred;
       struct opaque_auth ah_verf;
       union des block ah key;
       struct auth ops {
         void (*ah nextverf)();
         int (*ah_marshal)();
         int (*ah_validate)();
         int (*ah_refresh)();
         void (*ah_destroy)();
       } *ah ops;
       char *ah_private;
} AUTH;
struct authsys parms {
       unsigned long aup time;
       char *aup machname;
       uid_t aup_uid;
       gid_t aup_gid;
       unsigned int aup_len;
       gid t *aup gids;
};
```

# Figure 6-38: c.h>, Part 3 of 12

Figure 6-39: c.h>, Part 4 of 12

```
enum clnt stat {
       RPC SUCCESS=0,
       RPC_CANTENCODEARGS=1,
       RPC CANTDECODERES=2,
        RPC CANTSEND=3,
       RPC CANTRECV=4,
       RPC TIMEDOUT=5,
       RPC_INTR=18,
       RPC_VERSMISMATCH=6,
       RPC AUTHERROR=7,
       RPC PROGUNAVAIL=8,
       RPC PROGVERSMISMATCH=9,
       RPC_PROCUNAVAIL=10,
       RPC CANTDECODEARGS=11,
       RPC SYSTEMERROR=12,
       RPC UNKNOWNHOST=13,
       RPC UNKNOWNPROTO=17,
       RPC UNKNOWNADDR=19,
       RPC NOBROADCAST=21,
       RPC RPCBFAILURE=14,
       RPC_PROGNOTREGISTERED=15,
       RPC_N2AXLATEFAILURE=22,
       RPC_UDERROR=23,
       RPC TLIERROR=20,
       RPC FAILED=16
#define RPC PMAPFAILURE RPC RPCBFAILURE
```

Figure 6-40: c.h>, Part 5 of 12

```
#define _RPC_NONE
                           0
#define RPC_NETPATH
#define _RPC_VISIBLE
#define RPC_CIRCUIT_V
#define RPC DATAGRAM V 4
#define RPC CIRCUIT N 5
#define _RPC_DATAGRAM_N
#define _RPC_TCP
#define RPC UDP
#define RPC ANYSOCK -1
#define RPC_ANYFD RPC_ANYSOCK
struct rpc err {
       enum clnt_stat re_status;
       union {
              struct {
                int errno;
                int t errno;
              RE err;
              enum auth stat RE why;
              struct {
               unsigned long low;
                unsigned long high;
               } RE vers;
               struct {
               long s1;
                long s2;
               } RE 1b;
       } ru;
};
```

Figure 6-41: c.h>, Part 6 of 12

```
struct rpc_createerr {
enum clnt stat cf stat;
 struct rpc err cf error;
};
typedef struct {
       AUTH *cl auth;
       struct clnt ops {
         enum clnt_stat (*cl_call)();
         void (*cl abort)();
         void (*cl geterr)();
         int (*cl freeres)();
         void (*cl destroy)();
         int (*cl control)();
        } *cl ops;
       char *cl_private;
       char *cl netid;
       char *cl_tp;
} CLIENT;
#define FEEDBACK REXMIT1
#define FEEDBACK OK
#define CLSET TIMEOUT
                            1
#define CLGET_TIMEOUT 2
#define CLGET_SERVER_ADDR 3
#define CLGET FD
#define CLGET_SVC_ADDR
                           8
#define CLSET FD CLOSE
#define CLSET FD NCLOSE
                           9
#define CLSET RETRY TIMEOUT 4
#define CLGET RETRY TIMEOUT 5
```

### Figure 6-42: <rpc.h>, Part 7 of 12

```
extern struct
rpc createerr rpc createerr;
enum xprt stat {
 XPRT DIED,
 XPRT MOREREQS,
 XPRT IDLE
};
typedef struct {
        int xp_fd;
        unsigned short xp port;
        struct xp_ops {
          int (*xp recv)();
           enum xprt stat (*xp stat)();
           int (*xp getargs)();
            int (*xp reply) ();
            int (*xp freeargs)();
            void (*xp_destroy)();
        } *xp ops;
        int
              xp addrlen;
        char *xp_tp;
        char *xp_netid;
        struct netbuf xp ltaddr;
        struct netbuf xp rtaddr;
        char xp raddr[16];
        struct opaque auth xp verf;
        char *xp p1;
        char *xp p2;
        char *xp_p3;
 } SVCXPRT;
```

Figure 6-43: <pc.h>, Part 8 of 12

```
struct svc req {
        unsigned long rq prog;
        unsigned long rq_vers;
        unsigned long rq proc;
        struct opaque_auth rq_cred;
        char *rq_clntcred;
        SVCXPRT*rq_xprt;
};
extern fd set svc fdset;
typedef struct fdset {
       long fds_bits[32];
} fd set;
enum msg_type {
       CALL=0,
        REPLY=1
enum reply_stat {
       MSG ACCEPTED=0,
        MSG_DENIED=1
};
enum accept stat {
       SUCCESS=0,
       PROG UNAVAIL=1,
        PROG MISMATCH=2,
        PROC_UNAVAIL=3,
        GARBAGE_ARGS=4,
        SYSTEM ERR=5
};
```

## Figure 6-44: <pc.h>, Part 9 of 12

```
enum reject stat {
       RPC_MISMATCH=0,
        AUTH ERROR=1
};
struct accepted reply {
        struct opaque auth ar verf;
        enum accept stat ar stat;
        union {
         struct {
         unsigned long low;
         unsigned long high;
         } AR versions;
         struct {
          char *where;
          xdrproc t proc;
         } AR results;
         } ru;
};
struct rejected_reply {
        enum reject stat rj stat;
        union {
         struct {
         unsigned long low;
          unsigned long high;
         } RJ versions;
         enum auth_stat RJ_why;
         } ru;
 };
```

```
struct reply body {
        enum reply_stat rp_stat;
        union {
          struct accepted_reply RP_ar;
          struct rejected_reply RP_dr;
};
struct call body {
        unsigned long cb rpcvers;
        unsigned long cb prog;
        unsigned long cb vers;
        unsigned long cb proc;
        struct opaque auth cb cred;
        struct opaque_auth cb_verf;
};
struct rpc_msg {
        unsigned long rm xid;
        enum msg_type rm_direction;
         struct call body RM cmb;
         struct reply_body RM_rmb;
        } ru;
};
struct rpcb {
       unsigned long r prog;
        unsigned long r_vers;
       char *r netid;
       char *r addr;
       char *r_owner;
};
```

Figure 6-46: c.h>, Part 11 of 12

```
struct rpcblist {
 struct rpcb rpcb_map;
 struct rpcblist *rpcb_next;
};
enum xdr_op {
 XDR_ENCODE=0,
 XDR DECODE=1,
 XDR FREE=2
struct xdr_discrim {
 int value;
 xdrproc_t proc;
enum authdes_namekind {
 ADN FULLNAME,
  ADN NICKNAME
struct authdes fullname {
  char *name;
  union des block key;
  u_long window;
 struct authdes cred {
  enum authdes namekind adc namekind;
  struct authoes fullname adc fullname;
   unsigned long adc nickname;
 };
```

Figure 6-47: c.h>, Part 12 of 12

```
typedef struct {
       enum xdr op x op;
        struct xdr ops {
         int
                  (*x getlong)();
         int
                     (*x putlong)();
                     (*x getbytes)();
         int
                (*x_putbytes)();
         int.
         unsigned int (*x_getpostn)();
         int
                (*x setpostn)();
         long *
                      (*x inline)();
         void
                      (*x_destroy)();
        } *x ops;
        char x public;
        char
              x private;
        char x_base;
        int x handy;
} XDR;
typedef int (*xdrproc t)()
#define NULL xdrproc t ((xdrproc t)0)
#define auth destroy(auth)
                            ((*((auth)->ah ops->ah destroy))(auth))
#define clnt call(rh, proc, xargs, argsp, xres, resp, secs) \
        ((*(rh)->cl ops->cl call)(rh, proc, xarqs, arqsp, xres, resp, secs))
#define clnt freeres(rh, xres, resp) ((*(rh)->cl ops->cl freeres)(rh, xres, resp))
#define clnt_geterr(rh, errp) ((*(rh)->cl_ops->cl_geterr)(rh, errp))
#define clnt_control(cl, rq, in) ((*(cl)->cl_ops->cl_control)(cl, rq, in))
#define clnt_destroy(rh) ((*(rh)->cl_ops->cl_destroy)(rh))
#define svc_destroy(xprt) (*(xprt)->xp_ops->xp_destroy)(xprt)
#define svc_freeargs(xprt, xargs, argsp) \
        (*(xprt)->xp_ops->xp_freeargs)((xprt), (xargs), (argsp))
#define svc getargs(xprt, xargs, argsp) \
       (*(xprt)->xp ops->xp getargs)((xprt), (xargs), (argsp))
#define svc_getrpccaller(x) (&(x)->xp_rtaddr)
#define xdr getpos(xdrs)
                            (*(xdrs)->x ops->x getpostn)(xdrs)
#define xdr_setpos(xdrs, pos) (*(xdrs)->x ops->x setpostn)(xdrs, pos)
#define xdr inline(xdrs, len) (*(xdrs)->x ops->x inline)(xdrs, len)
#define xdr destroy(xdrs)
                             (*(xdrs)->x ops->x destroy)(xdrs)
```

### Figure 6-48: <search.h>

```
typedef struct entry { char *key; void *data; } ENTRY;
typedef enum { FIND, ENTER } ACTION;
typedef enum { preorder, postorder, endorder, leaf } VISIT;
```

Figure 6-49: <sys/sem.h>

```
#define SEM UNDO
                    010000
#define GETNCNT
#define GETPID
#define GETVAL
#define GETALL
                     6
#define GETZCNT
                     7
#define SETVAL
                   8
#define SETALL
struct semid ds {
       struct ipc perm
                         sem perm;
       struct sem *sem_base;
                   sem_pad[2];
       char
       unsigned short sem nsems;
       time_t sem_otime;
                 sem_ousec;
sem_ctime;
sem_cusec;
       long
       time t
       long
                    pad[4];
       long
};
struct sem {
       unsigned short semval;
              sempid;
       pid t
       unsigned short semncnt;
       unsigned short semzent;
};
struct sembuf {
       unsigned short sem_num;
       short
              sem op;
       short
                    sem flg;
};
```

Figure 6-50: <setjmp.h>

```
#define _JBLEN 13
#define _SIGJBLEN 64
typedef int jmp_buf[_JBLEN];
typedef int sigjmp_buf[_SIGJBLEN];
```

### Figure 6-51: <sys/shm.h>

```
struct shmid ds {
        struct ipc_perm shm perm;
        int shm_segsz;
        struct anon map
                              *shm amp;
        unsigned short shm lkcnt;
       char pad[2];
pid_t shm_lpid;
pid_t shm_cpid;
        unsigned long shm nattch;
        unsigned long shm cnattch;
        time_t
                  shm_atime;
                  shm_ausec;
shm_dtime;
shm_dusec;
        long
        time_t
        long
                    shm_ctime;
shm_cusec;
        time t
        long
        long
                     pad1[4];
};
#define SHMLBA
                      sysconf (31)
#define SHM RDONLY
                       010000
#define SHM RND
                       020000
```

# Figure 6-50: <setjmp.h>

```
#define _JBLEN 13
#define _SIGJBLEN 64
typedef int jmp_buf[_JBLEN];
typedef int sigjmp_buf[_SIGJBLEN];
```

### Figure 6-51: <sys/shm.h>

```
struct shmid ds {
        struct ipc_perm shm_perm;
        int
               shm_segsz;
        struct anon_map
                           *shm amp;
        unsigned short shm_lkcnt;
        char
                pad[2];
        pid_t shm_lpid;
pid_t shm_cpid;
        unsigned long shm nattch;
        unsigned long shm cnattch;
        time_t shm_atime;
                  shm_ausec;
shm_dtime;
shm_dusec;
shm_ctime;
shm_cusec;
        long
        time t
        long
        time_t
        long
        long
                     pad1[4];
#define SHMLBA
                      sysconf (31)
#define SHM RDONLY
                      010000
#define SHM RND
                       020000
```

## Figure 6-52: <sigaction.h>

Figure 6-53: <sys/siginfo.h>, Part 1 of 3

```
#define SI_FROMUSER(sip) ((sip)->si_code <= 0)</pre>
#define SI_FROMKERNEL(sip) ((sip)->si_code > 0)
#define SI_USER 0
#define ILL ILLOPC 1
#define ILL ILLOPN 2
#define ILL ILLADR 3
#define ILL ILLTRP 4
#define ILL PRVOPC 5
#define ILL PRVREG 6
#define ILL COPROC 7
#define ILL_BADSTK 8
#define NSIGILL
#define FPE_INTDIV 1
#define FPE_INTOVF 2
#define FPE_FLTDIV 3
#define FPE_FLTDIV
#define FPE FLTUND 5
#define FPE FLTRES 6
#define FPE FLTINV 7
#define FPE FLTSUB 8
#define NSIGFPE
                     8
#define SEGV MAPERR 1
#define SEGV ACCERR 2
#define NSIGSEGV
#define BUS_ADRALN 1
#define BUS_ADRERR
                      2
#define BUS_OBJERR
                      3
#define NSIGBUS
                      3
```

Figure 6-54: <sys/siginfo.h>, Part 2 of 3

```
#define TRAP BRKPT
                       1
#define TRAP TRACE 2
#define NSIGTRAP
#define CLD EXITED 1
#define CLD KILLED 2
#define CLD DUMPED 3
#define CLD TRAPPED 4
#define CLD STOPPED 5
#define CLD CONTINUED 6
#define NSIGCLD
#define POLL_IN
#define POLL OUT
#define POLL MSG
                         3
#define POLL ERR
                        4
#define POLL_PRI
                         5
#define POLL HUP
#define NSIGPOLL
#define SI_MAXSZ 128
#define SI PAD
                       ((SI MAXSZ/sizeof(int))-3)
#define si_pad
                       _data._proc._pid
#define si_uid data.proc.pdata.cld.status data.proc.pdata.cld.stime data.proc.pdata.cld.utime data.proc.pdata.cld.utime data.proc.pdata.kill.uid data.fault.addr
                         _data._proc._pdata._cld._status
#define si band
                         data. file. band
```

Figure 6-55: <sys/siginfo.h>, Part 3 of 3

```
typedef struct siginfo {
        int si_signo;
        int si_errno;
int si_code;
        union {
                       _pad[SI PAD];
                struct {
                       pid_t _pid;
union {
                               struct {
                                   uid_t _uid;
                                } _kill;
                                struct {
                                      clock_t _utime;
int _status;
clock_t _stime;
                                } _cld;
                        } _pdata;
                } proc;
                struct {
                       char * addr;
                struct {
                       int _fd;
long _band;
                } _file;
        } _data;
} siginfo_t;
```

Figure 6-56: <signal.h>, Part 1 of 2

```
#define SIGHUP
#define SIGINT
#define SIGOUIT
                 3
#define SIGILL
                  4
#define SIGTRAP
#define SIGIOT
                 6
#define SIGABRT
                7
#define SIGEMT
#define SIGEMT
#define SIGKILL
#define SIGBUS
                 10
#define SIGSEGV
                 11
#define SIGSYS
#define SIGPIPE
                  12
                 13
#define SIGALRM
                 14
#define SIGTERM
                 15
#define SIGUSR1
                 16
#define SIGUSR2
                 17
#define SIGCLD
                 18
#define SIGCHLD
                 18
#define SIGPWR
#define SIGWINCH
#define SIGPOLL
                 22
#define SIGSTOP
                 23
#define SIGTSTP
                 24
#define SIGCONT
                 25
#define SIGTTIN
                 26
#define SIGTTOU
                 27
#define SIGURG
                  33
#define SIGIO
                  34
#define SIGXCPU
                  35
#define SIGXFSZ
                   36
#define SIGVTALRM
                   37
#define SIGPROF
                   38
#define SIGLOST
```

Figure 6-57: <signal.h>, Part 2 of 2

```
65
#define NSIG
#define MAXSIG
                              64
#define SIG BLOCK
                              0
#define SIG UNBLOCK
                              1
#define SIG_SETMASK 2
#define SIG ERR (void(*)())-1
#define SIG IGN (void(*)())1
#define SIG_HOLD (void(*)())2
#define SIG_DFL (void(*)())0
#define SS ONSTACK 0x00000001
#define SS DISABLE 0x00000002
struct sigaltstack {
          char *ss_sp;
          int ss_size;
int ss_flags;
typedef struct sigaltstack stack t;
typedef struct sigset {
         unsigned long s[2];
} sigset t;
#define SIGNO MASK 0xFF
#define SIGDEFER 0x100
#define SIGHOLD 0x200
#define SIGRELSE 0x400
#define SIGIGNORE 0x800
#define SIGPAUSE 0x1000
```

Figure 6-58: <sys/stat.h>, Part 1 of 2

Figure 6-59: <sys/stat.h>, Part 2 of 2

```
#define S IFMT
                      0xF000
#define S IFIFO
                      0x1000
#define S IFCHR
                      0x2000
#define S IFDIR
                      0x4000
#define S IFBLK
                      0x6000
#define S IFREG
                      0x8000
#define S IFLNK
                      0xA000
#define S ISUID
                      04000
#define S ISGID
                      02000
#define S ISVTX
                      01000
#define S IRWXU
                      00700
#define S_IRUSR
                      00400
#define S IWUSR
                      00200
#define S IXUSR
                      00100
#define S IRWXG
                      00070
#define S IRGRP
                      00040
#define S IWGRP
                      00020
#define S IXGRP
                      00010
#define S IRWXO
                      00007
#define S IROTH
                      00004
#define S IWOTH
                      00002
                      00001
#define S IXOTH
#define S ISFIFO(mode) ((mode & S IFMT) == S IFIFO)
#define S ISCHR(mode) ((mode & S IFMT) == S IFCHR)
#define S ISDIR (mode) ((mode & S IFMT) == S IFDIR)
#define S ISBLK (mode) ((mode & S IFMT) == S IFBLK)
#define S_ISREG(mode) ((mode & S_IFMT) == S_IFREG)
```

# Figure 6-60: <sys/statvfs.h>

```
#define FSTYPSZ
                    16
typedef struct statvfs {
       unsigned long f_bsize;
       unsigned long f frsize;
       unsigned long f blocks;
       unsigned long f_bfree;
       unsigned long f_bavail;
       unsigned long f_files;
       unsigned long f ffree;
        unsigned long f_favail;
        unsigned long f fsid;
                  f_basetype[FSTYPSZ];
        unsigned long f_flag;
        unsigned long f namemax;
        char f_fstr[32];
        unsigned long f_filler[16];
} statvfs t;
#define ST_RDONLY 0x01
#define ST_NOSUID
                    0x02
```

## Figure 6-61: <stddef.h>

```
#define NULL 0
typedef int ptrdiff_t;
typedef unsigned int size_t;
typedef long wchar_t;
```

## Figure 6-62: <stdio.h>

```
typedef unsigned int size_t;
typedef long
              fpos_t;
#define NULL
                1024
#define BUFSIZ
#define EOF
                   (-1)
#define stdin
                   (&__stdinb)
#define stdout
                   (&__stdoutb)
#define stderr
                   (&__stderrb)
                   __stdinb;
extern FILE
extern FILE extern FILE
                   __stdoutb;
extern FILE __stderrb;
#define getchar() getc(stdin)
#define putchar(x) putc((x), stdout)
#define SEEK SET
                   0
                   1
#define SEEK_CUR
#define SEEK END
                   2
#define L ctermid
                    9
#define L cuserid
```

### Figure 6-63: <stdlib.h>

```
typedef struct {
    int quot;
    int rem;
} div_t;

typedef struct {
    long int quot;
    long int rem;
} ldiv_t;

typedef unsigned int size_t;

#define NULL 0
#define EXIT_FAILURE 1
#define EXIT_SUCCESS 0
#define RAND_MAX 32767

extern unsigned char __ctype[];
    #define MB_CUR_MAX __ctype[520]
```

Figure 6-64: <stropts.h>, Part 1 of 4

```
#define RNORM
                       0x000
#define RMSGD
                       0x001
#define RMSGN
                       0 \times 002
#define RMODEMASK
                       0x003
#define RPROTDAT
                       0x004
#define RPROTDIS
                       0x008
#define RPROTNORM
                       0x010
#define FLUSHR
                       0x01
#define FLUSHW
                       0x02
#define FLUSHRW
                       0x03
#define S_INPUT
                       0 \times 0001
#define S_HIPRI
                       0x0002
#define S OUTPUT
                       0 \times 0004
#define S MSG
                       0x0008
#define S ERROR
                       0x0010
#define S HANGUP
                       0x0020
#define S RDNORM
                       0 \times 0040
#define S WRNORM
                       S OUTPUT
#define S_RDBAND
                       0x0080
#define S WRBAND
                       0x0100
#define S BANDURG
                       0x0200
#define RS HIPRI
#define MSG HIPRI
                       0 \times 01
#define MSG ANY
                       0x02
#define MSG BAND
                       0 \times 04
#define MORECTL
                       1
#define MOREDATA
#define MUXID ALL
                      (-1)
```

Figure 6-65: <stropts.h>, Part 2 of 4

```
#define STR
                                          ('S' << 8)
#define I_NREAD
                                          (STR | 01)
#define I_PUSH
                                          (STR | 02)
#define I POP
                                          (STR | 03)
#define I_LOOK
                                          (STR | 04)
#define I_FLUSH
                                          (STR | 05)
#define I SRDOPT (STR|06)
#define I GRDOPT (STR|07)
#define I_STR (STR|010)
#define I_SETSIG (STR|011)
#define I_GETSIG (STR|012)
#define I_FIND (STR|013)
#define I_LINK (STR|014)
#define I_LINK (STR|015)
#define I_RECVFD (STR|016)
#define I_PEEK (STR|017)
#define I_FDINSERT (STR|020)
#define I_SWOPT (STR|021)
#define I_GWROPT (STR|023)
#define I_LIST (STR|025)
#define I_PLINK (STR|026)
#define I_PUNLINK (STR|026)
#define I_PUNLINK (STR|027)
#define I_PUNLINK (STR|027)
#define I GRDOPT
                                          (STR | 07)
  #define I_FLUSHBAND (STR|034)
 #define I_CKBAND (STR|035)
#define I_GETBAND (STR|036)
#define I_ATMARK (STR|037)
  #define I_SETCLTIME (STR|040)
  #define I GETCLTIME (STR | 041)
  #define I CANPUT
                                         (STR | 042)
```

Figure 6-66: <stropts.h>, Part 3 of 4

```
struct strioctl {
       int ic_cmd;
int ic_timout;
            ic len;
       int
       char *ic_dp;
};
struct strbuf {
      int maxlen;
       int
             len;
       char *buf;
};
struct strpeek {
      struct strbuf ctlbuf;
       struct strbuf databuf;
       long
             flags;
1;
struct strfdinsert {
       struct strbuf ctlbuf;
       struct strbuf databuf;
       long flags;
                fildes;
       int
       int
                  offset;
};
struct strrecvfd {
       int fd;
       uid t uid;
       gid_t gid;
       char fill[8];
};
```

# Figure 6-67: <stropts.h>, Part 4 of 4

Figure 6-68: <termios.h>, Part 1 of 6

```
#define NCC 8
#define NCCS 19
#define CTRL(c) ((c)&037)
#define IBSHIFT 8
#undef POSIX VDISABLE
typedef unsigned long tcflag t;
typedef unsigned char cc t;
typedef unsigned long speed t;
                  0
1
2
3
4
#define VINTR
#define VQUIT
#define VERASE
#define VKILL
#define VEOF
#define VEOL
#define VEOL2
#define VMIN
                    4
                    5
#define VTIME
                    7
#define VSWTCH
#define VSTART
#define VSTOP
#define VSUSP 10
#define VDSUSP 11
#define VREPRINT
                    12
#define VDISCARD
                    13
                    14
#define VWERASE
                 15
#define VLNEXT
```

Figure 6-69: <termios.h>, Part 2 of 6

```
0
#define CNUL
                  0377
'\\'
0177
#define CDEL
#define CESC
#define CINTR
                 034
'#'
'@'
#define CQUIT
#define CERASE
#define CKILL
                 04
#define CEOT
                 0
#define CEOL
#define CEOL2
#define CEOF
                   04
                  021
023
#define CSTART
#define CSTOP
#define CSWTCH
                   032
#define CNSWTCH 0
#define CSUSP CTRL('z')
#define CDSUSP CTRL('y')
#define CRPRNT
                   CTRL('r')
#define CFLUSH
                    CTRL('o')
#define CWERASE
                    CTRL('w')
#define CLNEXT
                    CTRL('v')
                    0000001
#define IGNBRK
 #define BRKINT
                    0000002
 #define IGNPAR
                     0000004
 #define PARMRK
                    0000010
 #define INPCK
                    0000020
                    0000040
 #define ISTRIP
                    0000100
 #define INLCR
                    0000200
0000400
 #define IGNCR
 #define ICRNL
 #define IUCLC
                    0001000
                    0002000
 #define IXON
 #define IXANY
                    0004000
0010000
 #define IXOFF
 #define IMAXBEL 0020000
```

Figure 6-70: <termios.h>, Part 3 of 6

```
#define OPOST
                     0000001
#define OLCUC
                     0000002
#define ONLCR
                    0000004
#define OCRNL
                   0000010
#define ONOCR
                   0000020
#define ONLRET
                   0000040
#define OFILL
                   0000100
#define OFDEL
                   0000200
#define NLDLY
                   0000400
#define NLO
#define NL1
                   0000400
#define CRDLY
                   0003000
#define CRO
                   0001000
#define CR1
#define CR2
                   0002000
#define CR3
                    0003000
                    0014000
#define TABDLY
#define TABO
#define TAB1
                   0004000
#define TAB2
                   0010000
#define TAB3
                   0014000
#define XTABS
                   TAB3
                   0020000
#define BSDLY
#define BSO
#define BS1
                   0020000
                  0040000
#define VTDLY
#define VTO
                   0040000
#define VT1
#define FFDLY
                    0100000
#define FF0
                    0
#define FF1
                    0100000
```

Figure 6-71: <termios.h>, Part 4 of 6

#define CBAUD	077600000
#define B0	0
#define B50	00200000
#define B75	00400000
#define B110	00600000
#define B134	01000000
#define B150	01200000
#define B200	01400000
#define B300	01600000
#define B600	02000000
#define B1200	02200000
#define B1800	02400000
#define B2400	02600000
#define B4800	03000000
#define B9600	03200000
#define B19200	03400000
#define EXTA	03400000
#define B38400	03600000
#define EXTB	03600000
#define CSIZE	00000060
#define CS5	0
#define CS6	0000020
#define CS7	0000040
#define CS8	0000060
#define CSTOPB	0000100
#define CREAD	0000200
#define PARENB	0000400
#define PARODD	0001000
#define HUPCL	0002000
#define CLOCAL	0004000
#define LOBLK	0010000
#define RCV1EN	0020000
#define XMT1EN	0040000
#define CIBAUD	037700000000
#define PAREXT	04000000

Figure 6-72: <termios.h>, Part 5 of 6

```
0000001
0000002
0000004
0000010
0000020
0000040
0000100
#define ISIG
                         0000001
#define ICANON
#define XCASE
#define ECHO
#define ECHOE
#define ECHOK
#define ECHONL
#define NOFLSH
                     0000400
0001000
0002000
0004000
#define TOSTOP
#define ECHOCTL
#define ECHOPRT
#define ECHOKE
#define FLUSHO
                      0020000
                       0040000
#define PENDIN
#define IEXTEN
                         0100000
#define IOCTYPE
                         0xff00
```

Figure 6-73: <termios.h>, Part 6 of 6

```
#define TIOC ('T'<<8)
#define TCSANOW (TIOC|14)
#define TCSADRAIN (TIOC|15)
#define TCSAFLUSH (TIOC|16)

#define TCIFLUSH 0
#define TCIOFLUSH 1
#define TCIOFLUSH 2
#define TCOOFF 0
#define TCOON 1
#define TCIOFF 2
#define TCION 3

struct termios {
    tcflag_t c_oflag;
    tcflag_t c_oflag;
    tcflag_t c_lflag;
    char c_padl;
    cc_t c_c[NCCS];
};
```

Figure 6-74: <sys/time.h>, Part 1 of 2

```
#define CLK TCK
#define CLOCKS PER SEC 1000000
#define NULL
typedef long clock t;
typedef long time t;
struct tm {
       int tm_sec;
       int tm min;
       int tm hour;
       int tm mday;
       int tm mon;
       int tm year;
       int tm wday;
       int tm_yday;
       int tm isdst;
};
struct timeval {
       time t tv sec;
       long tv usec;
};
extern long timezone;
extern int daylight;
extern char *tzname[2];
/* starred values may vary and should be
       retrieved with sysconf() of pathconf() */
```

## Figure 6-75: <sys/time.h>, Part 2 of 2

### Figure 6-76: <sys/times.h>

## Figure 6-77: <sys/tiuser.h>, Service Types

```
#define T_CLTS 3
#define T_COTS 1
#define T_COTS_ORD 2
```

## Figure 6-78: <sys/tiuser.h>, Transport Interface States

```
#define T_DATAXFER 5
#define T_IDLE 2
#define T_INCON 4
#define T_INREL 7
#define T_OUTCON 3
#define T_OUTREL 6
#define T_UNBND 1
#define T_UNINIT 0
```

Figure 6-79: <sys/tiuser.h>, User-level Events

```
#define T_ACCEPT1 12
#define T_ACCEPT2 13
#define T_ACCEPT3 14
#define T_BIND 1
#define T_CONNECT1 8
#define T_CONNECT1 8
#define T_CONNECT2 9
#define T_OPEN 0
#define T_OPEN 0
#define T_OPEN 1
#define T_PASSCON 24
#define T_RCV 16
#define T_RCVOINECT 10
#define T_RCVDIS1 19
#define T_RCVDIS2 20
#define T_RCVDIS3 21
#define T_RCVDER 23
#define T_RCVDER 6
#define T_RCVDER 7
#define T_RCVDER 7
#define T_SNDDIS1 17
#define T_SNDDIS1 18
#define T_SNDUDATA 5
#define T_SNDUDATA 5
#define T_UNBIND 3
```

Figure 6-80: <sys/tiuser.h>, Error Return Values

```
#define TACCES 3
#define TBADADDR 1
#define TBADADATA 10
#define TBADF 4
#define TBADF 4
#define TBADFLAG 16
#define TBADOPT 2
#define TBADSEQ 7
#define TBUFOVFLW 11
#define TLOW 12
#define TLOW 9
#define TNOADDR 5
#define TNOADDR 5
#define TNOOIS 14
#define TNOOIS 14
#define TNOTSUPPORT 18
#define TNOUDERR 15
#define TOUTSTATE 6
#define TSTATECHNG 19
#define TSYSERR 8
```

Figure 6-81: <sys/tiuser.h>, Transport Interface Data Structures, 1 of 2

```
struct netbuf {
      unsigned int maxlen;
       unsigned int len;
       char
                  *buf:
struct t bind {
       struct netbuf addr;
       unsigned int qlen;
};
struct t_call {
       struct netbuf addr;
       struct netbuf opt;
       struct netbuf udata;
       int sequence;
struct t discon {
        struct netbuf udata;
       int reason; int sequence;
};
struct t info {
        long addr;
        long options;
        long tsdu;
        long etsdu;
        long connect;
        long discon;
        long servtype;
};
```

Figure 6-82: <sys/tiuser.h>, Transport Interface Data Structures, 2 of 2

```
struct t_optmgmt {
    struct netbuf opt;
    long    flags;
};

struct t_uderr {
    struct netbuf addr;
    struct netbuf opt;
    long    error;
};

struct t_unitdata {
    struct netbuf addr;
    struct netbuf opt;
    struct netbuf opt;
    struct netbuf udata;
};
```

# Figure 6-83: <sys/tiuser.h>, Structure Types

```
#define T_BIND 1
#define T_CALL 3
#define T_DIS 4
#define T_INFO 7
#define T_OPTMGMT 2
#define T_UDERROR 6
#define T_UNITDATA 5
```

# Figure 6-84: <sys/tiuser.h>, Fields of Structures

# Figure 6-85: <sys/tiuser.h>, Events Bitmasks

### Figure 6-86: <sys/tiuser.h>, Flags

#### Figure 6-87: <sys/types.h>

```
typedef long time_t;

typedef long daddr_t;

typedef unsigned long dev_t;

typedef long gid_t;

typedef unsigned long ino t;

typedef int key_t;

typedef long pid_t;

typedef unsigned long mode_t;

typedef unsigned long typedef unsigned long typedef long off_t;

typedef long uid_t;
```

#### Figure 6-88: <ucontext.h>

#### Figure 6-89: <uio.h>

```
typedef struct iovec {
    char *iov_base;
    int iov_len;
} iovec_t;
```

```
Figure 6-90: <ulimit.h>
```

```
#define UL_GETFSIZE 1
#define UL_SETFSIZE 2
```

### Figure 6-91: <unistd.h>, Part 1 of 3

```
#define R OK
#define W OK
#define X OK
#define F OK
                  0
#define F ULOCK
                   1
#define F LOCK
#define F_TLOCK
                   2
#define F_TEST
                    3
                   0
#define SEEK SET
#define SEEK_CUR
#define SEEK END
#define POSIX_JOB_CONTROL
#define POSIX SAVED IDS
#undef _POSIX_VDISABLE
#define POSIX VERSION
#define XOPEN VERSION
/* starred values may vary and should be
       retrieved with sysconf() of pathconf() */
```

LIBRARIES 6-79

Figure 6-92: <unistd.h>, Part 2 of 3

```
#define SC_ARG_MAX 1
#define SC_CHILD_MAX 2
#define SC_CLK_TCK 3
#define SC_OPEN_MAX 4
#define SC_OPEN_MAX 4
#define SC_OPEN_MAX 5
#define SC_JOB_CONTROL 6
#define SC_VERSION 8
#define SC_VERSION 8
#define SC_WAXUMEMV 12
#define SC_MAXUMEMV 12
#define SC_MAXUMFOC 13
#define SC_MAXUMSGSZ 14
#define SC_SHMMAXSZ 16
#define SC_SHMMINSZ 17
#define SC_SHMMINSZ 17
#define SC_SHMMINSZ 17
#define SC_SHMSSGS 18
#define SC_NSEYSEM 19
#define SC_NSEYSEM 20
#define SC_NSEYMAP 21
#define SC_NSEYMAP 25
#define SC_NSEYMAP 26
#define SC_NSEYMAP 27
#define SC_NSEYMAP 27
#define SC_NIMER_VIRT 24
#define SC_NIMER_VIRT 24
#define SC_NIMER_VIRT 25
#define SC_NIMER_VIRT 26
#define SC_NIMER_VIRT 27
#define SC_NIMER_VIRT 29
#define SC_NIMER
```

#### Figure 6-93: <unistd.h>, Part 3 of 3

```
#define SC SHMLBA 31
#define SC_SVSTREAMS 32
#define SC_CPUID 33
#define SC_PASS_MAX 34
#define SC_PASS_MAX 34
#define SC_PAGESIZE 36
#define PC_LINK_MAX 1
#define PC_MAX_CANON 2
#define PC_MAX_INPUT 3
#define PC_NAME_MAX 4
#define PC_PATH_MAX 5
#define PC_PATH_MAX 5
#define PC_PO_TPE_BUF 6
#define PC_CHOWN_RESTRICTED 7
#define PC_NO_TRUNC 8
#define PC_VDISABLE 9
#define PC_VDISABLE 10
#define STDIN_FILENO 0
#define STDOUT_FILENO 1
#define STDOUT_FILENO 1
#define STDERR_FILENO 2
```

#### Figure 6-94: <utime.h>

```
struct utimbuf {
     time_t actime;
     time_t modtime;
};
```

Figure 6-95: <utsname.h>

#### Figure 6-96: <wait.h>

```
0177
#define WSTOPPED
#define WCONTINUED
                     0010
#define WUNTRACED 0004
#define WNOHANG
                     0100
#define WNOWAIT
                    0200
                    0001
#define WEXITED
#define WTRAPPED
                    0002
                    WTRAPPED
#define WTRACED
                      0177
#define WSTOPFLG
#define WCONTFLG
                      0177777
#define WSIGMASK
                      0177
#define WLOBYTE(stat) ((int)((stat)&0377))
#define WHIBYTE(stat) ((int)(((stat)>>8)&0377))
#define WWORD(stat) ((int)((stat))&0177777)
#define WCOREFLG
                    0200
#define WCOREDUMP(stat)
                             ((stat)&WCOREFLG)
                            (((s) \& 0xff00) >> 8)
#define WEXITSTATUS(s)
#define WIFCONTINUED(stat) (WWORD(stat) = = WCONTFLG)
                             (WTERMSIG(s) = = 0)
#define WIFEXITED(s)
                           (!WIFEXITED(s)&&!WIFSTOPPED(s))
#define WIFSIGNALED(s)
                             ((WTERMSIG(s) = -0x7f) && (((s)&0x80) = -0))
#define WIFSTOPPED(s)
                             (WIFSTOPPED(s)?WEXITSTATUS(s):0)
#define WSTOPSIG(s)
#define WTERMSIG(s)
                             ((s) & 0x7f)
```

6-83



# Index

68000 1: 1, 3: 22, 40 68000 family 3: 1 68008 1: 1 68010 1: 1 68020 1: 1, 3: 1 68030 1: 1, 3: 1 68040 1: 1, 3: 1, 10–11 68851 3: 1 68881 3: 1, 10–11, 24 68882 3: 1, 10–11, 24  A  %a0, pointer return value 3: 16 %r0, structure return value 3: 17 %r2, structure return value 3: 17 ABI conformance 3: 1, 22 see also undefined behavior 3: 1 see also unspecified property 3: 1 absolute code 3: 31, 5: 3 see also position independent code 3: 31 address	stack frame 3: 12, 41 structure and union 3: 3 allocation, dynamic stack space 3: 41 ANSI, C (see C language, ANSI) architecture implementation 3: 1 processor 3: 1 argc 3: 24 arguments bad assumptions 3: 40 exec(BA_OS) 3: 24 floating-point 3: 15 integer 3: 14 main 3: 24 pointer 3: 14 sign extension 3: 14 stack 3: 12, 14 structure and union 3: 16 variable list 3: 40 argv 3: 24 array 3: 3 automatic variables 3: 40 auxiliary vector 3: 26, 29
stack 3: 26 virtual 5: 1 address error exception 3: 23 address registers 3: 10 addressing, virtual (see virtual addressing) aggregate 3: 3 alignment array 3: 3 bit-field 3: 6 doubles in structures or unions 3: 2–3 executable file 5: 1 scalar types 3: 1	B base address 3: 29 behavior, undefined (see undefined behavior) bit field packing 3: 9 bit-field 3: 5 alignment 3: 6 allocation 3: 6 plain 3: 6 boot parameters (see tunable parameters) bra instruction 5: 7

branch instructions 3: 37	data registers 3: 10
breakpoint trap exception 3: 23	data representation 3: 1
bsr instruction 3: 35–36	diskettes, floppy 2: 1
	distribution media 2: 1
	double 3: 2
C	double-precision 3: 2, 15
Clanguage	dynamic linking 3: 19, 5: 5
	environment 5: 8
ANSI 3: 1, 24, 40	lazy binding 5: 8
fundamental types 3: 1	LD_BIND_NOW 5:8
main 3: 24	relocation 5: 5, 7
portability 3: 40	see also dynamic linker 5:5
switch statements 3: 37	dynamic segments 3: 20, 5: 4
calling sequence 3: 10	dynamic stack allocation 3:41
function epilogue 3: 17	signals 3: 41
function prologue and epilogue 3: 33	
%ccr (see condition code register)	
char 3: 2	E
sign of 3: 2	
chk, chk2 instruction exception 3: 23	emulation, instructions 3: 1
code generation 3: 31	environment 3: 29, 5: 8
code sequences 3: 31	exec(BA_OS) 3: 24
condition code register 3: 24	envp 3: 24
initial value 3: 24	exceptions
configuration parameters (see tunable	interface 3: 23
parameters)	signals 3: 23
coprocessor protocol exception 3: 23	exec(BA_OS) 3: 32
cptrapcc, trapcc, trapv exception 3: 23	interpreter 3: 28
	paging 5: 1
5	process initialization 3: 24
D	executable file, segments 5:3
9/ 20 mag also ratum college 0 44	execution mode (see processor execution
%a0, see also return value 3: 14 %d0	mode)
	extended-precision 3: 2, 15
see also return value 3: 14	external memory fault exception 3: 23
integer return value 3: 16	, , , , , , , , , , , , , , , , , , , ,
data	
process 3: 19	

uninitialized 5: 2

F	Н
faults (see traps) file, object (see object file)	heap, dynamic stack 3: 41
file offset 5: 1	
float 3:2	
Floating Point Coprocessor 3: 1, 10–11, 24  floating-point 3: 2, 14 arguments 3: 15 IEEE 3: 24 return value 3: 16 floating-point data registers 3: 24 floating-point exception 3: 23 fmovm.x instruction 3: 13, 16–17 format error exception 3: 23 formats array 3: 3	IEEE floating-point 3: 24 illegal instruction exception 3: 23 initialization, process 3: 24 installation, software 2: 1 instructions allowable 3: 1 emulation 3: 1 int 3: 2 integer arguments 3: 14 integer zero-divide exception 3: 23
structure 3: 3	J
union 3: 3 %fp (see frame pointer) %fp0, floating-point return value 3: 16 FPCP 3: 1, 10–11 frame pointer 3: 12–13, 34	jmp instruction 5: 7 jsr instruction 3: 35
frame size, dynamic 3: 41	L
function, void 3: 16 function call, code 3: 35 function linkage (see calling sequence)	lazy binding 5: 8 LD_BIND_NOW 5: 8 Id(SD_CMD) (see link editor) lea instruction 3: 33 Level 1 1: 3
G	Level 2 1: 3
global offset table 3: 32, 4: 2, 4, 6-8, 5: 5 _GLOBAL_OFFSET_TABLE_ 3: 33 %a5 3: 33 relocation 3: 32 _GLOBAL_OFFSET_TABLE_ (see global offset table)	libsys 6: 1 line 1010 emulator exception 3: 23 line 1111 emulator exception 3: 23 link editor 4: 6-7, 5: 5 link.l instruction 3: 13, 33 local variables 3: 40 long 3: 2

iong double 3. 2	Section 4: 2
long word 3: 12	see also archive file 4: 1
longjmp(BA_LIB) (see	see also dynamic linking 5:5
setjmp(BA_LIB))	see also executable file 4:1
	see also relocatable file 4:1
	see also shared object file 4:1
M	segment 5: 1
	shared object file 3: 32
main	special sections 4:2
arguments 3: 24	offset table, global (see global offset
declaration 3: 24	table)
malloc(BA_OS) 3: 21	optimization 3: 13
media, distribution 2: 1	5, 10
memory allocation, stack 3: 40-41	
memory fault exception 3: 23	Р
memory management 3: 19	
mmap(KE_OS) 3: 21	padding 3: 3–4
modes, processor (see processor execu-	requirements: structure, padding 3:3
tion mode)	requirements for 3:3
Motorola 68000 Family 5: 1	structure and union 3:3
Motorola 68000 family, generic term	page size 3: 19, 29, 5: 1
1: 1	Paged Memory Management Unit 3:1
mov.l instruction 3: 17, 34, 5: 7	paging 3: 19, 5: 1
movm instruction 3: 13	performance 5: 1
movm.l instruction 3: 13, 16–17, 33	parameters, system configuration (see tun-
	able parameters)
	PC-relative 3: 32, 37
N	performance 3: 1
	paging 5: 1
null pointer 3: 2, 20, 24	physical addressing 3: 19
dereferencing 3: 20	plain bit field 3: 6
	PMMU 3: 1
$\cap$	pointer 3: 2
O	function argument 3: 14
object file 4: 1	null 3: 2, 20, 24
ELF header 4: 1	portability
executable 3: 32	C program 3: 40
executable file 3: 32	instructions 3: 1
relocation 4: 3	position-independent code 3: 31–33, 5: 4
	p. 3.11311 maependem eode

see also absolute code 3: 31 see also global offset table 3: 31 see also procedure linkage table 3: 31 privileged opcode exception 3: 23	data 3: 10 description 3: 10–11, 13 floating-point 3: 10 global 3: 10
procedure linkage table 3: 32–33, 4: 2, 5, 7, 5: 5–6	initial values 3: 24, 26 local 3: 14
relocation 3: 32	saving 3: 12
procedures (see functions)	scratch 3: 14
process	signals 3: 14
dead 3: 41	relocation
entry point 3: 24	global offset table 3: 32 procedure linkage table 3: 32
initialization 3: 24	procedure initial
segment 3: 19	see object file 4: 3
size 3: 19	resources, shared 3: 19
stack 3: 25	return address 3: 16
virtual addressing 3: 19	return value
processor architecture 3: 1	floating-point 3: 16
processor execution mode 3: 22, 24	integer 3: 14, 16
processor supervisor mode 3: 22	pointer 3: 14, 16
processor-specific information 3: 1, 10,	structure and union 3: 17
19, 31, 5: 1, 5–6, 6: 1	rts instruction 3: 16–17
program counter, relative addressing (see	
PC-relative)	C
program loading 3: 28, 5: 1	3
Programmer's Reference Manual 3: 1	sbrk() 6: 1
purpose of ABI 1: 1	scalar types 3: 1
	scratch registers 3: 14
0	secondary storage 3: 19
Q	section, object file 5: 1
QIC cartridge 2: 1	segment
	dynamic 3: 20
	permissions 5: 2
R	process 3: 19–20, 5: 1, 6
0.47	segment permissions 3: 21
re-entrancy 3: 17	setjmp(BA_LIB) 3: 41
registers	setrlimit(BA_OS) 3: 21
address 3: 10	shared object file 3:32
calling sequence 3: 13	segments 3: 20, 5: 4

Index I-5

short 3: 2	see also <b>libsys</b> 6:1
SIGBUS exception 3: 23	system load 3: 19
SIGEMT exception 3: 23	3y 3tem 10ad - 5. 15
sign extension	
arguments 3: 14	T
bit-field 3: 6	
sign of char 3: 2	tape
	QIC cartridge 2: 1
signal(BA_OS) 3: 14, 23	reel-to-reel 2:1
signals 3: 14, 41	termination, process 3: 41
signed 3: 2, 6	text
SIGSEGV exception 3: 23	process 3: 19
single-precision 3: 2, 15	sharing 3: 32
sizeof 3: 1–2	trace exception 3: 23
software installation 2: 1	trap #2-15 exception 3: 23
%sp (see stack pointer)	trap instruction 3: 22
stack	traps (see exceptions)
address 3: 26	traps, access exception 3: 20
dynamic allocation 3: 41	tunable parameters
initial process 3: 26	process size 3: 19
process 3: 19–20	stack size 3: 21
system management 3: 21	
stack frame 3: 12, 40	
alignment 3: 12, 41	U
organization 3: 11–12, 39	
size 3: 12, 40	undefined behavior 3: 1, 17, 26, 28, 5: 2
stack pointer 3: 13, 26, 34	see also ABI conformance 3: 1
<stdarg.h> 3: 40</stdarg.h>	see also unspecified property 3:1
structure 3: 3	uninitialized data 5:2
function argument 3: 16	union 3: 3, 5
padding 3: 3	function argument 3: 16
return value 3: 17	return value 3: 17
structures, functions returning 3: 17	unions, functions returning 3: 17
supervisor mode (see processor supervisor	unlk instruction 3: 16–17
mode)	unsigned 3: 2, 6
switch code 3: 38	unspecified property 3: 1, 14, 16–17,
switch statements 3: 37	24–27, 5: 1, 3
sysconf(BA_OS) 3: 19, 29	see also ABI conformance 3:1
system calls 6: 1	see also undefined behavior 3:1

user mode (see processor execution mode)
User's Manual 1:2

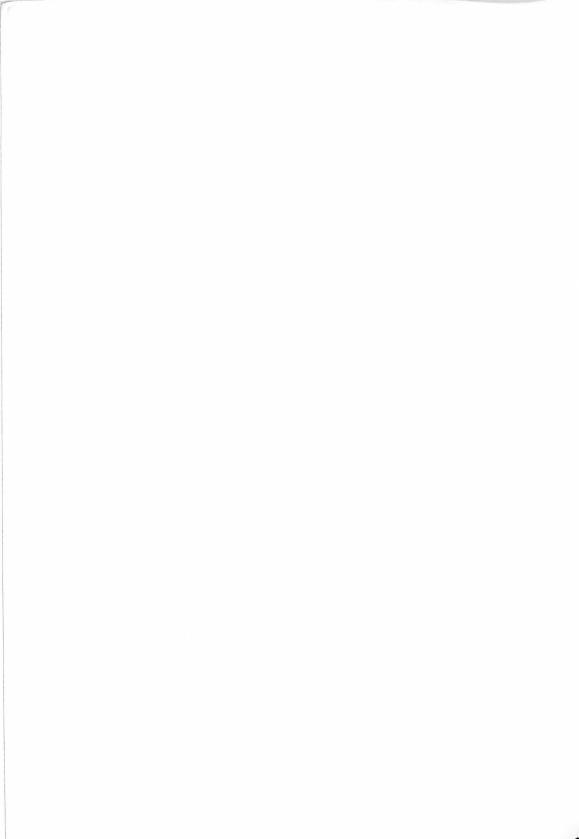


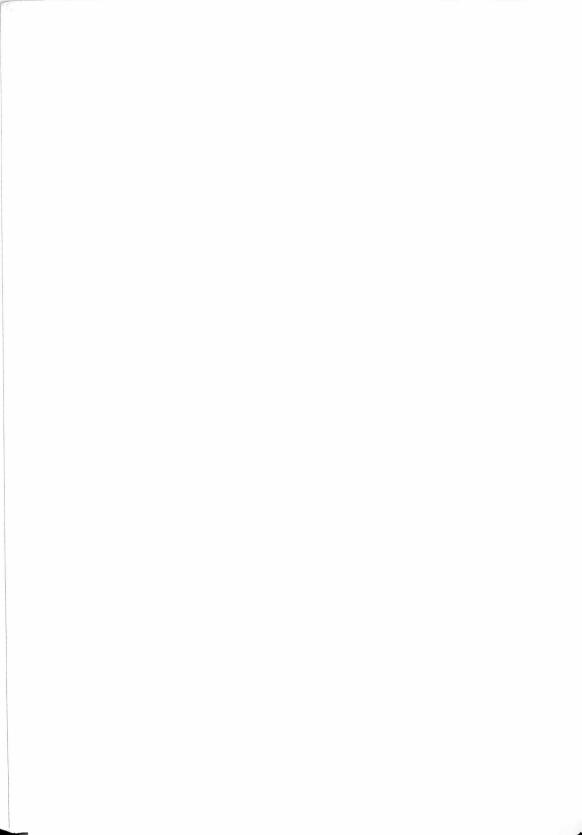
<varargs.h> 3: 40
variable argument list 3: 40
variables, automatic 3: 40
virtual addressing 3: 19, 32
bounds 3: 21
invalid 3: 20
void functions 3: 16

## Z

zero
null pointer 3: 2, 20
uninitialized data 5: 2
virtual address 3: 20
zero fill 3: 6

Index I-7

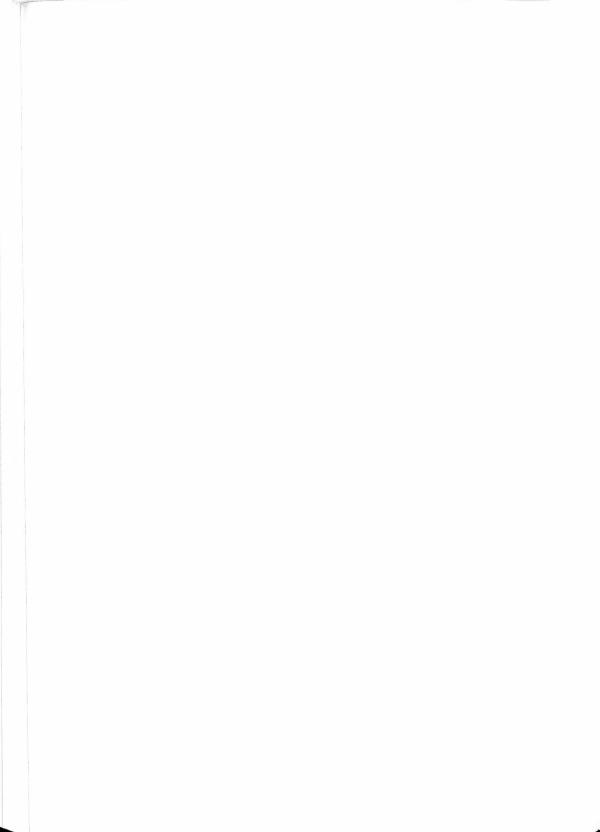




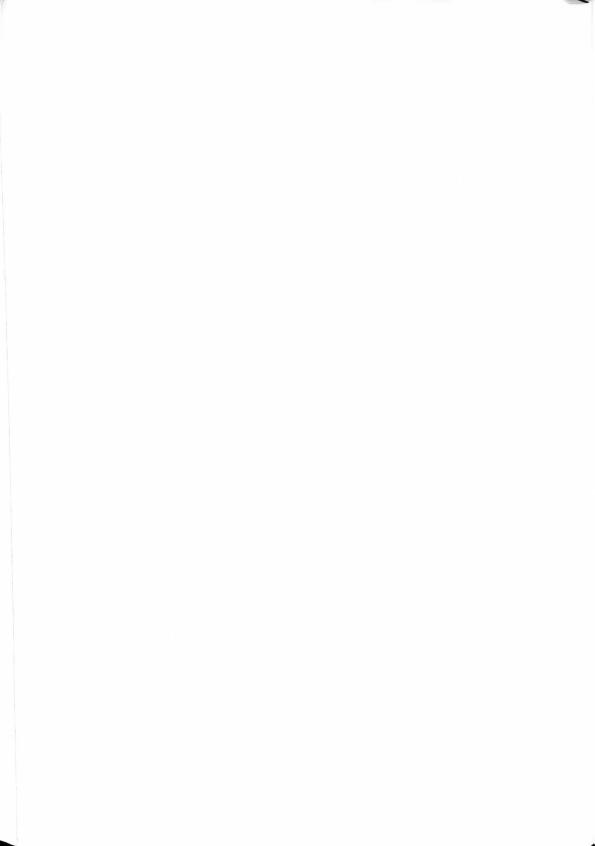




















Rückge're letum

# 277628 +02 NOV 1993

29. 11. 93

17/MAI95 A 005615

and the same