SOFTWARE MANUAL

AMOS
MONITOR CALLS

DWM-00100-42

REV. B01

alpha micro

NOTE: This printing of the manual contains the contents of Change Page Packet #1 for the "AMOS Monitor Calls Manual" (DSS-10000-12), which may be ordered separately from Alpha Micro.

First Printing: 1978
Second Printing: 1979
Third Printing: 30 April 1981
Fourth Printing: 1 July 1981

This book reflects AMOS Versions 4.5 and later

ALPHA MICROSYSTEMS
17881 Sky Park North
Irvine, CA 92714

PREFACE

One of the major features of the AMOS operating system is the large number of monitor calls available to the assembly language programmer. By making most common routines available in the monitor, AMOS frees the programmer from having to repetitively write the same routine. This manual describes these monitor calls.

We assume that the reader of this manual is familiar with assembly language programming and the AM-100 instruction set. We also assume that the reader is familiar with the AM-100 macro assembly system described in the AMOS Assembly Language Programmer's Reference Manual (DWM-00100-43). This reference manual is most emphatically NOT a tutorial on assembly language programming. Many such tutorials exist; if you are just learning assembly language, you should consult such a book before reading this manual.

Table of Contents

CHAPTER 1 COMMUNICATING WITH THE AM-100 MONITOR

1.1 MONITOR CALL CALLING FORMAT 1-1
1.1.1 Arguments 1-1
1.1.2 Standard Address Arguments 1-2

CHAPTER 2 JOB SCHEDULING AND CONTROL SYSTEM

2.1 THE JOB CONTROL BLOCK (JCB) 2-1
2.1.1 Example - Scanning The Job Control Area 2-2
2.2 ACCESSING YOUR JCB 2-3
2.2.1 Calling Sequence 2-3
2.3 JOB SCHEDULING CALLS 2-3
2.3.1 SLEEP - PUT JOB TO SLEEP 2-3
2.3.2 WAKE - WAKE UP JOB 2-4
2.4 JOB CONTROL BLOCK FORMAT 2-4
2.4.1 JOBSTS - The Job Status Word 2-4
2.4.2 JOBSPR - The Stack Pointer Reset Address 2-5
2.4.3 JOBNAM - The Job Name 2-5
2.4.4 JOBBAS - The Memory Base Address 2-5
2.4.5 JOBSIZ - The Memory Partition Size 2-5
2.4.6 JOBUSR - The Current PPN 2-6
2.4.7 JOBPRV - The Privilege Word 2-6
2.4.8 JOBPRG - The Current Program Name 2-6
2.4.9 JOBCMZ - The Command File Size 2-6
2.4.10 JOBCMS - The Command File Status 2-6
2.4.11 JOBERC - The Error Control Address 2-7
2.4.12 JOBTYP - The Job Type 2-7
2.4.13 JOBBPT - The Breakpoint Address 2-7
2.4.14 JOBBNK - The Memory Bank Pointer 2-7
2.4.15 JOBDEV - The Default Device 2-7
2.4.16 JOBDRV - The Default Drive 2-8
2.4.17 JOBTRM - The Terminal Block Pointer ... 2-8
2.4.18 JOBRBK - The Run Control Block 2-8
2.4.19 JOBFPE - The Floating-Point Trap Address 2-9
2.4.20 JOBRNQ - The Scheduling Area 2-9
2.4.21 JOBDYS - The DYSTAT Address 2-9
2.4.22 JOBWAT - Semaphore Wait Chain Link 2-9
2.4.23 JOBEXT - Job Exit-Trap Stack Pointer .. 2-10
2.4.24 JOBMSG - Job Message System Area Pointer 2-10
2.4.25 JOBIAL - Job Impure Area Link 2-10
2.4.22 JOBSTK - The Job's Stack Area 2-10

CHAPTER 3 MEMORY CONTROL SYSTEM CALLS

3.1 MEMORY PARTITION FORMAT 3-2
3.2 MEMORY MODULE FORMAT 3-5
3.3 MANIPULATING MEMORY MODULES 3-6

3.3.1 Allocating a Memory Module 3-8
3.3.2 Changing a Memory Module 3-8
3.3.3 Deleting a Memory Module 3-8
3.3.4 Permanent and Temporary Modules 3-8
3.4 MEMORY MAPPING SYSTEM 3-9
3.4.1 Internal Table Format 3-10
3.4.1.1 The MEMDEF Word 3-10
3.4.1.2 The JOBBNK Word 3-11
3.4.2 The Bank Switching Process 3-12
3.4.3 The BNKSWP Monitor Call 3-12
3.4.4 The DMADDR Monitor Call (For Memory Partition Controller) 3-13

CHAPTER 4 LOADING AND LOCATING MEMORY MODULES

4.1 THE SRCH AND FETCH CALLS 4-1
4.1.1 Specifying the Module Name 4-1
4.1.2 The Module Address 4-2
4.1.3 Flags 4-2
4.1.3.1 F.FCH - Fetch Module From Disk 4-2
4.1.3.2 F.USR - Bypass System Memory Search 4-3
4.1.3.3 F.ABS - Bypass Memory Search . 4-3
4.1.3.4 F.FIL - Mark Module as Permanent 4-3
4.1.4 Completion Codes 4-3

CHAPTER 5 MONITOR QUEUE SYSTEM CALLS

5.1 INCREASING THE AVAILABLE QUEUE LIST SIZE 5-1
5.2 QUEUE BLOCK USAGE BY THE SYSTEM 5-2
5.3 QUEUE SYSTEM MONITOR CALLS 5-3
5.3.1 QGET - Obtain a Free Queue Block 5-3
5.3.2 QRET - Return a Queue Block 5-3
5.3.3 QADD, QINS - Manipulating Queue Blocks 5-4

CHAPTER 6 THE FILE SERVICE SYSTEM

6.1 THE DATASET DRIVER BLOCK 6-1
6.1.1 DDB Format 6-2
6.1.1.1 Error Code 6-2
6.1.1.2 Flags 6-4
6.1.1.3 Buffer Address 6-4
6.1.1.4 Record Size 6-4
6.1.1.5 Buffer Index 6-4
6.1.1.6 Record Number 6-5
6.1.1.7 Queue Chain Link 6-5
6.1.1.8 JCB Address 6-5
6.1.1.9 Job Priority 6-5
6.1.1.10 Device Code 6-5
6.1.1.11 Drive 6-5
6.1.1.12 Call Level 6-5
6.1.1.13 Filename and Extension 6-6

6.1.1.14 PPN 6-6
6.1.1.15 Open Code 6-6
6.1.1.16 Driver Work Area 6-6
6.1.2 Device Transfer Buffers 6-6
6.1.3 Error Handling 6-7
6.1.3.1 Error Codes 6-7
6.2 FILE SERVICE MONITOR CALLS 6-8
6.2.1 FSPEC - Process an ASCII Filespec 6-8
6.2.2 INIT - Initialize the DDB 6-9
6.2.3 LOOKUP - Find the File 6-10
6.2.4 OPENI - Open a File for Input 6-10
6.2.5 OPENO - Open a File for Output 6-10
6.2.6 OPENA - Open and Append to Existing File 6-10
6.2.7 OPENR - Open a File for Random Processing 6-11
6.2.8 CLOSE - Close a File 6-11
6.2.9 READ - Perform a Physical Transfer 6-11
6.2.9.1 Sequential Devices 6-11
6.2.9.2 Random Devices 6-11
6.2.9.3 Interrupt Structure 6-12
6.2.10 WRITE - Perform a Physical Write 6-12
6.2.10.1 Sequential Devices 6-12
6.2.10.2 Random Devices 6-12
6.2.10.3 Interrupt Structure 6-12
6.2.11 INPUT - Perform a Logical Read 6-13
6.2.11.1 Sequential File Processing ... 6-13
6.2.11.1.1 Example 6-13
6.2.11.2 Random File Processing 6-13
6.2.11.3 Special Devices 6-14
6.2.12 OUTPUT - Perform a Logical Write 6-14
6.2.12.1 Sequential File Processing ... 6-14
6.2.12.1.1 Example 6-14
6.2.12.2 Random File Processing 6-15
6.2.12.3 Special Devices 6-15
6.2.13 DELETE - Delete a File 6-15
6.2.14 RENAME - Rename a File 6-15
6.2.15 ASSIGN - Assign a Device 6-15
6.2.16 DEASGN - Deassign a Device 6-16
6.3 DISK SERVICE MONITOR CALLS 6-16
6.3.1 Calling Sequence 6-16
6.3.2 The Bitmap Area 6-17
6.3.2.1 The Status Word 6-17
6.3.2.2 The Bitmap DDB 6-17
6.3.2.3 The Bitmap Buffe 6-17
6.3.2.4 The Bitmap 6-18
6.3.2.5 Altering the Bitmap 6-18
6.3.3 DSKCTG - Allocate a Contiguous Area ... 6-18
6.3.4 DSKALC - Allocate a Record 6-18
6.3.5 DSKDEA - Deallocate a Record 6-19
6.3.6 DSKBMR - Read the Bitmap 6-19

6.3.7 DSKBMW - Write the Bitmap 6-19
6.3.8 DSKDRL - Lock the Directory 6-19
6.3.9 DSKDRU - Unlock the Directory 6-20
6.4 MAGNETIC TAPE DRIVER MONITOR CALLS 6-20
6.4.1 REWIND Arg 6-20
6.4.2 WRTFM Arg 6-20
6.4.3 FMARK Arg 6-21
6.4.4 FMARKR Arg 6-21
6.4.5 TAPST Arg1,Arg2 6-21

CHAPTER 7 TERMINAL SERVICE SYSTEM

7.1 TERMINOLOGY 7-1
7.2 THE TERMINAL LINE TABLE 7-2
7.2.1 The Terminal Status Word 7-2
7.3 THE TERMINAL SERVICE CALLS 7-2
7.3.1 KBD {label} - Fetch a Line of Data 7-2
7.3.2 TTY - Output One Character 7-3
7.3.3 TIN - Get an Input Character 7-3
7.3.4 TOUT - Output One Character 7-3
7.3.5 TAB - Output One Tab 7-3
7.3.6 CRLF - Output a Carriage-Return / Line-Feed 7-4
7.3.7 TTYI - Output a String of Characters .. 7-4
7.3.8 TTYL - Output a String of Characters Indexed 7-4
7.3.9 PTYIN - Place Character in Input Buffer 7-4
7.3.10 PTYOUT - Fetch Character from Output Buffer 7-5
7.3.11 TTYIN - Fetch Another Job's Input 7-5
7.3.12 TTYOUT - Place a Character in Another Job's Output 7-5
7.3.13 TRMICP - Process Input Character Within Interface Driver 7-5
7.3.14 TRMOCP - Process Output Character Within Interface Driver 7-5
7.3.15 TRMBFQ - Process Output Characters Within Terminal Driver 7-6
7.3.16 TBUF - Output Large Amounts of Data ... 7-6
7.3.17 TCRT - Call Special Terminal Driver Routines 7-6
7.3.17.1 Standard Functions 7-7
7.3.17.1.1 Cursor Addressing 7-7
7.3.17.1.2 Other Functions . 7-7
7.3.18 Message Calls 7-8

CHAPTER 8 CONVERSION MONITOR CALLS

8.1 NUMERIC CONVERSION CALLS 8-1
8.1.1 Calling Format 8-1
8.1.1.1 Size Byte 8-1
8.1.1.2 Flags 8-2

8.2 RAD50 CONVERSION MONITOR CALLS 8-2
8.2.1 RAD50 Packing Algorithm 8-3
8.2.2 Packing and Unpacking Calls 8-3
8.2.2.1 PACK - Pack Three ASCII Characters into RAD50 . 8-3
8.2.2.2 UNPACK - Unpack Three RAD50 Characters into ASCII 8-4
8.3 PRINTING CONVERSION CALLS 8-4
8.3.1 PFILE - Output a Filespec From a DDB .. 8-4
8.3.2 PRNAM - Output a Filename 8-4
8.3.3 PRPPN - Output a PPN 8-4
8.4 ALPHABETIC CONVERSION--LCS AND UCS 8-4

CHAPTER 9 INPUT LINE PROCESSING CALLS

9.1 ALF - TEST A CHARACTER FOR ALPHABETIC 9-1
9.2 NUM - TEST A CHARACTER FOR NUMERIC 9-2
9.3 TRM - TEST A CHARACTER FOR TERMINATOR 9-2
9.4 LIN - TEST A CHARACTER FOR LINE TERMINATOR ... 9-2
9.5 BYP - BYPASS BLANKS 9-2
9.6 GTDEC - INPUT A DECIMAL NUMBER 9-2
9.7 GTOCT - INPUT AN OCTAL NUMBER 9-2
9.8 GTPPN - INPUT A PROJECT-PROGRAMMER NUMBER 9-3
9.9 FILNAM - INPUT A FILENAME 9-3

CHAPTER 10 MISCELLANEOUS MONITOR CALLS

10.1 EXIT - RETURN TO AMOS COMMAND LEVEL 10-1
10.2 CTRLC - BRANCH ON CONTROL-C 10-1
10.3 JLOCK, JUNLOK - PREVENT CONTEXT SWITCHING ... 10-2
10.4 RQST - REQUEST CONTROL OF A SEMAPHORE 10-2
10.5 RLSE - RELEASE CONTROL OF A SEMAPHORE 10-2
10.6 PCALL - INVOKE PROGRAM AS SUBROUTINE 10-3
10.7 AMOS - EXECUTE AMOS COMMAND AS SUBROUTINE ... 10-3

APPENDIX A DISK STRUCTURE FORMAT

A.1 PHYSICAL RECORD FORMAT A-1
A.2 DISK RECORD TYPES A-2
A.2.1 The Disk ID Record A-2
A.2.2 The Bitmap A-2
A.2.3 The Master File Directory A-3
A.2.4 The User File Directory A-3
A.2.5 Sequential File Data Records A-3
A.2.6 Contiguous File Data Records A-3
A.3 FILE STRUCTURE A-3
A.4 MFD ITEM FORMAT A-5
A.5 UFD ITEM FORMAT A-5

APPENDIX B SYSTEM COMMUNICATION AREA

B.1 SYSTEM - SYSTEM ATTRIBUTES WORD B-1
B.2 DEVTBL - ADDRESS OF THE DEVICE TABLE B-1

B.3 DDBCHN - ACTIVE DDB CHAIN B-1
B.4 MEMBAS & MEMEND - USER MEMORY POINTERS B-2
B.5 SYSBAS - BASE OF SYSTEM MEMORY B-2
B.6 JOBTBL - ADDRESS OF THE JOB TABLE B-2
B.7 JOBCUR - JCB ADDRESS OF THE CURRENT JOB B-2
B.8 JOBESZ - JOB TABLE ENTRY SIZE B-2
B.9 TIME - THE TIME OF DAY B-3
B.10 DATE - THE SYSTEM DATE B-3
B.11 HLDTIM - THE HEAD LOAD TIMER B-3
B.12 CLKFRQ - LINE CLOCK FREQUENCY B-3
B.13 SPXSAV - STACK POINTER SAVE LOCATION B-4
B.14 SPXINT - INTERNAL STACK B-4
B.15 LPTQUE - LINE PRINTER SPOOLER QUEUE B-4
B.16 TRMDFC - BASE OF THE TERMINAL DEFINITION TABLE B-4
B.17 TRMIDC - ADDRESS OF FIRST INTERFACE DRIVER ... B-4
B.18 TRMTDC - ADDRESS OF FIRST TERMINAL DRIVER B-4
B.19 TRMSCN - THE NON-INTERRUPT TERMINAL QUEUE B-4
B.20 CLKQUE - THE CLOCK QUEUE B-5
B.21 SCNQUE - THE IDLE SCAN QUEUE B-5
B.22 RUNQUE - THE JOB SCHEDULING QUEUE B-5
B.23 DRVTRK - THE DRIVE/TRACK TABLE B-5
B.24 MEMDEF & MEMBNK - MEMORY MANAGEMENT CONTROL .. B-5
B.25 ZSYDSK - ADDRESS OF SYSTEM DISK DRIVER B-5
B.26 SYSMEM - SYSTEM MEMORY LINK B-6
B.27 MSGQUE - SYSTEM LINK COMMUNICATION B-6
B.28 MSGDAT - MESSAGE SYSTEM DATA AREA B-6
B.26 QFREE - QUEUE SYSTEM CONTROL B-6

APPENDIX C ALPHABETIC LISTING OF AMOS MONITOR CALLS

INDEX

CHAPTER 1

COMMUNICATING WITH THE AM-100 MONITOR

The AM-100 monitor contains over 70 routines available for use by assembly language programs running in user or monitor memory. These routines are called by the supervisor calls SVCA and SVCB, which have been coded into macro form to make them easy to incorporate into user programs. The macros are included as a part of the system library file SYS.MAC in account [7,7] of the system disk. These calls have been grouped according to the function they perform and are described in this chapter and the following chapters.

1.1 MONITOR CALL CALLING FORMAT

The general format for all monitor calls is:

{label:} opcode {arguments} {;comments}

As the format shows, the only required item in all calls is the opcode itself, which is the name of the monitor call. A label may be used if desired, in which case it is assigned the address of the SVCA or SVCB instructions which start all monitor call sequences. The total number of words generated by any monitor call depends upon the call itself. Some calls generate up to four words of code to perform the function. Those calls which incorporate an ASCII message (such as the TYPE call) generate a string of bytes varying in length depending on the message involved. As in machine instructions, you may also place comments at the end of the line; each line of comments is identified by a preceding semi-colon.

1.1.1 Arguments

Some calls require one or more arguments to specify parameters for the execution of the monitor call function. These arguments most normally are source and/or destination address items for the data being manipulated by the monitor call. Some calls allow you to specify the location of data parameters, while other calls operate with predefined registers that you must set up beforehand. The following sections define each call and detail

the required arguments. Normally you define the arguments as expression values, standard addresses, or ASCII strings. An expression value may be any valid source expression which, after full evaluation, results in a value within the range of the argument definition. ASCII strings are just that; a string of characters typically used as a message to be displayed. Standard addresses are so important and complex that we devote the next entire section (1.1.2) to explaining them.

1.1.2 Standard Address Arguments

NOTE

> The following section is one of the most important, and most frequently misunderstood, sections of this manual. The concept of standard arguments is fundamental to understanding the monitor call calling sequences.

Standard addresses form the heart of many of the more complex monitor calls; you should therefore thoroughly understand them in order to gain maximum flexibility from the system. A standard address argument is coded exactly the same as a standard source or destination operand for a machine instruction such as ADD or MOV. Some restrictions should be noted, however, due to the method used in processing the standard address. Standard addresses are only used with those monitor calls that are coded as SVCB instructions. The SVCB pushes all user registers onto the stack, and it is from these stored values on the stack that the monitor call processor gains access to the address calculations using those registers. Standard addresses may take the form of any of the valid WD16 addressing modes; however, all autoincrement and autodecrement processing is done on a word basis, even though the monitor call may be requesting only one byte of data. In addition, the value used for SP register references is a dummy value which is not reloaded into SP when the monitor call exits, so the autoincrementing and autodecrementing modes will be ignored if used with the stack pointer register.

The monitor call processing software within the monitor actually duplicates the hardware, calculating the target address from the stored register value on the stack and the data from the extra word, if the address mode uses one. This target address then becomes the address of the data to be manipulated by the specific monitor call routine itself. This data may be only one byte, or it may be several words or more. The target address calculated by the processing of the standard address argument always points to the first byte of the data if more than one byte is required by the monitor call. A special case occurs when the standard address argument specifies the direct register address mode. In the WD16 hardware instructions, there is never more than one full word of data involved for the standard source and destination address modes, so direct register works on either the low byte or the full word in the target register. In the processing of monitor call standard addresses, however, this is not always the case since, as we

pointed out, some calls require several words of data to be manipulated. When direct register mode is used, the target address is actually the address of the stored register on the stack, which was a direct result of the SVCB hardware instruction processing. If more than one word is used by the call, it merely sequences right on through the stored words on the stack. In simple terms this means that if a monitor call wants three words of data for an argument and you specify the register R2 as the standard address argument, the three words that are used are actually those in R2, R3 and R4, in sequence. This is often very useful when writing re-entrant code.

CAUTION: If you specify a register for a call that wants more words than you have registers (most I/O calls want a 20-word DDB argument), the monitor call will walk right on through your stack and most likely crash the entire system.

One of the more common errors is forgetting that a standard argument needs a pound-sign (#) in front of a literal argument. For example, if you want the program to sleep for 20 clock ticks, the code reads:

```
SLEEP   #20.
```

Note that without the pound-sign, the program would sleep for the number of ticks contained in program-relative location 20.

It is very important that you understand the concepts outlined above. Think of the standard address arguments as source or destination addresses, as in the machine instructions. When you use them incorrectly, you will definitely find out about it quickly, since the usual result is a system crash.

CHAPTER 2

JOB SCHEDULING AND CONTROL SYSTEM

The AMOS timesharing monitor allocates jobs and schedules CPU time and resources for their operation. In order to properly write assembly language programs which make use of some of the more complex features of the system, you must have a basic understanding of how jobs are scheduled and controlled. The theory behind job-handling is too encompassing to cover in one section of this manual, but we can explain the fundamentals of job control by user programs.

Each job running in the system has two dedicated components which are not shared by any other job in the system: a monitor job control block and a user memory partition. In the monitor memory area itself, a job control table contains one area for each job that has been allocated to the system. One job is allocated for each JOBS command in the system initialization command file, which gives the job name and the terminal to which it is connected. The area allocated for each job in the job control table contains specific information about that job. This area is called the job control block and will be referred to from now on as the JCB.

2.1 THE JOB CONTROL BLOCK (JCB)

The format of the JCB is defined as a series of equate statements in the system library file SYS.MAC on DSK0:[7,7]. Each equate statement has the name JOBxxx, where xxx is a 3-character code for the specific item of the JCB being defined. The value of this symbol is actually the offset in bytes from the base of the JCB to the item itself. You may, during the course of your program, wish to read the current data in your own JCB or in some instances modify it. References to the JCB items should be made in one of two ways:

1. Use the system monitor calls JOBGET, JOBSET, and JOBIDX; which is the preferred method.

2. Locate the JCB for your job by moving @#JOBCUR into a register and then referencing all JCB items via JOBxxx(Rx).

Three words in the system communication area define the entire job control system during time-sharing operation. These three words are not part of the JCB areas but rather are non-sharable parameters set up during system initialization and not part of any one job. We point this out because the names of these three words are JOBTBL, JOBCUR and JOBESZ; which appear to be part of a user JCB but really are not. JOBTBL contains the base of the JCB table where all JCB's are stacked sequentially. This address is set up at system initialization time and is never changed. JOBCUR always contains the address of the JCB which has control of the CPU and is updated to point to the new JCB each time the job scheduler switches to a different job. Therefore, @#JOBCUR always points to your JCB if you reference it, because the reference is only executed while you have control of the CPU. JOBESZ contains the size of the JCB in bytes and is used by the system and by user programs for scanning through the JCB table. Since the size of the JCB may expand as new features are added to the system, JCB table scans must be made by setting an index to the base of the table (MOV @#JOBCUR,Rx) and then adding the size to the index to get to the next entry (ADD @#JOBESZ,Rx). In a JCB table scan, the first word of each JCB is guaranteed to be non-zero and the table is terminated by a null (zero) word. Again, these three words are a part of the master system communication area and not in the job table itself.

2.1.1 Example - Scanning The Job Control Area

The following is a brief example of how to scan the JCB table and process each JCB entry (such as for a system status report):

```
        MOV     @#JOBTBL,R0     ;set JCB table index R0 to table base
;Loop here to process each job table entry (JCB)
LOOP:   ...     ...             ;process JCB entry which is indexed by R0
        ...     ...             ;references to JCB items are via JOBxxx(R0)
        ...     ...
        ADD     @#JOBESZ,R0     ;advance R0 to next JCB entry
        TST     @R0             ;is this end of JCB table?  (null word)
        BNE     LOOP            ;  nope - go process valid JCB entry
;At this point we have finished the job table scan
        ...     ...
        ...     ...
```

2.2 ACCESSING YOUR JCB

You use three monitor calls to gain access to your own JCB when necessary. Two of the calls are used to transfer a single word of data to and from a specific word in the JCB; the other sets an index to a specific spot in the JCB area so that multiple words may be transferred, or so that faster access may be obtained when needed. These calls are as follows:

```
JOBGET  tag,item        ;Transfers one word from JCB item to tag
JOBSET  tag,item        ;Transfers one word from tag to JCB item
JOBIDX  tag,item        ;Sets absolute address of JCB item into tag
```

Since the locations may change, always use these calls as shown above.

2.2.1 Calling Sequence

All calls share the same basic format, where tag is a standard argument used for the transfer of one word of data in the JOBGET and JOBSET calls or to receive the index address in the JOBIDX call. The item argument is one of the JCB item tags (JOBSTS, JOBNAM, etc.), which identifies the item to be used in the transfer or to have the index set to. These items are equated to their relative offset value in SYS.MAC. Section 2.4 below explains how to use these items and points out their importance to the user.

2.3 JOB SCHEDULING CALLS

Three calls are used by various routines within the system monitor for controlling the job scheduling processes. These calls are JWAIT, JWAITC, and JRUN. JWAIT sets any job into the wait state. JWAITC sets your job into the wait state. JRUN then reactivates a job to the run state. If the J.NXT flag is specified, the job is placed at the beginning of the run queue; when J.NXT is not specified along with other JRUN flags, the job is placed at the end of the run queue. JWAIT and JRUN require that the job being controlled be indexed by R0 (which must point to the base of the JCB for that job), and that the argument specify one of the status control bits (in JOBSTS) to be used as the control flag. JWAITC assumes the current user.

The syntax for JWAIT, JWAITC and JRUN is as follows:

```
JWAIT flags
JWAITC flags
JRUN flags
```

2.3.1 SLEEP - PUT JOB TO SLEEP

SLEEP is a simple call that puts the user job to sleep for the number of line clock ticks you specify in the argument. After the specified amount of time has elapsed, the job is automatically awakened and execution continues with the instruction following the SLEEP call. The Z-flag is set if the job slept for the specified number of clock ticks. The Z-flag is reset if the job woke up prematurely because another job used the WAKE call.

CAUTION: A sleep call with an argument of zero clock ticks puts the job to sleep for about 18 minutes (65536 clock ticks).

The normal AM-100 system runs with a clock frequency of 60 Hz; each clock tick, therefore, has a value of 16.7 milliseconds. Also, the first clock tick may occur any time within the first 16.7 milliseconds (not necessarily a full clock tick).

Remember that SLEEP takes a standard argument; therefore, to cause the job to sleep for one minute, you would execute:

```
SLEEP   #3600
```

not

```
SLEEP   3600
```

Leaving off the pound sign (#) is a frequent coding error.

2.3.2 WAKE - WAKE UP JOB

This call wakes a specified job. R0 must point to the base of the JCB of the job you want to wake out of the sleep state. The Z-flag is set if the call is successful. If the specified job was already awake, the Z-flag is reset.

2.4 JOB CONTROL BLOCK FORMAT

The following is a list of the entries contained in your JCB. Each of these entries may be accessed via JOBGET, JOBSET, or JOBIDX by using the tag defined in each entry.

2.4.1 JOBSTS - The Job Status Word

The first word in each JCB is the job status flag word. Each bit in this word indicates a particular state in which the job may reside. Some legal states are defined by more than one bit being on at a time. The system and some of the system programs set and reset these bits as the current state of the job changes, but you should not alter this word without extreme caution. Following is a brief list of the bits and the mneumonics assigned to them, along with a basic description of the function of the bit when it is set.

```
J.ALC=1         ;Job entry is allocated (guarantees JOBSTS non-zero)
J.TIW=2         ;Job is in Terminal Input Wait state
J.TOW=4         ;Job is in Terminal Output Wait state
J.SLP=10        ;Job is in Sleep state
J.IOW=20        ;Job is in I/O Wait state
J.EXW=40        ;Job is in External Event Wait state
J.SMW=100       ;Job is waiting on a semaphore
J.CCC=200       ;A control-C abort is waiting to be processed
```

```
J.RUN=400       ;Job is running
J.MON=1000      ;Job is in monitor command mode (no program active)
J.LOD=4000      ;Program is being loaded for execution
J.SUS=10000     ;Job is in Suspend state
J.LOK=20000     ;Job has CPU locked (by user program command)
J.NXT=100000    ;Is always 0 in JOBSTS
```

If any of the following flags are on, the job will not be scheduled for CPU run time until the flag has been cleared: J.TIW, J.TOW, J.SLP, J.IOW, J.EXW, or J.SUS.

2.4.2 JOBSPR - The Stack Pointer Reset Address

One word, JOBSTR, is used to store the stack pointer reset address which is calculated when the system is initialized. This address is then used to reset the stack pointer each time the job exits back to monitor command mode. The user may allocate a larger stack area within his own partition by reloading this address if desired. The JOBSTK field is no longer used; see explanation of JOBSTK later in this chapter.

2.4.3 JOBNAM - The Job Name

Two words, JOBNAM, contain the 6-character job name packed RAD50. This name is set up by the JOBS command in the system initialization file. If a user program alters this word, it effectively alters the name of the job.

2.4.4 JOBBAS - The Memory Base Address

JOBBAS, one word, contains the base address of the user memory partition if one has been allocated for this job. This address is altered only by the MEMORY program which allocates and deallocates user memory partitions. We advise against altering this address unless you thoroughly understand the memory allocation process.

2.4.5 JOBSIZ - The Memory Partition Size

One word, JOBSIZ, contains the size of the user memory partition in bytes if one has been allocated for this job. This size word together with the above JOBBAS address word define the current user memory partition. JOBSIZ is altered only by the MEMORY program and the monitor command processor.

2.4.6 JOBUSR - The Current PPN

JOBUSR, one word, contains the current user PPN (account number) if the user is logged in. Zero indicates that no user is currently logged into this job. JOBUSR is modified by the LOG and LOGOFF programs and is tested by various protection schemes in the system to allow user access to files, etc.

2.4.7 JOBPRV - The Privilege Word

JOBPRV, one word, is used to store the privileges associated with the job. This word is not currently used but is allocated for future implementations of the security system. Further documentation will be provided when the system is completed.

2.4.8 JOBPRG - The Current Program Name

Two words, JOBPRG, contain the 6-character program name which is currently running or was the last job run if in monitor command mode. JOBPRG is loaded with the program name (packed RAD50) by the command processor when the program is loaded or located for execution. Currently, the only significance of this program name is in the displays created by the SYSTAT program (user terminal status display) and the DYSTAT program (video monitor).

2.4.9 JOBCMZ - The Command File Size

JOBCMZ is one word containing the size of the current command file area in the user memory partition if a command file is being processed. If this word is zero, no command file is currently in effect. This word is set to the initial size of a command file when that file is loaded into the top of the user partition and is decreased as each line is extracted from the area and sent to the monitor command processor. When it gets to zero, the command file is finished and the system returns to normal command mode input from the user terminal. The user should not alter this word.

2.4.10 JOBCMS - The Command File Status

JOBCMS is one word containing flags used by the command file processor when a command file is being processed. These flags should never be altered by the user, so they are not detailed here. JOBCMS works in conjunction with JOBCMZ to affect the command file processing scheme.

2.4.11 JOBERC - The Error Control Address

One word, JOBERC, controls the processing of WD16 hardware bus errors as described in the WD16 Programmer's Reference Manual. If JOBERC is zero a bus error causes a message to be printed on the user terminal, and the job is aborted. If JOBERC is non-zero a jump is made to the address specified in JOBERC, which should contain a valid routine for shutting down the program. Note that the bus error is fatal for this user only and does not normally kill the whole time-sharing system.

2.4.12 JOBTYP - The Job Type

JOBTYP, one word, specifies the type of job which is assigned to this jobstream. The following flags are currently implemented:

```
J.USR=1         ;Job is a user partition
J.NUL=2         ;Job is currently running the null subroutine
J.NEW=4         ;Job is processing a new memory allocation
J.LPT=10        ;Job is running the line-printer spooler (LPTSPL)
J.HEX=20        ;Binary inputs and outputs are in hex (not octal)
J.DER=40        ;Print disk error retry messages
J.VER=100       ;Activate auto-verify mode for disk writes
J.GRD=400       ;Terminal is guarded against SEND and FORCE commands
```

2.4.13 JOBBPT - The Breakpoint Address

JOBBPT is one word specifying the address to jump to if a breakpoint is encountered during the execution of a user program. JOBBPT is used by the DDT debug program for breakpoint handling and not normally used by user programs.

2.4.14 JOBBNK - The Memory Bank Pointer

JOBBNK is one word used by the memory management system to define the bank in which the job's current memory partition resides. It is actually a pointer to the control item within the memory mapping table which is used for turning the bank on and off when the job is allocated CPU time. This word must not be modified by the user.

2.4.15 JOBDEV - The Default Device

JOBDEV, one word, contains the RAD50 device code for the default device to be used if the file specification being processed by the FSPEC call does not explicitly specify a device. Normally this default device is DSK.

2.4.16 JOBDRV - The Default Drive

One word, JOBDRV, contains the drive number in binary for the default drive number to be used if the file specification being processed by the FSPEC call does not explicitly specify a drive number. Only used if the device code matches the code in JOBDEV or if the device code is left to default also. JOBDEV and JOBDRV normally contain the device and drive number set by the LOG program when a user logs in. They specify the disk device and drive which you usually use for processing.

2.4.17 JOBTRM - The Terminal Block Pointer

JOBTRM is one word containing a pointer to the terminal definition block for the terminal which is currently attached to this job. If no terminal is currently attached, this word contains a zero. The first word in the terminal definition block is the terminal status word, which is available to you for modification to set various terminal parameters such as echo control, image mode and lower-case processing. The old monitor call TIDX would deliver the address of this status word back to you in register R0. The TIDX call is no longer supported and must be replaced by the more general call:

```
        JOBGET  R0,JOBTRM       ;Get status word index
```

As with all of the JOBxxx calls, the destination may be any valid address and not just R0 as in the example above. The above example will replace the TIDX call exactly in performance, since TIDX used R0 as its destination.

For further information on the format of the terminal definition block and its use, refer to the source listing of the terminal service routine (TRMSER) which is made available to users on a special source diskette, as well as on the standard system disk pack. The terminal definition block is defined at the beginning of this routine.

2.4.18 JOBRBK - The Run Control Block

JOBRBK, a 14-word area, is the run control block for the jobstream. It is used for the loading of programs and overlays during job execution and is set up by the user program with the parameters needed to fetch the next program or overlay segment prior to the execution of a FETCH call. Refer to the description of the FETCH monitor call in section 4.1 for more details on the use of this item.

2.4.19 JOBFPE - The Floating-Point Trap Address

JOBFPE, one word, contains the address to jump to if a floating point error, such as a divide by zero, is executed. A user program which executes floating point instructions should enter its error trap address into JOBFPE and not into the vector at memory location 76, since this would destroy the sharable resource of that vector.

2.4.20 JOBRNQ - The Scheduling Area

JOBRNQ, a 7-word area, maintains the parameters for job scheduling and context switching of this job. The first four words are dynamically changing links used during the job scheduling process to place the job into the active run queue for future processing. Any altering of these four link words should be done with caution.

The fifth and sixth words are used to determine the job's run priority. The fifth word (at JOBRNQ+10) is the time counter which is decremented once for each clock interrupt whenever the job is running. When this count goes to zero, the job is put into the wait state and another job is activated. The sixth word (at JOBRNQ+12) is the actual priority of the job (set up by the JOBPRI command) and is used to initialize the above time counter each time the job is given control of the CPU for running.

The seventh word is used for storage of the current stack pointer value when the job is not in the active run state. The scheduler restores the stack pointer from this word each time the job is reactivated.

2.4.21 JOBDYS - The DYSTAT Address

JOBDYS, one word, contains the address of the byte in the VDM screen memory area for the job execution arrow. It is set by the DYSTAT program and referenced by the monitor job scheduler. The user should not alter this address.

2.4.22 JOBWAT - Semaphore Wait Chain Link

RQST and RLSE use this field to maintain a chain of JCBs waiting on a particular semaphore. This field contains the JCB address of the next job in this wait chain.

2.4.23 JOBEXT - Job Exit-Trap Stack Pointer

This field contains the value of the stack pointer the last time you executed an AMOS or PCALL monitor call. It is reset to the previous value each time you go through the exit-trap system.

2.4.24 JOBMSG - Job Message System Area Pointer

This field contains the address of the Link system message area for this job.

2.4.25 JOBIAL - Job Impure Area Link

Reserved for future use. Do not use.

2.4.26 JOBSTK - The Job's Stack Area

JOBSTK is a 100-word area that acts as the stack for this job. SP is set to the top of this area when a new program is initiated. You may reset your own stack pointer by moving the address of a larger area within your own partition, if the program needs more stack area. Be sure to allow at least 20 extra words or so for possible real-time interrupt handling which needs a valid stack area for register saves. The job scheduler also saves all user registers and processor status on the user stack during job context switching.

The label "JOBSTK" is not defined explicitly in SYS.MAC, but the area exists as the last 100 words in the JCB. The area has not been labeled because the JCB may be increased in size as the need arises, and the JOBSTK area should not be referenced by a label which will change value in future releases.

CHAPTER 3

MEMORY CONTROL SYSTEM CALLS

The AM-100 system contains a fairly sophisticated memory control system, even though there is no memory protection or mapping hardware associated with it. In order to make maximum use of the memory resources available and minimize system crashes due to memory violations, the assembly language programmer should understand how the monitor allocates memory and the rules under which memory should be accessed. This section describes the memory allocation scheme and the monitor calls that assist you in using memory in the proper way.

At present, memory may be handled in three different ways. Your AM-100 or AM-100/T system may have only one 64-kilobyte memory board; you can add memory boards to your configuration, jumpering them in such a way as to allow bank-switching; or you can use the AM-700 Memory Partition Controller (MPC) board to expand memory without requiring bank-switching. (For details on MPC, refer to Memory Management With The Memory Partition Controller in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.)

On the 64-kilobyte system, the top 256-byte portion is unavailable because it is mapped to the I/O ports. On a bank-switching system, 256 bytes at each jumpering junction are unavailable. On the MPC system, all memory may be allocated for either the monitor or user partitions.

In all three memory systems, the AMOS monitor resides in low memory beginning at location zero and extending upward as far as the monitor requires (typically around 14K bytes). The remaining memory above the monitor up to the end of the total amount of memory in your system is available for assignment as user memory partitions for each of the jobs. User partitions can be of varying sizes; none, however, may exceed 64K bytes minus the monitor size. The amount of memory a user program has available is therefore defined as the single contiguous memory partition which has been assigned to his job by the appropriate operator command (MEMORY, JOBMEM, or JOBSIZ).

This memory partition block is then allocated into smaller defined blocks called "modules," which are used by the system and the user to contain programs and data areas. Monitor calls exist which allow the user program to locate the absolute boundaries of its own memory partition and also to allocate, change, and delete memory segments in the form of defined modules.

(Changed 1 July 1981)

These modules can be named just like files (filename.extension), so they may be located by that name. Any program loaded for execution will be in the form of a module. During execution, some programs create other modules for device buffers, data tables, etc.

3.1 MEMORY PARTITION FORMAT

The memory partition assigned to a job may be located anywhere in memory depending on the memory that was available when the job assigned it using the appropriate allocation command. The user program may not count on any specific location for this partition. Within the partition, memory modules are allocated upward beginning at the base of the defined partition and building modules on top of each other as long as space permits. Modules may not be built that will extend past the top boundary of the user partition. As modules are deleted from memory, all modules above them are automatically shifted downward to fill up the space that the deleted module left. Also, when any module is changed in size, the modules above it are shifted in position accordingly. This method insures that all available memory is always at the top of your partition in one contiguous block. This method of grabbing the first portion of free memory to load a program into is the main reason that all programs must be written in totally relocatable code.

Figure 3-1 shows a typical memory layout for three users operating in a 64K system. The free memory at the 56K boundary could be used by a fourth job or by a current job that needs to expand.

Three monitor calls return information about your memory partition as it happens to be allocated. These three calls all take a single standard argument into which is delivered the absolute address of the base, end, or free base of the user memory partition. The three calls and the addresses that they return are listed below:

```
USRBAS  arg  - absolute base of user memory partition (last word)
USREND  arg  - absolute end of user memory partition (last word)
USRFRE  arg  - current base of remaining free memory (last module+2)
```

Since modules must always occupy an even number of bytes, the above calls always return an even address. If no modules are allocated in the current partition, the USRFRE address equals the USRBAS address. Otherwise, the USRFRE address is the word following the last currently allocated module in the memory partition. The remaining free user memory may be calculated by subtracting the USRFRE address from the USREND address.

Figure 3-2 shows a typical user job partition during the execution of a program which was loaded automatically by the operating system. The program itself was the first module to be allocated in the user partition and then was executed after being loaded. It remains in memory until it completes its task and exits to the monitor, at which time it is deleted by the operating system monitor. During execution, the program allocates a 1K data table module which may be used for storage of symbols or some similar function. Two I/O files are then opened on disk which causes the operating system file service routine to allocate the two disk buffer modules. The remaining memory in the partition has not yet been allocated in our example.

Note: Memory sizes are typical

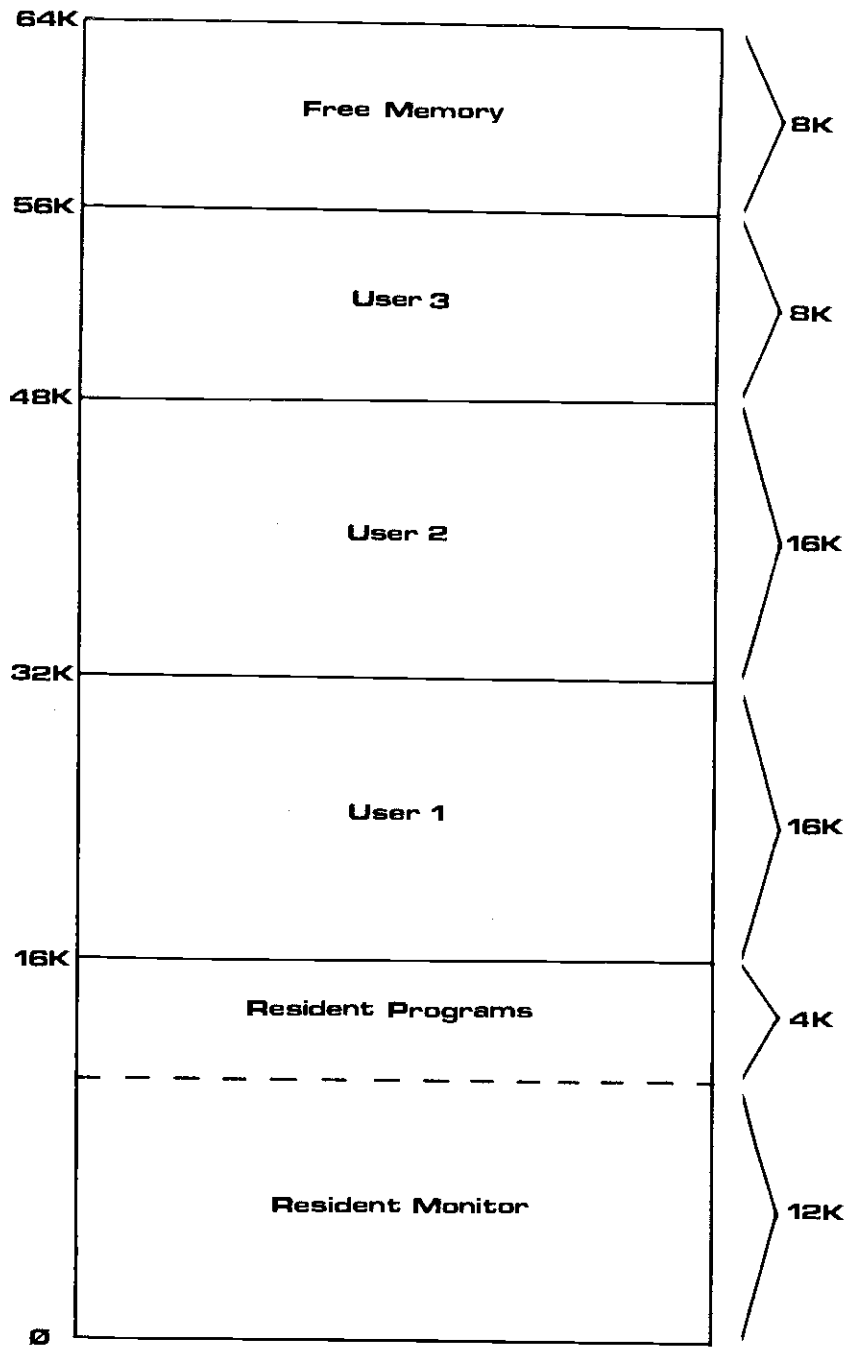

Total resident monitor size is 16K, leaving 48K for user partitions

Memory Map for a Typical 64K System (3 users)

Fig 3-1

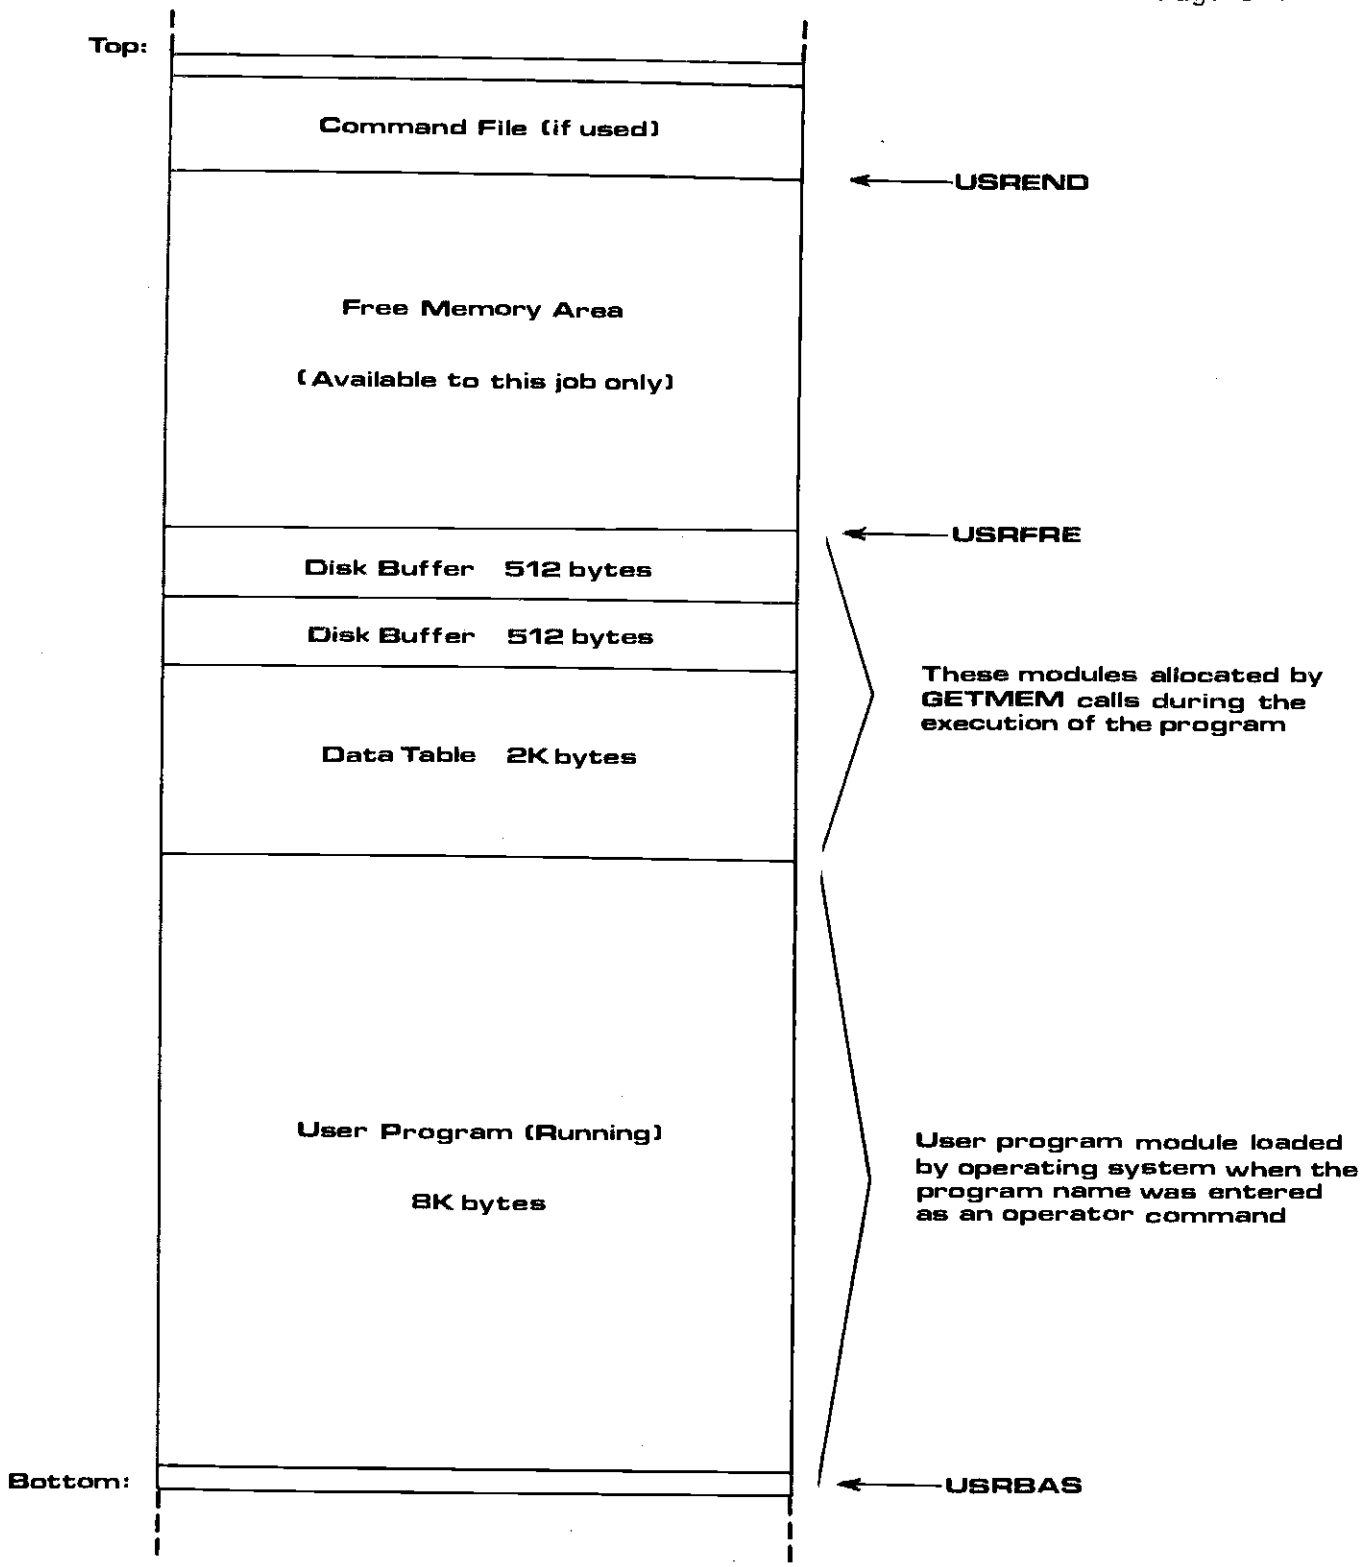

Memory Map for a Typical User Job Partition

Fig 3-2

Note that the USREND call does not actually return the absolute end of the partition but rather the end of the available free memory at the time of the call. If a command file is in progress, it occupies the upper part of the partition which we do not wish to alter during the execution of a program. In fact, the program should not have to take into consideration whether or not it was called by direct command or from a command file. Use of the USREND call insures that the user program may use all of free memory without having to compensate for the remaining part of any command file module.

Although the standard use of memory by the operating system is through the use of the memory management system calls (to be described next), you may find it easier to use free memory without regard to module boundaries, especially for use in variable length tables or hashing techniques. For this reason, the free memory space is always defined as the area between the addresses returned by the USRFRE and USREND calls. Note that the initialization of files normally results in the allocation of a buffer module; the operating system allocates this buffer at the current setting of the USRFRE address, then updates that USRFRE address. Therefore, you must be sure that all I/O buffers and any work modules are allocated before freely using the memory above the USRFRE address. The INIT and FETCH calls both cause the indirect allocation of a memory module in addition to the direct allocation or alteration of modules by the GETMEM, CHGMEM and DELMEM calls.

3.2 MEMORY MODULE FORMAT

Memory modules are the basic unit of formal data structure within the user memory partition. They are always allocated on word boundaries and must contain an even number of bytes to maintain this format. The monitor calls automatically pad an odd-sized module with a null byte to even it up. All modules contain five housekeeping words followed by any number of data words from zero to the maximum size left in the user memory partition. The five housekeeping words are always allocated, so a single-word module really takes up six words of memory.

The module format is as follows:

Word 1 - total size of module in bytes including the housekeeping words
Word 2 - module flag word
Word 3 - module filename packed RAD50
Word 4 - module filename packed RAD50
Word 5 - module extension packed RAD50
Words 6 thru n - module data area

Figure 3-3 gives a pictorial view of the above standard module format. The data area is usually the only area with which the user is concerned and so all references are made from the base of this area. The SRCH and FETCH calls (described in section 4.1) return this absolute address when locating or loading the requested module, instead of the address of the base of the housekeeping words. References to the housekeeping words should therefore be made via negative offsets relative to the data base address.

When scanning for a specific module or locating the end of the current module string, you may set your index using the USRBAS call, which returns the address of the size word of the first allocated module. You can then merely check the housekeeping words for the correct module name or other determining parameters and, if the module is to be bypassed, add the size word to the index. This bumps the index to the next module allocated. The last module always has a zero word following it, and you must be careful not to destroy this zero word if you are manipulating free memory directly without allocating it using the memory calls.

The module filename and extension follow the same format as the filenames on disk if the module in memory is named. The name is optional and need be used only if the module is to be located by name at a later time.

Modules may be either temporary or permanent depending on the method used to load them into memory. A module is made permanent by setting the file bit on in the housekeeping flag word when the module is allocated. Temporary modules are automatically deleted by the monitor when the program finishes and executes the EXIT call. Permanent modules are not automatically deleted but may be deleted by either the operator DELETE command or the monitor DELMEM call. Forcing a zero into the size word of the module is another way of deleting it, but this is not the recommended way since it also deletes all modules above it (the zero is the module area termination word).

3.3 MANIPULATING MEMORY MODULES

Three monitor calls are used to create, alter and delete these memory modules. All three calls take a single standard argument which must be the address of a 2-word block called a memory control block (MCB). The first word of this MCB contains the absolute memory address of the data area in the allocated module (past the housekeeping words). The second word contains the size of the data area in bytes (ten bytes less than the total module size since the housekeeping words are not included). The MCB therefore is the user's block, which defines a contiguous area in memory by its base address and size in bytes. You need not be concerned with the housekeeping words unless you need to access them directly; such a necessity should be rare.

The following three calls are used to manipulate memory modules:

```
GETMEM  MCB   - allocates a new memory module at current USRFRE
CHGMEM  MCB   - changes the size of the module defined by MCB
DELMEM  MCB   - deletes the memory module defined by MCB
```

The Z-flag is reset if GETMEM and/or CHGMEM fail (i.e., there is insufficient memory).

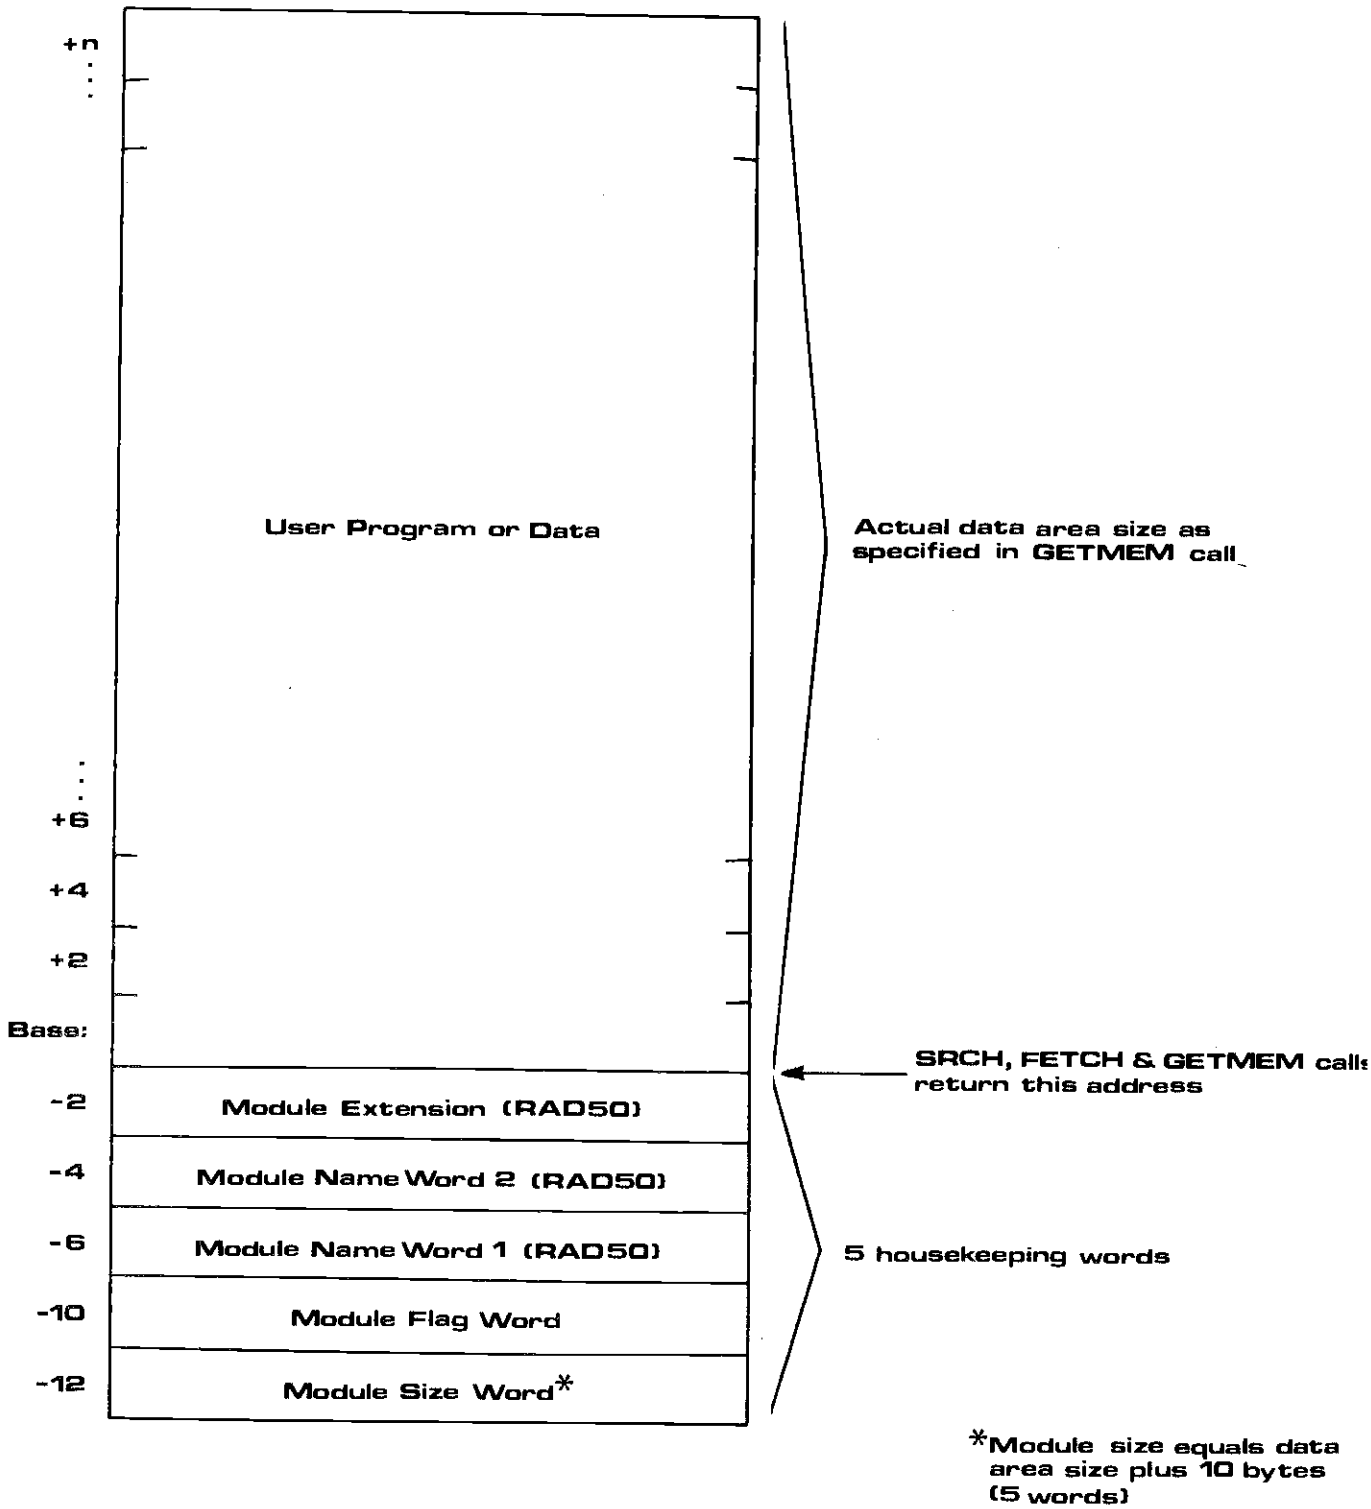

Standard Memory Module Format

Fig 3-3

3.3.1 Allocating a Memory Module

The following example shows the allocation of a 100-byte module

```
        MOV     #100.,MCB+2     ;set module size as 100 (decimal) bytes
        GETMEM  MCB             ;allocate module (MCB gets its address)
        BNE     NOMEM           ;no memory available
        ...
        ...
        ...
MCB:    WORD    0               ;receives address of module data area
        WORD    0               ;size of module data area in bytes

NOMEM:  EXIT
```

3.3.2 Changing a Memory Module

You may increase the size of the same module by:

```
        ADD     #20.,MCB+2      ;increase size word by 20 bytes
        CHGMEM  MCB             ;change its size
        BNE     NOMEM           ;not enough memory available
```

The above code causes the monitor to adjust the module housekeeping size word to reflect the new size. The address of the module does not change. However, note that the USRFRE address advances by 20 bytes and that any modules allocated after the one at MCB are shifted up in memory; but their corresponding addresses in their MCB are not adjusted by the monitor. I/O buffers allocated after the MCB module will therefore be erroneously addressed after the change, so the CHGMEM call must be used with care.

3.3.3 Deleting a Memory Module

To delete the above module we use the code,

```
        DELMEM  MCB             ;delete the module
```

3.3.4 Permanent and Temporary Modules

Recall that all temporary modules are automatically deleted by the monitor when the program exits. You may force the module to be permanently left in memory by giving it a name and setting the file bit (defined in SYS.MAC as "FIL") in the flag word. The following example illustrates the allocation of a 200-word module which is made permanent with the name "TABLE1.DAT":

```
        MOV     #200.,TBL1+2    ;set size as 200 bytes
        GETMEM  TBL1            ;allocate the module
        BNE     NOMEM           ;no memory available
        MOV     TBL1,R0         ;set R0 to index the data area base
        MOV     #[DAT],-(R0)    ;set the module name and extension (RAD50)
        MOV     #[LE1],-(R0)    ;  into the housekeeping words
        MOV     #[TAB],-(R0)    ;  in reverse order for efficient use of R0
        BIS     #FIL,-(R0)      ;set permanent file bit on in flag word
        ...
        ...
        ...
TBL1:   WORD    0               ;receives address of module
        WORD    0               ;size of module in bytes
```

Permanent memory modules may be saved onto disk using the operator SAVE command, or they may be deleted from memory when done by the operator DEL command. Refer to the AMOS User's Guide (DWM-00100-35) for details on these commands.

3.4 MEMORY MAPPING SYSTEM

The AMOS system is capable of supporting memory in excess of 64K by a simple bank switching technique which turns selected memory boards on and off under control of the operating system. This section defines some of the technical aspects of that system. It is assumed that you are already familiar with the operational aspects of the memory management system from the standpoint of setting up the SYSTEM.JNI file commands and operating procedures.

You must define for your own application the normal 64K memory as two general areas called sharable and switchable memory. Sharable memory always starts at location zero and extends upward far enough to totally contain the resident operating system and any system programs or sharable memory area needed for the application. Switchable memory then may occupy the remainder of the memory area up to the 64K address (octal 177376 inclusive).

There is only one sharable memory area that is always active. The switchable area, however, may be occupied by multiple memory boards referred to as "banks." Banks are defined to the operating system during system startup with the MEMDEF statements. Each MEMDEF statement defines the memory board (or boards) which are to be activated when that bank is selected by the operating system. Selection of the bank for activation is done when one of the user jobs which resides within that bank is granted CPU time by the AMOS job scheduling system. This action is automatic and transparent to the user. Only one bank may be active at a time, since all banks effectively respond to the same memory addresses (the area defined as switchable memory).

3.4.1 Internal Table Format

The memory bank switching system is controlled by a table which is built by the MEMDEF statements during system startup time. The table is basically a linked list of multi-word entries that resides within the monitor area. One entry defines the sharable memory area, and there is one entry for each bank defined by a MEMDEF statement. Two words that reside in the monitor system communication area are used to control the memory management system. These words are labled "MEMDEF" and "MEMBNK"; MEMDEF stores the base address of the table just defined, and MEMBNK stores the memory bank which is currently active. If memory management is not in use (no MEMDEF statements appeared in the SYSTEM.INI file) both of these words contain a zero value.

A system configured with an AM-700 or Memory Partition Controller (MPC) has a different controlling data structure than one using traditional bank swapping. (For information on the MPC, refer to the "System Operator's Information" section of the AMOS Software Update Documentation Packet.) The data structure is a linked list of queue elements, each containing four words. One element is allocated for the sharable memory area, one for each job on the system, one for each piece of switchable system memory, and one to indicate the end of physical memory. These elements are created by JOBS, BITMAP, and SYSTEM during the system initialization procedure. The queue is pointed to by the word labeled "MEMDEF" residing in the system communications area.

3.4.1.1 The MEMDEF Word - The MEMDEF word in the system communication area contains the address of the first entry in the table, which is always the entry defining the sharable memory boundaries. The format for this entry is:

Word 1 - link to next entry
Word 2 - base address of sharable memory (0)
Word 3 - top address of sharable memory plus 1

The remaining entries define the switchable memory banks in use and have the format:

Word 1 - link to next entry (0 if this is last entry)
Word 2 - base address of this switchable bank
Word 3 - top address of this switchable bank plus 1
Words 4 through n - hardware control codes for bank switching

The hardware control codes are one or more entries used to turn the memory boards on and off during bank switching. There is one control code for each physical board which has been defined as part of this bank. Each control code is two words in length, with the first word containing the address of the hardware port for the memory board and the second word containing the switch-on and switch-off bytes (low and high bytes, respectively) that are sent to that port. Note that in the MEMDEF statements you can specify more than one board per bank (even different types of boards) by separating the

board definitions with slashes. The final hardware code is followed by a single word of zero to indicate the end of the codes for this bank.

On a Memory Partition Controller (MPC) system (see reference in section 3.4.1 above), the word MEMDEF points to the data structure used by the operating system to control memory partitions. Each entry has the following format:

Word 1 - Link to the next entry
Word 2 - JCB pointer
Word 3 - Base address of partition
Word 4 - Limit address of partition

The element describing the sharable memory area has a 0 in word 2. An element describing a switchable system memory module has a -2 in word 2. The last element has a 0 in word 1. The base and limit addresses contained in words 3 and 4 are magnitude 256; that is, the real memory address shifts right eight bits. The sharable memory element has a 0 in word 3, and word 4 contains the end of the system area. The first job on the system has 0 in word 3, as the base address is an offset from the end of the sharable area. The element for the sharable memory area is first in the queue, the elements for jobs are next, occurring in the sequence that the JOBS statement lists them. Next are the elements for switchable system memory, occurring in reverse order of the BITMAP statements that generated them. The last element indicates the end of physical memory. For more details on exactly what base and limit addresses are and how they work, refer to the hardware documentation for the AM-700.

3.4.1.2 The JOBBNK Word - The JOBBNK word in each job's JCB contains the address of the word 4 in the above definition for the bank in which the job currently resides. This address is the base of the control codes for the hardware switching operation. The MEMBNK word in the system communication area always contains the same address as the JOBBNK word for the job that is currently running. This is used by the scheduling and switching system to turn off the current job and turn on the next job for running.

For a Memory Partition Controller (MPC) system (see reference in section 3.4.1 above), the JOBBNK word in a job's JCB points to word 3 in the corresponding MEMDEF queue element. The MEMBNK word in the system communications area always points to word 3 of the element corresponding to the memory partition currently mapped in by the AM-700.

3.4.2 The Bank Switching Process

Memory bank switching is performed by the job scheduler by a simple sequence of steps:

1. Use the MEMBNK word to locate the currently active bank entry.
2. Send the switch-off byte to the port address for each control code.
3. Use the JOBBNK word for the next job to be run to locate the bank entry for that job.
4. Send the switch-on byte to the port address for each control code.
5. Store the new job's JOBBNK data into the MEMBNK word for next time.

3.4.3 The BNKSWP Monitor Call

Under normal operation of the AMOS system each user is confined to an area that resides totally within any one defined memory bank. The BNKSWP call may be used by a more sophisticated assembly language routine to allow one user to access more than one bank of memory. The BNKSWP monitor call expects register R1 to contain the address of word 4 of the bank which is to be activated (similar to the automatic operation which uses the address within the JOBBNK word). The currently active memory bank is switched off and the new bank (per R1 address) is switched on. The MEMBNK word is updated properly to reflect the newly activated memory bank. Register R1 is also changed to contain the index to the previously operating bank, thereby allowing a convenient return to reactivate the previous bank if R1 is not altered.

Note that since the current bank is switched off, the BNKSWP call must be executed from somewhere in sharable memory to prevent the return from executing instructions in the new bank. This can be accomplished in one of several different ways, including pushing the routine onto your stack (within the JCB) or executing a special subroutine which has been loaded into system memory.

On a Memory Partition Controller (MPC) system (see reference in section 3.4.1 above), the BNKSWP call functions the same as it does on a bank swapped system, except that R1 is expected to point to word 3 of the MEMDEF queue element describing the memory partition the caller wants to map in. The same restrictions that existed before still apply. The user must check bit 15 in the SYSTEM word residing in the system communications area. If it's on, he must realize that the MEMDEF queue is structured differently than it would be on a bank swapped system.

3.4.4 The DMADDR Monitor Call (For Memory Partition Controller)

The AM-700 or Memory Partition Controller translates memory addresses for DMA devices as well as for the AM-100/T processor. This feature allows DMA activity to occur in one job's partition concurrently with another job running in another partition. On bank swapping systems, only the job that is doing DMA activity can be running. All other jobs are locked out for the duration of the DMA operation. Device drivers for DMA I/O devices (e.g., the magnetic tape) must include a DMADDR monitor call when executing on an MPC system. The one argument passed to the DMADDR is the DMA level of the device. When called, DMADDR sets up the appropriate base address and limit address registers on the MPC. If DMADDR is called on a system configured without the MPC, nothing is done at all.

In order to utilize the advantages of the MPC, the driver should test the word SYSTEM in the system connumications area; if bit 15 is set, other jobs should be allowed to run while DMA activity is ongoing. If bit 15 is not set, the normal bank-swapping code should be executed. The calling sequence for DMADDR appears as follows:

```
DMADDR  DMALEV          ; Set up MPC hardware for this DMA activity.
```

DMALEV is the DMA level of the device, which is constant for any particular device but changes from one device to another. There are no return arguments from DMADDR.

(For a more complete explanation of the Memory Partition Controller, refer to the "System Operator's Information" section of the AMOS Software Update Documentation Packet.)

CHAPTER 4

LOADING AND LOCATING MEMORY MODULES

Memory modules may contain an optional filename and extension, which may be used to locate modules, both in memory and on the disk. This chapter deals with locating and loading modules via these optional filenames and extensions. Normally, when you enter a command from the terminal, AMOS first searches for the requested program in the resident system memory area, then in your own memory partition. If the program is resident in either of these places, it need not be loaded in from disk, and execution begins immediately using the resident program in system or user memory.

4.1 THE SRCH AND FETCH CALLS

The user may make use of two monitor calls (FETCH and SRCH) for locating and loading modules in memory by name. In actuality, the SRCH call is a specialized version of the FETCH call and is included only for convenience and compatibility with older programs that are still in the system. Basically, the SRCH call only locates a module if it is in memory, while the FETCH call automatically loads a module into memory from the disk if it is not found to be in memory already.

Both calls have the same basic format:

```
SRCH    nameblock,index,control-flags
FETCH   nameblock,index,control-flags
```

4.1.1 Specifying the Module Name

Nameblock is a standard argument used in the SRCH and FETCH calls to specify the name of the module to be located or loaded. The format of the actual nameblock referenced is different in each case, however. In the case of the SRCH call, nameblock refers to a 3-word block of memory (or 3 contiguous registers) containing the filename and extension of the desired module in RAD50 packed form. For the FETCH call, nameblock refers to a full file Dataset Driver Block (DDB) which allows the user to specify a full disk file

specification to load the module from in case it is not located in memory. The DDB has not yet been introduced and is defined and explained in section 6.1.1. In brief, the DDB is a 24 (octal) word area in memory which contains all the information and work areas to define and manipulate a specific disk file in any area on any defined disk device. The DDB is normally set up by processing an ASCII file specification with the FSPEC call (more on this later).

4.1.2 The Module Address

The second argument is the index which is to receive the absolute memory address of the located (or loaded) memory module data area. Refer to figure 3-3 in the preceding chapter for the layout of the memory module and the place that this index is set to. The index argument is also a standard argument, although the normal mode is to receive the module address in a general register (R0-R5). If the index argument is not specified in the call, the default used is register R0 which is compatible with older versions of this system.

4.1.3 Flags

The third argument is the optional control flags which may be used to control the operation of the SRCH and FETCH calls. This argument is any valid expression which evaluates down to a value in the range of 0-17 (octal). Only the low order four bits are significant and they have been given the following mnemonic definitions in the system library SYS.MAC:

```
F.FCH=1         ;Fetch module from disk if not in memory
F.USR=2         ;Search user memory only
F.ABS=4         ;Load absolute segment from disk
F.FIL=10        ;Set module permanent file flag after load from disk
```

4.1.3.1 F.FCH - Fetch Module From Disk - F.FCH is the flag that actually differentiates the SRCH call from the FETCH call, since they both technically are the same SVCB supervisor call. The SRCH call forces this bit off while the FETCH call forces this bit on. When set, the F.FCH bit causes the nameblock to be interpreted as a full file DDB and the module to be loaded from disk if not located in memory first. Since the use of this bit is controlled by specifying either SRCH or FETCH as the calling opcode, you should not include this bit in the control-flags argument of your call.

4.1.3.2 F.USR - Bypass System Memory Search - F.USR is the flag used to specify bypassing the searching of the resident system memory area for the module and proceed directly to searching the user area only. This allows specific versions of modules to be loaded and used even though they may be duplicated in the system memory area. This flag is not normally used by programs other than system software.

4.1.3.3 F.ABS - Bypass Memory Search - When set, F.ABS forces a direct search to the disk for the requested module, bypassing all memory searches that would normally occur. The module is then loaded into memory at the absolute address specified by the index argument in the calling sequence. No housekeeping words are allocated, and the first word of the module gets loaded into the first word specified by the index argument. Note that this form is the only time the index argument is used to pass an address to the FETCH processor instead of being used to receive the address of the located module. The F.ABS form of the FETCH call is used to load program segment overlays.

4.1.3.4 F.FIL - Mark Module as Permanent - F.FIL is used to force the permanent file flag bit on in the module flag word after the module has been loaded from disk. The FETCH call always places the filename and extension into the housekeeping words 3-5 so even if the module is only temporary, it may still be located by name as long as the program which loaded it is still active. This is useful for dynamic loading of subprograms and/or data modules. Setting the F.FIL flag on in the control-flags argument means that the module will not be deleted from memory by the operating system when the calling program finally exits. The operator LOAD command uses this method to load a program into memory and leave it there to be called by name.

4.1.4 Completion Codes

When the SRCH or FETCH call returns, the user must test the status of the Z-bit to see if the module was located or loaded successfully. If the Z-bit is set (tested by BEQ), the operation was successful. If the Z-bit is not set (tested by BNE), the module was either not located or would not fit into the remaining free memory within the user's partition.

CHAPTER 5

MONITOR QUEUE SYSTEM CALLS

The monitor queue is a list of blocks in system memory which are linked to each other in a forward chain. The base of this chain, and the count of the blocks in the chain, are contained in the QFREE monitor communications words (see Appendix B). Each queue block in the chain links to the next one by storing its address in the first word of the queue block. The last queue block in the chain contains a zero link word to flag it as the end. Each queue block is currently 8 words (16 bytes) in size, although this value may increase with the next release of the file system. The monitor initially contains 20 blocks in the available queue list.

During normal monitor operation various functions use these queue blocks to perform certain tasks. When a routine needs a queue block, it issues a QGET monitor call, which delivers the first available queue block by returning its base address in register R3. The routine then uses this area to temporarily store information during processing. When the routine no longer requires the block, it issues a QRET monitor call, which returns the queue block to the available list for later re-use.

The monitor queue system is necessary to provide storage for interrupt driven hardware (AM-300 board) and for storage during memory management operations. The queue blocks always reside in sharable system memory and therefore may be used by interrupt routines without regard to memory management context switching. The monitor queue system will be used more and more as the monitor is improved but is also available to the user if desired. The XLOCK subroutine (for multi-user locks in AlphaBasic) uses the monitor queue system to store the lock parameters.

5.1 INCREASING THE AVAILABLE QUEUE LIST SIZE

It is apparent that the number of queue blocks in use at any one time varies with system loading, number of users, and tasks being performed. Some applications may demand a larger available list of queue blocks to insure safe system operation. A check is performed to see if the available queue is exhausted. However, you can increase the size of the available queue list during system startup time.

The monitor is initially generated with 20 free blocks in the available queue. At any time in the SYSTEM.INI file prior to the final SYSTEM command

you may execute the QUEUE nnn command which allocates "nnn" more queue blocks for general use. A typical increase for a large system with several users running extensive applications might be 100 more blocks for a total of 120.

Once the system is up and running no more queue blocks may be added to the list, so you must give your best guess at your total requirements. The QUEUE command takes on a new life once the system is running. If you type the QUEUE command, the system responds by typing back the current number of free queue blocks in the available queue list. It is by this method that you may keep a close eye on the relationship between the system operation and queue block usage.

5.2 QUEUE BLOCK USAGE BY THE SYSTEM

This section lists the areas of the monitor which currently make use of the queue system, to give you a better idea on how to estimate your particular needs. Remember that this list will probably expand in future releases of the monitor. Also, add to this any applications that you may write which include the QGET and QRET calls (described in section 5.3).

The terminal service system makes frequent use of the queue system during output operations. A typical terminal driver may have up to four or five queue blocks in use at any one time, for linking buffers and storing immediate data values.

The monitor SLEEP call uses one queue block during the time the job is asleep.

The Persci disk driver uses one queue block while the head is loaded.

The XLOCK AlphaBasic subroutine uses one queue block for each separate system lock that is currently active by any job. This block is not returned to the available list until the lock is released by the job that has it locked.

The FLOCK AlphaBasic subroutine uses a number of queue blocks that varies with the number of jobs accessing files, the number of files open at one time, and the number of records open for each file. Currently, at any given moment during the use of FLOCK, the number of queue blocks being used equals:

twice the number of different files open using FLOCK, plus
the number of different records open using FLOCK, plus
the number of jobs with files open using FLOCK, plus
the total number of FLOCK opens (i.e., # of Action 0's)
that haven't been closed, plus
the total number of record uses (i.e., # of Action 3's)
that haven't been released

If FLOCK changes in the future, the above formula may also require modification.

The last two factors of the above equation anticipate circumstances where the same file and/or the same record is being accessed by more than one job at a time. If two jobs are reading the same file, that is two opens or two Action 0's.

The line printer spooler, as of version 4.1, uses the queue system to store the printer queue as well as a list of printers connected to the system.

5.3 QUEUE SYSTEM MONITOR CALLS

You can utilize the monitor queue system by using one of the four monitor queue management calls (QGET, QRET, QADD, QINS). These calls are fast for use in interrupt level routines. All calls work through register R3 and no other registers are disturbed. Since most queue blocks will be used in some form of sharable resource chain or interrupt level routine, the processor must be locked before executing any of the queue management calls. Violating this rule could destroy the available queue list or result in inter-job errors. None of the calls require any arguments to be passed except for the address in R3.

5.3.1 QGET - Obtain a Free Queue Block

This call obtains the first free queue block from the available list and returns its base address in R3. The Z-flag is set if the queue block was available, and is reset if no queue blocks were available. The queue block is first removed from the available list, and then all words in the block are cleared to zeros.

5.3.2 QRET - Return a Queue Block

This call returns a queue block to the available queue list in the monitor. The address which was in the first word of the block (usually a link to the next block in your chain) is returned in R3 after the block has been linked back into the available queue list. All queue blocks that have been allocated by QGET, QADD or QINS should eventually be returned to the monitor by the QRET call when they are no longer needed.

5.3.3 QADD, QINS - Manipulating Queue Blocks

Similar to the QGET call, these two calls obtain the first free queue block from the available list. The Z-flag is set if the queue block was available, and is reset if no queue blocks were available. If available, the queue block is linked into your own specific list whose address is in R3. This is because most system calls use queue blocks as elements of some specific list, depending on the application. The XLOCK subroutine, for

instance, maintains a list of all active system locks and adds or deletes queue blocks from this list as locks are set and reset.

The standard format of these individual lists follows the format of the free list. Each block links to its successor by storing its address in the first word of the block. All other words in the queue block are available for the storage of specific data. The last block in the list contains a zero in word 1 to mark the end of the list. The QADD call scans down the chain marked by the address in R3 and then inserts the new queue block at the end of the existing list. The QINS call inserts the new queue block in the chain at the point indexed by R3 and links the remaining list elements (if any) to the newly inserted block. Both calls then return the address of the second word of the new queue block in R3. This is the base of the data area of the queue block where you may store the data.

Remember that the current size of each queue block is eight words in length. The QADD and QINS calls place a link in the first word, leaving seven words of data storage for your application. The QRET call always requires the address of the first word when returning the queue block to the available list, regardless of the call used to obtain the block.

SEE ALSO
B.29 GFREE

CHAPTER 6

THE FILE SERVICE SYSTEM

The AMOS monitor has a simple yet powerful device-independent file service system which relieves the programmer of the task of I/O coding for each device with which he wishes his program to interface. In addition to this device independence, the monitor contains all routines to manage the disk file system on a logical-call basis. The programmer need not be concerned with the exact physical placement of files on the disk except in rare instances where the system software is being developed or tested. The monitor also contains an efficient means for developing new device drivers to be incorporated into the system when unsupported devices must be interfaced. This section gives a general overview of the file service system and describes the Dataset Driver Block (abbreviated as DDB) which is the descriptor link for all I/O and file calls to the monitor.

6.1 THE DATASET DRIVER BLOCK

All I/O operations and file operations are accomplished by monitor calls with reference to a DDB, which defines the device or file being operated upon. Whether the operation is to a unit-record device such as a printer, or to a specific file within a file-structured device such as a disk, depends upon the parameters passed to the monitor through the referenced DDB. There is no limit to the number of devices or files that may be active at any given time, but there must be one separate DDB for each device or file in use concurrently. There are no internal channel numbers or device numbers to limit the number of concurrently active devices or files. The general sequence of events for the complete processing of a device or file operation can be summed up as follows:

1. The DDB is set up with the defining parameters such as device name, drive number, filename and extension, project-programmer number, etc. This data normally comes from the processing of an ASCII file specification such as DSK1:FILTST.MAC[101,1] by an FSPEC call.

2. The I/O buffers are allocated either directly by the user program or by an INIT call referencing the DDB in use.

3. The logical opening processes for the device or file are performed, normally by an OPEN call referencing the DDB.

4. Data transfers to or from the device are performed by either READ and WRITE calls for physical transfers or INPUT and OUTPUT calls for logical transfers.

5. The logical closing processes for the device or file are performed, normally by a CLOSE call referencing the DDB.

The monitor contains complete error processing routines which allow the programmer to specify (by flags in word 1 of the DDB) whether any uncorrectable errors are to result in an automatic error message to the operator on his terminal, an aborting of the program and return to monitor, or both. You may also elect to process the errors yourself by checking the error code returned in word 1 of the DDB.

6.1.1 DDB Format

Figure 6-1 shows the format of the DDB which must be allocated within the user program area and set up by the user before any I/O operations can take place. The DDB is 24 (octal) words in size and is usually allocated by a BLKW 24 statement. The DDB can be assigned any tag which will then become the reference tag for all subsequent operations to that dataset. Some of the items in the DDB you must set up before certain operations may be called for, while other items are set up and used by the monitor file service routines. The following descriptions explain the use of each item.

6.1.1.1 Error Code - This byte is set to a non-zero code at the completion of an I/O operation that was unsuccessful for various reasons. A zero indicates the operation was successful. You need to test this byte only if the error control flag in the flags byte (DDB+1) specifies returning to the user on an error condition or if the operation allowed a non-fatal error condition to occur. The error codes are listed at the end of this section.

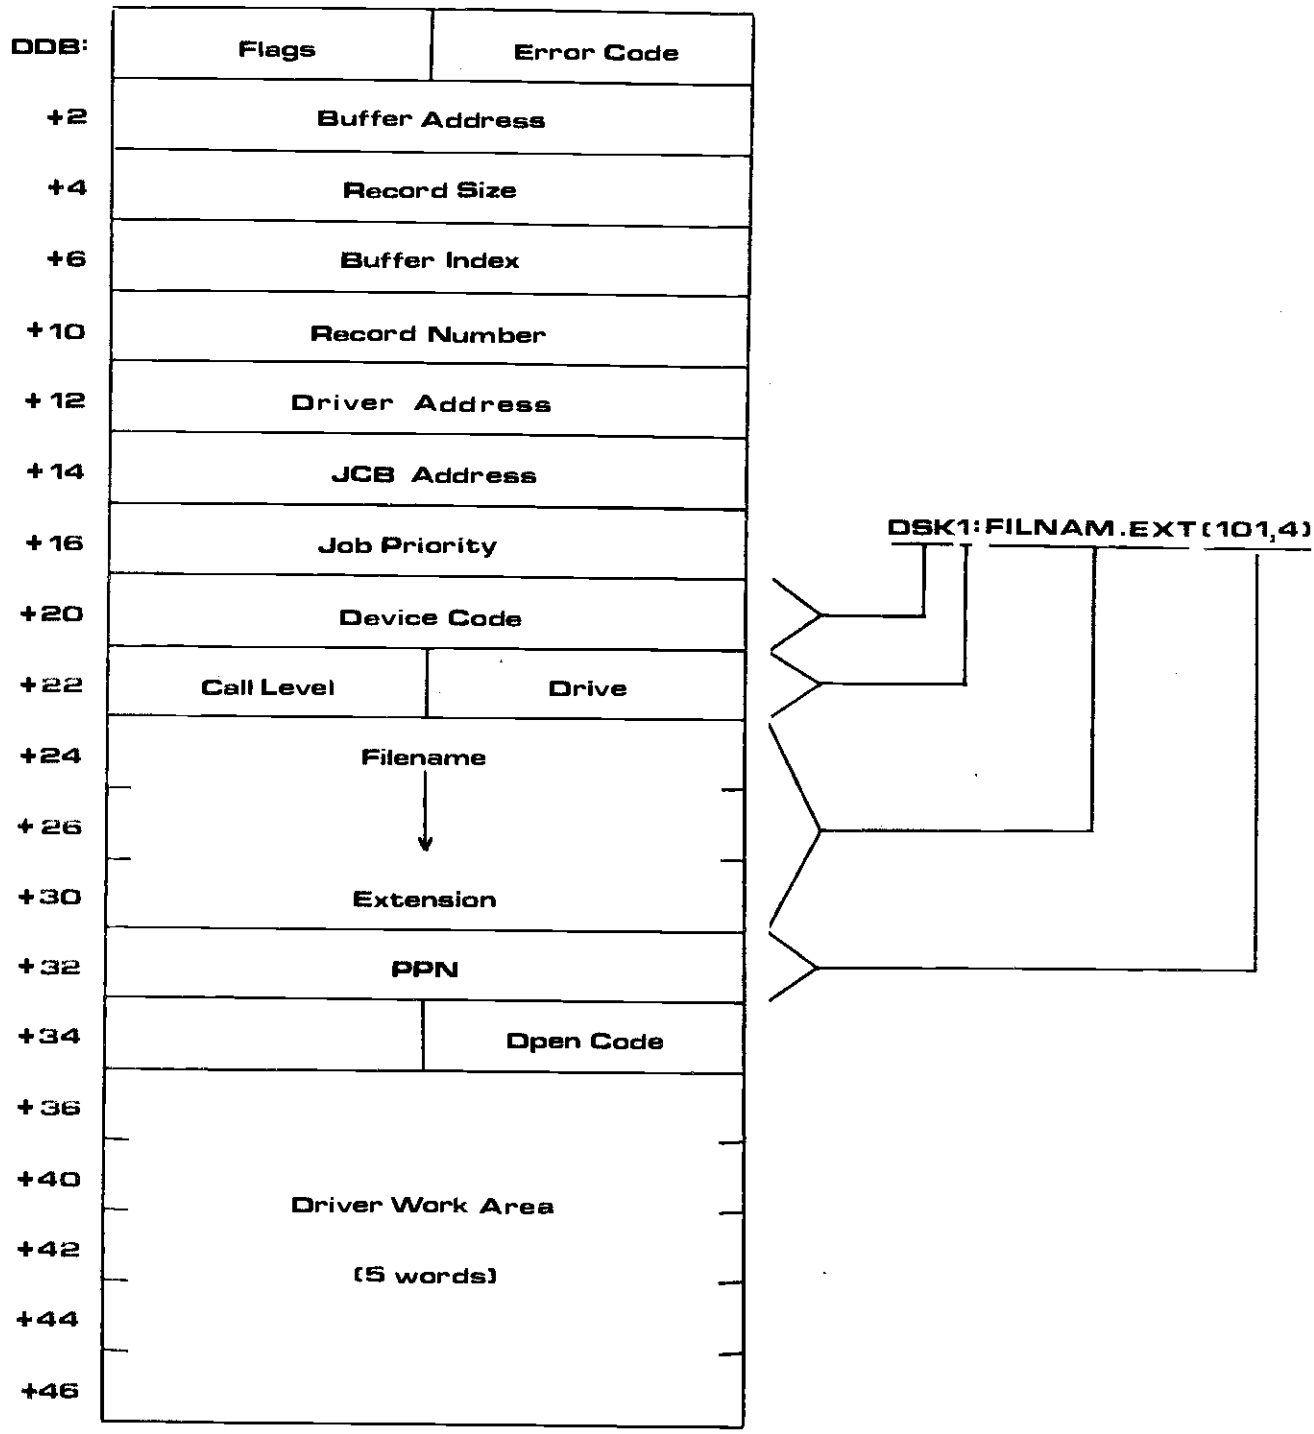

Dataset Driver Block

Fig 6-1

6.1.1.2 Flags - This byte is used to control the flow of the I/O operation and the handling of error codes by the file service routines. The following functions are controlled by the eight flag bits:

```
  1 - set by user to force a return on error condition (abort if clear)
  2 - set by user to bypass printing of error messages on error conditions
  4 - real-time transfer flag (currently not implemented)
 10 - spare
 20 - transfer initiated (for internal file service use only)
 40 - read if 0 or write if 1 (for internal file service use only)
100 - device INITed - set by INIT call or user if explicit buffer in use
200 - dataset busy (transfer initiated or queued)
```

6.1.1.3 Buffer Address - This is the 16-bit absolute address of the base of the buffer to be used for all dataset transfers (read and write). It is set by the INIT call which allocates a buffer, or by the user program if it is allocating its own buffer and not using the INIT call. This address is used in conjunction with the flag bit 6 above, which indicates that a buffer has been allocated either by the INIT call or by the user. No transfers can take place without a buffer.

6.1.1.4 Record Size - This is the size in bytes for the physical transfer to use. The READ call transfers this number of bytes from the device to the user buffer beginning with the address in DDB+2. The WRITE call transfers this number of bytes from the user buffer to the user device. The INIT call sets this size to the standard buffer size, or you can set the size if you are doing your own buffering. You may modify the size for transferring records of variable sizes as long as it does not exceed the buffer size of the capacity of the device or driver in use. Various logical file service routines set this size word during processing, such as the OPEN call for the disk which must perform directory operations on a 512-byte buffer at all times.

6.1.1.5 Buffer Index - This is a byte counter which is used by logical routines (INPUT and OUTPUT calls) for keeping track of bytes transferred into and out of the user buffer. Various calls reset this value, and you then use it and increment it as bytes are transferred into and out of the buffer. Details are given in later sections where the calls themselves are described. This buffer index word is normally not a true buffer pointer but rather an offset from the buffer base (per DDB+2) to the current byte being manipulated.

6.1.1.6 Record Number - You set the record number to read or write a specific random record from a random access device such as disk. The first record on the device is considered record zero, and the record numbers increment sequentially from there. This record number is actually used only by the physical driver routines for READ and WRITE calls, but other logical calls set this word to perform transfers to specific disk areas such as directory operations on disk. Most non-disk devices are not random access, in which case this record number is ignored by the respective drivers.

6.1.1.7 Queue Chain Link - This word is for internal use only, It is the link used by the I/O queueing routines for interrupt driven transfers. You should not alter this word.

6.1.1.8 JCB Address - File service routines store the address of the controlling job's JCB so that interrupt driven drivers can locate the corresponding job for activation on transfer complete status. This word is also for internal use only.

6.1.1.9 Job Priority - The current software job priority is set here by file service routines to specify the priority of the transfer in queued operations. This byte is for internal use only. The top byte of this word (DDB+17) is currently not used.

6.1.1.10 Device Code - The 3-character device code (packed RAD50) must be set here by an FSPEC call or directly by the user before any I/O operations may be performed.

6.1.1.11 Drive - Used only by drivers for devices with multiple drives, this byte must be set to specify the drive to be used for the transfer. A -1 byte (octal 377) may be used to indicate the current default drive number. If the device is DSK, the default drive used is the drive onto which you are currently logged. Other devices may have different defaults.

6.1.1.12 Call Level - For internal use only, this byte is used to keep track of the level of nesting of the file service calls for proper error recovery handling. This byte must be zero before the first file call is executed.

6.1.1.13 Filename and Extension - These are three words which contain the RAD50 packed filename and extension for file-structured devices. These words are ignored by drivers for devices which are not file-structured, but they may cause inaccurate error messages if they are not set to zero values.

6.1.1.14 PPN - This is the octal project-programmer bytes for the area to be used to locate the file. It is used only on file-structured devices which are multi-user based such as disk. A zero causes the default value to be the current PPN which the job is logged in under. To prevent inaccurate error messages, this word should be zero, if not used.

6.1.1.15 Open Code - This byte is set by the OPEN call to indicate the mode of the open statement for future processing operations. It is normally ignored by drivers for devices which are not file-structured. It is for internal use only and should not be modified by the user. The corresponding top byte of the word (DDB+35) is currently not used. The following open codes are in use:

0 - file is not open
1 - file is open for sequential input (OPENI call)
2 - file is open for sequential output (OPENO call)
10 - file is open for appending (OPENA call)
4 - file is open for random input/output (OPENR call)

6.1.1.16 Driver Work Area - The remaining five words are for internal use by the device drivers for links, record counts, etc., and should not be modified by the user during processing. Not all drivers make use of the work area, but it must be there if device independence is to be preserved.

6.1.2 Device Transfer Buffers

Each dataset must have an associated transfer buffer to handle input and output operations. This buffer must be allocated either directly or through use of the INIT call which allocates the buffer as a memory module by using a GETMEM call. The INIT call allocates a standard size buffer for the device being used (the size of the buffer is defined within the driver itself). If you do not wish to use the INIT call, you may allocate any size buffer you wish (must be large enough for any logical calls to be performed) and then set its address in DDB+2. Refer to the section detailing the I/O calls themselves for more details on the use of these buffers.

6.1.3 Error Handling

When an error occurs during any file service call, the file service routines normally perform typical error correction procedures. If the error is fatal (uncorrectable), two operations may or may not take place depending on the setting of bits 0 and 1 in the flags byte at DDB+1. First, bit 1 is tested and if it is not set, the monitor outputs a standard error message to the user terminal, giving the type of call that failed, the file specification for the device that the error occurred on, and the reason for the error. The appropriate error code is also placed in the error byte at DDB+0 for later testing by the user. Second, bit 0 of the flags byte is tested and if it is not set, the user program is aborted by the file service system and you are returned to monitor mode. You normally set these bits on before any I/O calls are made, if you wish to process the errors within the user program itself.

6.1.3.1 Error Codes - The following list gives the error code (in octal) returned in the DDB error byte by the file service system, along with the reason for the error:

01 - file specification error (FSPEC)
02 - insufficient free memory for buffer allocation (INIT)
03 - file not found (OPENI, OPENR, OPENA, DELETE, RENAME)
04 - file already exists (OPENO)
05 - device not ready (all calls)
06 - device full (OUTPUT)
07 - device error (all calls)
10 - device in use (ASSIGN)
11 - illegal user code (all file calls)
12 - protection violation (OPENO, OPENR, DELETE, RENAME)
13 - write protected (all output calls)
14 - file type mismatch
15 - device does not exist (all calls)
16 - illegal block number (READ, WRITE)
17 - buffer not initiated (all calls except INIT)
20 - file not open (READ, WRITE, INPUT, OUTPUT, CLOSE)
21 - file already open (all OPEN calls)
22 - bitmap kaput (all disk bitmap calls)
23 - device not mounted (all calls)
24 - invalid filename (OPENO, FSPEC, DSKCTG)

At the conclusion of every file service monitor call, the error byte at the base of the DDB is tested for the convenience of the user program. This allows you to test for an error status directly after the call with a BNE instruction without having to first explicitly test the byte with a TSTB instruction. This, of course, only applies if you have the error trapping bit set in the DDB status word to prevent the job from being aborted on a file error.

6.2 FILE SERVICE MONITOR CALLS

This section describes the file service calls which are available to the user program for both logical and physical I/O operations. All calls have the same general format, which uses a single argument representing the dataset driver block (DDB) to be used for the operation. See the preceding chapter for a complete description of the DDB format. In brief, the calls described in this section are:

FSPEC	process a device specification
INIT	initialize a dataset driver block buffer
LOOKUP	lookup a file to see if it exists
OPENI	open a file for sequential input
OPENO	open a file for sequential output
OPENA	open a file for appending
OPENR	open a file for random input/output
CLOSE	close a file to further processing
READ	read a physical record
WRITE	write a physical record
INPUT	read a logical record
OUTPUT	write a logical record
DELETE	delete a file
RENAME	rename a file
ASSIGN	assign a device to a job
DEASGN	deassign a device from a job

6.2.1 FSPEC - Process an ASCII Filespec

The FSPEC call is used to process an ASCII file specification from a command line (or any other ASCII buffer) and set up the parameters in the DDB according to the results of the processing. The ASCII file specification must be indexed by R2 and must be in the standard format of dev:filnam.ext[p,pn] with a valid termination character, if a short default specification is used.

The FSPEC call is slightly different from the rest of the I/O calls in that it allows you to use a second argument if you wish. This argument must be the default extension for the filename parameter to be used in the event that the file specification does not contain an explicit extension (identified by a period after the filename). If the second argument does not exist, the FSPEC processor does not process the input file specification past the colon which terminates the device/drive parameters.

The device code (3 characters) is packed RAD50 and stored in DDB+20 if it exists as marked by the terminating colon. The drive number is stored in the byte at DDB+22 if it exists. If the device code does not exist, the current default device (stored in the job's JCB item JOBDEV) is stored in DDB+20. If the drive number is not in the input specification an octal 377 is stored in DDB+22 to flag the default drive number to the device driver.

The filename and extension are then processed unless no second argument was used in the call, in which case the FSPEC processor returns to the user at this point. The filename and extension are packed RAD50 and stored in the three words at DDB+24 through DDB+30. If no filename is entered in the input specification, the word at DDB+24 is cleared to zero to flag the absence of the filename parameter. If a filename is entered but no extension is entered, then the default extension specified in the second argument of the FSPEC call is stored as the extension in DDB+30.

If a project-programmer number is in the file specification (marked by a left square bracket "["), it is processed and stored in DDB+32. If no p,pn is entered, DDB+32 is cleared to zero to flag its absence.

At the conclusion of the processing of the input file specification, the index R2 is pointing to the termination character (the first character following the file specification string). If an error in the input string is detected, the FILE SPECIFICATION ERROR message is printed (unless suppressed by bit 1 in DDB+1) and the program is aborted (unless suppressed by bit 0 in DDB+1). The error code 01 is set in DDB+0 error code byte.

No other modifications take place to the DDB area except that the error byte at DDB+0 is cleared at the start of the FSPEC processing. If you do not use the FSPEC call to set up your DDB, you must use some other form of explicit code to insure that the DDB is set up properly to define the device and file for any subsequent I/O operations.

6.2.2 INIT - Initialize the DDB

The INIT call is the normal means for allocating the dataset buffer and initializing the DDB for processing. The INIT call locates the device driver (searching [1,6] on DSK0: if not in memory), then allocates a standard size buffer based on the size specified in the driver. Bit 6 of the flag byte at DDB+1 is set to indicate the initialization. The address of the buffer is set into DDB+2, and the size in bytes is set into DDB+4.

No calls deallocate the buffer once it has been allocated by the INIT call. Multiple OPEN-CLOSE processes may be performed on the DDB once the INIT has been done. The buffer is temporary and is deallocated automatically when the program exits to monitor, or it can be explicitly deallocated by using the DELMEM call with the address stored in DDB+2. Recall that the buffer is allocated as a standard memory module with a GETMEM call.

NOTE

All file service calls with the exception of the FSPEC call require the use of a disk buffer, and therefore must be preceded by the INIT call for processing.

6.2.3 LOOKUP - Find the File

This is a form of the OPEN call which does nothing except search for the file and return an error code if it is not found. The file is not actually opened for processing, and an OPENI call must be used if the file is to be subsequently read from. The LOOKUP call is useful for determining if a file that is about to be opened for output already exists, so that it can first be deleted by the DELETE call. The LOOKUP call is ignored for devices which are not file-structured.

The LOOKUP call is also useful for some system programming techniques since it returns parameters about the file in the DDB work area. The work area is located in the last five words of the DDB. The first three words of this work area are loaded with the three words of the directory item if the file is found. These three words are the number of records in the file, the number of active data bytes in the last record, and the record number of the first data record in the file. Refer to Appendix A, "Disk Structure Format," for complete details on the directory format.

6.2.4 OPENI - Open a File for Input

The OPENI call locates a file in a file-structured device and sets up the DDB parameters (work area) for subsequent INPUT processing. An error results if the file is not found. The code 01 is set into DDB+34 to flag the OPENI operation. The OPENI call is normally followed by a series of INPUT calls which deliver sequential records from the file to the user buffer. The OPENI call is ignored for devices which are not file-structured.

6.2.5 OPENO - Open a File for Output

The OPENO call first searches the specified device in the specified user area and returns an error if the file already exists. If it does not, the DDB is set up for OUTPUT processing. The code 02 is set into DDB+34 to flag the OPENO operation. The OPENO call is normally followed by a series of OUTPUT calls which transfer data from the user buffer to sequential records in the file. The OPENO call is ignored for devices which are not file-structured.

6.2.6 OPENA - Open and Append to Existing File

The OPENA call is similar to OPENO, except that it allows you to append data to an existing file. The code 10 is set into DDB+34 to flag the OPENA operation. The OPENA call is normally followed by a series of OUTPUT calls which transfer data from the user buffer to the end of the file. This call is ignored for devices which are not file-structured.

6.2.7 OPENR - Open a File for Random Processing

The OPENR executes basically the same as the OPENI call, but the code stored in DDB+34 is 04 to flag random processing. The file located for random processing must be a contiguous file. The OPENR call is normally followed by a series of INPUT and OUTPUT calls which transfer data between specific records in the file and the user buffer in both directions. The OPENR call is also ignored for devices which are not file-structured.

6.2.8 CLOSE - Close a File

The CLOSE call finishes up logical processing of a file and clears the open code in DDB+34. No further INPUT or OUTPUT operation may occur once a file has been closed. No action is normally done on a file which is open for input. For files open for output, the final record is written out and the file is added to the directory system on the specific device. The CLOSE call is ignored for devices which are not file-structured.

6.2.9 READ - Perform a Physical Transfer

This is the physical transfer call for reading input data from a device. No check is made for file open status since the READ call is not a logical file call.

6.2.9.1 Sequential Devices - For sequential access devices such as a paper tape reader, the READ call delivers one record from the device to the user buffer. The size of this record is normally the number of bytes specified in DDB+4, but this may not necessarily be true if the driver does not transfer under the rules of the system. If the device is not capable of generating the requested number of bytes per DDB+4 (such as a tape reader which runs out of tape), a lesser number may be transferred in which case the count in DDB+4 is adjusted to reflect the true number actually transferred to the user buffer.

6.2.9.2 Random Devices - For random access devices such as disk, you must specify the record number to be located and read, by placing that number into DDB+10 before executing the READ call. Most random access devices always transfer the requested number of bytes per DDB+4 into the user buffer. (If the buffer is larger than the physical block, the system reads multiple contiguous blocks to fill up the buffer.) An error results if the record number is not within the range of the specific device. For example, the standard AMOS floppy disk is structured as 500 (decimal) records of 512 bytes each. The legal record numbers therefore range from 0 through 499, decimal. Similar range restrictions apply for each random device.

6.2.9.3 Interrupt Structure - The system allows interrupt driven devices to be queued and processed in a priority fashion. Normally, the execution of a READ call suspends the running of the user program until the transfer has been completed, at which time the user job is reactivated. You must then either test the dataset busy bit (bit 7) of the flag byte or use the WAIT call to stall until the transfer has been completed. The dataset busy flag is reset when the transfer has been completed. You must then check for errors. The realtime bit is ignored for devices which are not interrupt driven or whose drivers do not run under the I/O queue system.

6.2.10 WRITE - Perform a Physical Write

This is the physical transfer call for writing data to a device. No check is made for file open status, since the WRITE call is not a logical file call.

6.2.10.1 Sequential Devices - For sequential access devices such as a printer, the WRITE call delivers one record to the device from the user buffer. The size of this record is the number of bytes specified in DDB+4. The driver is responsible for the correct transfer count, and you may alter the number in DDB+4 for each new WRITE call to the same device for the writing of variable length records.

6.2.10.2 Random Devices - For random access devices such as disk, you must specify the record number to be located and read, by placing that number into DDB+10 before executing the WRITE call. Most random access devices always transfer the requested number of bytes per DDB+4 into the user buffer. An error results if the record number is not within the range of the specific device. The standard AMOS floppy disk is structured as 500 (decimal) records of 512 bytes each. The legal record numbers, therefore, range from 0 through 499, decimal.

6.2.10.3 Interrupt Structure - The system allows interrupt driven devices to be queued and processed in a priority fashion. Normally, the execution of a WRITE call suspends the running of the user program until the transfer has been completed, at which time the user job is reactivated. You must then either test the dataset busy bit (bit 7) of the flag byte or use the WAIT call to stall until the transfer has been completed. The dataset busy flag is reset when the transfer has been completed. You must then check for errors. The realtime bit is ignored for devices which are not interrupt driven or whose drivers do not run under the I/O queue system.

6.2.11 INPUT - Perform a Logical Read

The INPUT call is the logical equivalent of the READ call for logical processing of datasets. The INPUT call reads a logical record within a file or device dataset under the control of the specific driver in use. A dataset must be opened for input (OPENI) or random access (OPENR) before INPUT calls are performed. The INPUT call first sets the standard buffer size into DDB+4, so you may not use this call to transfer non-standard record sizes. The number of bytes actually read may be less than the standard record size due to the driver processing or due to an end-of-file condition. The actual number of bytes transferred is set into DDB+4 by the driver routine.

6.2.11.1 Sequential File Processing - The INPUT call is mainly used in logical sequential file processing; it sets up the buffer index value in DDB+6 to direct the processing of the data by the user routines. This index value is actually the offset to the first byte of valid data within the user buffer, whose base address is at DDB+2. For unit record devices, the value is zero since all data within the buffer is user data. For sequential disk files, however, the first word in each record within the file is a link word to the next record; therefore, the value set into DDB+6 by the disk driver is 2, so that processing starts with the third byte in the user buffer.

6.2.11.1.1 Example - The following subroutine is normally used to get each byte of data from a sequential file:

```
;Subroutine to get next byte from file defined as INDDB and leave it in R1
;
INBYTE: CMP     INDDB+6,INDDB+4 ;is the buffer empty?
        BLO     INBG            ;  no - get next byte
        INPUT   INDDB           ;read next logical record into buffer
        CMP     INDDB+6,INDDB+4 ;check for end of file (no data transferred)
        BEQ     INEOF           ;  go to end of file routine
INBG:   PUSH    INDDB+2         ;stack the buffer base address
        ADD     INDDB+6,@SP     ;  and add the index offset to get position
        MOVB    @(SP)+,R1       ;pick up the next byte from user buffer
        AND     #377,R1         ;insure upper byte is cleared in R1
        INC     INDDB+6         ;increment the buffer index for next time
        RTN                     ;subroutine return
```

6.2.11.2 Random File Processing - A special situation arises for files opened for random access by the OPENR call. Instead of the next sequential record being read, the specific relative record whose number is in DDB+10 is read into the user buffer. You first set this number up and then execute the INPUT call. The record number is actually relative to the base of the file and has no direct relationship to the physical record on the device as would be returned by a READ call.

6.2.11.3 Special Devices - For devices that do not implement special processing of logical calls, the INPUT call performs a READ call instead.

6.2.12 OUTPUT - Perform a Logical Write

The OUTPUT call is the logical equivalent of the WRITE call for logical processing of datasets. The OUTPUT call writes a logical record to a file or device dataset under the control of the specific driver in use. A dataset must be opened for output (OPENO) or random access (OPENR) before OUTPUT calls are performed. The OUTPUT call transfers the number of bytes in DDB+4, but it normally does it as a standard record (depends on the driver in use). We discourage attempts to use the OUTPUT call for transferring non-standard record sizes.

6.2.12.1 Sequential File Processing - The main use of the OUTPUT call is in logical sequential file processing. The OUTPUT call sets up the buffer index value in DDB+6 to direct the processing of the data by the user routines. This index value is actually the offset to the first byte position for valid data within the user buffer whose base address is at DDB+2. For unit record devices this value is zero, since all data within the buffer is user data. For sequential disk files, however, the first word in each record within the file is a link word to the next record; therefore, the value the disk driver sets into DDB+6 is 2, so that processing starts with the third byte in the user buffer.

6.2.12.1.1 Example - The following subroutine is normally used to put each byte of data to a sequential file:

```
;Subroutine to put next byte from R1 into file defined as OTDDB
;
OUTBYT: CMP     OTDDB+6,OTDDB+4 ;is the buffer full now?
        BLO     OUBYT           ;no - add this byte
        OUTPUT  OTDDB           ;yes - write it
MOUBYT: PUSH    OTDDB+2         ;stack the buffer base address
        ADD     OTDDB+6,@SP     ;  and add index offset to get position
        MOVB    R1,@(SP)+       ;move data byte to user buffer
        INC     OTDDB+6         ;increment the buffer index offset value
        RTN                     ;subroutine return
```

6.2.12.2 Random File Processing - A special situation arises for files opened for random access by the OPENR call. Instead of the next sequential record being written, the specific relative record whose number is in DDB+10 is written out from the user buffer. You first set this number up and then execute the OUPUT call. The record number is actually relative to the base of the file and has no direct relationship to the physical record on the device as would be written by a WRITE call.

6.2.12.3 Special Devices - For devices that do not implement special processing of logical calls, the OUTPUT call performs a WRITE call instead.

6.2.13 DELETE - Delete a File

The DELETE call deletes a specific file from a file-structured device. The filename, extension and p,pn (if used) must be set in the DDB before executing the call. An error results if the file is not found. The DELETE call is ignored for devices which are not file-structured.

6.2.14 RENAME - Rename a File

The RENAME call renames a specific file on a file-structured device. The filename, extension and p,pn (if used) must be set in the DDB before executing the call. The new filename and extension must be packed RAD50 into the three words immediately following the DDB in memory. The RENAME call merely locates the directory item for the file and replaces the three words which store the filename and extension. The RENAME call is ignored for devices which are not file-structured.

6.2.15 ASSIGN - Assign a Device

The ASSIGN call is used to assign a non-sharable device (such as a printer) to the current user's job by setting a flag in the device's entry in the device table in monitor memory. Once a device has been assigned by this call, any attempt to assign it by another job results in an error. The device stays assigned to this job until deassigned by the DEASGN call. The ASSIGN call performs no action if the specified device is sharable, such as a disk.

6.2.16 DEASGN - Deassign a Device

The DEASGN call is used to deassign a device which has been assigned to the user's job by the ASSIGN call. Once deassigned, the device becomes available for assignment by other jobs. The DEASGN call performs no action if the specified device is sharable or if it is not currently assigned to the user's job. All devices are deassigned when the program exits to the monitor.

6.3 DISK SERVICE MONITOR CALLS

In the previous section we covered the file-oriented monitor calls. Those calls allow you to access data files without regard to the actual structure of the data on the device. Internally, of course, AMOS does have to deal with the structure of the data. This section deals with the monitor calls used to manipulate that structure. A description of the data structures used to maintain files on a device can be found in Appendix A, "Disk Structure Format."

The disk presents special problems which require the use of special monitor calls to control the accessing of the directory and bitmap records. These records have a non-sharable attribute associated with them, even though the disk in general is a sharable device. For instance, two user programs may not both be updating the same directory records at the same time. The same holds true for the bitmap records. The following monitor calls are used to control the access to these non-sharable records:

DSKCTG - allocates a contiguous file for random processing
DSKALC - allocates the next available record on disk
DSKDEA - deallocates a specific record on disk
DSKBMR - reads disk bitmap and sets re-entrant lock flag
DSKBMW - rewrites disk bitmap after user modification
DSKDRL - sets re-entrant directory lock for a specific user
DSKDRU - clears re-entrant directory lock for a specific user

The access to these records is normally done by the monitor routines as a direct result of normal I/O processing by file service calls. It is a somewhat tricky process and the disk calls should not be used except with extreme caution, since misuse could violate the integrity of the file structure on the disk. The following descriptions are directed at those system programmers who are familiar with shared file techniques.

6.3.1 Calling Sequence

All calls use a standard argument which is the address of the associated DDB to be used for the call. In addition to the first argument which is the DDB, some calls use an optional second argument for processing. The second argument is detailed in the description of the call.

6.3.2 The Bitmap Area

The bitmap area is an area in monitor memory which is allocated by the BITMAP program run at system startup time by the BITMAP command in the system initialization command file. This area consists of a status word, a DDB for bitmap reads and writes, and a buffer for the actual bitmap including the hash total words. The format of the bitmap area is as follows:

```
BLKW    0               ;Bitmap status word
BLKW    12              ;Partial DDB for bitmap I/O
BLKW    Bitmap-size     ;Bitmap buffer (size depends on device)
BLKW    2               ;Hash total words
```

The device table entry for each drive has the address of the corresponding bitmap area to be used for that drive. More than one drive may share the same bitmap area, forcing a rewrite each time a different drive is referenced. This is not efficient with regard to time but can save some memory for larger devices where the bitmap buffer may be several hundred words or more.

6.3.2.1 The Status Word - The status word (first word in bitmap area) contains two flags which are used to control bitmap access. Bit 0 is the bitmap lock flag and is set to flag that the bitmap is locked and being read or modified by some user job. The DSKBMR call sets this flag on, and it is up to you to clear it after you have finished the bitmap access and modification. Bit 1 is the bitmap rewrite flag which is set to indicate that one or more modifications have been made to the bitmap in memory, and that it must be rewritten to disk before being discarded. If the user program modifies the bitmap in memory, it must set the rewrite flag to insure that the bitmap is rewritten.

6.3.2.2 The Bitmap DDB - The bitmap DDB is a partial DDB because no files are ever referenced, and the rest of the DDB is not needed. The bitmap is normally allocated as record 2 of each disk, and it extends across successive records for those devices which overflow one record.

6.3.2.3 The Bitmap Buffer - The bitmap buffer area is the exact size required to contain the entire bitmap from the disk. Two extra words are allocated to contain the hash total which is used to insure the integrity of the bitmap in memory and on disk. Each time the bitmap is read, or before the bitmap is rewritten, this hash total is checked and an error results if it is bad. The hash total is merely the double-word binary sum of the entire bitmap buffer. You must update this hash total each time you modify the bitmap, or else an error results when it is time to rewrite the bitmap to disk.

6.3.2.4 The Bitmap - The bitmap itself contains one bit for each logical record on the disk structure. This bit is off if the record is free, and on if the record is in use by anyone, including the system structure records themselves. Each word in the bitmap can define up to 16 records. The first word in the bitmap defines records 0 through 17 (octal) with bit 0 defining record 0 and proceeding upward throughout the word. The second word defines records 20 through 37, and so on. To define the 500 decimal records in a standard IBM-compatible AMOS floppy disk, we need 32 words (32 times 16 = 512) with the last word not being totally used. The bitmap itself therefore takes up 34 words, including the two hash total words.

6.3.2.5 Altering the Bitmap - Altering the bitmap is tricky but the sequence recommended is:

1. Read the bitmap using the DSKBMR call
2. Alter the bitmap as necessary (recompute the hash total)
3. Set the rewrite flag (status word bit 1)
4. Clear the bitmap lock (status word bit 0)
5. Rewrite the bitmap using the DSKBMW call

6.3.3 DSKCTG - Allocate a Contiguous Area

The DSKCTG call is used to allocate a contigous file on a random access device. A standard argument is used as the second argument which represents the number of records to be allocated in the file. A search is made to find the first available area on the disk which can fully contain the requested number of records. These records are marked as in-use on the disk bitmap, and a file descriptor item is added to the user directory. The word which gives the number of bytes in the last record is set negative to flag this file as contiguous, distinguishing it from the normal sequential files. A device-full error results if no area can be found on the disk which is large enough to contain the file.

6.3.4 DSKALC - Allocate a Record

The DSKALC call is used to allocate one record for use by this user as a directory record or as a file record. A standard argument is used as the second argument, which represents the word that is to receive the record number of the allocated record. An error results if there are no free records left on the specified disk. A DSKBMR call is first performed to insure that the current job has access to the bitmap, and then the first free record is located and marked in use. The bitmap record is flagged as modified, causing it to be rewritten at the next DSKBMW call or if it must be swapped out to make room for another bitmap sharing the same area in memory.

6.3.5 DSKDEA - Deallocate a Record

The DSKDEA call is used to deallocate a specific record on a disk and make it immediately available for use by another user (or the same user). A standard argument is used as the second argument, which represents the address of the word containing the record number of the record to be deallocated. No check is made to insure that this record is allocated to either the current user or any other user. A DSKBMR call is first performed to insure that the current job has access to the bitmap, then the specified record's bit is set to zero to indicate that the record is free. The bitmap record is flagged as modified to force a rewrite.

6.3.6 DSKBMR - Read the Bitmap

The DSKBMR call locates the bitmap area in monitor memory for the specified disk and insures that it is not locked by another job. If it is locked, a stall is made until it is released. It is then locked for this job and a return is made to the user. The address of the bitmap area is set into the word specified by the second argument in the calling sequence. The second argument is a standard argument in format. Refer to the description of the bitmap area above and note that the second argument receives the address of this area and not the address of the bitmap itself. You may locate the bitmap itself because its address is in the second word of the bitmap area (second word of the bitmap DDB).

6.3.7 DSKBMW - Write the Bitmap

The DSKBMW call locates the bitmap area in monitor memory for the specified disk and insures that it is not locked by another job. If it is locked, a stall is made until it is released. It is then locked for this job and rewritten to disk from memory unless the hash total is bad. After the rewrite is complete both the rewrite and lock flags are cleared and a return is made to the user.

6.3.8 DSKDRL - Lock the Directory

The DSKDRL call locks the directory for the specified drive for modification by the user program. It is used by such file service routines as CLOSE for output files, DELETE and RENAME calls. If the directory is already locked by another job, a stall is made until it is released. The user program or routine must unlock the directory via the DSKDRU call after the modifications have been made.

6.3.9 DSKDRU - Unlock the Directory

The DSKDRU call unlocks the directory for the specified drive after it has been locked by the DSKDRL call for modification. No action is performed if the directory is not locked by the current job.

6.4 MAGNETIC TAPE DRIVER MONITOR CALLS

Some monitor calls allow your assembly language programs to access the magnetic tape unit driver, MTU.DVR. For information on using the magnetic tape utility programs, refer to Using the Magnetic Tape Unit in the "User's Information" section of the AMOS Software Update Documentation Packet. That document also defines some of the terms we use in the following discussion.

Before you begin use of MTU.DVR, make sure your magnetic tape units are defined in your system device table, and that the program MTSTAT.SYS has been included in the monitor (via the SYSTEM command in the system initialization command file).

In addition to the magnetic tape drive monitor calls detailed below, you can use the READ and WRITE calls to input and output data to and from the magnetic tape unit, in the same way you would use them to perform disk I/O.

6.4.1 REWIND Arg

This call issues a rewind command to the specified tape unit. REWIND accepts a standard argument that represents a DDB on which you have already performed an FSPEC, an INIT, and an OPEN monitor call.

The DDB selects the device to which you want to issue a REWIND command. If an error results from this call, you see the standard system file operation error messages (e.g., ?Cannot INIT Devn: - device does not exist).

6.4.2 WRTFM Arg

This call issues a write-file-mark command to the specified tape unit. WRTFM accepts a standard argument that represents a DDB on which you have already performed an FSPEC, an INIT, and an OPEN monitor call.

The DDB selects the device to which you want to write a file mark. If this call results in an error, you see the standard system file operation error messages.

6.4.3 FMARK Arg

This call issues a find-file-mark command to the specified tape unit. FMARK accepts a standard argument that represents a DDB on which you have previously performed an FSPEC, an INIT, and an OPEN monitor call. The DDB selects the device to which you want to issue a find-file-mark command. The FMARK call causes the MTU driver to read forward on the specified tape until it finds a file mark. Any errors resulting from this call are indicated by standard file operation error messages.

6.4.4 FMARKR Arg

FMARKR causes the MTU driver to read in reverse on the tape until it finds a file mark. The call accepts a standard argument that represents a DDB on which you have previously performed an FSPEC, an INIT, and an OPEN monitor call. The DDB selects the device to which you want to issue the FMARKR command. Any resulting error is indicated by standard file operation error messages.

6.4.5 TAPST Arg1,Arg2

This call issues a read-tape-status command to the specified tape unit. TAPST accepts two standard arguments. The first, Arg1, represents a DDB on which you have previously performed an FSPEC, an INIT, and an OPEN monitor call. The DDB selects the device whose status you want to return. The returned status code appears in Arg2. The staus bits TAPST returns are as follows:

BIT	FUNCTION	COMMENTS
0	7-track	Indicates that unit is in 7-track mode.
1	NRZI mode	Indicates that unit is in NRZI recording mode.
2	End-of-tape	Indicates that end-of-tape was detected during the previous command.

CHAPTER 7

TERMINAL SERVICE SYSTEM

The AMOS monitor has several calls which deliver data to and from both the user terminal and other terminals connected to the system. A terminal is defined as an ASCII character-oriented device which is capable of both output and input. This is the formal definition and does not preclude the use of output-only devices on terminal designated ports. Also, the system includes software terminals known as "pseudo terminals," which can be used to control jobs that are not actually associated with a hardware interface on a designated port address. The calls listed here normally input from or output to the terminal which is controlling the job that is executing the call. Some calls (as specified) will input from or output to another terminal not connected to the current job or to a pseudo terminal controlling another job.

Programs which make use of the standard terminal service calls that communicate with the user terminal can be run without modification in a job controlled by a pseudo terminal. Keyboard input calls and terminal output calls always go to the controlling terminal, regardless of which job they are running in. Therefore, you need not be concerned with the physical port address or attributes of the terminal which is controlling the job. The monitor routines handle all this automatically.

7.1 TERMINOLOGY

Due to a holdover from older system terminology, most terminal output calls reference the device name of "TTY," which used to define the teletype device on systems that normally used teletypes as terminals. The input device of the teletype was then called the keyboard, and the calls reference the device name of "KBD." These are strictly mnemonics and do not necessarily reflect the physical attributes of the terminals, which now are more commonly the higher speed video display terminals.

7.2 THE TERMINAL LINE TABLE

Each terminal has associated with it a terminal line table which is a work area in monitor memory set up to contain the parameters and work areas associated with the control of the terminal device. Most of the items in this terminal line table are for internal use only, and you need not be concerned with them. The JOBGET Rx,JOBTRM call may be used to set an index to the associated terminal line table, so that you can inspect or modify the items within.

7.2.1 The Terminal Status Word

Normally, you need to be concerned only with the terminal status word, which is the first word in the terminal line table. This word has certain flags in it that you may modify to alter the operation of your terminal calls. The terminal status word has the following flag positions defined:

```
    1 - user sets to force image mode input (see KBD call)
    2 - user sets to suppress echoing of input characters
    4 - user sets to allow escapes to be processed (as in EDIT)
   10 - engages data mode to allow complete data transparency on input
        (^C, nulls, bit 8 characters).
   20 - user sets to allow lower-case input (disables conversion)
  200 - internal flag used to indicate output is in progress
 1000 - flag used to indicate "hog" mode for terminal (set by TRMDEF)
 2000 - user sets to indicate terminal runs in local mode (no echo)
```

The terminal status word is cleared each time the user program exits back to monitor mode upon program completion, thereby restoring normal terminal operation regardless of program operation.

7.3 THE TERMINAL SERVICE CALLS

AMOS includes 17 monitor calls to perform input and output between the system and any of its connected terminals.

7.3.1 KBD {label} - Fetch a Line of Data

The KBD call accepts one full line of input from the user terminal into a monitor line buffer, then sets index R2 to the base of that buffer for the user reference. During the inputting of the line, the user job is set into the terminal input wait state, thereby consuming no CPU time until the line is finished. All normal line editing features are active (rubout, control-U, tab, etc.) and a control-c input aborts the job unless the user has set up control-c trapping via the JOBICP item in the JCB for the job. If you specify a label with the KBD call, the program automatically branches to that label. The line is terminated when a carriage-return or a line-feed is entered. The monitor automatically appends a line-feed to the carriage-return, and a null byte is set after the line-feed character.

If the echo-suppress flag is set in the terminal status word, normal echoing of the input characters is suppressed, such as when the password is being entered for the LOG command. If the image-mode input flag is set, the KBD command has a different effect. No editing is performed and instead of one line being accepted, only one character is accepted and it is delivered back to you in register R1 instead of register R2 being set to the monitor line buffer. Image-mode input echoing is still under control of the echo-suppress flag as in normal line mode.

7.3.2 TTY - Output One Character

The TTY call outputs one character from register R1 to the controlling terminal and then returns. Tabs are echoed as spaces up to the next modulo-8 carriage position, unless the image-mode output flag is set in the terminal status word. If the job is running under the control of a command file, the character will only be output to the terminal if the output suppress command is in the normal state (:R revives it, :S silences it).

7.3.3 TIN - Get an Input Character

TIN gets the next input character from either the terminal input buffer or from the command string if the job is controlled by a command file. The character is delivered in R1. This call is normally only used within the operating system itself and not by user programs.

7.3.4 TOUT - Output One Character

TOUT outputs one character to the controlling terminal of the job or to the job which has this job attached (by the address in the JOBATT item). This call differs from the general TTY call in that the command file status is not checked by the TOUT call. The TOUT call, like the TIN call, is normally only used within the operating system itself.

7.3.5 TAB - Output One Tab

This convenience call outputs a single tab character to the user terminal. In effect, it is the same as the code sequence:

```
MOVI    11,R1
TTY
```

7.3.6 CRLF - Output a Carriage-Return / Line-Feed

This convenience call outputs a carriage-return and line-feed pair to the user terminal. In effect, it is the same as the code sequence:

```
MOVI    15,R1
TTY
MOVI    12,R1
TTY
```

7.3.7 TTYI - Output a String of Characters

The TTYI call outputs a string of characters which follows the call itself up to but not including a null byte. The call could be used as follows to output two lines of data to the user terminal:

```
TTYI
ASCII   /LINE 1 DATA/
BYTE    15
ASCII   /LINE 2 DATA/
BYTE    15,0
EVEN
```

The TTYI call also automatically appends a line-feed to all carriage-returns included in the string.

7.3.8 TTYL - Output a String of Characters Indexed

The TTYL call is similar to the TTYI call in that it outputs a string of ASCII characters up to a null byte. The string of characters for the TTYL call may be anywhere in memory and not in line with the call itself in the program flow. TTYL takes one standard argument--the address of the message to be output. It is therefore useful for outputting from a table of messages by setting an index to the specific message within the table (per some numeric director code), and then using that register as the argument to the TTYL call. The TTYL call also appends a line-feed to each carriage-return in the string.

7.3.9 PTYIN - Place Character in Input Buffer

The PTYIN call allows one job to force a character into the input buffer of another job which is probably controlled by a pseudo terminal. This call takes two standard arguments. The first is the data byte to be sent to the other job and the second argument is the address of the JCB of the job into which the character is to be forced. PTYIN is the call through which the FORCE operator command functions.

7.3.10 PTYOUT - Fetch Character from Output Buffer

The PTYOUT call allows one job to get a character from the terminal output buffer of another job which is controlled by a pseudo terminal. If no output is available from the specified job, the calling job is put to sleep until a character is available. The PTYOUT call takes two standard arguments. The first argument is the address of the byte which will receive the data character, and the second argument is the address of the JCB from which the character is to be taken.

7.3.11 TTYIN - Fetch Another Job's Input

The TTYIN call allows one job to get waiting input data from the terminal input buffer of another job. This call has not yet been fully inplemented.

7.3.12 TTYOUT - Place a Character in Another Job's Output

The TTYOUT call allows one job to put data into another job's terminal output buffer. This call, like the TTYIN call, is not yet fully implemented.

7.3.13 TRMICP - Process Input Character Within Interface Driver

The TRMICP call is executed from within a terminal interface driver to process one character which has just been received from the terminal by the hardware interface. R1 must contain the input character to be processed, and R5 must index the terminal definition table entry for the specific terminal being serviced. TRMSER then takes the character and passes it to the terminal driver input routine for pre-processing if desired. When the terminal driver passes it back to TRMSER, it is then edited for control codes and other special characters and then added to the terminal input buffer. All the pertinent flags are set automatically to indiciate actions to be taken by the application program when it requests the input data. If the input character is a break character (line-feed), or if image mode is active, the associated job is awakened to process the available data.

7.3.14 TRMOCP - Process Output Character Within Interface Driver

The TRMOCP call is executed from within a terminal interface driver to get from TRMSER the next output character to be sent to the terminal. This is usually in response to an interrupt from the interface board, indicating that the prior character has been fully output and the board is ready to transmit the next character. R5 must index the terminal definition table entry for the specific terminal being serviced, and R1 gets the next available character upon return from TRMSER processing of the call. If

there is no more output available in the output buffer, R1 is set to -1 as a flag, and the associated job is awakened to fill the output buffer again.

7.3.15 TRMBFQ - Process Output Characters Within Terminal Driver

The TRMBFQ call is a physical output call usually executed from within a terminal driver or a monitor routine. There are, however, times when it can be used by an assembly language application program. The TRMBFQ call effectively adds a buffer full of data characters to the output buffering system for a specific terminal. It does this by linking the buffer into the dynamic output queue list used by TRMSER for this terminal. When this call is used, R2 must index the buffer to be queued, R3 must contain the number of characters in the buffer, and R5 must index the terminal definition table entry for the specific terminal. The TRMBFQ call performs the output initiation function if the output system for the terminal is currently idle.

7.3.16 TBUF - Output Large Amounts of Data

The TBUF call is the normal call for user programs to use for queueing up large amounts of data into the terminal output system of a terminal where the single character calls are considered inefficient. It is a buffered call in that it works through the two output buffers for the terminal, as opposed to going directly into the output queue system. If you try to output more data via the TBUF call than there is currently room for in the output buffers, the user job is suspended while the output buffers are unloaded to the terminal. Each time one of the output buffers is emptied, the job is awakened and the TBUF call proceeds to fill that buffer. This continues until the original amount of data is exhausted, at which time the call returns to the user program. When the call is executed, R2 must index the buffer to be output and R3 must contain the number of characters to be output (similar to the TRMBFQ call). R5 need not index the terminal definition table entry since this is a user level call.

7.3.17 TCRT - Call Special Terminal Driver Routines

The TCRT call is the linkage into the special processing routine portion of a terminal driver. R1 usually contains a 2-byte code which is interpreted by the terminal driver routine as a special function, such as cursor positioning or special editing action. The only action actually performed by the TCRT call within TRMSER is to locate the terminal driver for the attached terminal and call the driver control routine within it. You must refer to the actual driver listing to determine the action performed relative to the code passed to it in R1.

7.3.17.1 Standard Functions - The TCRT call is most commonly used for controlling such special CRT functions as cursor addressing and screen clearing. To maintain compatibility between terminal drivers, Alpha Micro has defined the following functions within the terminal drivers it supports.

7.3.17.1.1 Cursor Addressing - To perform cursor addressing, R1 is loaded with a 2-byte argument defining the screen row and column to which the cursor is to be moved. The high-order byte is loaded with the row, and the low-order byte is loaded with the column. The uppermost-leftmost (Home) position is column 1, row 1.

7.3.17.1.2 Other Functions - To perform other special CRT functions, the high-order byte of R1 should be loaded with 377 (octal). The low-order byte is then loaded with one of the special function codes listed below.

0	Clear Screen and set normal intensity
1	Cursor Home (move to 1,1)
2	Cursor Return (move to column 1 without line-feed)
3	Cursor Up one row
4	Cursor Down one row
5	Cursor Left one column
6	Cursor Right one column
7	Lock Keyboard
8	Unlock Keyboard
9	Erase to End of Line
10	Erase to End of Screen
11	Enter Background Display Mode (reduced intensity)
12	Enter Foreground Display Mode (normal intensity)
13	Enable Protected Fields
14	Disable Protected Fields
15	Delete Line
16	Insert Line
17	Delete Character
18	Insert Character
19	Read Cursor Address
20	Read Character at Current Cursor Address
21	Start Blinking Field
22	End Blinking Field
23	Start Line Drawing Mode (enable alternate character set)
24	End Line Drawing Mode (disable alternate character set)
25	Set Horizontal Position
26	Set Vertical Position
27	Set Terminal Attributes

Not all terminal drivers have all of the above functions, simply because all terminals do not have all of the functions. If your terminal has additional features, Alpha Micro recommends starting at 100 (octal) when assigning function codes.

7.3.18 Message Calls

Three calls have been defined in SYS.MAC as macros using the TTYI call. These calls are for the convenience of the programmer and to make the program more readily understandable. They all take a single argument which is an ASCII message string to be output to the user terminal. Due to the way that macro arguments are processed, if the message has leading or trailing spaces, or if it has imbedded commas, it must be enclosed in angle brackets or part of it will be lost. The three calls are:

```
TYPE    msg     ;Types the message on the user terminal as is
TYPESP  msg     ;Types the message and appends one space to it
TYPECR  msg     ;Types the message and appends a CRLF pair to it
```

The macros are defined in SYS.MAC as follows:

```
DEFINE  TYPE    MSG
        TTYI
        ASCII   /MSG/
        BYTE    0
        EVEN
        ENDM

DEFINE  TYPESP  MSG
        TTYI
        ASCII   /MSG' /
        BYTE    0
        EVEN
        ENDM

DEFINE  TYPECR  MSG
        TTYI
        ASCII   /MSG/
        BYTE    15,0
        EVEN
        ENDM
```

It should be noted that the message may not contain any slashes, since these are used as delimiters for the ASCII statement in the macros.

CHAPTER 8

CONVERSION MONITOR CALLS

8.1 NUMERIC CONVERSION CALLS

The AMOS monitor contains two calls which perform conversions from a single binary word value to an ASCII formatted decimal or octal string. Options for the conversion allow the string to be sent to the user terminal, to an output file or to a buffer in memory. Options also allow control of the result format.

8.1.1 Calling Format

Both calls have the same general format and take two arguments, each of which must be an expression that evaluates down to a byte value within the specified range. The two calls are:

```
DCVT    size,flags      ;Convert binary number in R1 to decimal
OCVT    size,flags      ;Convert binary number in R1 to octal
                        ; (hexadecimal if J.HEX is set for this job)
```

8.1.1.1 Size Byte - The size byte determines the number of digits in the output result. A zero size specifies a floating format in which the number of digits used is just enough to fully contain the result. A non-zero size specifies a fixed number of digits for the result with leading zeros being replaced by blanks. In either form, if the R1 value is zero, at least one zero digit will be output as the result.

8.1.1.2 Flags - The flags byte contains six flags which control the destination of the result string and also some other formatting options. The following list gives the flag bit positions and the action taken when the flag is set:

1 - disables leading zero blanking
2 - outputs the result to the user terminal
4 - outputs the result to the file whose DDB is indexed by R2
10 - puts result in memory at buffer indexed by R2 and updates R2
20 - adds one leading space to the result
40 - adds one trailing space to the result

Note that the maximum value which can be displayed using these calls is the maximum value of a 16-bit word. All numbers are considered unsigned so the largest decimal number is 65535, the largest octal number is 177777, and the largest hex number is FFFF.

If the size byte is non-zero, the sense of the leading zero blanking flag described below is reversed. In other words, when the size byte is zero, the conversion calls default to leading zero blanking, with bit 0 turning that blanking off. When the size byte is non-zero, the calls default to leading zeroes, with bit 0 specifying that leading zeroes are to be blanked.

The following examples may clarify things a bit. All examples assume the value in R1 is 964 (decimal), and the letter "b" in the result field indicates a blank.

```
DCVT    0,2     prints 964
DCVT    0,22    prints b964
DCVT    0,42    prints 964b
DCVT    5,2     prints 00964
DCVT    5,3     prints bb964
DCVT    5,43    prints bb964b
DCVT    5,62    prints b00964b
DCVT    2,2     prints 64 (the 9 is lost)
```

8.2 RAD50 CONVERSION MONITOR CALLS

Radix-50 packing is used throughout the system where the packing of filenames and other data entities lends itself. Radix-50 (RAD50) packing is a system by which three ASCII characters may be packed into a single 16-bit word using a special algorithm based on the value of octal 50. The character set that may be packed RAD50 is limited in scope to the alphanumeric characters, the period, the dollar sign, and the blank. The following list gives the legal characters that may be packed RAD50 and their equivalent octal codes:

Character	RAD50 code
blank	0
A-Z	1-32
a-z	1-32
$	33
.	34
0-9	36-47

There is no character for the RAD50 code 35.

8.2.1 RAD50 Packing Algorithm

The packing algorithm for a 3-character input to a 16-bit RAD50 result is:

1. The first character code is multiplied by 3100 octal (50x50).
2. The second character code is multiplied by 50 and added to the first.
3. The third character code is added to the above to form the result.

The unpacking algorithm merely reverses the above sequence to get the triplet.

8.2.2 Packing and Unpacking Calls

There are two monitor calls which perform the above packing and unpacking algorithms. Both calls use registers R1 and R2 as indexes to the components and require no calling arguments.

8.2.2.1 PACK - Pack Three ASCII Characters into RAD50 - The triplet (3 ASCII characters) indexed by R2 is packed into RAD50 form and the result is left in the word indexed by R1. R1 is incremented by 2 to receive the next result word for multiple packing. R2 is left indexing the first character, which was not included in the packing of this triplet. The PACK call terminates packing and forces blank fill for any input which does not contain three valid RAD50 characters. For the PACK call, a blank is considered an illegal input character and terminates packing.

8.2.2.2 UNPACK - Unpack Three RAD50 Characters into ASCII - The word in the address indexed by R1 is unpacked, and the triplet is left in the three bytes beginning with the byte currently indexed by R2. R1 is incremented by 2 for the next word, and R2 is incremented by 3 for the next triplet result. Blanks are legal in unpacking and are placed into the result if they are decoded from the input word.

8.3 PRINTING CONVERSION CALLS

There are three calls in the monitor which accept a system unit input and convert the unit to standard printable form and then output it to the user terminal. These calls are used to print out file specifications, filenames, and project-programmer numbers. Each call takes one standard argument which addresses the system unit to be converted and printed.

8.3.1 PFILE - Output a Filespec From a DDB

The argument addresses a file DDB, and the PFILE call extracts the parameters in the file specification words. It then prints them on the user terminal in the standard format of dev:filnam.ext[p,pn].

8.3.2 PRNAM - Output a Filename

The argument addresses a 3-word filename.extension block (packed RAD50), and the PRNAM call prints the converted result on the user terminal in the standard format of filnam.ext.

8.3.3 PRPPN - Output a PPN

The argument addresses a 1-word project-programmer code, and the PRPPN call prints the converted result on the user terminal in the standard format of proj,prog. The p,pn is output in octal, regardless of the setting of J.HEX.

8.4 ALPHABETIC CONVERSION--LCS AND UCS

The AMOS monitor inlcudes two calls that switch between upper- and lower-case alphabetic characters. LCS converts one character in R1 to lower case. UCS converts one character in R1 to upper case. The Z-flag is set if the call is successful.

CHAPTER 9

INPUT LINE PROCESSING CALLS

When a program is executed by an operator command, register R2 is left pointing to the first non-blank character on the command line which follows the command name itself. The remainder of the line is normally interpreted by the particular program and used to determine the files to be acted on, the record number to be dumped, the devices to be accessed, etc. For example, the MACRO call requires the name of the program and any switch options to follow the MACRO command name on the same line. The macro assembly program then processes the program name and the switch options by way of the R2 index which was left indexing the rest of the command line. This command line is actually the user's terminal input buffer.

In addition to the command input line, the KBD monitor call also leaves R2 set to the input line buffer which contains the user input data. Also, various translators and file processing programs may read in a line of data and then set index R2 to the base of that line for scanning. For this reason, there exists a number of monitor calls which perform scanning and conversion functions based on an input line which is indexed by R2. Some of the calls merely test the character indexed by R2 for a specific condition and return with flags set, based on the result of the test. In these instances R2 is not modified. In calls which perform scan conversions, R2 is updated to point to the character which terminated the conversion. With the exception of the FILNAM call, none of these calls require any arguments. Conversion results are always delivered back to the user in register R1.

9.1 ALF - TEST A CHARACTER FOR ALPHABETIC

The character indexed by R2 is tested for alphabetic (A-Z; a-z); the Z-flag is set if it is, and cleared if it is not. R2 is not changed.

9.2 NUM - TEST A CHARACTER FOR NUMERIC

The character indexed by R2 is tested for numeric (0-9); the Z-flag is set if it is, and cleared if it is not. R2 is not changed.

9.3 TRM - TEST A CHARACTER FOR TERMINATOR

The character indexed by R2 is tested for a legal terminator defined as a blank, tab, comma, semicolon, carriage-return, line-feed, or null. The Z-flag is set if the character is a terminator, and cleared if it is not. R2 is not changed.

9.4 LIN - TEST A CHARACTER FOR LINE TERMINATOR

The character indexed by R2 is tested for a legal end-of-line defined as a semicolon, carriage-return, line-feed, or null. The Z-flag is set if the character is an end-of-line character, and cleared if it is not. R2 is not changed.

9.5 BYP - BYPASS BLANKS

Index R2 is advanced past all characters which are blanks or tabs and left indexing the first non-blank, non-tab character it finds.

9.6 GTDEC - INPUT A DECIMAL NUMBER

Index R2 is used to process a decimal number whose value may be from 0 to 65535 in the input line (leading zeros are legal), and to deliver the resultant binary value back in R1. The N-flag is set if there is an error (i.e., result is greater than 65535). R2 is updated to point to the character following the decimal input number. In the case of an error, R2 is left indexing the digit that would have caused the overflow past 65535 for double-word processing techniques.

9.7 GTOCT - INPUT AN OCTAL NUMBER

Index R2 is used to process an octal number whose value may be from 0 to 177777 in the input line (leading zeros are legal), and to deliver the resultant binary value back in R1. The N-flag is set if there is an error (i.e., result is greater than 177777). R2 is updated to point to the character following the octal input number. If J.HEX is set for this job (via the SET HEX command), this call processes input in hexadecimal instead of octal.

9.8 GTPPN - INPUT A PROJECT-PROGRAMMER NUMBER

Index R2 is used to process a project-programmer number in the standard format of proj,prog, and to deliver the resultant binary code back in R1. The format dictates that project numbers be octal numbers with a value between 1 and 377, and programmer numbers be octal numbers with a value between 0 and 377. The N-flag is set if the PPN was not in valid format. R2 is updated to point to the character following the PPN.

9.9 FILNAM - INPUT A FILENAME

Index R2 is used to process a filename.extension input string, leaving the RAD50 packed 3-word result in the three words starting with the address specified as the first argument of the call. In format, this argument is a standard monitor call argument. The second argument is a 1- to 3-character extension to be used in case no explicit extension is entered in the input string. R2 is updated to index the terminating character. The Z-bit is set if there was no filename to process (i.e., the first character was not a legal RAD50 character).

CHAPTER 10

MISCELLANEOUS MONITOR CALLS

This section deals with the monitor calls which do not fit into any of the categories treated thus far.

10.1 EXIT - RETURN TO AMOS COMMAND LEVEL

This is the normal means that a program uses to terminate processing and return to monitor command mode. The EXIT call takes no arguments. The monitor, upon executing the EXIT call, deletes all temporary memory modules in the user partition and resets any parameters that are program dependent such as JOBICP, JOBBPT, etc. All assigned devices are also released at this time. The user terminal is then placed in the monitor command mode, ready to process another operator command.

10.2 CTRLC - BRANCH ON CONTROL-C

Whenever a control-C is entered on a terminal keyboard (usually to abort a program), no action takes place immediately, but rather a flag is set in the JCB status word which must be tested later by the program. The CTRLC call is used within an application program to check the status of the control-C flag (in the JCB status word) and branch to a specific address if the flag is set. This call is a convenience since the user could perform the same task with a few instructions by locating his own JCB status word and checking the J.CCC flag within it. The format of this call is:

```
CTRLC   routine-address
```

where routine-address is the address to branch to within the program if the control-C flag is set.

The CTRLC call does not reset the J.CCC flag but merely indicates that it is set (this allows nested routines to unwind themselves correctly). The user program must then reset the flag explicitly by clearing it in the JCB status

word or implicitly by performing the EXIT call, which kills the program and returns to monitor mode, clearing J.CCC.

10.3 JLOCK, JUNLOK - PREVENT CONTEXT SWITCHING

The JLOCK call prevents context switches from occurring and allows the current user to run. JUNLOK reverses the effect of JLOCK.

10.4 RQST - REQUEST CONTROL OF A SEMAPHORE

RQ points to a 2-word semaphore which may conventionally be associated with any type of resource (disk, buffer, queue block, etc.). When a job requires access to a resource, it should RQST the semaphore associated with that resource. RQST decrements the semaphore count (representing the number of available resources) by 1. If the resulting count is greater than or equal to 0, the RQST returns, allowing access to that resource. If the difference is less than 0, the job is placed in a wait chain until the resource is available.

To illustrate, suppose a job needs to access one of 20 available queue blocks. A semaphore with an initial value of 20 (to represent the available queue blocks) could be set up and accessed prior to any attempts to allocate a queue block. A RQST call decrements the count from 20 to 19, confirms that 19 is greater than or equal to 0, then returns control of the job so it can get a queue block. If none of the 20 queue blocks were available (i.e., the semaphore count < 0), the job would be placed in a wait state until a queue block was identified as freed via a RLSE call (see section 10.5 below).

10.5 RLSE - RELEASE CONTROL OF A SEMAPHORE

If, upon execution of the RQST call (see section 10.4 for explanation), the semaphore count is less than or equal to 0 (i.e., none of the resources requested is available), the requesting job is put to sleep in a wait chain. When one job is finished with one of those resources, a RLSE call on the semaphore associated with that resource increments the count by 1 and determines if the result is less than or equal to 0. If it is, the next job in the wait chain is awakened and allowed to finish the RQST.

For example, if none of 20 queue blocks is currently available, the count is less than or equal to 0--let's say it's 0. Before a job tries to get a queue block, a RQST on the semaphore decrements the count from 0 to -1 and places the job in a wait chain. After a job frees a queue block, it uses the RLSE call on the semaphore associated with "queue blocks." This call increments the semaphore count by 1, resulting in 0, and wakes the first job in the wait chain, which allows it to continue on and allocate a queue block. The following diagram illustrates the semaphore:

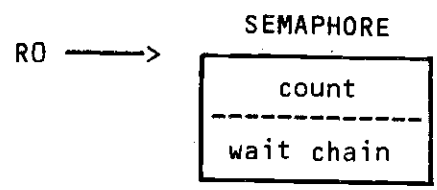

10.6 PCALL - INVOKE PROGRAM AS SUBROUTINE

PCALL is similar to the standard machine instruction call (JSR), except return is not done via the RTN instruction but is accomplished via the EXIT supervisor call. The format is:

```
PCALL subroutine-address
```

where the subroutine address is the address of the program you wish to call.

If you wish to use the PCALL monitor call to execute a program that locates its stack area in the user partition (e.g., VUE), you must first place your own stack area in your user partition.

10.7 AMOS - EXECUTE AMOS COMMAND AS SUBROUTINE

When AMOS is used as a monitor call, the character string pointed to by R2 is treated as a monitor command line, and the AMOS command in this command line is executed without leaving the current program.

If you wish to use the PCALL monitor call to execute a program that locates its stack area in the user partition (e.g., VUE), you must first place your own stack area in your user partition.

APPENDIX A

DISK STRUCTURE FORMAT

The AMOS monitor supports a flexible disk file system which relieves you of the task of keeping track of files, links and record counts. The structure of the standard disk format used in the AMOS system is described here for those programmers who wish to do some disk file manipulation or system software programming.

A.1 PHYSICAL RECORD FORMAT

The logical record size for all disks used within the AMOS file structure, regardless of type, is 512 bytes. For efficiency, the hard-disk structures (such as the AM-500 or Trident subsystems), and the AMS floppy format all define the physical record size to be this 512-byte logical record size. To maintain compatibility with other systems, the standard IBM-compatible floppy disk format is somewhat different and will be expained in more detail here.

The standard IBM-compatible floppy disk has 2002 128-byte physical records on 77 tracks, each track having 26 sectors numbered 1 through 26. The AMOS system uses a logical record size of 512 bytes (256 words) for each record, so the actual record is made up of four standard size 128-byte records on the floppy disk itself. The disk driver routine is responsible for translating the AMOS record number (0-499) to the proper four physical records on the disk. There are only 500 records of 512 bytes each, as far as the programmer is concerned, and the last two 128-byte records on the floppy disk are lost to his use.

The driver translates the AMOS record number into a starting record number, which is four times as great. In addition, a physical sector interleave factor is used so that a 512-byte record requires only one rotation of the disk instead of four, which would be the case if an attempt was made to access four physically contiguous sectors on the floppy disk. The interleave factor is 5, meaning that there are four sectors between each logically contiguous pair of sectors.

A.2 DISK RECORD TYPES

There are six different record types in use in the AMOS system, categorized by their use in the logical processing of files. Each record is 512 bytes long, but their internal structure differs due to different usage in the system. The six record types are:

1. Disk ID record
2. Bitmap records
3. Master File Directory record (MFD)
4. User directory records
5. Sequential file data records
6. Contiguous file data records

The following three record types take care of records 0-2, which are the same on all disks. Initializing the disk by using the "I" command in the SYSACT program writes out record 1 (empty MFD of all zeros) and record 2 (bitmap with records 0-2 allocated), logically clearing the disk of all users and files and making all remaining records (3-499) available. These records are then allocated as either user directory records or file data records.

A.2.1 The Disk ID Record

The Disk ID record is always record 0 and is not currently used by the AMOS system. It has been reserved for use by user routines which may want to store disk identification information in it. It is permanently allocated, so it will not accidently be used as a data record by any system routine. Since this record is reserved for the disk ID, you should not attempt to use it for other purposes.

A.2.2 The Bitmap

The bitmap is one or more records which always begin with record 2 and extend into as many sequential records as necessary to represent the entire disk. Each word in the bitmap is capable of representing the state of 16 logical records with one bit being used for each record. The bit is set if the record is in use and cleared if it is free. The last two words of every bitmap are a double-word hash total used to maintain bitmap integrity during processing. Any remaining words in the last bitmap record are unused. The bitmap itself is permanently allocated but contains no links to other system disk records. If you destroy the bitmap, you can run the DSKANA program to recover it.

A.2.3 The Master File Directory

The master file directory record is always record 1 and forms the root of the file structure tree. It contains one entry of four words for each user PPN which is allocated to this disk by the SYSACT program. A maximum of 63 users may be allocated on any one disk, since only one MFD record is available.

A.2.4 The User File Directory

User directory records contain up to 42 entries of six words each to describe user files in the corresponding PPN. The first word of each directory record is a link word to the next directory record in the event that more than 42 files are allocated in the current user area. The final directory record has a zero link word indicating that no more directory records follow.

A.2.5 Sequential File Data Records

Sequential file data records have a link word and 255 data words. The link word is the record number of the next record in the file. A zero link word indicates this is the last record in the file. The last record in the file may have anywhere from 0 to 509 active data bytes in its data area. The directory record item contains this number. Sequential files are normally processed as one long string of bytes from start to finish.

A.2.6 Contiguous File Data Records

Contiguous file data records have 256 data words and no links. Contiguous files must be allocated as a block of records with no intervening records belonging to other files. They must be allocated before their use while sequential files are allocated one record at a time as they are required. Contiguous files allow random access processing, since any record may be located as a direct offset relative to the base record.

A.3 FILE STRUCTURE

The file structure is depicted in figure A-1 and resembles a tree with the MFD record as its root. The MFD record has one item for each allocated user on this disk. Each MFD item then contains the record number of the first user directory record for that PPN number. The user directory record has one item for each data file in this user's area. Each directory item then contains the record number of the first data record in he file. Sequential

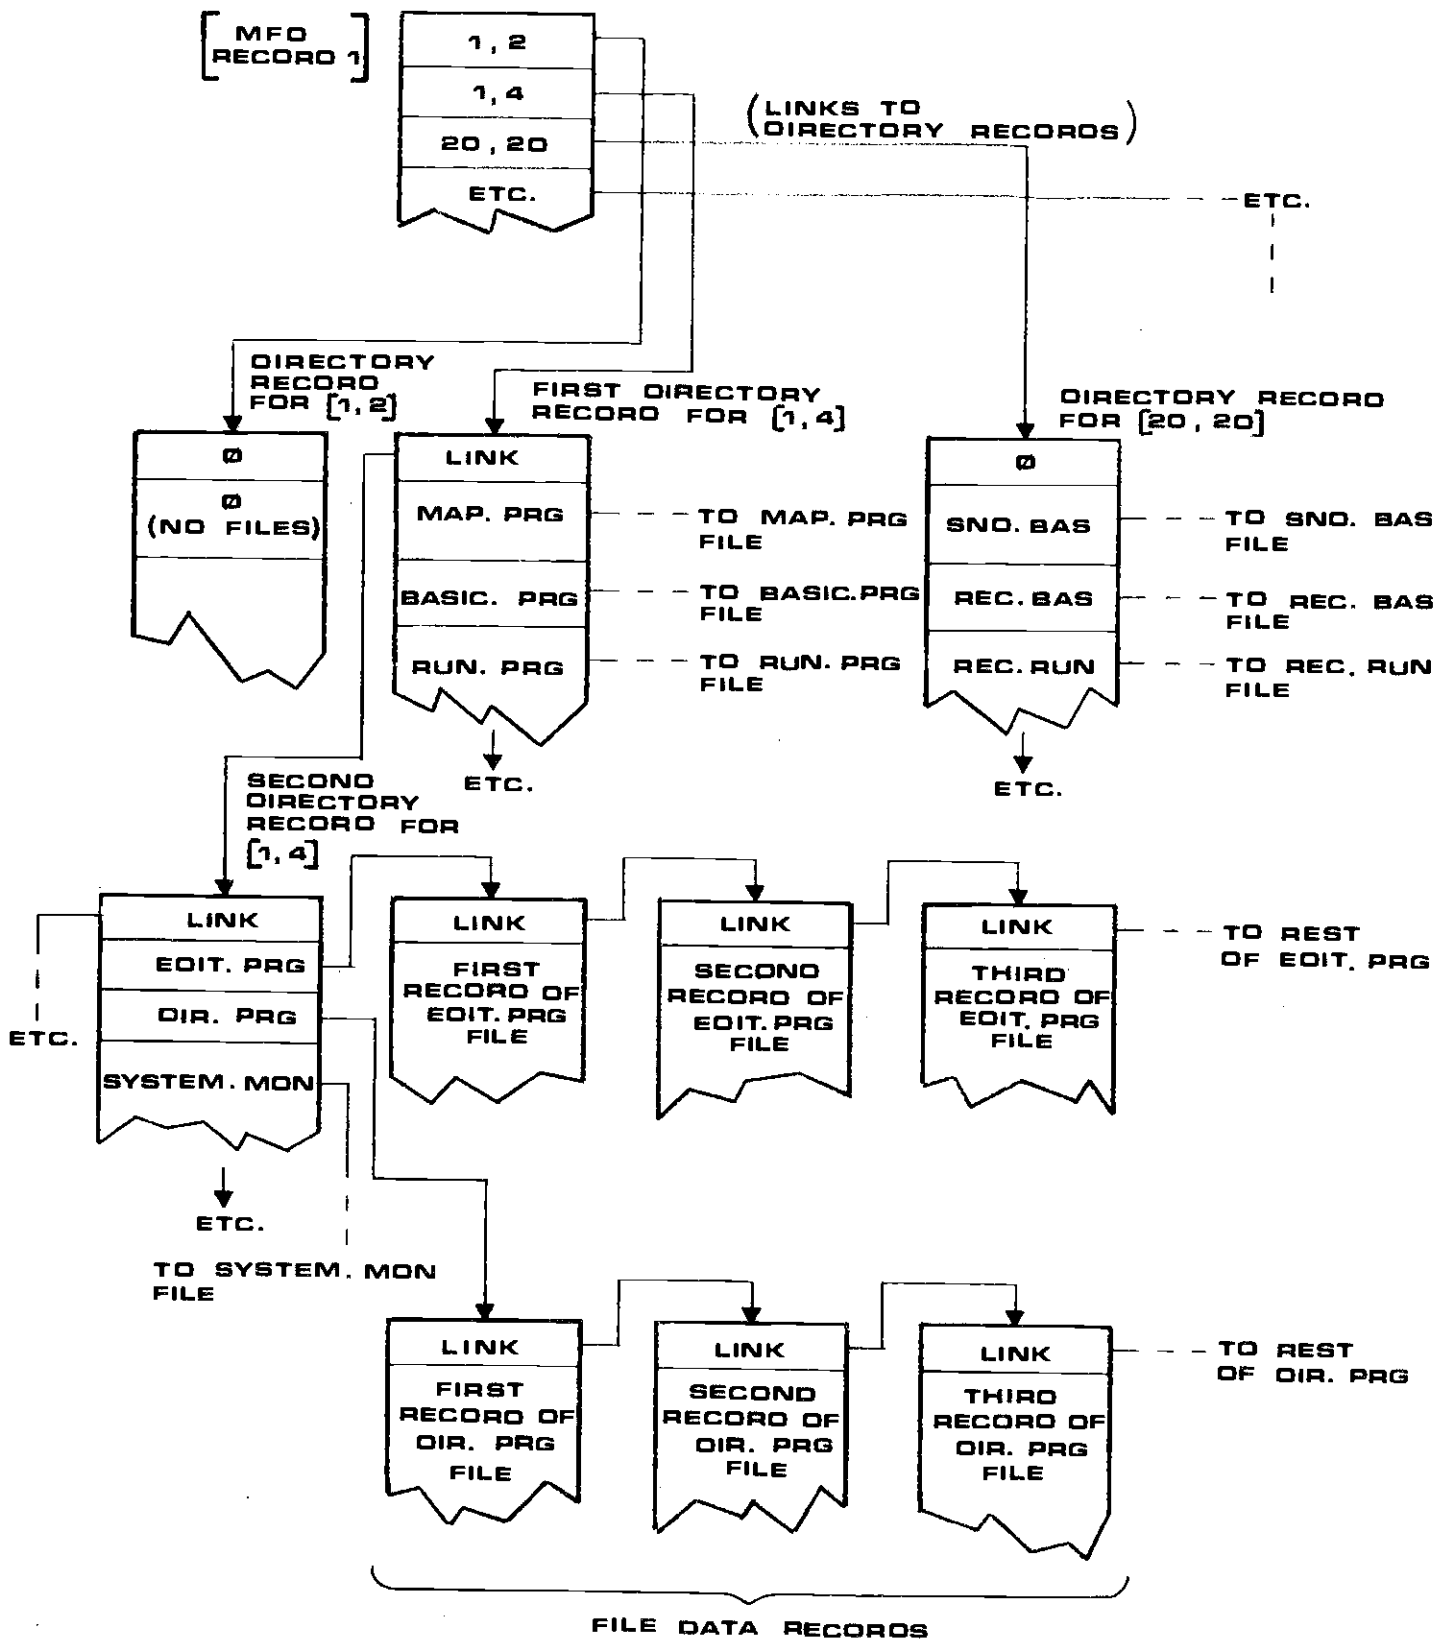

Disk File Structure

Fig A-1

files then chain through the data records by link words as shown in the diagram. The two files that are partially depicted are EDIT.PRG and DIR.PRG in user area [1,4] which happens to be the system program area. Contiguous files have no link words and must occupy physically adjacent records beginning with the first record as addressed in the directory item. Contiguous files are not depicted in the diagram since they are so straightforward in organization.

A.4 MFD ITEM FORMAT

Each MFD item is four words long and contains the PPN specification, user directory link, and password. The format of the item is:

Word 1 - user PPN (proj and prog are each one byte)
Word 2 - record number of first user directory record
Words 3-4 - password packed RAD50 (up to 6 characters)

Word 2 is zero if no files have been allocated to this user yet, meaning no directory records have yet been allocated. Words 3-4 are zero if no password is required to gain access to this user account when logging on via the LOG command.

MFD items are added, deleted, and changed by the SYSACT program.

A.5 UFD ITEM FORMAT

Each user directory item is six words long and contains information about the data file which it defines. The format of the item is:

Words 1-3 - filename.extension of the file packed RAD50
Word 4 - number of data records in this file
Word 5 - number of active data bytes in last record
Word 6 - record number of first data record in file

Word 1 is -1 (octal 177777) if this file has been erased and the directory item is available for another file definition. Word 1 is zero, to mark the logical end of the user directory. The byte count in word 5 is negative if this is a contiguous file. It also represents the negative active byte count of the file if the contiguous file has been opened for output and written into sequentially.

APPENDIX B

SYSTEM COMMUNICATION AREA

One area in monitor memory starting at location 100 (octal) is called the system communication area. It is defined mnemonically in SYS.MAC and contains specific parameters that deal directly with singular system resources and root addresses. They are briefly defined here for those users who wish to carefully reference them; but such action should be rare and must be undertaken with great caution. All references to these parameters should be made symbolically in the absolute addressing mode. For example, the instruction MOV @#JOBTBL,R0 should be used to set the base of the user job table into index register R0.

B.1 SYSTEM - SYSTEM ATTRIBUTES WORD

This word contains system attribute and status flags. Currently it is only used to indicate that the system has been properly loaded when bit 0 is set on.

B.2 DEVTBL - ADDRESS OF THE DEVICE TABLE

Set up by the DEVTBL program in the system initialization command file, this word contains the absolute address of the device table in monitor memory.

B.3 DDBCHN - ACTIVE DDB CHAIN

This is the base of the active DDB chain for interrupt driven routines. It is set up and altered by the file service routines as new I/O DDB's are queued for transfer requests, and goes to zero each time there are no requests pending. It is not used for non-interrupt driven devices.

B.4 MEMBAS & MEMEND - USER MEMORY POINTERS

These two words define the beginning and end of the complete user memory area. MEMBAS is the address of the first word following the complete resident monitor, including the system memory area for user resident programs. MEMEND is the address of the last word in the total physically contiguous RAM memory in the machine. It is set up by the INITIA program when the monitor first starts up, by a memory scan technique which locates the last available 1K bank. If memory management is active, MEMEND can only reflect the end of switchable memory within bank 0, and its use in the system diminishes.

B.5 SYSBAS - BASE OF SYSTEM MEMORY

This is the address of the system memory area which is used to contain any user programs set up by the SYSTEM command in the system initialization command file. It is zero if no system memory area exists.

B.6 JOBTBL - ADDRESS OF THE JOB TABLE

This is the address of the user job table which contains one JCB entry for each user allocated via the JOB command in the system initialization command file. For a complete description of the job table and JCB entries, refer to Chapter 2, "JOB SCHEDULING AND CONTROL SYSTEM."

B.7 JOBCUR - JCB ADDRESS OF THE CURRENT JOB

This word always contains the address of the JCB for the job that is currently running and has control of the CPU. For the user program, it always points to your own JCB as long as you are running. Obviously if you are referencing this word you must be running. JOBCUR is updated only by the job scheduler in the time-sharing monitor.

B.8 JOBESZ - JOB TABLE ENTRY SIZE

This word is set up when the monitor is built and contains the size in bytes of the JCB entry in the job table. This way, when the JCB item expands, the programs which scan the job table will not have to be reassembled since they get the JCB size dynamically from JOBESZ. This includes routines within the monitor itself.

B.9 TIME - THE TIME OF DAY

THIS 2-word field is incremented each time the line clock interrupts. It represents the current time of day, stored as the number of ticks since midnight. You can reference this parameter to keep track of the time it takes to do something on the machine. Remember, TIME is used to count clock ticks and not seconds or milliseconds. To calculate the actual time in seconds, divide the elapsed time in ticks by the clock frequency which is stored in the CLKFRQ constant described further on. This, of course, assumes that the CLKFRQ command has been used in the system initialization command file to properly set up the constant for your particular frequency (50 Hz overseas, remember?).

B.10 DATE - THE SYSTEM DATE

This 2-word field is used by various date routines to store the current date in some specific format. Its use depends upon the applications which are defining the format. The DATE field is not accessed or altered by the system monitor itself.

B.11 HLDTIM - THE HEAD LOAD TIMER

This 2-word area controls the head-load timing for the AM-200 floppy disk system when used with the Persci Floppy Disk Drive. The second word (at HLDTIM+2) is set up by the HEDLOD program, in the system initialization command file, to the number of clock ticks desired to wait before unloading the disk heads during periods of inactivity. Each time the head is loaded or another disk transfer is initiated, the count in the second word is transferred to the first word. Each time the clock interrupts, the count in the first word is decremented, and if it ever gets to zero the head is unloaded.

B.12 CLKFRQ - LINE CLOCK FREQUENCY

This word is set up by the CLKFRQ command in the system initialization command file to contain the frequency at which the line clock is running. It is used by routines which compute elapsed time by counting the clock ticks in the TIME constant. It is normally set to 60 for systems in North American countries and to 50 for systems running overseas.

Remember that CLKFRQ specifes only the local line frequency. Changing CLKFRQ has no effect on the execution speed of the computer.

B.13 SPXSAV - STACK POINTER SAVE LOCATION

This word is used by the clock interrupt routine for saving the user stack pointer just prior to switching to the internal stack.

B.14 SPXINT - INTERNAL STACK

This is the address of the internal work stack used for processing clock interrupts. It is set up by the initial load routine and used by the clock interrupt processor.

B.15 LPTQUE - LINE PRINTER SPOOLER QUEUE

This is the dynamic link address to the base of the line printer spooler queue. The format of the spooler queue is subject to frequent change, so it is not detailed here.

B.16 TRMDFC - BASE OF THE TERMINAL DEFINITION TABLE

This is the link to the base of the terminal definition table. There is one entry in this table for each terminal defined at system startup by a TRMDEF statement in the SYSTEM.INI file.

B.17 TRMIDC - ADDRESS OF FIRST INTERFACE DRIVER

This is the link to the first terminal interface driver defined in the system. Each driver then links to the next one in the chain.

B.18 TRMTDC - ADDRESS OF FIRST TERMINAL DRIVER

This is the link to the first terminal driver defined in the system. Each driver then links to the next one in the chain.

B.19 TRMSCN - THE NON-INTERRUPT TERMINAL QUEUE

TRMTSC is the link to the chain of queue blocks for all terminals which are defined as non-interrupt driven and which require terminal scan service each clock tick.

B.20 CLKQUE - THE CLOCK QUEUE

CLKQUE is the link to the clock queue which gets scanned every clock interrupt. This queue has some entries that remain constant and some that are continuously added and deleted (such as SLEEP command queue blocks). CLKQUE is actually the base entry in the queue chain and therefore is two words in size.

B.21 SCNQUE - THE IDLE SCAN QUEUE

This is the link to that point within the clock queue chain which defines the idle scan queue or that portion of the clock queue which will be continuously scanned when the system is idle. SCNQUE is actually the base entry in the queue chain and therefore is two words in size.

B.22 RUNQUE - THE JOB SCHEDULING QUEUE

This 5-word block forms the base and end entries for the job scheduling and run queue, along with the necessary control information. Its format is unimportant to the user, and you should never alter it.

B.23 DRVTRK - THE DRIVE/TRACK TABLE

DRVTRK is a 4-byte block that stores head track positioning information for floppy disks used in the system. It is used only by the head unload and head positioning routines in various floppy disk drivers.

B.24 MEMDEF & MEMBNK - MEMORY MANAGEMENT CONTROL

These two words are used by the memory management system (when active) to store the base of the memory bank definition table and the currently active bank index. They are explained in detail in Chapter 3, "MEMORY CONTROL SYSTEM CALLS."

B.25 ZSYDSK - ADDRESS OF SYSTEM DISK DRIVER

This word contains the base address of the system disk driver within the monitor. It is used by MONGEN to overlay the disk driver with another one when changing the resident disk type.

B.26 SYSMEM - SYSTEM MEMORY LINK

This word contains the address of the bank-switched system memory, typically containing bitmaps.

B.27 MSGQUE - SYSTEM LINK COMMUNICATION

This word is the base address of the Link system interrupt queue.

B.28 MSGDAT - MESSAGE SYSTEM DATA AREA

This word points to the system data communication area for the Link system.

B.29 QFREE - QUEUE SYSTEM CONTROL

QFREE consists of two words, the first containing the number of queue blocks currently available, the second pointing to the first available queue block. Queue blocks are allocated and deallocated by getting and returning them from the front of the list controlled by this address, automatically incrementing or decrementing the free count in the process. The operation of the queue system is more fully explained in Chapter 5, "MONITOR QUEUE SYSTEM CALLS."

APPENDIX C

ALPHABETIC LISTING OF AMOS MONITOR CALLS

The following is a quick reference to all AM-100 monitor calls:

ALF	tests the character indexed by R2 for alphabetic
AMOS	executes AMOS command without exiting current program .
ASSIGN	assigns a non-sharable device to a job
BNKSWP	changes banks when running under memory management system
BYP	bypasses all spaces and tabs in the string indexed by R2
CHGMEM	changes the size of a user memory module
CLOSE	closes a logical dataset
CRLF	prints a carriage-return line-feed pair on the user terminal
CTRLC	checks for a control-c pending
DCVT	converts a binary value to decimal and prints it on the user terminal
DEASGN	deassigns a non-sharable device from a job
DELETE	deletes a file from a file-structured device
DELMEM	deletes a user memory module from his partition
DSKALC	allocates next available record on disk and returns block number
DSKBMR	reads disk bitmap and sets re-entrant lock for user modification
DSKBMW	rewrites disk bitmap after user modification
DSKCTG	allocates a contiguous file for random processing
DSKDEA	deallocates a record on disk and makes it available for use again
DSKDRL	sets re-entrant directory lock for a specific user's directory
DSKDRU	clears re-entrant directory lock for a specific user's directory
EXIT	exits from user program and returns to monitor command mode
FETCH	fetches a module from disk into user memory unless already in memory
FILNAM	processes a filename specification indexed by R2 into RAD50 format
FMARK	find file mark on specified magnetic tape unit
FMARKR	read in reverse to find file mark on specified magnetic tape unit
FSPEC	processes a complete file specification indexed by R2 and sets up DDB
GETMEM	allocates a user memory module in his partition
GTDEC	converts a decimal number indexed by R2 into binary and returns it in R1
GTOCT	converts an octal number indexed by R2 into binary and returns it in R1
GTPPN	converts a p,pn format indexed by R2 into binary and returns it in R1
HTIM	sets up the diskette head unload timer function
INIT	initializes a dataset driver block (DDB) for I/O processing
INPUT	performs a logical record input I/O function on an open dataset
JLOCK	prevents context switches and allows current user to run
JOBGET	retrieves a job control block item for the current job

JOBIDX	set an index to a job control block item for the current job
JOBSET	sets data into a job control block item for the current job
JRUN	restores a waiting job to the run request state
JUNLOK	enables context switches (reverses effect of JLOCK)
JWAIT	sets an active job into the wait state
JWAITC	sets your job into the wait state
KBD	accepts input from user terminal keyboard (character or line mode)
LCS	converts one character in R1 to lower case
LIN	tests the character indexed by R2 for valid end-of-line character
LOCK	locks the processor against interrupts (performs IDS instruction)
LOOKUP	looks for a specific file on disk and returns information about it
NUM	tests the character indexed by R2 for numeric
OCVT	converts a binary value to octal and prints it on the user terminal
OPEN	general form of the I/O logical dataset open calls
OPENA	opens a logical dataset for appending
OPENI	opens a logical dataset for input
OPENO	opens a logical dataset for output
OPENR	opens a logical dataset for random access
OUTPUT	performs a logical record output I/O function on an open dataset
PACK	packs an ASCII triplet into its RAD50 code
PCALL	invokes program as subroutine
PFILE	prints a complete file specification on user terminal from a DDB
PRNAM	prints a filename specification on user terminal from its packed format
PRPPN	prints a p,pn specification on user terminal from its packed format
PTYIN	forces one character into another job's terminal input buffer
PTYOUT	retrieves one character from another job's terminal output buffer
QADD	adds a queue block to the end of a queue list
QGET	gets a queue block from the free list and clears it for use
QINS	inserts a queue block into a queue list at a defined point
QRET	removes a queue block from a queue list and returns it to the free list
READ	performs a physical record read I/O function on a dataset
RENAME	renames a file on a file-structured device
REWIND	rewind magnetic tape on specified magnetic tape unit
RLSE	releases control of a semaphore and allows waiting job to access source
RQST	requests control of a semaphore to access source or to wait in wait chain
SCAN	forces a single scan of the idle scanner queue (SCNQUE)
SLEEP	puts the user job to sleep for a specified number of line clock ticks
SRCH	searches for a named memory module and returns its address
TAB	sends a tab character to the user terminal
TAPST	read tape status of specified magnetic tape unit
TBUF	queues up a variable length data buffer for output to a terminal
TCRT	executes the special function CRT routine in the active terminal driver
TIN	reads one character from the user terminal input buffer
TOUT	sends one character to the user terminal output buffer
TRM	tests the character indexed by R2 for a valid termination character
TRMBFQ	adds a data buffer to the active output queue of a terminal
TRMICP	processes one input character (used within terminal drivers)
TRMOCP	processes one output character (used within terminal drivers)
TTY	outputs one character to the user terminal
TTYI	outputs an in-line message to the user terminal
TTYIN	retrieves one character from any job's terminal input buffer
TTYL	outputs a message to the user terminal
TTYOUT	forces one character into any job's output buffer

TYPE	types an ASCII message on the user terminal
TYPECR	types an ASCII message on the user terminal with appended CRLF pair
TYPESP	types an ASCII message on the user terminal with one appended space
UCS	converts one character in R1 to upper case
UNLOCK	unlocks the processor for interrupts (performs IEN instruction)
UNPACK	unpacks a RAD50 code word into its equivalent ASCII triplet
USRBAS	returns the address of the current user's memory partition base
USREND	returns the address of the current user's memory partition end
USRFRE	returns the address of the current user's free memory area
WAKE	wakes a job out of sleep state
WRITE	performs a physical record write I/O function on a dataset
WRTFM	write a file mark to specified magnetic tape unit

Index

ALF 9-1
Alphabetic conversion 8-4
AM-100 3-13
AM-100/T 3-13
AM-700 3-10
AMOS 10-3
ASSIGN 6-15

Bitmap Format 6-17
Bitmaps A-2
BNKSWP 3-12
BYP 9-2

CHGMEM 3-6
CLKFRQ B-3
CLKGDE B-5
Clock Frequency B-3
CLOSE 6-11
Contiguous Files A-3
Control-C 10-1
Convenience Macros 7-8
CRLF 7-4
CTRLC 10-1
Cursor Addressing 7-6

DATE B-3
DCVT 8-1
DDB 6-1
  Buffer Address 6-4
  Buffer Index 6-4
  Buffers 6-6
  Call Level 6-5
  Device Code 6-5
  Drive 6-5
  Driver Work Area 6-6
  Error Code 6-2
  Error Handling 6-7
  Extension 6-6
  Filename 6-6
  Flags 6-4

JCB Address 6-5
Job Priority 6-5
Open Code 6-6
PPN 6-6
Queue Chain Link 6-5
Record Number 6-5
Record Size 6-4
DDB Format 6-2
DDBCHN B-1
DEASGN 6-16
Decimal Input 9-2
Decimal Output 8-1
DELETE 6-15
DELMEM 3-6
DEVTBL B-1
DEVTBL program B-1
Disk File Structure A-3
Disk ID Record A-2
Disk Record Types A-2
Disk Service Monitor Calls . . . 6-16
Disk Structure A-1
DSKALC 6-18
DSKBMR 6-19
DSKBMW 6-19
DSKCTG 6-18
DSKDEA 6-19
DSKDRL 6-19
DSKDKU 6-20

EXIT 10-1

FETCH 4-1
Flags 4-2
File marks 6-21
File Service Monitor Calls . . . 6-8
File Service System 6-1
File Structure A-3
Filenames 8-4, 9-3
Filespecs 8-4
FILNAM 9-3
FMARK 6-21
FMARKR 6-21
FORCE command 7-4
FSPEC 6-8

GETMEM 3-6, 6-6
GTDEC 9-2
GTOCT 9-2
GTPPN 9-3

Head Load Time B-3
HEDLOD program B-3
Hexadecimal Input 9-2

Hexadecimal Output 8-1
HLDTIM B-3

INIT 6-6, 6-9
INPUT 6-13
Input Line Processing Calls . . . 9-1
Interface Drivers 7-5, B-4

JCB 2-1, B-2
  Size B-2
JCB Entries
  JOBBAS 2-5
  JOBBNK 2-7, 3-11
  JOBBPT 2-7
  JOBBRK 2-8
  JOBCMS 2-6
  JOBCMZ 2-6
  JOBCUR 2-1
  JOBDEV 2-7
  JOBDRV 2-8
  JOBDYS 2-9
  JOBERC 2-7
  JOBEXT 2-10
  JOBFPE 2-9
  JOBIAL 2-10
  JOBMSG 2-10
  JOBNAM 2-5
  JOBPRG 2-6
  JOBPRV 2-6
  JOBRNQ 2-9
  JOBSIZ 2-5
  JOBSPR 2-5
  JOBSTK 2-9
  JOBSTS 2-4
  JOBTRM 2-8
  JOBTYP 2-7
  JOBUSR 2-6
JLOCK 10-2
Job Control Block 2-1, B-2
  Size B-2
Job Table B-2
JOBBAS 2-5
JOBBNK 2-7, 3-11
JOBBPT 2-7
JOBBRK 2-8
JOBCMS 2-6
JOBCMZ 2-6
JOBCUR 2-1, B-2
JOBDEV 2-7
JOBDRV 2-8
JOBDYS 2-9
JOBERC 2-7
JOBESZ B-2

JOBEXT 2-10
JOBFPE 2-9
JOBGET 2-1, 2-3
JOBIAL 2-10
JOBIDX 2-1, 2-3
JOBMSG 2-10
JOBNAM 2-5
JOBPRG 2-6
JOBPRV 2-6
JOBRNQ 2-9
JOBSET 2-1, 2-3
JOBSIZ 2-5
JOBSPR 2-5
JOBSTK 2-9
JOBSTS 2-4
JOBTBL B-2
JOBTRM 2-8
JOBTYP 2-7
JOBUSR 2-6
JRUN 2-3
JUNLOK 10-2
JWAIT 2-3
JWAITC 2-3

KBD 7-2, 9-1

LIN 9-2
Line Printer Spooler B-4
LOOKUP 6-10
LPTQUE B-4

Magnetic tape drivers . . . . . 6-20 to 6-21
Master File Directory A-3, A-5
MEMBAS B-2
MEMONK 3-11, B-5
MEMDEF B-5
MEMDEF Program 3-9
MEMEND B-2
Memory Management 3-9, B-5
Memory Mapping 3-9
Memory Modules 3-5, 4-1
Memory Partition Controller . . . 3-10 to 3-13
Memory Partitions 3-2
MFD A-3, A-5
Miscellaneous Monitor Calls . . . 10-1
Monitor Calls
  ALF 9-1
  Alphabetic conversion 8-4
  AMOS 10-3
  Arguments 1-1
  ASSIGN 6-15
  BNKSWP 3-12
  BYP 9-2

Calling Format 1-1
CHGMEM 3-6
CLOSE 6-11
CRLF 7-4
CTRLC 10-1
DCVT 8-1
DEASGN 6-16
DELETE 6-15
DELMEM 3-6
Disk Service 6-16
DSKALC 6-18
DSKBMR 6-19
DSKBMW 6-19
DSKCTG 6-18
DSKDEA 6-19
DSKDRL 6-19
DSKDRU 6-20
EXIT 10-1
FETCH 3-5, 4-1
File Service 6-8
FMARK 6-21
FMARKR 6-21
FSPEC 6-8
GETMEM 3-6, 6-6
GTDEC 9-2
GTOCT 9-2
GTPPN 9-3
INIT 6-6, 6-9
INPUT 6-13
Input Line Processing 9-1
JLOCK 10-2
JOBGET 2-1, 2-3
JOBIDX 2-1, 2-3
JOBSET 2-1, 2-3
JRUN 2-3
JUNLOK 10-2
JWAIT 2-3
JWAITC 2-3
KBD 7-2, 9-1
LIN 9-2
LOOKUP 6-10
Magnetic tape drivers . . . . . 6-20 to 6-21
Memory Control 3-1
Miscellaneous 10-1
NUM 9-2
Numeric Conversion 8-1
OCVT 8-1
OPENA 6-10
OPENI 6-10
OPENO 6-10
OPENR 6-11
OUTPUT 6-14
PACK 8-3

PCALL 10-3
PFILE 8-4
Printing Conversion 8-4
PRNAM 8-4
PRPPN 8-4
PTYIN 7-4
PTYOUT 7-5
QADD 5-3
QGET 5-3
QINS 5-3
QRET 5-3
RAD50 Conversion 8-2
READ 6-11
RENAME 6-15
REWIND 6-20
RLSE 10-2
RUST 10-2
SLEEP 2-3
SRCH 3-5, 4-1
Standard Address Argument . . . 1-2
TAB 7-3
TAPST 6-21
TBUF 7-6
TCRT 7-6
Terminal Service 7-1
TIDX (obsolete) 2-8
TIP 7-3
TPOS 7-3
TRP 9-2
TSPEED 7-6
TTCCOP 7-5
TTMCOP 7-5
TTY 7-3
TTYI 7-4, 7-8
TTYIN 7-5
TTYI 7-4
TTYOUT 7-5
TYPE 1-1
UNPACK 8-4
USRBAS 3-2
USREND 3-2
USRFRE 3-2
WAKE 2-4
WRITE 6-12
WRTFM 6-20
MFC . . . . . 3-10 to 3-13
MSGDAT B-6
MSGQUE B-6
MTU.DVR . . . . . 6-20 to 6-21

NUM 9-2
Numeric Conversion Monitor Calls 8-1
Numeric Input 9-2

Octal Input 9-2
Octal Output 8-1
OCVT 8-1
OPENA 6-10
OPENI 6-10
OPENO 6-10
OPENR 6-11
OUTPUT 6-14

PACK 8-3
PCALL 10-3
PFILE 8-4
Physical Disk Record Format . . . A-1
PPNs 8-4, 9-3
Printing Conversion Monitor Calls 8-4
PRNAM 8-4
Project-Programmer Numbers . . . 8-4, 9-3
PRPPN 8-4
Pseudo Terminals . . . . . 7-4 to 7-5
PTYIN 7-4
PTYOUT 7-5

QADD 5-3
QFREE B-6
QGET 5-3
QINS 5-3
QRET 5-3
QUEUE command 5-2
Queue System 5-1, B-6
  Manipulating Queue Blocks . . . 5-3
  Obtaining a Free Queue Block . 5-3
  Returning a Queue Block 5-3

RAD50 Conversion Monitor Calls . 8-2
Random File Processing 6-11, 6-13
Random Files A-3
READ 6-11
RENAME 6-15
REWIND 6-20
RLSE 10-2
RQST 10-2
RUNQUE B-5

SCNQUE B-5
Semaphores 10-2
Sequential Files A-3
SLEEP 2-3
SPXINT B-4
SPXSAV B-4
SRCH 3-5, 4-1
  Flags 4-2
Standard Address Argument 1-2
SYS.MAC 1-1, 2-1, 7-8

SYSBAS B-2
SYSMEM B-6
SYSTEM B-1
System Communication
QFREE 5-1
System Communication Area B-1
CLKFRQ B-3
CLKQUE B-5
DATE B-3
DDBCHN B-1
DEVTBL B-1
DRVTRK B-5
HLDTIM B-3
JOBCUR B-2
JOBESZ B-2
JOBTBL B-2
LPTQUE B-4
MEMBAS B-2
MEMBNK 3-10, B-5
MEMDEF 3-10, B-5
MEMEND B-2
MSGDAT B-6
MSGQUE B-6
QFREE B-6
RUNQUE B-5
SCNQUE B-5
SPXINT B-4
SPXSAV B-4
SYSBAS B-2
SYSMEM B-6
SYSTEM B-1
TIME B-3
TRMDFC B-4
TRMIDC B-4
TRMSCN B-4
TRMTDC B-4
ZSYDSK B-5
System Date B-3

TAB 7-3
TAPST 6-21
TBUF 7-6
TCRT 7-6
Terminal Definition Table B-4
Terminal Drivers 7-6, B-4
Terminal Input 7-2
Terminal Service Monitor Calls . 7-1
Terminal Status Word 7-2
TIDX (obsolete) 2-8
TIME B-3
Time of Day B-3
TIN 7-3
TOUT 7-3

TRM 9-2
TRMBFQ 7-6
TRMDFC B-4
TRMICP 7-5
TRMIDC B-4
TRMOCP 7-5
TRMSCN B-4
TRMTDC B-4
TTY 7-3
TTYI 7-4, 7-8
TTYIN 7-5
TTYL 7-4
TTYOUT 7-5
TYPE 1-1, 7-8
TYPECR 7-8
TYPESP 7-8

UFD A-3, A-5
UNPACK 8-4
User File Directory A-3, A-5
USRBAS 3-2
USREND 3-2
USRFRE 3-2

WAKE 2-4
WRITE 6-12
WRTFM 6-20

ZSYDSK B-5