AMOS
VERSION 4.6
SOFTWARE UPDATE
DOCUMENTATION

DSS-10000-20

REV. A00

alpha micro

AMOS VERSION 4.6 SOFTWARE UPDATE DOCUMENTATION PACKET

'Alpha Micro', 'AMOS', 'AlphaBASIC'
'AlphaPASCAL', 'AlphaLISP', 'AlphaSERV',
'AlphaVUE', and 'AlphaACCOUNTING'

are trademarks of

ALPHA MICROSYSTEMS
Irvine, CA   92714

MASTER TABLE OF CONTENTS
FOR THE AMOS SOFTWARE UPDATE DOCUMENTATION PACKET

(For a complete list of Alpha Micro software documentation, see
<u>A Guide to the Alpha Micro Software Documentation Library</u>.)

October 1981

AMOS VERSION 4.6 RELEASE NOTES

The AMOS 4.6 software release contains a number of new programs that support several new hardware products. We have also released new versions of existing programs that have additional capabilities and are easier to use.

This document describes the changes and additions that have been made for the AMOS Version 4.6 software release. It describes only those changes made since the last release (AMOS 4.5A); you may find general system documentation in the appropriate manuals. Please refer to the document A Guide to the Alpha Micro Software Documentation Library in this documentation packet for a list of all Alpha Micro software documentation.

Please read the documentation supplied with this release before converting your system from AMOS 4.5 or 4.5A to AMOS 4.6.

1.0 NEW PRODUCTS

This release supports several new hardware products:

   The AM-610 Video Cassette Recorder Interface provides high-capacity file backup via a standard consumer-grade Video Cassette Recorder.

   The AM-620 1/4" Streaming Tape Drive provides convenient tape cartridge file backup.

   Several new Winchester technology disk-based systems are available that use a 32-megabyte Winchester disk or a 60-megabyte Winchester disk. Several of these systems are optionally configured with the AM-610 Video Cassette Recorder Interface or the AM-620 1/4" Streaming Tape Drive as backup and warm boot devices. (NOTE: The 8.5-megabyte Winchester disk based system was made available with AMOS Release 4.5A.) These disks run under the control of the AM-420 Hard Disk Controller.

This release also provides support for Winchester disk subsystems as well as Winchester disk-based integrated systems. Several new programs have been created to support the Winchester disks; before using your Winchester disk, it is important that you read the Software Notice for AM-420 Users in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.

2.0 THE RELEASE MECHANISM

The AMOS 4.6 release software is available on a variety of media. Please contact your Alpha Micro dealer for a list of the part numbers associated with the release media. You will receive a complete set of system software, not just the programs that have been changed or newly created for the AMOS 4.6 Release.

Note that because of space restrictions, users who purchase their software update on single-density, single-sided floppy diskettes, may also want to purchase several additional, miscellaneous diskettes in order to receive the complete set of software which is available on other media.

For more information, see the document Disks Available from Alpha Micro in this documentation packet.

3.0 DOCUMENTATION INCLUDED WITH THE AMOS 4.6 RELEASE

Included with this release is the standard AMOS Software Update Documentation Packet, (DSS-10000-20).

We have issued the following new documents and manuals since the AMOS 4.5A Release:

Change Page Packet #1 for the "ISAM System User's Guide", (DSS-10000-22)

Change Page Packet #4 for the "AMOS System Commands Reference Manual", (DSS-10000-18)

Change Page Packet #3 for the "AlphaBASIC User's Manual", (DSS-10000-19)

Software Information for the AM-610 Video Cassette Recorder Interface, (DSS-1000-24)

Software Information for the AM-620/AM-621 1/4" Streaming Tape Drive, (DSS-10000-23)

We have revised the following manual for this release:

AlphaVUE User's Manual, (DWM-00100-15, Revision C00)

4.0   NEW SOFTWARE FOR THE AMOS 4.6 RELEASE

The AMOS 4.6 Release contains all drivers, certification programs, and file backup programs necessary for using the products discussed in Section 1.0.

FIX420
FIX420 configures a disk driver for a Winchester technology disk based on the size of the drive (8.5mb, 32mb, or 60mb), and the number of logical units per physical device (1 to 8 on the 8.5mb and 32mb drives; 2 to 8 on the 60mb drive).

A major enhancement made to support Winchester disks is that you may use the FIX420 program to define multiple logical devices on a single Winchester physical device. For information on using FIX420, refer to the <u>Software Notice for AM-420 Users</u> in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.

WRMGEN
One of the features provided by the new backup devices (the video cassette recorder and the 1/4" streaming tape drive) is the ability to generate a "warm boot" tape from which you can boot your Winchester system if you should need to restore your Winchester System Disk. (A "warm boot" brings your system up under a pre-initialized monitor with just one job and one system terminal, and just enough software so you can restore your System Disk from another tape cassette or cartridge.)

Please refer to the documents <u>Software Information for AM-610 Video Cassette Recorder Interface,</u> (DSS-10000-24) <u>and Software Information for the AM-620/AM-621 1/4" Streaming Tape Drive,</u> (DSS-10000-23) for information on using the WRMGEN program to generate a warm boot cassette or cartridge. (NOTE: The command files VCRGEN and STRGEN have been included with this release to run WRMGEN for you and to answer many of its questions for you.)

VCR/
Streamer
Backup
Several new programs have been created to support the AM-610 Video Cassette Recorder Interface and the AM-620 1/4" Streaming Tape Drive. These programs allow you to perform file-oriented backup by saving files from disk to a tape cassette or cartridge and restoring the files from the cassette or cartridge back to the disk. You may also see a directory of the files on the data cassette or cartridge.

Please refer to the documents <u>Software Information for AM-610 Video Cassette Recorder Interface,</u> (DSS-10000-24), <u>and Software Information for the AM-620/AM-621 1/4" Streaming Tape Drive,</u> (DSS-10000-23), for information on using video cassette recorder and 1/4" streaming tape drive file backup programs.

5.0 ENHANCED SOFTWARE FOR THE AMOS 4.6 RELEASE

We have made changes to several existing programs:

AlphaVUE   This release includes Version 2.5 of AlphaVUE, the
             screen-oriented text editor. This version of AlphaVUE is
             faster and smaller than previous versions. Note that we
             have also completely rewritten the AlphaVUE manual for this
             release to make it easier to use for the new user who is
             unfamiliar with AlphaVUE.

ISAM   A new feature has been added to ISAM-- when searching a
             data file sequentially by symbolic key, you may ask ISAM to
             return the symbolic key as well as the relative record
             number of the data record. See Change Page Packet #1 for
             the "ISAM System User's Guide", (DSS-10000-22), for
             details. The Change Page Packet #3 for the "AlphaBASIC
             User's Manual", (DSS-10000-19), also contains information
             on this new feature.

FIXTRM   We have rewritten the NEWTRM program to make it easier to
             use. NOTE: We renamed this program to FIXTRM for
             compatibility with our other program names.

MEMPLT   This program has been enhanced so that it now displays the
             system memory map for non-memory management and bank
             switching systems as well as for systems that use the
             AM-700 Memory Partition Controller.

TAPE   A new option has been added that allows you to append
             carriage return/line feeds at the end of records being
             transferred from tape to disk. The purpose of this feature
             is to help you convert your data so that it can be accessed
             via AlphaBASIC programs and the AlphaVUE text editor.

SYSACT   SYSACT has been changed so that when you use a SYSACT
             command, any resulting changes to the disk MFD are made at
             that time, and not when you exit SYSACT. This avoids any
             conflict with multiple users accessing the MFD at the same
             time, but does mean that you cannot change your mind after
             using several SYSACT commands and Control-C out of SYSACT,
             assuming that the MFD modifications have not been made.

AM120.DVR   The driver for the AM-120 Auxiliary I/O Controller has been
             changed to allow greater compatibility with new and
             existing products. Because of this change, you may need to
             reset the date and time on the AM-120 the first time you
             boot under AMOS 4.6 if you have two AM-120 boards in your
             system.

6.0 PROGRAMS INCLUDED WITH THE AMOS 4.6 RELEASE

In the past, this document included a list of all programs contained in the current AMOS release with a description of their function. Since the release documentation is printed a considerable amount of time before the actual release media masters are created, to give you more precise and accurate information, we have decided to omit the list of programs from the Release Notes. Instead, we direct your attention to the file DIRECT.LST in account [1,4] of your release media for a list of all files (with hash totals) on that media. By generating hash totals for the files and then comparing the hash totals with those in the DIRECT.LST file, you can verify that all files on the media are intact.

DISKS AVAILABLE FROM ALPHA MICRO

1.0 DISKS AVAILABLE

The disks currently available from Alpha Micro are listed below. Listed
after the disks' descriptions are their part numbers. To obtain the prices
for these disks, please contact your Alpha Micro Dealer.

        WANGCO AM-200 System Diskettes (STD format)   PDB-00104-06
        WANGCO AM-200 System Diskettes (AMS format)   PDB-00104-07
        PERSCI AM-200 System Diskettes (STD format)   PDB-00104-02
        PERSCI AM-200 System Diskettes (AMS format)   PDB-00104-03
        CDC AM-210 System Diskette (Double Density, AMS)   PDB-00104-11
        AM-500 System Disk (5-Mbyte disk pack)   PDB-00104-12
        AM-410 System Disk (15-Mbyte disk pack)   PDB-00104-13

        Miscellaneous Program Diskette   PDB-00104-19
        Driver Source Diskette   PDB-00104-20
        AlphaLISP/AlphaPASCAL Diskette   PDB-00104-21

Because of the large amount of software now available from Alpha Micro, it
is no longer possible to fit all of the System programs onto one
single-density floppy diskette. Thus, floppy diskettes other than just the
System Diskette are available for some of the standard floppy devices (e.g.,
Wangco and Persci). These additional diskettes contain software and sources
not included in the single-density System Diskettes.

The AM-410 and AM-500 System Disks and the CDC AM-210 System Diskette
contain all of the single-density System Diskette software, plus all of the
software contained on the other non-System Diskettes. The three non-System
Diskettes are supplied in STD (128 byte/sector) format only.

Trident users wishing to update their systems from a System Diskette should
choose the System Diskette for their type of floppy drive, in whichever
format is most convenient.

The following sections discuss the overall organization of the Alpha Micro
software by disk account, and discuss what types of software and sources are
available on the various single-density floppy diskettes.

2.0 SOFTWARE ORGANIZATION

The general system release software is organized into thirteen different categories. Each category of software resides in a separate disk account:

[1,2]   This is the System Operator's account, OPR:. It is the account from which the System Operator runs various system management and maintenance programs. Only the AM-410 System Disk has a file in this account: BADBLK.SYS (which contains disk certification data for that disk).

[1,4]   This is the System Program Library account, SYS:. It is the account that contains all of the system software. You execute a program in this account by typing its name at AMOS command level.

[1,6]   This is the Device Driver Library account, DVR:. It contains all terminal, interface, and general device drivers.

[2,2]   This is the Command File Library account, CMD:. It contains a number of useful command and DO files.

[7,0]   This is the Library account, LIB:. It contains auxiliary files for system software (such as the AlphaFIX overlay files).

[7,1]   This is the Help File Library account, HLP:. It contains the HELP files. Any text file with a .HLP extension that resides in this account can be accessed via the HELP command from anywhere on the system.

[7,2]   This is the AlphaMAIL account, BOX:. The AlphaMAIL Operator uses this account in maintaining and managing the AlphaMAIL system.

[7,4]   This is the AlphaLISP Library account, LSP:. It contains AlphaLISP programs.

[7,5]   This is the AlphaPASCAL Library account, PAS:. It contains AlphaPASCAL programs, as well as the AlphaPASCAL compiler, linker, and standard library.

[7,6]   This is the AlphaBASIC Library account, BAS:. It contains various assembly language subroutines callable by your BASIC programs, and contains compiled BASIC programs. Any .RUN or .SBR file in this account can be accessed by BASIC from anywhere on the system.

[7,7]   This is the MACRO Library account, MAC:. It contains various assembly language files that can be accessed by your assembly language programs.

[10,1]   This account contains source files to various miscellaneous programs.

(Changed 31 October 1981)

[10,2]   This account contains source files for a number of terminal
           and device drivers.

Each one of the accounts above is included in the AM-500 and AM-410 System
Disks and the double-density floppy System Diskettes.

3.0 SINGLE-DENSITY FLOPPY DISKETTES

These paragraphs discuss what accounts listed above are included on which
single-density floppy diskettes:

3.1 The Single-Density System Diskette

The System Diskette is actually several diskettes that contain all of the
system software, drivers, and command files distributed with the current
release. (That is, these diskettes contain accounts [1,4], [1,6], [2,2],
[7,0], and [7,7].)

We have taken care to make sure that the diskette labeled "Part 1 of X"
contains the software necessary to boot your system under the minimal system
initialization command file contained on the diskette. There is also enough
room on the first diskette to allow you to copy files and to edit your
SYSTEM.INI file.

The four different types of System Diskettes differ only in the device
driver the monitor has been generated with. Each System Diskette contains a
SYSTEM.MON file that has been generated to access a particular System
Device.

3.2 The Single-Density Non-System Diskettes

The other three single-density floppy diskettes contain the rest of the
accounts listed in Section 2.0 above that are not included in the
single-density System Diskette.

The Miscellaneous Programs Diskette - This floppy diskette contains all of
the files in accounts [7,1], [7,2], and [10,1].

The Driver Sources Diskette - This floppy diskette contains all of the
files in account [10,2].

The AlphaLISP/AlphaPASCAL Diskette - This floppy diskette contains all of
the files in [7,4] and [7,5]. It also contains the system software and
command files necessary to use AlphaLISP and AlphaPASCAL.

A GUIDE TO THE ALPHA MICRO SOFTWARE DOCUMENTATION LIBRARY

Alpha Micro software documentation is grouped into two categories: 1) individual manuals that you can purchase separately; and 2) documents that are packaged together in the AMOS Software Update Documentation Packet.

A number in parentheses after a document title (e.g., DWM-00100-20) indicates that the document is a manual that can be purchased separately; the number is the part number by which you may order that document.

Refer to Section 3.0 for an alphabetic list of all Alpha Micro software documentation. (In that list we indicate those documents that are part of the AMOS Software Update Documentation Packet by following their titles with the code "AMOS PKT." You may only order those documents by ordering the entire AMOS Software Update Documentation Packet, part number DSS-10000-20.)

1.0 THE ALPHA MICRO SOFTWARE MANUALS

This section lists those documents not included in the AMOS Software Update Documentation Packet or as documentation for the AlphaAccounting system.

NOTE: We are in the process of writing an AMOS System Operator's Guide, which contains information necessary to the person in charge of system management (e.g., system software installation, modifying the system initialization command file, formatting disks, memory management, performing diagnostic tests on disks and memory, creating user accounts, etc.).

TITLE: AlphaBASIC User's Manual

            (DWM-00100-01, Revision B02)

READER: Aimed at all BASIC users. Assumes that you have had some prior experience with BASIC and with programming.

TOPICS: Discusses the operation of AlphaBASIC in both interactive and compiler modes. Lists all AlphaBASIC commands and functions, and discusses the various data formats supported by AlphaBASIC. Also discusses the use of MAP statements for mapping data structures into memory. Talks about advanced BASIC programming techniques (e.g., the AlphaBASIC file I/O system, calling external assembly language routines, chaining to other programs, print using, error trapping and using the ISAM system). Contains many program examples.

TITLE: AlphaFIX User's Manual

            (DWM-00100-69, Revision A01)

READER: Advanced assembly language programmer, who already fully
            understands the use of MACRO, LINK, and DDT.

TOPICS: Explains the use of FIX, the Alpha Micro screen-oriented
            debugger program for machine language programs. Lists all
            the FIX commands and modes.

TITLE: AlphaLISP User's Manual

            (DWM-00100-05)

READER: This manual is written for the LISP programmer who is
            already quite familiar with the LISP language.

TOPICS: Discusses the functions available in AlphaLISP, along with
            a description of the data types recognized. Instructions
            are included on operating the language processor, along
            with a group of sample programs written in AlphaLISP.

TITLE: AlphaMAIL User's Manual

            (DSS-10000-06)

READER: Aimed at all users of the system who are authorized users
            of AlphaMAIL. Assumes no prior experience with electronic
            mail systems but some familiarity with the Alpha Micro
            system is required. Further instructions for an
            experienced user acting as the AlphaMAIL Operator are also
            provided.

TOPICS: Gives full operating instructions for AlphaMAIL, along with
            a summary of all AlphaMAIL commands. Discusses the general
            user features of sending, receiving, printing and
            forwarding messages. Further discusses the maintenance
            system available to the Operator controlling AlphaMAIL.

TITLE: <u>AlphaPASCAL User's Manual</u>

(DWM-00100-08, Revision B00)

READER: This manual is aimed at the PASCAL programmer who is already familiar with standard PASCAL (i.e., Wirth and Jensen PASCAL).

TOPICS: Describes the AlphaPASCAL Version 2.0 compiler and linker. This implementation of PASCAL is fully compatible with the AMOS file system and the AlphaVUE text editor. The manual contains an introduction to AlphaPASCAL, information on compatibility with previous versions of AlphaPASCAL, and complete operating instructions for the AlphaPASCAL compiler, linker, and run-time package. The book also contains a complete summary of AlphaPASCAL, including information on all functions, procedures, data types, file-handling techniques, assembly language subroutines, and information on writing and modifying an external library.

TITLE: <u>AlphaVUE/TXTFMT Training Guide</u>

(DSS-10000-03)

READER: Written for all users who are new to AlphaVUE and TXTFMT, the Alpha Micro text processors. Assumes no prior experience with computers or text editors.

TOPICS: A tutorial containing exercises and demonstrations that teach the new user how to use AlphaVUE and TXTFMT. The emphasis of the book is on creating documents in a business environment. For complete information on AlphaVUE and TXTFMT, turn to the reference manuals: <u>AlphaVUE User's Manual</u> and <u>TXTFMT User's Manual</u>.

TITLE: <u>AlphaVUE User's Manual</u>

(DWM-00100-15, Revision D00)

READER: Aimed at all users of the system who want to create text files. Assumes no prior experience with text editor programs.

TOPICS: Gives full operating instructions for VUE, along with a summary of all VUE commands. Discusses the creation of the VUE initializaton file, VUE.INI, and describes those display features of your CRT terminal that are required by VUE.

TITLE: AMOS Assembly Language Programmer's Reference Manual

                (DWM-00100-43, Revision B00)

READER: Written for the experienced assembly language programmer
                who wants to become familiar with the AM-100 assembly
                language programming system.   Does NOT teach assembly
                language. For information on the instruction set used by
                the AM-100 processor, refer to the WD16 Microcomputer
                Reference Manual, (DWM-00100-04).

TOPICS: Discusses assembly language programming on the AMOS system.
                Describes operation of the AMOS macro-assembler, MACRO, the
                linkage editor, LINK, and the symbol-table file program,
                SYMBOL.   Discusses use of the object file library
                generator, LIB and the global cross reference program,
                GLOBAL.   Discusses types of expressions, data and
                statements recognized by the assembler, and gives
                information on pseudo-operations, macros and writing
                relocatable code. Also gives detailed operating
                instructions for the symbolic debugger program, DDT.

TITLE: AMOS Monitor Calls Manual

                (DWM-00100-42, Revision B01)

READER: Aimed at the advanced systems programmer, who wants to
                interface assembly language programs to the monitor via
                those monitor calls available to user programs.

TOPICS: Describes in detail the 70+ monitor routines resident in
                the operating system that are available to user programs.
                The manual also discusses the file service system, the
                terminal service system, the structure of a memory
                partition, and the format of various data structures used
                by the system (e.g., Job Control Blocks, Dataset Driver
                Block, Master File Directory, etc.).

TITLE: AMOS System Commands Reference Manual

                (DWM-00100-49, Revision A03)

READER: Written for the experienced user of the AMOS system who
                wants a quick reference guide to all commands on the
                system. We assume that the reader is already very familiar
                with the AMOS system software.

TOPICS: Contains two- or three-page summaries of all AMOS commands.
                Also contains a chart of the ASCII character set.

TITLE: AMOS User's Guide

           (DWM-00100-35, Revision A01)

READER: This manual is written for the beginning user of the system
           who has had some prior experience with computers, but who
           is new to the AMOS system.

TOPICS: Written in two parts: Part I deals with turning the system
           on and off, typing commands, logging in and specifying
           files. Part II covers the major system commands, and
           discusses command files and DO files, special wildcard
           commands and file backup procedures. Command descriptions
           are ordered by type and function.

TITLE: Introduction to AMOS

           (DWM-00100-65)

READER: This manual was written for the person who wants some
           background information on the AMOS system in particular and
           computers in general. Assumes little computer experience.

TOPICS: NOT an operations or demonstration manual.   Written in
           three parts: Part I deals with elementary computer
           concepts, defining terms such as "files," "hexadecimal,"
           "CPU," and "program."   Part II talks about some of the
           programs available on the AMOS system, such as AlphaBASIC,
           VUE, TXTFMT, and ISAM. Part III gives an overview of the
           Alpha Micro operating system, giving a very general
           introduction to the Terminal Service System, Job Scheduler,
           Command Processor, and other major portions of AMOS.

TITLE: ISAM System User's Guide

             (DWM-00100-06, Revision A02)

READER: This manual is aimed at the assembly language programmer
             who wants to use the ISAM (Indexed Sequential-files Access
             Method) package from within an assembly language program.
             Assumes thorough knowledge of assembly language, and prior
             experience with ISAM. NOTE: For complete information on
             using ISAM from within BASIC programs, refer to the
             AlphaBASIC User's Manual.

TOPICS: Describes the ISAM calls available to the assembly language
             programmer for modifying and using ISAM files, as well as
             the various ISAM programs that the user can run from the
             monitor level for building and dumping ISAM files. Also
             discusses the process of building and dumping ISAM files.
             Contains brief discussion of using ISAM from within a BASIC
             program.

TITLE: TXTFMT User's Manual

             (DWM-00100-07, Revision B00)

READER: Written for all users of the system who want to create
             formatted documents. Assumes experience with one of the
             text editing programs, AlphaVUE or EDIT.

TOPICS: Lists all of the text formatting commands recognized by
             TXTFMT. Also discusses TXTFMT operation and error
             messages.

TITLE: WD16 Microcomputer Reference Manual

             (DWM-00100-04)

READER: This manual is aimed at the experienced assembly language
             programmer who wants to become familiar with the
             instruction set used by the AM-100 processor.

TOPICS: The manual does NOT contain information on assembly
             language programming, or on using the AMOS assembler. It
             does list all of the instructions and addressing modes used
             by the WD16 microprocessor. For information on assembly
             language programming on the AMOS system, refer to the AMOS
             Assembly Language Programmer's Reference Manual.

2.0 THE AMOS SOFTWARE UPDATE DOCUMENTATION PACKET

The documents in the AMOS Software Update Documentation Packet (DSS-10000-20) are organized into four major groups:

      1. User's Information

      2. System Operator's Information

      3. System Programmer's Information

      4. BASIC Programmer's Information

For a list of the documents included in the AMOS Software Update Documentation Packet, refer to the Master Table of Contents; this document is the first document in the AMOS Software Update Documentation Packet.

3.0 AN ALPHABETIC LIST OF ALL ALPHA MICRO SOFTWARE DOCUMENTATION

NOTE: The code AMOS PKT following a title indicates that the document is part of the AMOS Software Update Documentation Packet.

AlphaAccounting Release Notes - Version 1.4 (DSS-10000-17), available only to users licensed for the AlphaAccounting software package.

AlphaBASIC User's Manual (DWM-00100-01, Revision B02)

AlphaFIX User's Manual (DWM-00100-69, Revision A01)

AlphaLISP User's Manual (DWM-00100-05)

AlphaMAIL User's Manual (DSS-10000-06)

AlphaPASCAL User's Manual (DWM-00100-08, Revision B00)

The AlphaVUE/TXTFMT Training Guide (DSS-10000-03)

AlphaVUE User's Manual (DWM-00100-15, Revision D00)

AMOS 4.6 Software Update Documentation Packet (DSS-10000-20)

AMOS Assembly Language Programmer's Reference Manual (DWM-00100-43, Revision B00)

AMOS Monitor Calls Manual (DWM-00100-42, Revision B01)

AMOS System Commands Reference Manual (DWM-00100-49, Revision A03)

AMOS User's Guide (DWM-00100-35, Revision A01)

AMOS Version 4.4 Method of Handling Bad Disk Blocks (AMOS PKT)

AMOS 4.6 Release Notes (AMOS PKT)

BASORT - BASIC Subroutine for Sorting Random and Sequential Files (AMOS PKT, Revision A03)

Building a Terminal Driver (The FIXTRM Program) (AMOS PKT, Revision A01)

Change Page Packet #3 for AlphaBASIC User's Manual (DSS-10000-19)

Change Page Packet #1 for AlphaPASCAL User's Manual (DSS-10000-10)

Change Page Packet #1 for the AlphaVUE/TXTFMT Training Guide (DSS-10000-16)

Change Page Packet #4 for AMOS System Commands Reference Manual (DSS-10000-18)

Change Page Packet #1 for the ISAM System User's Guide, (DSS-10000-22)

Change Page Packet #1 for the TXTFMT User's Manual (DSS-10000-13)

COMMON - BASIC Subroutine to Provide Common Variable Storage (AMOS PKT)

Configuring Floppy Disk Drivers (AMOS PKT, Revision A01)

Defining Non-system Disk Devices (AMOS PKT, Revision A01)

Defining Switchable System Memory (AMOS PKT, Revision A03)

Disk Drivers and Formats (AMOS PKT, Revision A05)

Disk Labeling Procedures (AMOS PKT)

Disk Maintenance Procedures for the System Operator (AMOS PKT, Revision A06)

Disks Available from Alpha Micro (AMOS PKT, Revision A06)

The DUMP Command (AMOS PKT, Revision B00)

EDIT: A Character-oriented Text Editor (AMOS PKT, Revision A01)

FLOCK - BASIC Subroutine to Coordinate Multiuser File Access (AMOS PKT, Revision A01)

Generating System Monitors (AMOS PKT, Revision A03)

A Guide to the Alpha Micro Software Documentation Library (AMOS PKT, Revision A06)

Important Notice for ISAM Users (AMOS PKT)

Important Notice for LISP Users (AMOS PKT)

Introduction to AMOS (DWM-00100-65)

I/O Programming for the Alpha Micro Computer (AMOS PKT, Revision A01)

ISAM System User's Guide (DWM-00100-06, Revision A02)

The Magnetic Tape File Backup Programs (AMOS PKT, Revision A01)

Memory Management Option (AMOS PKT, Revision A01)

Memory Management with the Memory Partition Controller (AMOS PKT, Revision A01)

New Command File and DO File Features (AMOS PKT, Revision A01)

Program Design Language Formatting System (AMOS PKT)

Setting Up the Line Printer Spooler (AMOS PKT, Revision A03)

Software Information for the AM-610 Video Cassette Recorder Interface (DSS-10000-24)

Software Information for the AM-620 1/4" Streaming Tape Drive (DSS-10000-23)

Software Information for the Users of the AM-960 Status Display (AMOS PKT)

Software Installation Instructions for the AM-120 (AMOS PKT, Revision A01)

Software Installation Instructions for the AM-710 Memory Board (AMOS PKT)

Software Notice for AM-410 Users (AMOS PKT, Revision A03)

Software Notice for AM-420 Users (AMOS PKT, Revision A01)

SPOOL - BASIC Subroutine for Spooling Files to the Line Printer (AMOS PKT, Revision A02)

The System Initialization Command File (AMOS PKT, Revision A06)

Terminal Service System (AMOS PKT)

TXTFMT User's Manual (DWM-00100-07, Revision B00)

Using the Magnetic Tape Utility Programs (AMOS PKT, Revision A01)

WD16 Microcomputer Reference Manual (DWM-00100-04)

The Winchester/Floppy Backup Programs (AMOS PKT)

XLOCK - BASIC Subroutine for Multiuser Locks (AMOS PKT, Revision A03)

XMOUNT - BASIC Subroutine to Mount a Disk (AMOS PKT, Revision A01)

AMOS Software Update Documentation
AMOS Release 4.6
October 1981

USER'S INFORMATION

This section contains the following documents:

Important Notice for ISAM Users

New Command File and DO File Features, Revision A01

The DUMP Command, Revision B00

Important Notice for LISP Users

EDIT: A Character-oriented Text Editor, Revision A01

Program Design Language Formatting System

AMOS VERSION 4.6 SOFTWARE UPDATE DOCUMENTATION PACKET

'Alpha Micro', 'AMOS', 'AlphaBASIC'
'AlphaPASCAL', 'AlphaLISP', 'AlphaSERV',
'AlphaVUE', and 'AlphaACCOUNTING'

are trademarks of

ALPHA MICROSYSTEMS
Irvine, CA   92714

ALPHA MICROSYSTEMS
17881 Sky Park North
Irvine, CA   92714

NEW COMMAND FILE AND DO FILE FEATURES

Command and DO files play an important part in extending the power and range of the AMOS command language. This document describes several new features and programs that greatly increase the flexibility of your command files. Among other things, these new command file features allow you to: change the command file trace flag as many times as you wish within a command file; conditionally transfer control to various portions of a command file based on whether a specific file exists; and, exit a command file, perform various AMOS commands, and then resume execution of that command file. We also describe several new AMOS level commands that are particularly useful as command file elements.

For detailed information on the concepts of command files and DO files, refer to Chapter 8, "Command Files and DO Files," in the AMOS User's Guide, (DWM-00100-35). In brief, a command file is a special kind of text file which can contain a series of AMOS commands, specifications of other command files, test data, etc. To execute all of the commands and data in a DO file or a command file, just enter the name of the file at AMOS command level. Command files are extremely useful for performing entire sequences of commands that you use frequently (e.g., commands that do disk backup or that compile a series of BASIC programs).

DO files are a type of command file that allow you to specify arguments when you invoke the DO file that are substituted into that DO file where special parameter symbols appear. This allows you to use one DO file with a wide range of files or programs. Since DO files are just one type of command file, when we mention "command files" in the rest of this document, we will be talking about both standard command files and DO files.

Special symbols may appear in command files that allow you to ask the user of the file for input, and to display messages to that user. In addition, you may use the :T, :S and :R symbols to choose whether or not users of the command file see program output that result from the actions of your command file. Except for these special symbols, when AMOS processes a command file, it treats most of the elements of the file as if you were entering each command file line from a terminal, line by line.

For more information on any command discussed below, refer to the reference sheet for that command in the AMOS System Commands Reference Manual, (DWM-00100-49, Revision A01 and later).

1.0 THE COMMAND FILE TRACE FLAG (:T AND TRACE)

Command files have always allowed you to choose whether or not the user of
your command file sees the contents of the file while it is being processed
by AMOS. AMOS checks the status of the command file "trace flag" to
determine whether or not to display the contents of the command file. Until
the advent of the TRACE command, the only way to affect the trace flag was
by including the :T symbol at the place in the command file where you wanted
users to see command file contents. The main disadvantage to this was that
you could only use the :T symbol once to turn the trace flag on in a
command file, and then could not turn it off again. (Of course, if you
wanted the trace flag off, you could simply omit the :T from the command
file.)

NOTE: When the trace flag is off, you can still allow users to see program
output and messages by use of the special :S, :R, and :<> symbols; however,
they will not see comments or command lines in the command file. For
example, if the line:

    .ERASE *.BAK (RET)

appears in the command file when the trace flag is off, if no :R precedes
the command line, the user of the file does not see the command line OR the
information output by the ERASE program. If a :R precedes the command line
and the trace flag is off, the user does not see the command line, but does
see the output produced by ERASE. NOTE: Remember to always include a :R in
your command file if the trace flag is off but you want the user of the
command file to see any output generated by the commands in the command
file.

1.1 The TRACE Command

The TRACE command allows you to turn the trace flag on or off at any point
within the command file as many times as you wish. You may also use TRACE
to switch the current state of the trace flag, regardless of whether it is
on or off. You may only use TRACE in a command file.

To use TRACE, enter it in your command file in one of the following three
formats:

    TRACE ON   Turns the command file trace flag ON. Users see all of
                    the contents of your command file, including comments and
                    all program command lines. :S and :R are ignored when
                    TRACE ON is in effect.

    TRACE OFF   Turns the command file trace flag OFF. Users see only
                    messages enclosed with the :<> symbols. If you include
                    the :R symbol, users also see program output. (The :S
                    symbol turns off the :R symbol.) (NOTE: :S and :R work
                    with TRACE OFF as if no TRACE or no :T were in the
                    command file.)

(Changed 30 April 1981)

TRACE SWITCH   Switches the current status of the command file trace flag. If the flag is OFF, this command turns it ON and vice versa.

If you do not include an argument on the TRACE command line, AMOS ignores the command.

As an example of the use of the TRACE command, let's create a command file that does disk backup:

```
; Command file to backup HAWK System Disk onto another disk.
LOG DSK0:1,2
:< Make sure nobody else is running on DSK0: or your backup disk>
ERASE *.BAK[]
DIRSEQ
;
TRACE ON
DSKCPY
DSK0:
DSK1:
:<
All done. Remove cartridge and label it. Log back into your
account.>
```

Not having a :T or TRACE OFF makes sure that the user of the command file won't see the output of the cleanup functions we perform. (The user still sees those messages bracketed by the :<> symbols.) The TRACE ON ensures that the user will see that we are running the DSKCPY program.

2.0 LOOKING UP FILES FROM WITHIN A COMMAND FILE (LOOKUP)

There are many occasions when a command line in a command file causes AMOS to search for a file. If AMOS fails to find the specified file, execution of the command file does not abort, but continues even though the necessary file was not found. Sometimes this can be a severe inconvenience if several subsequent command lines assume that the nonexistent file exists.

2.1 The LOOKUP Command

The LOOKUP command allows you to search for a file from within a command file, and then to perform several actions based on the results of that search. If a file is not found, you may resume execution of the command file at the command line following the lookup, or you may cause the command file to terminate execution. You may also choose whether to display your own error message or a standard AMOS error message if the file is not found.

By using LOOKUP in combination with the GOTO command, you can choose which portions of the command file to execute based on the results of the file lookup. This allows you to perform conditional branching within a command file.

You may only use the LOOKUP command within a command file. To use LOOKUP, enter it into your command file in one of the following formats:

LOOKUP Fspec   where Fspec specifies the file you want to search for. If the file is found, LOOKUP continues execution of the command file; if it is not found, LOOKUP displays the appropriate AMOS error message (e.g., ?Cannot OPEN STDMOD.BAD - file not found), and returns the user of the command file to AMOS command level.

LOOKUP Fspec Msg   where Fspec specifies the file you want to search for, and Msg is a message supplied by you. If the file is found, LOOKUP just continues execution of the command file. If the file is not found, LOOKUP displays your message (instead of the appropriate AMOS error message) and returns the user of the command file to AMOS command level.

LOOKUP Fspec/   where Fspec specifies the file you want to search for, and the "/" symbol tells LOOKUP not to abort command file execution if the file is not found. If the file is found, LOOKUP skips over the next line in the command file and resumes execution at the line past it. If the file is not found, LOOKUP continues execution at the next line after the LOOKUP command, and displays the appropriate AMOS error message.

LOOKUP Fspec/Msg   where Fspec specifies the file you want to search for, "/" tells LOOKUP not to abort command file execution if the file is not found (see the paragraph above), and Msg is a message supplied by you. If the file is found, LOOKUP skips over the next line in the command file and resumes execution at the line past it. If the file is not found, LOOKUP resumes execution at the next line in the command file, and displays the specified message. (If you include both the "/" symbol and a message, the message must follow the slash on the LOOKUP command line.)

If you omit portions of the file specification, LOOKUP assumes the device and account the user of the command file is logged into and a .PRG extension.

The LOOKUP "/" option becomes especially useful when you use the GOTO and EXIT statements to select certain portions of the command file to be

(Changed 30 April 1981)

executed as a result of the LOOKUP operation.   (See Section 4.0,
"Transferring Control Within a Command File (GOTO, EXIT)," for an example of
conditional branching within a command file.)

As an example of the use of the LOOKUP command, let's create a DO file that
installs a BASIC program into an account. (The user of the DO file provides
the name of the program to install as an argument when he or she invokes the
DO file. This argument gets substituted into the DO file for the $0 symbol,
which is a special DO file parameter symbol.)

```
; Command file to install $0.BAS from project library account
; into the account user is logged into.
;
:R
LOOKUP $0.BAS[300,0] ?That BASIC program doesn't exist. Try again.
COPY =[300,0]$0.BAS
;
TRACE ON
COMPIL $0.BAS
RUN $0.RUN
```

If the user of the DO file doesn't specify an argument, or gives an
incorrect file specification, the LOOKUP command above catches it and ends
command file execution.

NOTE: The command file above provides a good example of why you often need
to check to see if a proper file specification has been given.   If the
example above did not use the LOOKUP command, and if the user of the command
file supplied no argument at all (thus causing a space to be substituted for
the $0 symbol), the COPY command would copy ALL .BAS files from [300,0]
over to the account the user is logged into.

3.0 TEMPORARILY INTERRUPTING COMMAND FILE EXECUTION (PAUSE, CONT)

It would often be convenient to temporarily exit a command file, perform
various cleanup or housekeeping functions, and then resume use of the file.
The PAUSE and CONT commands allow you to do so.

The PAUSE command causes the temporary interruption in the execution of the
command file in which it appears.   You may then execute AMOS commands,
invoke other command files, use a text editor, etc. To resume execution of
the command file, use the CONT command.

3.1   The PAUSE Command

You may only use PAUSE within a command file. Enter the PAUSE command in
your command file where you want to temporarily interrupt execution of the
file.   You may optionally include a message on the PAUSE command line which
is displayed when the command file pauses. For example:

  PAUSE Type a K; COPY old .BAS files to [40,1]; type CONT to resume backup

This is what happens when a PAUSE statement is processed in a command file
you are using:

    1.   If a message has been included on the PAUSE command line,
         PAUSE displays that message to you.

    2.   Whether or not it has displayed a message, PAUSE now stops
         and waits for you to type a character. If you type a RETURN,
         PAUSE resumes execution of the command file. If you type
         anything but a RETURN, PAUSE returns you to AMOS command
         level.

    3.   Once PAUSE returns you to AMOS command level, you can now run
         any programs or command files you want. To resume execution
         of the command file (at the point after the PAUSE command),
         just type CONT at AMOS command level. (The CONT command may
         also appear within a command file.)

When PAUSE interrupts execution of a command file, it saves the elements of
the command file past the PAUSE command line in a special disk file named
CNT.CMD. (This file appears in the device and account the user of the
command file is logged into.) When you use the CONT command, it loads into
your memory partition the CNT.CMD file that is in your account and resumes
executing that command file.   (This means, of course, that to resume
execution of a PAUSEd command file, the user of the file must be logged into
the device and account where the command file was originally interrupted.)
The CNT.CMD file always contains the most current command file that has been
PAUSEd in that account.   For example, if you use a command file that is
interrupted because of a PAUSE command, the rest of that file is stored as
CNT.CMD in your account.

If, after performing various commands, you do not use a CONT command before
invoking another command file that also uses PAUSE to interrupt execution,
the contents of CNT.CMD (the first command file) are replaced by the
contents of the second command file. If you use the CONT command now, you
resume execution of the second command file, not the first. When a second
PAUSE causes the current contents of CNT.CMD to be written over (because a
CONT has not been used to resume execution of the previous command file),
you see the following message:

         %Supersedes existing file

(Changed 30 April 1981)

No harm is done if you do not resume command file execution by using a CONT
command   before   another   PAUSE command occurs, but you do lose the previous
contents of CNT.CMD when the new command file replaces them.

As an example of the use of PAUSE, let's create a command file that performs
"housekeeping" functions on a disk:

```
; This command file cleans up the disk and performs a backup.
:R
:<
If disk labeled BACKUPA is in drive, type a RETURN; otherwise,
type anything else, go put pack in drive, and then type CONT
when you are ready to resume backup.>
;
PAUSE
; Proper backup pack (DSK5:) is now in drive.
MOUNT DSK5:
:<
If you want to save everything on the disk, type a RETURN.
If there are any scratch files you don't want to save, type
anything but a RETURN.  Then erase your old, working files
from all accounts on the disk.  To resume backup, type CONT.

>
PAUSE Enter a character:
;
; Disk is ready to back up.  Chain to another command file
; that does actual backup.
BACKUP
:<
All done.  You may remove BACKUPA from drive.>
```

3.2 The CONT Command

If a command file has been temporarily interrupted as the result of a   PAUSE
command (that   is,   if a CNT.CMD file appears in the account you are logged
into), you may resume execution of that   command   file   by using   the   CONT
command.   For example:

        .CONT (RET)

If there is a CNT.CMD file, CONT tells AMOS to process it.   When the entire
command file has been processed, CONT erases the CNT.CMD file from the disk.
If no CNT.CMD file exists in the account, CONT displays the message:

        ?Can't continue

and you are returned to AMOS command level.

4.0 TRANSFERRING CONTROL WITHIN A COMMAND FILE (GOTO, EXIT)

Since the LOOKUP command allows you to choose which line of a command file
to execute, it now becomes possible to use a transfer command, GOTO, to
select which portion of a command file to execute based on the results of a
file lookup.   Used in combination with one another, the GOTO, EXIT, and
LOOKUP commands allow your command files to perform conditional branching.

4.1 The GOTO Command

You may only use the GOTO command within a command file. The GOTO command
allows you to transfer control from one portion of your command file to
another. The GOTO command line must contain both the GOTO command and an
argument which is the name of the label to branch to. For example:

```
:R
; Command file to compile BASIC programs.
;
LOOKUP TAXTBL.BAS/?Couldn't find file.  Are you in right account?
; If file not found, go to NOFILE.
GOTO NOFILE
TRACE ON
COMPIL TAXTBL.BAS
RUN TAXTBL.RUN
;
EXIT *That's all...Returning you to AMOS command level*
;
;NOFILE
:<
We're going to enter VUE so you can create TAXTBL.BAS.  Type
an X if you don't want to create the file; otherwise,>
PAUSE Hit RETURN when ready:
;
VUE TAXTBL.BAS
Y
```

In the example above, the GOTO command line contains the argument NOFILE.
NOFILE is the label of the portion of the command file to which the GOTO
command transfers control.

There are some conventions you must follow in setting up a GOTO, its
argument, and the label that designates the portion of the command file to
which you want to transfer.

> 1. GOTOs must precede the labels they branch to. That is, GOTO
>   statements may only transfer forward in the file.

2.   An argument may not contain trailing spaces. That is, the end of
     an argument must be either a RETURN or a semicolon (which
     identifies the start of a comment). That means that if you include
     a comment on the GOTO command line, it must begin directly after
     the argument.

3.   The command file label may either be a comment (that is, begin with
     a semicolon) or a valid, executable command file element.

     a.   If a label is a comment, the argument in the GOTO command line
          that refers to the label must <u>not</u> begin with a semicolon.
          (See the sample command file above.)

     b.   If a label is not a comment, it must be a valid command file
          element (e.g., a program name, a command file specification,
          etc.).

     If the GOTO statement directs you to a label that is not a comment
     or a valid command file element, the command file resumes execution
     after the label.

4.   You may begin a label with spaces, a semicolon, or spaces followed
     by a semicolon.   (There may be no spaces between a semicolon and
     the rest of the label.) These are ignored when GOTO compares an
     argument to the label it selects.

5.   Labels may be of any length (as long as they fit on one line), and
     must be the only thing on the line.

If GOTO cannot find the specified label, the user of the command file sees:

     <u>?Label not found</u>

and is returned to AMOS command level.

4.2   The EXIT Command

Whenever you create conditional branches within a command file, you face the
problem of separating portions of the command file so that users not
deliberately transferred to a labeled portion do not "fall into" that
section of the command file as they proceed through the file.   The EXIT
statement allows you to create one or more points in the file which cause
the user to gracefully be returned to AMOS command level. You may only use
EXIT within a command file. The sample command file above demonstrates the
use of the EXIT statement. As another example, consider the DO file below,
which does different things with a file, based on the extension of that
file:

```
; Command file that handles text files.
; If file .TXT file, format it.
:R
LOOKUP $0.TXT/?not .TXT file
GOTO NOTTXT
TRACE ON
TXTFMT $0.TXT
TRACE OFF
;
;NOTTXT
LOOKUP $0.LST/?not .LST file
GOTO NOTLST
PRINT $0.LST
EXIT *Your file is formatted and the .LST version is printed.*
;
;NOTLST
EXIT ?Couldn't find a .TXT or .LST file of that name.
```

5.0 ADDITIONAL USEFUL COMMANDS (BATCH, COM)

In addition to the commands we discussed above, two other new commands exist
that are particularly helpful when used within command files. You may also
use these commands at AMOS command level.

5.1 The BATCH Command

A command file executes faster if the programs it accesses are already
loaded into memory when that command file needs them.   The BATCH command
loads into your memory partition programs that are frequently used when
making use of the new command file features (GOTO, LOOKUP, EXIT, TRACE,
PAUSE, and LOAD).   (NOTE: GOTO, LOOKUP, EXIT, TRACE, and PAUSE are
re-entrant, and may be placed into system memory by the System Operator.) If
you use BATCH from within a command file rather than at AMOS command level,
you will want to put BATCH at the front of the command file. BATCH takes up
about 1K of your memory partition. To use BATCH, enter the command followed
by a RETURN:

        .BATCH (RET)

5.2 The COM Command

The COM command processes a file based on its extension. To use the
command, enter COM followed by the name of the file you want to affect. You
may not include the extension of the file.   In addition, the file must
appear within the account you are logged into, but you may specify a
different device. For example:

        _.COM STD1:MNMENU (RET)

COM begins looking for the disk file in this order:

    1.   .MAC file? Then assemble with MACRO.

    2.   .BAS file? Then compile with BASIC.

    3.   .PAS file? Then compile with PRUN CMPILR.

    4.   .TXT file? Then format with TXTFMT.

If you omit portions of the file specification, COM assumes the account and
device you are logged into. (Of course, if COM is used within a command
file, COM uses the account and device of the user of the command file as the
defaults.)

If COM can't find the file you have specified, or if the file does not have
one of the extensions listed above, you see:

        ?Filename is not a compilable file

where Filename is the file you specified on the COM command line.

THE DUMP COMMAND

1.0 INTRODUCTION

This document describes the DUMP utility program. (Also, see the reference sheet on DUMP in the AMOS System Commands Reference Manual, (DWM-00100-49), for a brief summary of DUMP command format.)

DUMP gives you a simple method for examining data, either in memory or on the disk. You can use DUMP to take a look at the contents of memory, the contents of a block on the disk, a disk Master File Directory, a disk User File Directory, or a disk bitmap. You can use DUMP to look at the contents of both random and sequential files. DUMP is re-entrant, and may be loaded into system memory by the System Operator.

Some uses of DUMP require that you give it one or more numeric arguments. These arguments must be in the number base the system is using for your numeric displays (usually octal). (You can use the SET command to change this number base from octal to hexadecimal, and vice versa. See the SET reference sheet in the AMOS System Commands Reference Manual (DWM-00100-49), for more information on SET.) Some uses of DUMP require that you supply keywords which select the DUMP function you want to use. You may abbreviate these keywords by giving just as many characters as will uniquely identify that keyword. (For example, you may enter DI instead of DIRECTORY.)

NOTE: Previous versions of this document used the term "disk records." The use of the word "record" can cause some confusion since it is sometimes used in other documentation to mean different things. In the interests of clarity, therefore, we have adopted the convention that the 512-byte groups of data into which AMOS organizes the disk are called "disk blocks," not "disk records." We have changed this document and the DUMP program accordingly. (We have retained the earlier DUMP format "DUMP RECORD"; it performs exactly the same function as the new "DUMP BLOCK" command.)

2.0 DUMP FUNCTIONS

DUMP allows you to select six different display functions. The DUMP format you use selects the specific function you want to perform:

2.1 Displaying Memory

FORMAT:

    .DUMP Address1 Address2 (RET)

To display memory, give DUMP two memory addresses in the number base the
system is using for your numeric displays. (If your system is a
bank-switched system, remember that you may only display memory addresses
that are in the memory bank within which your job resides.) For example:

    .DUMP 110000 120002 (RET)

DUMP displays the contents of memory from the first address to the second
address, inclusive.   If you do not supply a second memory address, DUMP
displays only the first 16 bytes of data. NOTE: DUMP rounds the starting
address down to the nearest multiple of 16.

A DUMP memory display looks something like this:

110000:006562 020056 052040 062550 066400 071557 020164 067543 rm. The most co
110020:066555 067157 072040 070171 020145 063157 062040 071551 mmon type of dis
110040:066160 074541 072040 065541 071545 072040 062550 063040 play takes the f

Let's take a look at the first line of this display:

1. The first number on the left (ending with a colon) is the memory
    address that contains the first byte of data on the line. In this
    example, memory addresses 110000 and 110001 contain the two bytes
    of data 006562.

2. Each group of six digits after the memory address represents two
    bytes (16 bits) of data in octal form. (If the system is using
    hexadecimal for your numeric displays, DUMP displays the data in
    groups of eight bits-- one byte.)

3. On the far right of the display is a field that gives the ASCII
    form of the data. It displays the same data as the numbers in the
    center of the display, but translated into their character
    representation.   Non-printing ASCII characters (such as
    Control-characters) appear as dots.

2.2 Displaying a Disk File

FORMAT:

    .DUMP Filespec (RET)

(Changed 1 May 1980)

To display a random or sequential disk file, type DUMP followed by the specification of the file you want to see. Then type a RETURN. For example:

.DUMP DSK1:PROJCT.OBJ[12,45] (RET)

If you omit portions of the file specification, DUMP assumes the device and account you are logged into and a .PRG extension. DUMP displays the entire file in the same form as the memory display (see above).

DUMP precedes the display with a message that tells you the number of the block you are seeing. For example:

Block number 12033 of DSK0:DATA.DAT[35,4]

If you are dumping a sequential file, DUMP also displays the next block link. For example:

Block number 784 of DSK1:PROJCT.OBJ [12,45], next block link is 11027.

2.3 Displaying a Disk Block

FORMAT:

.DUMP BLOCK Block-number1 {Block-number2} {Devn:} (RET)

To display the data in a disk block, type DUMP BLOCK. (You may also use the format DUMP RECORD.) Now enter the number of the block you want to see and the specification of the logical unit that holds that block.

If you want to see the contents of several blocks, enter a second block number. Type a RETURN. DUMP now displays the data from the first to the second block, inclusive. For example:

.DUMP BLOCK 1355 1360 HWK2: (RET)

If you do not specifiy a device, DUMP assumes the device you are logged into. The DUMP display looks much like the memory display above. The addresses on the left side of the display give the relative position of the data in each block.

2.4 Displaying a Bitmap

FORMAT:

.DUMP BITMAP {Devn:} (RET)

To display a disk bitmap, enter DUMP BITMAP. Now enter the specification of the logical unit whose bitmap you want to see. (If you omit the device

specification, DUMP assumes the device you are logged into.) Now type a
RETURN. For example:

    .DUMP BITMAP DSK5: (RET)

You now see a display of ones and zeros that represents the bitmap of the
disk. (A bitmap is a disk allocation map that marks which disk blocks are
in use and which are available.) Every block on the disk is represented in
the bitmap by a one (if in use) or a zero (if empty).   Each line of the
display begins with a number which ends in a colon; this octal or hex number
is the number of the disk block represented by the first one or zero on that
line.

At the end of the display, DUMP prints the total number of free blocks (in
decimal) on the disk. (An STD-format diskette has a total of 512 blocks; a
Hawk hard disk has 9696 blocks, and a Phoenix hard disk has about 29088
blocks.) For example:

    .DUMP BITMAP HWK1: (RET)

Bitmap Dump of HWK1:

```
000000:  1111111111111110111101111111111011111111101111111111111111111111
000100:  1111111111011111111101111111111111111111100111111111100111111111
000200:  1111011100011000011001000010101010111111110001000111111111111111111
000300:  1111111111111111111111110111111111111110111111111111110011111111111
                                   .
                                   .
                                   .
022500:  000000000000010000000000001110000001010111111000000000000110000000
022600:  000000000000000000000001000000000000000000000000000001110000000000000
022700:  000000000000000000000000000000000
```

    4850 free blocks

    .

2.5 Displaying a Disk Master File Directory

FORMAT:

    .DUMP MFD {Devn:} (RET)

Every disk contains a Master File Directory in Block 1 that contains a list
of all accounts on the disk and the starting block number of all individual
account directories (the User File Directories, or UFDs). To see the MFD of
a disk, type DUMP MFD. Now enter the specification of the logical unit
whose MFD you want to see. (If you omit the device specification, DUMP
assumes the device you are logged into.) Finally, type a RETURN.   For
example:

    .DUMP MFD HWK3: (RET)

The display you see looks something like this:

        Master File Directory Dump of HWK3:

        000000   [1,2]   002110
        000010   [11,5]   002105
        000020   [50,1]   000024
                     .
                     .
                     .
        000170   [0,0]   000000
        000200   [0,0]   000000
          .

Each line represents one User File Directory. The number on the left gives
the relative address of the MFD entry in the disk block. The characters in
the center give the account PPN. The number on the right gives the disk
block at which the UFD for that account begins.
Nonexistent accounts appear as:

    [0,0]

This is a normal part of the MFD display.

2.6 Displaying a Disk User File Directory

FORMAT:

    .DUMP DIRECTORY Block-number {Devn:} (RET)

or:

    .DUMP DIRECTORY [p,pn] {Devn:} (RET)

To display the directory for an individual account (i.e., the User File
Directory, or UFD), type DUMP DIRECTORY followed by the number of the block
at which the UFD starts. Or, you may simply enter the PPN associated with
that account, using the standard AMOS PPN format of [p,pn]. (If you use
[p,pn] format, and DUMP is not able to find the account you specified, you
see: ?Illegal user code.)

Now enter the specification of the logical unit that contains the UFD you
want to see. Type a RETURN. For example:

    .DUMP DIRECTORY 002105 SMD4: (RET)

or:

    .DUMP DIRECTORY [110,2] SMD4: (RET)

(Changed 1 May 1980)

If you omit the device specification, DUMP assumes the device you are logged into. The display may look something like this:

```
Directory dump of block 2105, next block link is 2564
Addr       Filename       Size      Active      Link
000002     SYSINO TXT     74        000166      002306
000016     :80RLD TXT     21        000034      002421
000032     PHONX  LST     34        000053      002446
000046     :80ST  MAC     27        000227      002650
000062     DOCMAN BAS     21        000413      002654
000076                    0         000000      000000
000112                    0         000000      000000
```

The first line of the display tells you the number of the block you are looking at. The next line tells you what disk block contains the next section of the directory. (A "next block link" of 0 indicates that the block you are displaying is the last block in the directory.) The rest of the display gives information about the directory entries.

You see this information for each directory entry:

Addr   The position (in bytes) of the directory entry relative to the start of the block.

Filename   The name and extension of the file.

Size   The number of disk blocks in the file.

Active   The number of active data bytes in the last block of the file.

Link   The address of the first disk block of the file.

NOTE: Directory entries in which the filename begins with the characters ":80" represent files that have been deleted from the directory. These are normal elements of the DUMP DIRECTORY display. The next time the system writes a new entry into the directory, it overwrites the first deleted entry.

(Changed 1 May 1980)

IMPORTANT NOTICE FOR LISP USERS

Several new functions and enhancements have been added to LISP in AMOS Release 4.4. These features include improved error reporting and the addition of functions to the Extended Library to handle breakpoints.

1.0 ERROR HANDLING

When LISP reports an error, it now displays the <u>user function</u> in which the error occurred. For example:

```
*(DE DOUBLE (X) (PLUS XX))
DOUBLE
*(DOUBLE 2)
UNBOUND VARIABLE - EVAL IN DOUBLE
                             =========

XX
*
```


2.0 NEW FUNCTIONS

Three new functions have been added to LISP: RETFROM, BREAK, and UNBREAK. In addition, we have added the variable BREAKFNS (which is maintained by BREAK and UNBREAK).

2.1 RETFROM

The call (RETFROM fn val) causes the most recent call of function <u>fn</u> to return with value <u>val</u>. If the specified function is not active, LISP generates an error message. For example:

```
*(DE F1 (X) (PROGN (F2) X))
F1
*(DE F2 () (RETFROM @F1 5))
F2
*(F1 7)
5
```

NOTE: The call (RETFROM PROG val) behaves exactly the same as (RETURN val).

2.2 BREAK (added to the Extended Library)

The call (BREAK fn1 fn2 ...) causes execution of a program to be interrupted if an attempt is made to call any of the specified functions. You may then single-step execution of the interrupted function by typing a line-feed, or resume execution by typing (RESUME). NOTE: fn1, fn2, ... are not evaluated.

2.3 UNBREAK (added to the Extended Library)

The call (UNBREAK fn1 fn2 ...) restores the specified functions so that they no longer interrupt program execution when called. (That is, this function clears breakpoints set via the BREAK function.) NOTE: fn1, fn2, ... are not evaluated.

2.4 BREAKFNS (added to the Extended Library)

BREAKFNS is a variable which contains a list of all functions which will interrupt program execution when called. BREAKFNS is maintained by BREAK and UNBREAK; therefore, you should not directly modify this variable. (See BREAK and UNBREAK, above.)

PROGRAM DESIGN LANGUAGE FORMATTING SYSTEM

1.0 INTRODUCTION

The Program Design Language Formatting System is a tool that helps you to produce a program-design document.

The first step in creating a program-design document is to use one of the Alpha Micro text editors (EDIT or VUE) to write a document. Write the document in the form set by the Program Design Language-- PDL (see Section 2.0, "Program Design Language").

When you exit the text editor, and are again at the AMOS command level, you may use the Program Design Language Formatting System (PDLFMT) to transform your text file into a finished program-design document.

PDLFMT produces a document which contains the following:

1. Table of Contents.

2. Formatted Design Listing - Each procedure in the program is listed on a separate, numbered page, with the page numbers that refer to other procedures in the margin. Each page indents the text to show control-structure nesting.

3. Reference Trees - Indented listing shows how procedure references are nested.

4. Cross Reference - Alphabetical listing of all sections and procedures. An index of where the sections and procedures appear, and where references to them appear (page and line numbers).

2.0 PROGRAM DESIGN LANGUAGE

A program design written in PDL has this form:

```
/T  Design Title
/S  Section Name
Text describing the section........
...................................
/P  Procedure Name
Text giving the procedure design...
...................................
/P  Procedure Name
.........  Text  ..................
...................................
/S  Section Name
........  Text  ..................
/R
Procedure name
Procedure name
...................................
```

NOTE: the slash (/) must be the first character on the line.

1. /T Design Title - Specifies the name of the program design. This title appears on every page of the finished document. (/T is called the title command.) The title command must always be the first command in the design, and must always be present.

2. /S Section Name - Specifies the start of a new section of procedure designs. The section name specified will appear on the pages of the finished design document as subtitles. After the section command is free-form text that describes the section. (This text may be any length or form that you want, but no line of it may begin with a /.)

3. /P Procedure Name - Specifies the start of a procedure design, and assigns Procedure Name as the name of that design. Any time the Procedure Name occurs as a statement within a procedure design, PDLFMT considers that occurrence as a reference to the procedure design.

4. /R Reference Tree - Specifies the start of a list of procedure names on successive lines. Each procedure name is a root of a reference tree listing. The Reference Tree command is optional, but if you do include it in your document, it must be the last command in the design.

2.1 Procedure Design

A procedure design consists of a sequence of statements. You may label each statement. (Labels are an alphanumeric identifier followed immediately by a colon.) PDLFMT will indent labels by -2 in the final document.

You may also precede each statement by one of the keywords: IF, ELSEIF, ELSE, ENDIF, DO, ENDO, or ENDDO. PDLFMT uses these keywords when it formats your design document. The rest of the line after an IF or ELSEIF is considered a condition; PDLFMT will not consider it a potential procedure reference.

Statements are sequences of text that end with carriage returns. They may contain embedded comments enclosed in parentheses. If the text of a statement (ignoring the comments) matches the name of a procedure or section, PDLFMT considers it a reference to that procedure or section.

You may continue statements on one or more lines by placing an ampersand (&) at the beginning of succeeding lines.

2.2 Control Structures

You can use the keywords IF, ELSEIF, ELSE, ENDIF, DO, and ENDDO (or ENDO) to indicate a variety of control structures. The paragraphs that follow give some idea of the possibilities.

2.2.1 The IF Construct
The IF construct provides the means for indicating condition execution. It corresponds to the classic IF...THEN...ELSE construct found in Algol-60 and PL/I, augmented by the ELSEIF of languages such as Algol-68. The ELSEIF is used to prevent excessive indentation of levels when cascaded tests are used.

The general form of the construct is:

        IF condition
         one or more statements
         ...
        ELSEIF condition
         one or more statements
         ...
        ELSEIF condition
         one or more statements
         ...
        ELSE
         one or more statements
         ...
        ENDIF

NOTE: You are allowed any number (including zero) of ELSEIFs, and you are allowed one ELSE at the most.

2.2.2 <u>The DO Construct</u> - Use the DO construct to indicate repeated execution, and for case selection. Indicate the iterative DO by:

```
DO iteration criteria
   one or more statements
   ...
ENDDO
```

You can choose the iteration criteria to suit the problem. Typical criteria begin with the words WHILE, UNTIL, or FOR. WHILE denotes a continuation criteria which is checked before each iteration. UNTIL denotes a termination criteria that is checked after each iteration. FOR denotes a range of items over which the one or more statements are to be applied. Examples:

```
DO WHILE THERE ARE INPUT RECORDS
DO UNTIL "END" STATEMENT HAS BEEN PROCESSED
DO FOR EACH ITEM IN THE LIST EXCEPT THE LAST ONE
```

Provision for premature exit from a loop and premature repetition of a loop are frequently useful. To accomplish this, you can take the statement UNDO to mean that control is to pass to the point following the ENDDO of the loop.   Likewise, CYCLE can be taken to mean that control is to pass to the iteration criteria test. If you want to apply UNDO or CYCLE to an outer loop in a nest of loops, you may label any DO and place the label after the UNDO or CYCLE.

You can indicate case selection by:

```
DO CASE selection criteria
```

In general, labels are used in the body of the DO to indicate where control passes for each case:

```
      DO CASE OF TRANSACTION TYPE
ADD:
      CREATE INITIAL RECORD
DELETE:
      IF DELETION IS AUTHORIZED
         CREATE DELETION RECORD
      ELSE
         ISSUE ERROR MESSAGE
      ENDIF
CHANGE:
      INCREMENT CHANGE COUNT
      CREATE DELETION RECORD
"OTHER":
      ISSUE ERROR MESSAGE
   ENDO
```

3.0 OPERATING INSTRUCIONS

Call PDLFMT from the AMOS command level by typing:

     .PDLFMT Filespec(RET)

where Filespec specifies a design file prepared with a text editor. If you
omit a filename extension, PDLFMT assumes a .PDL extension.   The formatted
design document is placed in the file Filespec.LST.

A demonstration file, TLGRAM.PDL, is included with PDLFMT.PRG on the dealer
distribution disk. For a demonstration of PDLFMT, type:

     .PRINT TLGRAM.PDL(RET)
     .PDLFMT TLGRAM(RET)
     .PRINT TLGRAM(RET)

4.0 ERROR MESSAGES

X IS AN ILLEGAL COMMAND - BYPASSING LINE

     PDLFMT found a command /X in the design document, but X is not a legal
     command. Use /T, /S, /P, or /R.

REPEATED DEFINITION: xxxxxxx, REFERENCES WILL BE TO LAST OCCURRENCE

     xxxxx is a section or procedure name that occurred in a section or
     procedure command more than once.

5.0 ACKNOWLEDGMENTS

PDLFMT is based on PDL and its processor, as described in "PDL - A Tool for
Software Design," by Stephen H. Caine and E. Kent Gordon of Caine, Farber,
and Gordon, Inc.

EDIT - A CHARACTER-ORIENTED TEXT EDITOR

1.0 INTRODUCTION

There are two text editing programs available on the AMOS system: VUE (a screen-oriented text editor) and EDIT (a character-oriented editor). For information on VUE, see the manual AlphaVUE User's Guide (DWM-00100-15). Unfortunately, a manual for EDIT does not yet exist. This document gives only a brief summary of the EDIT commands.

Character-oriented text editors were originally designed to be used on non-CRT terminals. Because these kinds of editors were designed to be used on hard copy terminals that do not permit fast display, the emphasis of such an editor is not on display, but on speed and power.

When you edit a text file, a text editor brings a copy of the file into memory and allows you to make your editing changes to the copy in memory; then the editor writes your changed file back out to the disk. EDIT maintains a pointer (called DOT) that points to your current position in the copy of the text file that is in memory. Most commands that you give to EDIT reference this pointer to see what text to affect.

You do not see any of the text in memory unless you explicitly ask EDIT to display one or more lines of text. You advance throughout the text in memory by using the various EDIT commands to move DOT. EDIT commands are one or two characters long, and some require arguments (e.g., you follow the search command with a string of text for which to search).

2.0 EDITING A NEW FILE

Before you can begin to enter text into an empty file, you must create the file by using the MAKE program. Type MAKE and the specification you want to assign to the new file. For example:

        .MAKE DSK1:NEWFIL.TXT[100,2] (RET)

You may only create a file in your own account or in an account within your own project.

After you have used the MAKE command, you can now use EDIT to enter text into the file (see below).

3.0 EDITING AN EXISTING TEXT FILE

You may use EDIT on any sequential file that contains ASCII characters. To
edit an existing file, type:

      .EDIT Filespec (RET)

where Filespec is the name of the file you want to edit. The default EDIT
extension is .MAC. After you hit RETURN, you see the EDIT prompt: *. You
are now ready to enter EDIT commands. DOT initially points to the first
character in your file. When you exit EDIT, the editor renames your
original disk file to a .BAK extension (for BACKUP), and saves your edited
copy under the original file's name and extension.

4.0 SPECIAL CHARACTERS

When entering text and commands to EDIT, you may use the RUB key (also
labeled RUBOUT, DEL, DELETE, etc.) to erase single characters, and a
Control-U to erase an entire line of input. EDIT itself has a group of
commands that you must use to delete those characters and lines already part
of the file you are editing.

EDIT uses the Escape key (labeled on your keyboard as ESC, ALT MODE, etc.)
rather than a RETURN as a command delimiter; this allows you to enter
carriage returns as part of your text. When you type an Escape to EDIT, you
see the character displayed on your terminal as a dollar sign, $. Whenever
you see a $ in this document, the symbol indicates an Escape.

5.0 THE COMMANDS

You may enter EDIT commands either in upper or lower case. You may enter
the commands one at a time, ending your input with two Escapes. For
example:

      *C$$
      *

After the EDIT prompt symbol, *, we entered the Character-advance command,
C, which moved DOT ahead in the file by one character-position. To tell
EDIT that the command line was complete, we entered two Escapes. EDIT
responds with another prompt to let us know that it is ready for another
command.

You may also enter commands as a group. For example:

      *LKT$$

The example above tells EDIT to move DOT to the beginning of the next line (L), kill (that is, delete) from DOT to the end of the current line, and type (that is, display) the characters from DOT to the end of the current line.

A line consists of all of the characters between two carriage-return/linefeed character pairs.

Some commands take numeric arguments (e.g., "3D" says delete the three characters after DOT); numeric arguments are decimal numbers, and always precede the command to which they apply. If a command takes a text argument (e.g., "Stext" says search for the word "text"), you must end the text argument with one or two Escapes. (One Escape tells EDIT that the last command is complete, but that it may not yet take action upon the current input line. Two Escapes tell EDIT to go ahead and act upon the current command line.) For example:

> *IThe "I" command tells EDIT to insert text$10T$$

The command line begins with an Insert command, I. EDIT will insert into your text file (at the current DOT position) all text (including carriage returns) following the I command up to a single or a double Escape. The single Escape above terminates the text entry string; next is a display command, 10T, that tells EDIT to display the 10 lines of characters that occur after the current position of DOT. The double Escapes tell EDIT to go ahead and act upon the entire string of commands. You may enter as many lines of commands and text as you wish; EDIT will not take action upon the input until you hit two Escapes.

If you want to cancel a string of input, you can do so by typing a Control-C (as long as you type the Control-C before entering the double Escapes).

5.1 A SUMMARY OF THE EDIT COMMANDS

Below is an alphabetical list of the EDIT commands. (Remember, DOT is the pointer that marks your position in the file.)

A   APPEND - Appends one or more records of the input file to the
            data buffer if there are at least 2000 free bytes of memory
            left, and DOT has not reached the end of the file.

C   CHARACTER ADVANCE - Moves DOT forward one character (e.g., C$$).

nC   CHARACTER ADVANCE - Moves DOT forward by "n" characters (e.g.,
            3C$$).

-nC   CHARACTER ADVANCE - Moves DOT backward by "n" characters (e.g.,
            -5C$$).

0C   CHARACTER ADVANCE - Moves DOT backward to the beginning of the
            current line.

D   DELETE - Deletes the first character after DOT.

nD   DELETE - Deletes the next "n" characters after DOT (e.g., 3D$$).

-D   DELETE - Deletes the character just behind DOT (e.g., -D$$).

-nD   DELETE - Deletes the the previous "n" characters behind DOT
             (e.g., -20D$$).

0D   DELETE - Deletes characters from the beginning of the line up to
             DOT.

HD   DELETE - Deletes entire buffer; that is, deletes as much of the
             file as is in memory.

E   EXIT - Exit to monitor. Outputs data buffer, and rest of input
             file.   Renames new file to original file's name and extension,
             and renames original file to a .BAK extension.

EG   EXIT AND GO - Exits to monitor and, if it is a .BAS or a .MAC
             file, processes the text file as is appropriate for its
             filetype. An EG$$ command used on a .BAS file tells the monitor
             to load in BASIC and compile the file; the EG command used on a
             .MAC file tells the monitor to load in MACRO and assemble the
             file.

EQ   EXIT AND QUIT - Exits to monitor, but doesn't make the editing
             changes you entered; the original file is left as is, untouched,
             and is not renamed to .BAK.

F   FREE MEMORY - Prints decimal number of free bytes left in your
             memory partition.

Gx   GET AUXILIARY - Gets auxiliary buffer "x" where the symbol x may
             be the letters A-Z. Inserts the buffer into the file at the
             current position of DOT.   DOT is moved forward the number of
             characters inserted.

Itext$ INSERT - Inserts specified text into the file at the current
             position of DOT.   You may insert carriage returns and other
             special symbols except for those Control-characters discarded by
             AMOS on input (See Special Insert, below).   Remember that the
             text is not actually inserted until you type two Escapes. For
             example:

                  *IThis is all one
                  ‾input; all of the
                  characters, even the
                  carriage returns, can
                  be entered with one
                  insert command$$

(Changed 1 July 1979)

nI   SPECIAL INSERT - You may insert special Control-characters not
              usually accepted by EDIT, by preceding the I command with the
              decimal ASCII code of the character you want to insert (e.g.,
              12I$$ inserts an ASCII character 12-- a form-feed).

0J   JUMP - Jumps DOT back to the beginning of the data buffer (i.e.,
              back to the beginning of the portion of your file that is in
              memory).

nJ   JUMP - Jumps DOT to immediately in front of the "nth" character
              in the data buffer.

ZJ   JUMP - Jumps DOT to the end of the buffer.

K   KILL - Kills the characters from DOT to the end of the current
              line.

nK   KILL - Kills the next "n" lines of text past DOT.

0K   KILL - Kills from the beginning of the current line to DOT.

-K   KILL - Kills from the beginning of the previous line to DOT.

-nK   KILL - Kills from the start of the "nth" line behind DOT up to
              DOT itself.

HK   KILL - Kills the entire data buffer.

L   LINE ADVANCE - Advances DOT to the beginning of the next line.

nL   LINE ADVANCE - Advances DOT forward "n" lines.   DOT is
              positioned at the start of the line.

0L   LINE ADVANCE - Moves DOT back to the start of the current line.

-L   LINE ADVANCE - Moves DOT back to the start of the previous line.

-nL   LINE ADVANCE - Moves DOT backward "n" lines from the current
              position of DOT, and positions DOT to the start of the line.

Linefeed   LINE ADVANCE AND TYPE - Typing a linefeed (Control-J) performs
              same function as LT$$; that is, advances to the front of the
              next line, and displays that line. A backspace (Control-H)
              performs the same function as the -LT$$ command; that is, moves
              DOT back to the start of the previous line, and displays that
              line.

Ntext$ WHOLE FILE SEARCH - Searches the current data buffer, beginning
              at DOT, for the first occurrence of "text". If the search
              within the current data buffer is not successful, EDIT writes
              that data buffer out to the disk, and brings in more text; DOT
              is reset to the beginning of that buffer, and the search begins
              again.   This process continues until "text" is found, or until

the end of the file has been reached. If "text" is found, DOT
is positioned just after it; if "text" is not found, EDIT
displays an error message: [SEARCH FAILED], and any commands
occurring in the input string after the search command are
aborted.

nNtext$ WHOLE FILE SEARCH - Same as above, except that search stops at
               the "nth" occurrence of "text" (e.g., 10Nsilicon$).

N$ WHOLE FILE SEARCH - Same as Ntext$, except that it uses the last
               text string that you entered to a search command.

nN$ WHOLE FILE SEARCH - Same as nNtext$ except that the last search
               string entered is used.

R   REVERSE - Same as -C command.

nR   REVERSE - Same as -nC command.

0R   REVERSE - Same as 0C command.

-R   REVERSE - Same as C command.

-nR   REVERSE - Same as nC command.

Stext$ SEARCH - Searches the data buffer beginning with DOT for the
               first occurrence of "text". Positions DOT just after "text" if
               it finds it; otherwise it displays the message: [SEARCH FAILED],
               DOT is positioned to the front of the buffer, and the rest of
               the commands in the input string are aborted.

nStext$ SEARCH - Same as Stext$, but EDIT searches for the "nth"
               occurrence of the search string "text".

S$ SEARCH - Same as Stext$ above, but EDIT uses the last search
               string entered.

nS$ SEARCH - Same as nStext$, but uses the last search string
               entered.

FSoldtext$newtext$$

               SEARCH-AND-REPLACE - Searches for "oldtext" and replaces it with
               "newtext". The command FNoldtext$newtext$$ performs the same
               function, but on the entire file, rather than on just the
               current data buffer.

T   TYPE - Displays the characters from DOT to the end of the line.

nT   TYPE - Displays "n" lines of characters starting from DOT.

0T   TYPE - Displays the characters from the beginning of the current
               line up to DOT.

-T   TYPE - Displays the characters from the start of the previous
                Line up to DOT.

-nT   TYPE - Displays from the "nth" Line behind DOT up to DOT.

Vx   VERIFY - Verifies auxiliary buffer contents, where the symbol
                "x" is a character from A-Z.   Lists the contents of the
                auxiliary buffer.

Xx   SAVE - Saves, in auxiliary buffer "x", the characters from DOT
                to the end of the current Line. (X is a character A-Z.) The
                previous contents of the auxiliary buffer are lost.

nXx   SAVE - Saves "n" Lines past DOT in auxiliary buffer "x".   ("x"
                is a character A-Z.)

0Xx   SAVE - Saves from the beginning of the current Line to DOT in
                auxiliary buffer "x".

-nXx   SAVE - Saves, in auxiliary buffer "x", from the start of the
                "nth" Line previous to the Line DOT is in up to DOT.

;text$ SEMICOLON INSERT - Performs same function as the Itext$ command,
                except that a semicolon is placed at the start of the inserted
                text.

TABtext$ TAB INSERT - Performs same function as the Itext$ command,
                except that a TAB character is placed at the start of the
                inserted text.

SPACEtext$ SPACE INSERT - Performs same function as the Itext$ command,
                except that a space is placed at the start of the inserted text.

n<...>   REPEATS - ALL of the commands within the angle brackets are
                repeated "n" times (e.g., 10<FSprimpt$print$>$$ tells EDIT to
                search for the word "primpt" and replace it with "print" ten
                times). ALL EDIT commands can be executed in a repeat,
                including other repeats. The maximum nesting level for repeats
                is eight. An error message and an abort occurs if you exceed
                the nesting limit; search failures also abort repeats.

                If you omit "n" the group of commands repeat endlessly until an
                error occurs, or until a Control-C is typed.   Often used to
                replace   ALL   occurrences   of   an   item (e.g.,
                <FSregistrar$register$>$$).

AMOS Software Update Documentation
AMOS Release 4.6
October 1981

SYSTEM OPERATOR'S INFORMATION

This section contains the following documents:

   The System Initializaton Command File, Revision A06

   Setting Up the Line Printer Spooler, Revision A03

   Memory Management Option, Revision A01

   Memory Management with the Memory Partition Controller, Revision A01

   Defining Switchable System Memory, Revision A03

   Configuring Floppy Disk Drivers, Revision A01

   Software Installation Instructions for the AM-120, Revision A01

   Software Installation Instructions for the AM-710 Memory Board

   Software Notice for AM-420 Users, Revision A01

   The Winchester/Floppy Backup Programs

   Software Notice for AM-410 Users, Revision A03

   Disk Labeling Procedures

   Disk Maintenance Procedures for the System Operator, Revision A05

   AMOS Version 4.4 Method of Handling Bad Disk Blocks

   Defining Non-system Disk Devices, Revision A02

   Disk Drivers and Formats, Revision A05

   Generating System Monitors, Revision A03

   Using the Magnetic Tape Utility Programs, Revision A01

   The Magnetic Tape File Backup Programs, Revision A01

   Building a Terminal Driver (The FIXTRM Program), Revision A01
      (NOTE: In the AMOS 4.5 Release, this document was titled "Building a
      Terminal Driver (The NEWTRM Program)."

DOCUMENTS CATEGORIZED BY CONTENT

SYSTEM SETUP

The System Initialization Command File, Revision A06
Setting Up the Line Printer Spooler, Revision A03
Memory Management Option, Revision A01
Memory Management with the Memory Partition Controller, Rev A01
Defining Switchable System Memory, Revision A03
Generating System Monitors, Revision A03

DISK MAINTENANCE

AMOS Version 4.4 Method of Handling Bad Disk Blocks
Disk Labeling Procedures
Disk Maintenance Procedures for the System Operator, Revision A05
Disk Drivers and Formats, Revision A05

DISK BACKUP

The Winchester/Floppy Backup Programs
Using the Magnetic Tape Utility Programs, Revision A01
The Magnetic Tape File Backup Programs, Revision A01

INSTALLING AND USING NEW DEVICES

Building a Terminal Driver (The FIXTRM Program), Revision A01
Software Installation Instructions for the AM-120, Revision A01
Software Installation Instructions for the AM-710 Memory Board
Software Notice for AM-420 Users, Revision A01
Software Notice for AM-410 Users, Revision A03
Defining Non-system Disk Devices, Revision A02
Configuring Floppy Disk Drivers, Revision A01

Purpose of these Standards

The Alpha Micro allows the use of a large variety of terminals, jobs, disks printers and devices. Each of these items is given a name in the SYSTEM.IN and San Luis Software (SLS) programs allow these names to be whatever yo like. However, each change in the names in the SYSTEM.INI requires corresponding change in the setup (not programming) of the SLS programs. Th fewer changes that are made from the typical SLS SYSTEM.INI the faster th installation time and the less likely that errors will show up at a late time. The names we have chosen are not going to match the names that th whole world is now using, but they are reasonable names and we simply had t pick a single name for each item. Consistency is a virtue!

San Luis Software is providing service to a large number of clients. Ou dial-up diagnostic baud rate is 300. It greatly increases our investigatio speed if we know that jobs are named JOB1, JOB2, etc., and terminals are name CRT1, CRT2, etc. Any time a name does not conform to the sample shown below we must do SYSTAT's or other time-consuming commands in order to determine th names used. Standard names save large amounts of time during these dial-u sessions and get the system back to the client sooner.

Typical SLS SYSTEM.INI

The following sample will be used to help explain the standards that must b used with SLS SYSTEM.INI's. As usual, there are a number of of ways to set u the SYSTEM.INI but SLS can't guarantee that .INI's that vary greatly from thi will work.

***** Start of Sample *****

```
; SLS1.INI - San Luis Software SYSTEM.INI
; Pete - 17 July 80
; Pete - 8 Jan 81 - change CRT1 to HAZEL driver

JOBS JOB1,JOB2,SPOOLA,SPOOLB,REMOTE

TRMDEF CRT1,    AM300=1:16,HAZEL,  80, 80, 20

SET DSKERR

XY=0
: San Luis Software, Inc
----------------------
SLS Version 2.1b
Amos 4.4a(2)

System Startup

TRMDEF CRT2,    AM300=3:16,HZ1500,  80, 80, 20
TRMDEF NEC,     AM300=4:10,HAZEL,  40, 40, 10
```

UTL005 (Rev. 26-Dec-81)

```
        TRMDEF TT,        AM300=2:16,HAZEL,  40,  40,  10
        TRMDEF REMOTE     DCH=120:234,HAZEL, 80,  80,  20

        :   Assign Devices

        DEVTBL DSK1,AMS0,AMS1,TRM,MEM,RES,IMG0,STD0,STD1

        QUEUE 40

        SET HEX
        MEMDEF 40,0,F                         ; Pi 64K
        MEMDEF 41,1,0/4C,FF,0                 ; Micro 32K, IMS 16K
        MEMDEF 42,1,0/4E,FF,0                 ; Micro 32K, IMS 16K

        SYSMEM 1:F900-FFFF

        :   Allocate Device Maps

        BITMAP AMS,39,0,1/S
        BITMAP DSK,605,0,1/S
        BITMAP STD,32,0,1/S

        :   Load Memory Modules

        SYSTEM FLTCNV.PRG[1,4]                ; All Basic programs call this
        SYSTEM SL0007.SBR[7,6]                ; All Basic programs call this
        SYSTEM SL0030.SBR[7,6]                ; All Basic programs call this
        SYSTEM AMS.DVR[1,6]                   ; Driver for diskette
        SYSTEM TRM.DVR[1,6]                   ; For direct printing
        SYSTEM SL0310.PRG[100,0]              ; Spooler control program
        SL0309 5 5 100                        ; All SLS systems require this area
        SYSTEM

        MOUNT DSK1:
        CLKFRQ 60
        HEDLOD 1800

        :   Assign Job Memory

        JOBMEM JOB2 1:8000-F4E4
        JOBMEM REMOTE 2:8000-FFFF
        JOBMEM SPOOLA 1:F4E6-F8FE

        :   Setup TI Printer

        ATTACH CRT2,SPOOLA
        KILL SPOOLA
        FORCE SPOOLA SL0310[100,0] TI NORMAL CRT2 /DFS

        :   Setup NEC Printer

        ATTACH CRT2,SPOOLB
        KILL SPOOLB
        FORCE SPOOLB MEMORY 1050
        FORCE SPOOLB SL0310[100,0] NEC NORMAL CRT2 /FS
```

UTL006 (Rev. 26-Dec-81)

San Luis Software SYSTEM.INI Standards

```
:    Startup Jobs

ATTACH CRT2,     JOB2
ATTACH TT,       SPOOLA
ATTACH NEC,      SPOOLB
ATTACH PHONE,    REMOTE

KILL JOB2
KILL REMOTE

FORCE JOB2 LOG 1,4
FORCE REMOTE LOG 1,4

: System Ready

Control Returned to AMOS
Enter MENU to start SLS systems


MEMORY 0

***** End of Sample *****
```

Overall Comments

We recommend that :T (trace) be used in all .INI's so that any errors in boot-up can be seen. The comments included in this sample will give a good trace of the bootup activity if the :T trace is not used.

The most recent changes should be listed in the comments at the top of file so that problems can be related back to the date of the changes.

The use of spaces, both vertically and within the statements, makes it easy to pick out important pieces of the .INI.

For the most part, like statements are grouped together, such as JOBS, TRMDEF's, JOBMEM's, etc.

After any modification to the SYSTEM.INI, boot and then use the SYSTEM command to check the monitor size. The monitor must be 32,767 bytes or less.

JOB Statement

Use JOB1, JOB2, JOB3, JOB4, JOB5, etc for user jobs. Use SPOOLA for the 1st matrix printer, SPOOLB for the 1st letter quality printer. Use SPOOLC for the 2nd matrix printer and SPOOLD for the 2nd letter quality printer. Use REMOTE for a dial-up job.

TRMDEF Statement

Use CRT1, CRT2, CRT3, CRT4, CRT5, etc., for user terminals. Use TI for Texas Instruments TI810, NEC for NEC Spinwriter, DIA for Diablo. Use TI2 for the 2nd TI, NC2 for the 2nd NEC, and DI2 for the second Diablo. Use PHONE for the dial-up job terminal.

UTL005 (Rev. 26-Dec-81)

Use baud rate 9600 (code 16) for all terminals and the TI. Use 1200 (code 10 for the NEC.

To conserve memory use the same terminal driver for each printer as for CRT1.

For each terminal, use 80, 80, 20 for buffers. For each printer, use 40, 40 10. Note that there is quite a bit of controversy over how much degredatio there is with these minimal sized buffers.

DEVTBL Statement

Use the Alpha Micro recommended names of DSK, AMS, STD, HWK, SMD and IMG Always include TRM and include MEM and RES if there is room.

QUEUE Statement

Use QUEUE size of 40 unless there is not enough room.

SET Statement

Set DSKERR on after the TRMDEF statement to catch any device errors. Set HE before the MEMDEF's so that the MEMDEF's and SYSMEM may be coded in he instead of octal (we needed to choose one or the other; we chose hex).

MEMDEF Statement

Add a comment to each line to show the Manufacturer and size of each board This helps the hardware person to understand which boards go with which banks

SYSMEM Statement

We always put bitmaps into switchable system memory.

BITMAP Statement

Overlay the bitmaps and put them into switchable system memory (/S).

SYSTEM Statement

You must include each non-system driver for Basic to access the device Include TRM.DVR for direct printing and many programmer functions. RUN an FLTCNV can be loaded into each job area but their are used by nearly ever program and really should be in the system area.

SL0310 is the spooler control program and can be included in system memory o can be part of each spooler. If SL0310 is included in system memory eac spooler requires 884 bytes. If not in system memory, each spooler require 1050 bytes.

Other SL subroutines can be loaded in if space permits. In the followin order, load in any subroutines that will fit: SL0007.SBR, , SL0020.SBR SL0026.OVR, SL0010.SBR, SL0030.SBR, SL0031.SBR, SL0005.SBR.

SL0309 Statement

UTL006 (Rev. 26-Dec-81)

SL0309, like the SYSTEM statement, loads modules into the system area. It requires TABLE.SYS in DSK0:100,0. It must be included for all SLS systems.

MOUNT Statement

Include a MOUNT for each device that is used regularly. Do not mount all Phoenix drives if they are not normally used.

JOBMEM Statement

Use a JOBMEM for each terminal job except JOB1. Include a JOBMEM for each printer that will not be in Bank 0. Each printer requires 884 bytes if SL0310 is in system memory and 1050 bytes if SL0310 is NOT in system memory.

Setting up each Printer

Attach CRT2 to the spooler job to avoid printout when booting; this printout always misaligns the paper and will usually mess up the first print job. KILL the spool job and FORCE the SL0310 statement. SL0310 logs itself to 100,0 and therefore does not require a log statement. SL0310 resides in DSK0:100,0 so it needs the PPN specification to be found.

The format of an SL0310 statement is:

SL0310[100,0] Queue Starting-Form System-Terminal /Switches

Use the same name for the Queue as for the printer terminal, i.e. TI, NEC, TI2, etc.

Use NORMAL for the Starting-Form.

The System-Terminal will receive any "change form" messages; specify whichever terminal you want. Note that if the /S switch (see below) is set, there will be no messages.

Valid Switches are: DFS, where D means default printer (use for 1 printer only); F, which means Form-Feed at end of file; and S, which means Suppress "change form" messages. S can be used since all SLS systems will tell the operator to mount a special form before starting a spooler.

ATTACH Statement

Always attach CRT1 to JOB1, 2 to 2, etc; TI to SPOOLA, NEC to SPOOLB, etc. Attach PHONE to REMOTE.

Finishing Up

Be sure to KILL each job (JOB1 and printers are already set) and log in any jobs that you wish. Leave jobs at the AMOS level and do not try to force the SLS MENU command to them.

Finish off with a MEMORY 0 so that JOB1 regains all excess Bank 0 memory.

UTL006 (Rev. 26-Dec-81)

THE SYSTEM INITIALIZATION COMMAND FILE

October 1981
Revision A06

This document reflects AMOS versions 4.6 and later

Table of Contents

1.0 INTRODUCTION

The Alpha Micro Operating System has been designed so that you can
"customize" it for your particular hardware set-up (that is, adapt it to run
with your terminals, your disks, etc.) In fact, you MUST do that initial
system software installation if your system software has not already been
set up so that it conforms to your machine's hardware configuration.
Because a user's hardware and system needs can keep changing as a system
grows, Alpha Micro has created a simple mechanism for adapting the operating
system to reflect those changes: the system initialization command file (the
SYSTEM.INI).

The SYSTEM.INI is a special kind of command file. (A command file is a text
file that contains system commands; the system reads and obeys the
instructions it reads from the file.) Whenever the system is powered up or
reset, it consults the SYSTEM.INI to find out what devices you use on the
system and what special programs and functions you want to add to the
operating system area of memory.

The rest of this document discusses the elements of the SYSTEM.INI and the
changes that you can make to the file to reflect changes in your system
configuration. (NOTE: While you read this document, keep in mind that the
terms "operating system," "monitor," and "system" are roughly
interchangeable in the pages that follow.)

1.1 Modifying the SYSTEM.INI

CAUTION: Let's assume that you have a SYSTEM.INI file that gets your
system up and running, and you want to change it to reflect some hardware
additions or changes. Before you edit the file, it's a good idea to make a
backup copy of your System Disk so that you can bring the system up again,
even if something goes wrong with your modified SYSTEM.INI.

If you change the file that works (and somehow your new SYSTEM.INI doesn't
work), you won't be able to get the system up and running off of that disk.
Before you can again use the disabled disk as a System Disk, you'll have to
bring the system up off of another System Disk, and transfer over a copy of
a good SYSTEM.INI. The system will not come up unless you have a valid
SYSTEM.INI file on your System Disk.

If you are running the system off physical drive zero (i.e., the
fixed-platter in a CDC Hawk hard-disk system, the first fixed disk in the
CDC Phoenix hard-disk system, or Drive Zero in a floppy-disk system) you can
make a copy of the good SYSTEM.INI under a different name, edit THAT version
of the file, and then use the MONTST command to test the new copy. This

procedure leaves you with a valid SYSTEM.INI under its original name, so that if you have to reset the computer, the system will be able to come up under the control of the original, valid SYSTEM.INI.

To use the MONTST command, type MONTST, SYSTEM.MON, a comma and the name of the modified command file:

       .MONTST SYSTEM.MON,NEWSYS.INI

If you are not running off physical drive zero, you will not be able to use the MONTST command, and will have to modify the SYSTEM.INI file itself; in that case, make sure that you have a valid System Disk that you can bring the system up on before you change your SYSTEM.INI file. (Or, in the case of a Winchester-disk based system that contains an AM-610 Video Cassette Recorder Interface or an AM-620 1/4" Streaming Tape Drive, make sure you have a valid warm boot cassette or cartridge.)

To change the SYSTEM.INI to reflect your own needs and hardware configuration, edit it with one of the text editors on the system (EDIT or VUE). The SYSTEM.INI resides in area [1,4] of the System Disk. After you change the file, reboot the system by pressing the reset button or by using the MONTST command (see above).

NOTE: Beginning with AMOS Release 4.6, a new tool is available for restoring a Winchester System Disk that runs under the control of the AM-420 Hard Disk Controller-- users who use the AM-610 Video Cassette Recorder Interface or the AM-620 1/4" Streaming Tape Drive may create a special tape that contains a "warm boot" monitor that may be used to reboot the system if the System Disk becomes damaged. For more information, see the manuals Software Information for the AM-610 Video Cassette Recorder Interface, (DSS-10000-24), and Software Information for the AM-620 1/4" Streaming Tape Drive (DSS-10000-23).

1.2 System Start-up

The monitor performs a certain set of procedures while it is coming up. If one of these steps fails, the system will not come up. These steps are:

    1.   When you press the reset button, the Alpha Micro CPU starts executing instructions at the address set up in its header. (If the CPU is an AM-100, this address is the address of the PROM on the disk controller board; if the CPU is an AM-100/T, the address is the address of a PROM on the CPU board itself.)

    2.   The program in the PROM transfers itself down into RAM between 31K-32K. (If your memory in these locations is bad or nonexistent, the system start-up will not proceed beyond this step.) If the phantom memory option is installed on the disk controller board, the phantom memory now becomes active.

(Changed 31 October 1981)

3. The PROM program (the bootstrap loader) is now in RAM and it begins
    to execute. It reads in the operating system skeleton monitor,
    which is a file called SYSTEM.MON in the [1,4] System Disk program
    area. The loader reads SYSTEM.MON into memory beginning at
    location zero and extending as far as necessary.

4. When SYSTEM.MON is in memory, it executes the initialization
    routine (INITIA) within the monitor itself. The purpose of INITIA
    is to scan memory to determine how much is available, and then to
    set up a user memory partition in the last 8K of memory.   (This
    user partition is temporary, and is just used to execute the system
    start-up functions under control of the SYSTEM.INI.)

5. Once INITIA is through setting up the user partition, the monitor
    reads in the system initialization command file, SYSTEM.INI.   Each
    line in the SYSTEM.INI represents one system function or parameter
    which determines the characteristics of the running monitor.

6. The monitor executes the commands in the SYSTEM.INI just as it
    would the commands in any other command file. Because this command
    file is the system initialization file, however, the monitor
    performs some of the commands in the SYSTEM.INI differently than it
    would the same commands after the system is completely up.   The
    execution of certain system commands (e.g., JOBS, TRMDEF, DEVTBL,
    etc.) performs the actual system generation.

During system start-up, certain programs cause the monitor to create new
areas at the end of itself; these areas include terminal definition blocks,
terminal drivers, job control blocks, device tables, memory bank tables,
system queues, and disk bitmaps. Your SYSTEM.INI may optionally specify a
list of programs to be added to the resident monitor area of memory during
system start-up.

The size of the monitor is not fixed, but is expanded during system
start-up. This is why INITIA allocates the initial user partition in the
top 8K portion of memory. This gives the monitor room as it expands so that
it won't overlap the user partition. (The monitor executes the SYSTEM.INI
in that user partition.)

NOTE FOR BANK-SWITCHED SYSTEMS: If your system uses the bank switching form
of memory management, the system uses the top 8K of Bank Zero to process the
SYSTEM.INI.   You can go ahead and allocate that portion of bank zero to a
job as a user memory partition (via the JOBMEM command), but DO NOT try to
use that job to run anything until after the SYSTEM.INI is fully processed
and the system is up and running. (That is, do not use a FORCE command
within the SYSTEM.INI to force input to the job.) NOTE: For information on
memory management via bank switching, see the Alpha Micro Integrated
Systems User's Guide, (DWM-00101-00).

NOTE FOR MEMORY PARTITION CONTROLLER SYSTEMS: On an MPC system, the monitor
uses the top 8K of the first 64K of memory to process the SYSTEM.INI.   For
information on the MPC (Memory Partition Controller) form of memory

management, see the document <u>Memory Management with the Memory Partition</u>
<u>Controller</u> in the "System Operator's Information" section of the AMOS
Software Update Documentation Packet.

2.0 A SAMPLE SYSTEM INITIALIZATION COMMAND FILE

This page and the next contain a typical SYSTEM.INI file for a system that
uses the bank switching form of memory management. The sections that follow
base their discussions on this sample SYSTEM.INI. (For information on how
to set up a system that uses the Memory Partition Controller to perform
memory management, see the document <u>Memory Management with the Memory</u>
<u>Partition Controller</u> in the "System Operator's Information" section of the
AMOS Software Update Documentation Packet.)

```
        :T
        JOBS  JOB1,JOB2,JOB3,JOB4,SPOOL
        ;
        TRMDEF  TRM1,AM300=1,ADM3,100,100,200
        TRMDEF  TRM2,AM300=2:6,SIL700,100,80,30
        TRMDEF  TRM3,AM300=3:16,SOROC,100,100,100
        TRMDEF  TRM4,AM300=5:16,SOROC,100,100,100
        TRMDEF  PRNTR,AM300=6:10,SOROC,100,100,20      ; Define printer
        TRMDEF  NULL,PSEUDO,NULL,25,25,2               ; Define
        ;                                                pseudo-terminal.
        ;
        MEMDEF  100,0,14            ; 32K switchable (Bank 0)
        MEMDEF  101,14,0            ; 32K switchable (Bank 1)
        MEMDEF  102,14,0            ; 32K switchable (Bank 2)
        MEMDEF  101,3,0             ; 32K switchable (Bank 3)
        MEMDEF  102,3,0             ; 32K switchable (Bank 4)
        ;
        SYSMEM  4:100000-160000     ; Define 24K as switchable
        ;                             system memory.
        ;
        MEMERR  250                                    ; Initialize memory
        ;                                                double-bit errors.
        DEVTBL  DSK1,DSK2,DSK3,DSK4,DSK5,HWK0,HWK1
        DEVTBL  AMS0,AMS1,TRM,RES,MEM,/MTM,STR          ; Define devices.
        ;
        BITMAP  HWK,606,0,1
        BITMAP  AMS,39,0,1
        BITMAP  DSK,1818,0/S        ; Put DSK bitmaps in
        BITMAP  DSK,1818,1/S        ; switchable system memory.
        BITMAP  DSK,1818,2/S
        BITMAP  DSK,1818,3/S
        BITMAP  DSK,1818,4/S
        BITMAP  DSK,1818,5/S
        ;
        QUEUE  20                   ; Add 20 more queue blocks.
        CLKFRQ 60                   ; Set clock frequency
```


(Changed 31 October 1981)

```
;
SYSTEM HWK.DVR[1,6]
SYSTEM VUE.PRG[1,4]
SYSTEM
;
SET DSKERR                          ; Enable full disk error reporting
SET GUARD                           ; for this job.
ATTACH TRM2,JOB2
JOBMEM JOB2 1:100000-177376          ; Give JOB2 32K of memory
KILL JOB2                           ; Initialize JOB2
FORCE JOB2 LOG DSK2:22,2            ; Log JOB2 into the system.
FORCE JOB2 SET DSKERR              ; Enable full disk error reporting
;
ATTACH TRM3,JOB3
JOBMEM JOB3 2:100000-177376          ; Give JOB3 32K of memory
KILL JOB3                           ; Initialize JOB3
;
ATTACH TRM4,JOB4
JOBMEM JOB4 3:100000-177376          ; Give JOB4 32K of memory
KILL JOB4                           ; Initialize JOB4
;
ATTACH NULL,SPOOL                    ; Attach line printer spooler
;                                      job to pseudo-terminal.
KILL SPOOL                          ; Initialize spooler job.
;
FORCE SPOOL                         ; Force input to spooler job.
MEMORY 4K                           ; Give spooler 4K sharable memory.
LOG 1,2                             ; Log job into system.
LPTINI PRNTR.INI                    ; Set up line printer spooler.

WAIT SPOOL                          ; Wait for SPOOL to finish before
;                                      proceeding.
MOUNT DSK1:                         ; Mount the six Phoenix
MOUNT DSK2:                         ; logical units.
MOUNT DSK3:
MOUNT DSK4:
MOUNT DSK5:
;
MEMORY 0                            ; Give JOB1 rest of sharable
                                    ; memory not used by monitor.
```

The following sections discuss the elements of the SYSTEM.INI file and their functions.

3.0 THE TRACE FUNCTION (:T)

The first line in a SYSTEM.INI is:

        :T

This turns on the trace function of the command file processor. That is, the ":T" tells the monitor to display the command file on a terminal while

(Changed 31 October 1981)

it is processing the SYSTEM.INI. When the monitor first finds the SYSTEM.INI, the process of reading and processing the file is initially under control of the first job defined (to which no terminal has yet been attached). As soon as a terminal is defined (by the first TRMDEF command), the monitor displays the remainder of the SYSTEM.INI on that terminal screen as it executes the file.

If you don't want the system to display the SYSTEM.INI as it processes it, omit the :T from the file.

NOTE: The semicolons in the example above indicate comment lines. The system does not process comment lines, but does display them as it processes the SYSTEM.INI if a :T appears in your file.

4.0   ALLOCATING JOBS (THE JOBS COMMAND)

The first command in the SYSTEM.INI (JOBS) tells the monitor what jobs to allocate in the system, and gives a name (1 to 6 characters) to each job. Each job named in the JOBS line causes one JCB (Job Control Block) area to be allocated in system memory. (Each job's JCB maintains information about that job for the system.)

The first job listed in the JOBS command line is the job that processes the system initialization command file when the system is reset. (We call this the Operator Job.)

IMPORTANT NOTE: You may define any job names you choose. However, the first character of a job name should not be a number. (On a bank switching system, a jobname that begins with a number will confuse the memory allocation command, JOBMEM, which will think the jobname refers to a memory bank rather than a job.)

If you wish to allocate more jobs than will fit on one line, you may have as many JOBS commands as you wish as long as they are before the first TRMDEF command in the file.

Each job allocated takes up about 150 words of system memory. NOTE: The first job defined on the JOBS command line and the terminal defined by the first TRMDEF statement are automatically attached; the system comes up under that job and you see the SYSTEM.INI file displayed on that terminal. However, except for this one case, the JOBS command does not automatically associate a terminal with a job; this is done using the TRMDEF command to define a terminal, and the ATTACH command to associate that terminal to a specific job. You must explicitly attach terminals to jobs in this way. In the sample SYSTEM.INI in Section 2.0, the JOBS command line looks like this:

          JOBS JOB1,JOB2,JOB3,JOB4,SPOOL

This line tells the system to allocate JCBs for five jobs.

5.0 DEFINING TERMINALS (TRMDEF)

After the JOBS command(s) must come one TRMDEF command for each terminal you want connected to your system. Every terminal has a name by which it is referenced by the monitor (1 to 6 characters), a specific hardware interface to which it is connected, and a terminal driver (a program that does any necessary character conversions). The TRMDEF command also specifies the size of the various buffers that are used in the data transfers between the terminal and the computer. The TRMDEF command takes this form:

        TRMDEF Name,Interface,Terminal,In-width,In-buffer,Out-buffer{,HOG}

When the monitor processes a TRMDEF command line, it builds a terminal definition unit in system memory which includes all of the elements above. The system then loads in the correct terminal driver and interface driver and links them to the definition unit; then it executes the interface driver which performs any necessary interface initialization.

NOTE: The buffer size values that you specify in your TRMDEF command lines affect the total size of the monitor.

After the monitor has finished processing the SYSTEM.INI, the TRMDEF command performs a different function. After the system is up and running, TRMDEF becomes a user command. At this time, TRMDEF displays on the screen the current terminal configuration of the system in a form similar to the original TRMDEF command lines in the SYSTEM.INI. The octal number that follows each terminal name is the absolute address in the monitor of the terminal definition unit for that terminal. (That information is sometimes useful when debugging the terminal service system; the general user can ignore it.)

In the sample SYSTEM.INI in Section 2.0, the first TRMDEF command line looks like this:

        TRMDEF TRM1,AM300=1,ADM3,100,100,200

Now we'll discuss the different elements of the TRMDEF command line:

5.1 Name

The terminal name consists of one to six alphanumeric characters chosen by you. Every terminal on the system must have a different name, although you may choose to use a terminal name that duplicates a job name or a program name. The system uses the terminal name to identify the terminal that you want to attach to a job or that you want to access using the TRM device driver.

5.2 Interface

The interface is the hardware board that connects the terminal to the system bus.   The interface statement gives the name of the terminal interface and its I/O port address on the system. (The I/O port address follows the name of the terminal interface, and is separated from it by an equals sign-- e.g., PS3=1.) As the system processes each TRMDEF command line, it loads the proper interface driver into system memory from area [1,6] of the System Disk.   (If the driver is already in memory because of a previous TRMDEF command line, the system does not load it in again.) Interface drivers are the programs that actually transfer data between the terminal and the terminal interface boards; these programs have the extension .IDV and must reside in account [1,6] of the System Disk. The interface drivers have the same name as the interface boards they work with.   The currently defined interface drivers available with the system are:

5.2.1 PS3

Noninterrupt driver for the Processor Technology 3P+S serial interface board. The interface statement must include the octal address of the control status port for the serial side of the board (e.g., PS3=0, PS3=20, etc.)

5.2.2 IMSIO

Noninterrupt driver for the IMSAI SIO-2 serial interface board.   The interface statement must include the octal address of the control status port for the selected side (A or B). The address of side A is 3 greater than the board I/O port address, while the address of side B is 5 greater than the I/O port address. Examples: IMSIO=3, IMSIO=5, IMSIO=23, IMSIO=25, etc.

5.2.3 AM100T

Interrupt-driven driver for the two serial ports contained on the AM-100/T CPU. The command format is identical to that for the AM310 driver, except that the I/O port must be either 0 or 1.

5.2.4 AM300

Full interrupt driver for the six-port Alpha Micro serial interface board. If you are using an AM-300 board, you may optionally include a code that selects the terminal baud rate. (We give these codes below.) The interface statement includes an I/O port address (1-6) and the optional baud rate code (separated from the I/O port address by a colon). The baud rate code is an octal number (0-17). If you omit the code, the AM-300 driver assumes a rate of 19200 baud. An AM300 interface statement takes the form:

        AM300=I/O port address{:baud rate code}

Some examples of AM-300 interface statements:

                AM300=1   (port 1 at 19200 baud)
                AM300=2   (port 2 at 19200 baud)
                AM300=3:6   (port 3 at 300 baud)
                AM300=5:12   (port 5 at 2400 baud)

The baud rate codes that you can specify in an interface statement are:

                :0   50   baud
                :1   75   baud
                :2   110   baud
                :3   134.5   baud
                :4   150   baud
                :5   200   baud
                :6   300   baud
                :7   600   baud
               :10   1200   baud
               :11   1800   baud
               :12   2400   baud
               :13   3600   baud
               :14   4800   baud
               :15   7200   baud
               :16   9600   baud
               :17   19200   baud

5.2.5 AM310

Fully interrupt-driven driver for the four-port Alpha Micro communications
controller.   If you are using the AM-310, you may optionally include a code
that selects the terminal baud rate. The interface statement includes an
I/O port address (0-3) and an optional baud rate code (separated from the
I/O port address by a colon). The baud rate code is actually a two-byte
command. The low byte is sent to Mode Register 1 (MR1) of the Programmable
Communications Interface on the AM-310 board. The high byte is sent to Mode
Register 2 (MR2).   (For more information, see the AM-310 Technical
Manual.)   If you omit the baud rate code, the AM-310 driver assumes a rate
of 19200 baud. The interface statement for the AM-310 board takes this
form:

            AM310=I/O port address{:baud rate code}

Some examples of AM-310 interface statements:

            AM310=0   (the first port at 19200 baud)
            AM310=3:37316   (the fourth port at 9600 baud)

The baud rate codes that you can specify are:

:30316	50	baud
:30716	75	baud
:31316	110	baud
:31716	134.5	baud
:32316	150	baud
:32716	300	baud
:33316	600	baud
:33716	1200	baud
:34316	1800	baud
:34716	2000	baud
:35316	2400	baud
:35716	3600	baud
:36316	4800	baud
:36716	7200	baud
:37316	9600	baud
:37716	19200	baud

5.2.6 AX310

Fully interrupt-driven driver. Allows you to run with multiple AM-310 interface boards in your system. The command format is identical to that of the AM310 driver, and the baud rate codes you may specify are the same.

5.2.7 AM120

The AM120 driver allows you to use the two serial I/O ports on the AM-120 Auxiliary I/O Controller for terminals or printers. You must also use this driver if you want to use the AM-120 parallel ports in interrupt driven mode.   For more information on the AM-120 board, see the document Software Installation Instructions for the AM-120 in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.

The interface statement portion of the TRMDEF for this interface board is identical to that of the AM-310 board, and uses the same optional baud rate codes as the AM-310 (see Section 5.2.5, above). For example:

    TRMDEF TERM1,AM120=1:37316,SOROC,100,100,100

5.2.8 PSEUDO

You normally use the PSEUDO driver with either the PSEUDO or NULL terminal specifications in a TRMDEF command line. This sets up a software interface driver that communicates with a pseudo terminal for those occasions when you have a job that doesn't need a real, hardware-controlled terminal for processing (e.g., a print spooler job). The PSEUDO interface driver is built into the monitor, and does not reside in area [1,6] of the System Disk.

(Changed 31 October 1981)

5.3 Terminal

The terminal statement tells the system what kind of terminal is connected
to the interface board, and thus what kind of terminal driver to load into
system memory from area [1,6] of the System Disk.   Different terminals
process characters differently. A terminal driver is the program that does
the necessary code conversion and character processing required by the
particular terminal that it supports. It is the terminal driver, then, that
takes care of the special functions (e.g., cursor control, Control-U,
rubout, null characters after RETURNs, etc.) that differ between terminal
types.

Terminal drivers have the extension .TDV and are sharable; that is, a given
driver is loaded only once into system memory, no matter how many terminals
of the same type are defined.   Some of the terminal drivers currently
available on the system are:

5.3.1 ADM3

The driver for the Lear Siegler ADM3 dumb terminal.   When this driver
receives a rubout, it tells the terminal to backspace and erase the
character in that position. It processes a Control-U by backspacing and
erasing the entire line. It also reverses the case condition of the RUB key
so that you do not need to shift rubouts, but you do need to shift
underscores. Most people prefer this convenience, but you can disable this
feature by setting the BIT 8 switch on your terminal to position "1."

5.3.2 SOROC

Similar to the ADM3 driver, but contains the codes for the SOROC
CRT-terminal.

5.3.3 HAZEL

Codes for the HAZELTINE 1500, 1510, and 1520 CRT-terminals.

NOTE: If you used earlier versions of the HAZEL driver, you were not able to
use the standard VUE commands because the terminal driver had to make some
special character translations to allow you to use the cursor control keys
of the 1510 and 1520 models. (For example, an end-of-line command was a
Control-G instead of the standard Control-N.) For the sake of convenience,
the current HAZEL driver allows you to use the standard VUE control
commands; however, you may not use the cursor control keys. (To move the
cursor, use the Control-H, Control-J, Control-K, and Control-L commands.)
The sources for the HAZEL driver are available on the "Driver Sources
Diskette" and on the Phoenix and Hawk System Disk packs.

5.3.4 ACTIV

Codes for the ACT-iv CRT-terminal.

5.3.5 DMEDIA

Codes for the Data Media ELITE 1520 CRT-terminal.

5.3.6 ADDS

Codes for the ADDS REGENT 100 terminal.

5.3.7 TELTYP

Driver for standard KSR and ASR Teletypes. Rubouts echo by typing the rubbed out characters between backslashes. A Control-U echoes as "^U" followed by a carriage return/line feed, which takes the printing mechanism to the next line. The line that was ended with a Control-U is ignored by the computer. This driver does no other code conversions, and no null characters are appended after a line feed.

5.3.8 SIL700

Driver for the Texas Instruments Silent-700 terminal. This driver is identical to the TELTYP driver above, except that it adds 3 null-characters at the end of every line-feed to prevent character overrun.

5.3.9 PSEUDO

Driver for the software-controlled pseudo terminals. This driver merely stops echoing of input characters and allows buffering of input to and from the controlled job. Use it only with the PSEUDO interface statement. The PSEUDO terminal driver is built into the monitor, and is not in the [1,6] area of the System Disk.

5.3.10 NULL

Driver identical to the PSEUDO driver above, except that it discards the terminal output from the job, instead of buffering it to wait for some other job to pick it up. Use this terminal driver when you want to control a job whose terminal output is of no importance (e.g., the print spooler). When using this driver for a job like the line printer spooler, you will usually use the FORCE command to send commands and data to that job; make sure that the buffer sizes you define in the pseudo-terminal's TRMDEF statement (see below for information on buffers) are large enough to accept the lines of data that you are going to FORCE to the job.

5.3.11 Other Terminal Drivers

Other terminal driver programs exist that support a variety of terminals and printers. Refer to the current AMOS System Disk, account DSK0:[1,6], to see what .TDV programs are available.

5.3.11.1 Building Your Own Terminal Driver (FIXTRM)

You may use the FIXTRM program outside of the SYSTEM.INI to build your own terminal driver for a particular terminal. See Building a New Terminal Driver (The FIXTRM Program) in the "System Operator's Information" section of the AMOS Software Update Documentation Packet for more information. (NOTE: In AMOS Release 4.5, this program was called "NEWTRM.")

5.4 In-width

The in-width statement specifies the maximum terminal line-width allowed before the system begins to discard input characters. Allowing a large width, such as 100, gives an added margin of safety when typing long lines.

5.5 In-buffer

There are times when the system cannot immediately process characters that you type from the keyboard. Instead, it stores the characters in an input buffer until it can get around to them. The in-buffer statement specifies the size of this buffer. The number that you specify, then, is also the number of characters that you can type ahead of the system before it starts to discard characters.   When you've reached the end of the type-ahead buffer, the system echoes any additional characters as bell codes and discards them. If you want to be able to type ahead a full line, make this parameter at least as large as the in-width value.

5.6 Out-buffer

The out-buffer statement specifies the size of the terminal output buffer. This is the buffer that holds the characters that the system sends to the terminal. The terminal empties this buffer at its own speed.   The system allocates two output buffers of the size specified in the out-buffer statement. The system allows a job to stay active until it fills these buffers; then the job is put into the terminal output wait state. In general, specify larger output buffers for faster terminals, and specify smaller output buffers (perhaps only 10 characters or so) for slower terminals.

For more details on these input and output buffers, see the document, Terminal Service System.   Remember: large buffers result in a larger resident monitor size.

(Changed 31 October 1981)

5.7 HOG

The last element of the TRMDEF statement is an option that you may wish to
include for terminals that run noninterrupt hardware, such as the 3P+S or
IMSAI SIO boards. The system limits the output of these boards to 60
characters per second whenever any job is demanding CPU time.

The HOG statement tells the system to use any time remaining in the job's
scheduled quantum to output characters to the terminal at maximum speed,
instead of giving that time over to another job for its task. This, of
course, reduces the total system throughput, since the time spent waiting
for the terminal interface to become ready is lost for any other tasks.

Interrupt-driven interface boards (such as the AM-300) ignore the HOG
statement; they always run at maximum speed regardless of CPU demands.

There is a bug in HOG. If the output buffer is less than 512 and you fill
the out-buffers up, printing will stop until you type a character. Because
of this problem, don't use an out-buffer of less than 513 when you include
the HOG statement in a TRMDEF command line.

6.0 BANK SWITCHING MEMORY MANAGEMENT (MEMDEF)

The sample SYSTEM.INI in Section 2.0 makes use of the bank switching memory
management option. For detailed information on memory management and the
use of the MEMDEF statement, see the Alpha Micro Integrated Systems User's
Guide, (DWM-00101-00) and the document Memory Management Option (in the
"System Operator's Information" section of the AMOS Software Update
Documentation Packet). Briefly, however-- memory management enables you to
expand the amount of memory the system as a whole can access by allowing you
to address more than 64K of memory. In one method of memory management
(called "bank switching"), each user is still limited to a maximum of 64K,
but the system can select between several different sets of memory banks
(i.e., it can bank switch memory). (Another form of memory management
that uses the Memory Partition Controller also allows you to use more than
64K of memory, but without bank switching. See the document Memory
Management with the Memory Partition Controller for more information.)

To tell the system that you are bank switching memory, use the MEMDEF
command in the SYSTEM.INI to define switchable memory banks. The example in
our SYSTEM.INI:

              MEMDEF 100,0,14   ; 32K switchable (Bank 0)
              MEMDEF 101,14,0   ; 32K switchable (Bank 1)
              MEMDEF 102,14,0   ; 32K switchable (Bank 2)
              MEMDEF 101,3,0   ; 32K switchable (Bank 3)
              MEMDEF 102,3,0   ; 32K switchable (Bank 4)

defines five switchable memory banks of 32K each, and a sharable,
non-switchable area of 32K which contains the monitor and programs that all
users can access.

(Changed 31 October 1981)

(The numbers that follow the MEMDEF statements depend on the type of memory boards you are using, and on the particular memory configuration you are setting up. For the example above, we used three Piiceon 32K-word memory boards. The arguments we supply above set up banks that use the following memory board rows: Bank 0-- rows C&D of board 100 (the monitor uses rows A&B); Bank 1-- rows C&D of board 101; Bank 2-- rows C&D of board 102; Bank 3-- rows A&B of board 101; and, Bank 4-- rows A&B of board 102. Refer to the memory management documentation mentioned above for instructions.)

If you are not bank switching memory (for example, if you do not have more than 64K of memory on your system or if your system is an MPC system), do not include the MEMDEF command in your SYSTEM.INI.

7.0 SWITCHABLE SYSTEM MEMORY (SYSMEM)

As the variety of devices that you can add to your system grows, it becomes more likely that you will have a greater number of DEVTBL entries and BITMAP commands in your SYSTEM.INI. Each new type of device that you add to the system increases the size of your monitor, because the software and control tables for that device must be incorporated into the monitor area of memory. Beginning with AMOS version 4.3, if your system uses bank switching memory management, you may now set aside an area of switchable memory for the use of the monitor. Currently, the only use you may make of this "switchable system memory" is to place bitmaps in it. (See Section 11.0, "The Disk Bitmap (BITMAP)," for information on placing bitmaps in switchable system memory on a bank switched system.) In the future, you may be able to allocate other sections of the monitor to switchable memory.

NOTE FOR MPC SYSTEMS: If your system is a Memory Partition Controller system, you may also place bitmaps in a memory area outside of the monitor, but do not use the SYSMEM command to do so. This area is called "supplemental system memory," and serves the same function as switchable system memory on a bank switching system. See the document Memory Management with the Memory Partition Controller in the "System Operator's Information" section of the AMOS Software Update Documentation Packet for information on using supplemental system memory on an MPC system.

SYSMEM tells the system what area of switchable memory you want to set aside for system use. SYSMEM takes this form:

         SYSMEM Bank#:StartAddress-EndAddress

where Bank# indicates the memory bank you want to allocate to system memory, and StartAddress and EndAddress give the beginning and ending memory addresses within that bank of the block you want to set aside. Be sure and allocate as much memory as you need for the bitmaps you want to place into switchable system memory. Place the SYSMEM command after the MEMDEF commands and before the BITMAP commands.

After the system is up and running, SYSMEM becomes a user command. SYSMEM followed by a RETURN tells you what switchable area of memory is set aside for the system.

(Changed 31 October 1981)

There are some important restrictions on the use of SYSMEM; for more information, refer to <u>Defining Switchable System Memory</u> in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.

8.0 MEMORY ERROR DETECTION (MEMERR)

The MEMERR command enables double-bit error detection for a Piiceon 32K-word memory board. It also initializes the Alpha Micro AM-710 128K byte memory board.

8.1 Piiceon 32K Word Memory Boards

MEMERR initializes the Piiceon board by instructing it to abort when a double-bit memory error occurs. Make sure that the Piiceon memory board is properly jumpered for the I/O error port you specify in the MEMERR command line. The error interrupt-enable jumper (jumper 54) must be installed on the memory board.

<u>NOTE:</u> MEMERR was designed to be used with the AM-100/T CPU. If you are using the 32K-word Piiceon memory boards as 64K-byte memory boards (that is, if you are using the AM-100 CPU), you can still use MEMERR if you use an unused error interrupt-enable jumper other than jumper 54 and enable the new line on the AM-100 board. (Jumper 54 is an AM-100/T interrupt line.)

You may assign the same I/O port number to more than one memory board because the system only issues write-status commands to the memory boards.

If you don't use MEMERR, the Piiceon 32K-word memory board automatically corrects single-bit errors, but ignores double-bit errors; if you use MEMERR, the memory board still corrects single-bit errors, but causes the system to halt on a double-bit error.

If the system halts, look at the error light on the Piiceon memory board (a red LED). If the light is on, the system halt occurred because of a double-bit memory error.

Put MEMERR after the TRMDEF commands in your SYSTEM.INI. Include the number of the I/O error port you have assigned to the memory board(s). (This I/O port is usually 250, octal or A8, hex.) For example:

          MEMERR 250

If double-bit errors are frequent on your system, you may want to replace the memory board on which the errors occur.

8.2 AM-710 128K Byte Memory Board

The AM-710 memory board also requires the use of MEMERR.   If your system
contains AM-710 and Piiceon 32K word memory boards, include on the MEMERR
command line the I/O error port assigned to the Piiceon boards (as in the
example above). If your system contains only AM-710 memory boards, you must
not supply any argument to the MEMERR command. For example:

              MEMERR   ; No Piiceon boards, only AM-710 boards.

Note that the AM-710 board also requires the use of the PARITY command
(discussed below).

9.0 PARITY ERROR DETECTION (PARITY)

The AM-710 memory board requires the presence of the PARITY command in your
SYSTEM.INI in order to enable parity error detection and reporting. Place
the PARITY command after the MEMERR command (also required).   PARITY takes
this form:

              PARITY I/O-port{,I/O-port2,...I/O-portN}

where I/O-port identifies the I/O port address of the one or more AM-710
boards in your system. (One I/O port address must appear for each address
selected by the AM-710 boards.) NOTE: Typically, all AM-710 memory boards
are addressed to the same address (100, octal; or 40, hex)-- therefore,
usually only one I/O port address will appear on the PARITY command line.
For information on parity error handling, see the document Software
Installation Instructions for the AM-710 Memory Board in the "System
Operator's Information" section of the AMOS Software Update Documentation
Packet.

10.0 THE DEVICE TABLE (DEVTBL)

Following the TRMDEF command lines is the DEVTBL command. It defines the
devices that your system can access. You must always have at least one
DEVTBL command line in the SYSTEM.INI file. (NOTE: The system already knows
that the System Disk, DSK0:, is present, so don't put DSK0: in the DEVTBL
command line. If your system has no devices other than DSK0:, enter DEVTBL
alone on a line, and do not follow it with any arguments.)

For each DEVTBL line, list all sharable devices before a slash; all private,
non-sharable devices after it. (A sharable device is one that all users can
access, such as a disk drive; a non-sharable device is one that only one
user at a time can access, such as a magnetic tape unit.)

If your system has more devices than will fit on one DEVTBL command line,
you can have as many DEVTBL command lines as you want, as long as they are
not separated by intervening commands.

As the system processes the DEVTBL command line, it builds a device table in
system memory. The file system consults the device table for device
assignments.   (If   your   devices   run   under   the control of the AM-410 or
AM-420, DEVTBL also builds a table in memory that keeps track of any
alternate tracks assigned for those devices.)

Here are some sample device names that the system recognizes:

        STD1   Floppy disk drive that uses IBM-standard format.
        DDA1   Floppy disk drive that uses double-density, double-sided AMS
                    format.
        DSK0   Logical unit zero of the System Device.
        HWK1   Logical unit 1 of a hard disk drive that runs under control
                    of the AM-500 controller.
        SMD5   Logical unit 5 of a hard disk drive that runs under control
                    of the AM-410 controller.
        MTU0   Magnetic tape unit that runs under control of the AM-600
                    Magnetic Tape Transport Formatter Interface.
        STR0   The AM-620 1/4" Streaming Tape Drive.
        VCR0   Video cassette recorder that runs under control of the AM-610
                    Video Cassette Recorder Interface.
        TRM   The generalized terminal service driver. Allows input and
                    output to any terminal connected to the system.
        RES   Driver that allows you to use system memory as a device:
                    e.g., .DIR RES:LOG.PRG.
        MEM   Driver that allows you to use user memory partition as a
                    device: e.g.,.COPY MEM: =WRKFIL.PRG.

NOTE: Several commands on the system require that you have both MEM and RES
defined in your DEVTBL as system devices.

After the system is fully up and running, the DEVTBL command becomes a user
command that tells you what devices are in the device table in system
memory; it also tells you which devices are sharable among users.

11.0 THE DISK BITMAP (BITMAP)

To   randomly   access   information on a disk, the AMOS file structure needs a
disk allocation map (a bitmap). The BITMAP command sets up these bitmaps.
If the disks on your system run under different disk controllers, each type
of disk must have its own device name and separate bitmap areas.   (NOTE:
Floppy disk drives, which may use disks in several different formats, must
have a different device defined for each type of format even though the
drives may run under the same controller. Also, several types of Winchester
drives can run under the control of the same AM-420 Hard Disk Controller--
however, because these devices are different sizes, each type must have its
own name and bitmap area.)

The BITMAP command specifies the device name, the number (in decimal) of
words that the bitmap buffer needs, and the list of drive numbers that are
to share this particular bitmap area. The Hawk hard disk drive requires 606
words per logical device; the Phoenix hard disk drive requires 1818 words

per surface. (For information on the number of words needed for the bitmap
of a particular floppy disk device, see Configuring Floppy Disk Drivers in
the "System Operator's Information" section of the AMOS Software Update
Documentation Packet.) You may have as many BITMAP commands as you want; as
many as one for each disk drive. Each BITMAP command creates a separate
bitmap area.

In the sample SYSTEM.INI in Section 2.0, one BITMAP command line looks like
this:
          BITMAP HWK,606,0,1

We could replace it with two BITMAP command lines to create two separate
bitmap areas for drives 0 and 1:

          BITMAP HWK,606,0
          BITMAP HWK,606,1

The monitor builds one sharable bitmap area in memory for each BITMAP
command line it encounters. The BITMAP command specifies the disks to be
accessed by SYSTAT when SYSTAT prints the number of free blocks left on the
devices on the system.

11.1 Switchable and Supplemental System Memory Option (/S)

Beginning with AMOS version 4.3, on a bank switching system you may use the
SYSMEM command to set aside part of switchable memory for use by the
monitor, thus reducing monitor size. (See Section 7.0, "Switchable System
Memory (SYSMEM)," for more information.) On a Memory Partition Controller
system, this area of memory is called "supplemental system memory" rather
than "switchable system memory," and you do not use the SYSMEM command to
define it. Instead, the monitor automatically defines the supplemental
system memory when it encounters the BITMAP command /S option.

For the present, bitmaps are the only monitor elements that can be placed in
switchable or supplemental system memory. Designate a specific bitmap as
one which is to be placed in switchable or supplemental system memory by
ending the BITMAP command line with a /S. For example:

          BITMAP DSK,1818,0,1,2,3,4,5/S

Make sure you have enough room in the area of additional system memory you
have allocated for the bitmaps you want to place there.

NOTE: If you place a device's bitmap in switchable or supplemental system
memory, and that device is not the System Device, you must remember to place
the driver for that device in system memory via the SYSTEM command;
otherwise, attempts to access that device will fail. (See Section 15.0 for
more information on the SYSTEM command.)

12.0 THE MONITOR QUEUE (QUEUE)

The monitor has a general purpose queue system that several commands use, and which is also available to user programs. The queue contains a fixed number of eight-word blocks which are assigned and then returned during the course of processing. The number of queue blocks that you need depends upon the size of your monitor, and the tasks that it performs.

The monitor initially contains 20 blocks; you may add more by using the QUEUE command in the SYSTEM.INI. Place the QUEUE command before any SYSTEM commands. The QUEUE command allocates additional queue blocks. So, the example in the sample SYSTEM.INI file, QUEUE 20, adds 20 blocks to the basic queue size of 20 to give a total queue size of 40 blocks.

For information on how your assembly language programs can access the monitor queue, refer to AMOS Monitor Calls Manual, (DWM-00100-42).

13.0 THE CLOCK FREQUENCY (CLKFRQ)

The AM-100 CPU board contains a real-time clock that several programs (e.g., the AlphaBASIC compiler and DYSTAT) refer to when they calculate time intervals; the system also uses this clock to perform job scheduling and timekeeping functions. The CPU board contains an external input for the line clock frequency that is connected to an AC line of approximately 10 volts.

The system has to know what frequency is being applied to the clock input line on the CPU board so that the programs that refer to the clock can know how to convert the clock tick count into actual time in seconds. This frequency is usually 60 Hz (the standard line frequency in the United States). Since many systems are shipped overseas, however, where the standard line frequency is 50 Hz, you must specify which frequency you are using.

The CLKFRQ command specifies the frequency (in Hz) that is being applied to the external clock input. The system stores this value so that programs that need to convert clock ticks to real time will be able to find out what frequency the clock is operating at. The CLKFRQ has nothing to do with the actual frequency at which the computer runs, and changing the CLKFRQ value does not affect the speed of the system. The CLKFRQ command in the sample SYSTEM.INI in Section 2.0 was:

        CLKFRQ 60

If you do not include the CLKFRQ command in the SYSTEM.INI, the system stores a zero in the monitor location reserved for the clock frequency. In this case, programs trying to convert clock ticks into actual time aren't able to do so.

(Changed 31 October 1981)

You may place the CLKFRQ command anywhere in the SYSTEM.INI. (NOTE: If you
are going to reset the system date and time from the AM-120 board
clock/calendar, you must place the CLKFRQ command before the DATE and TIME
commands in your SYSTEM.INI. See the next section.)

14.0   RESETTING THE SYSTEM DATE AND TIME FROM THE AM-120 (DATE AND TIME)

The AM-120 Auxiliary I/O Controller contains as one of its features a
clock/calendar with battery backup. If your system contains an AM-120, you
may include the DATE and TIME commands in your SYSTEM.INI to reset the
system date and time from the AM-120 clock/calendar. If you do so, you must
place the DATE and TIME commands before the final SYSTEM command in the
SYSTEM.INI. NOTE: Placing DATE and TIME within the SYSTEM.INI before the
final SYSTEM command if you do not have an AM-120 in the system could
cause problems when your system tries to boot under that SYSTEM.INI file.

You must also place the CLKFRQ command before the DATE and TIME commands.
For example:

        CLKFRQ 60
        DATE
        TIME
        SYSTEM HWK.DVR
        SYSTEM

For information on DATE and TIME, see the DATE and TIME reference sheets in
the AMOS System Commands Reference Manual, (DWM-00100-49). For
information on the AM-120, see Software Installation Instructions for the
AM-120 in the "System Operator's Information" section of the AMOS Software
Update Documentation Packet.

15.0   INCORPORATING PROGRAMS IN SYSTEM MEMORY (SYSTEM)

You may incorporate programs into the system monitor by using the SYSTEM
command in the SYSTEM.INI file. These programs may be system programs or
your own programs. When the system reads the SYSTEM.INI, it loads into
system memory the programs you've specified in the SYSTEM commands. These
programs actually become part of the monitor, and so dynamically increase
its size as they are loaded into memory.

CAUTION-- the programs to be included in the monitor must be re-entrant
(that is, sharable by more than one user). If they are not re-entrant,
there is a possibility of system failure when two users attempt to access
the same program. Many of the AMOS programs are re-entrant. Check with the
AMOS System Commands Reference Manual, (DWM-00100-49), to see if a
particular command program is re-entrant.

The most common use of the SYSTEM command in the SYSTEM.INI is to include
the AlphaBasic runtime package (RUN.PRG) in the monitor so that each user
does not need to load RUN into his own memory partition. You may also
include the interactive compiler (BASIC.PRG) itself if you expect heavy

development work by more than one user. If users on your system will be making extensive use of the screen-oriented text editor, VUE, you may want to load it into sharable memory via SYSTEM.

You might also want to use the SYSTEM command to include any realtime routines in the monitor that must be in memory at all times so that they can process asynchronous events (such as data collection from interrupting devices). For example, the DYSTAT program runs a continuously changing system display on a video monitor, and must be in memory at all times once it has begun execution so that it can update the display.

Another reason to use the SYSTEM command is if you want to include frequently-called user subroutines in the monitor. (You can locate such subroutines by name via the SRCH and FETCH monitor calls from assembly language programs.) Again, if these programs are to be shared by several users, they MUST be re-entrant.

NOTE: Before Release 4.5, you were required to place the generalized terminal driver, TRM.DVR[1,6], in system memory. In most cases this is no longer necessary. However, see Section 15.1 for a discussion of when it is necessary to include driver programs in system memory.

To include programs in the system monitor, use one SYSTEM command for each program. Follow each SYSTEM command with the file specification of the program you want to include. Example:

        SYSTEM VUE.PRG
        SYSTEM RUN.PRG
        SYSTEM AMS.DVR[1,6]

If your file specification does not include an extension, the monitor assumes a .PRG extension; if you don't supply an account specification, the monitor assumes account [1,4]. You must place any SYSTEM commands after all other commands in the SYSTEM.INI that expand the monitor size.

The SYSTEM command has another use beside the inclusion of programs in the system monitor. A SYSTEM command alone on a line in a SYSTEM.INI (that is, not followed by a file specification), tells the system that monitor expansion is finished. The system then flags the monitor as up and running. The system also sets a flag in the system communication area that indicates that the system initialization is done except for final cleanup. Various commands (including SYSTEM) test this flag to see which mode they may operate in. For example, before the system is up and running, the JOBS command allocates new jobs on the system; after system initialization, the JOBS command displays the jobname of the user that typed the JOBS command.

Whether or not you include any programs in the monitor, your SYSTEM.INI must have a SYSTEM command without a file specification to tell the operating system that the system has been initialized. This SYSTEM command must be after any other commands that expand the monitor size (and that includes any other SYSTEM commands that are followed by a file specification).

(Changed 31 October 1981)

After the system is up and running, the SYSTEM command performs a new function as a user command. After system initialization, the SYSTEM command tells you what programs are in system memory (that is, what programs are a part of the monitor), and the total size (in decimal words) of the monitor. For more information on the programs in system memory, you can use the MAP command. (For information on MAP, refer to the AMOS User's Guide, (DWM-00100-35).)

15.1 Including Device Drivers in System Memory

There are certain situations in which you must include device drivers in system memory:

1. If users on the system will be accessing a device that is not the System Device, and the bitmap for that device is in switchable or supplemental system memory, you must load the driver for that device into system memory or the device access will fail. (See Section 11.0, "The Disk Bitmap (BITMAP)," for information on placing bitmaps in switchable or supplemental system memory.)

2. If users on the system will be accessing a device that is not the System Device from within AlphaBASIC or AlphaVUE, you must preload the driver for that device into system or user memory before invoking BASIC or VUE or the device access will fail.

   The reason for this is that to perform their special functions, BASIC and VUE have to handle their own memory setup-- however, because of this, once BASIC or VUE is invoked, there is no room in memory to load a device driver.

   For example: users often access the generalized terminal driver, TRM.DVR, within BASIC (e.g., OPEN #100,"TRM:TRM6",OUTPUT) or VUE (e.g., > Unyank TRM:TRM6); if this is so on your system, remember to keep TRM.DVR in system memory or preload it into user memory before invoking BASIC or VUE.

16.0 SYSTEM INITIALIZATON CLEANUP

After the monitor processes the SYSTEM commands in the SYSTEM.INI, the system is technically up and running. There are a couple of things still left to do, however, before the initialization procedure is complete. You may now include any commands in the SYSTEM.INI that you want the monitor to perform automatically at the time of system start-up. These commands are all commands that you can enter from the keyboard for yourself, but it's sometimes convenient to have the monitor perform them automatically every time you power up or reset the system. For example, you can have the monitor mount the disks that you are going to use, connect terminals to jobs, etc. You can also use the FORCE command to force input to a particular job. (You can use this feature to login a user, run a business

program, etc., all without direct operator intervention.)   We discuss some
of these commands below.

After you've included the functions you want performed automatically, there
is one last thing to do before system initialization is complete--
deallocate the temporary 8K user memory partition in which the SYSTEM.INI
was processed. For bank switching systems or systems that do not use memory
management, use the MEMORY 0 command at the end of the SYSTEM.INI. For
Memory Partition Controller systems (MPC system), use a JOBSIZ MAX command
at the end of the SYSTEM.INI. Remember that the MEMORY 0 or JOBSIZ MAX
command must be the last command in the SYSTEM.INI.

Now we'll discuss some of the commands you may want the system to perform
automatically for you at system start-up. Remember that unlike some of the
commands discussed in earlier sections, all of these commands are legal user
commands, and may be used outside of the SYSTEM.INI.

16.1 Setting Options (SET)

The SET command can perform a variety of system functions. For example, the
SET BPI command sets the bits-per-inch data density value used by the
magnetic tape transport driver. The examples in our sample SYSTEM.INI:

          SET DSKERR
          SET GUARD

tell the system to report any soft disk errors that occur (SET DSKERR) and
to guard the terminal of the job processing the SYSTEM.INI from any messages
sent by other terminals (SET GUARD).

Note that the SET command only affects the job that used it.   For example,
the SET DSKERR command above only affects the job the system comes up under.
(For information on forcing commands to other jobs, see Section 16.5,
"Forcing Input to a Job (FORCE).")

See the SET reference sheet in the AMOS System Commands Reference Manual,
(DWM-00100-49), for more information on SET.

16.2 Attaching Jobs (ATTACH)

When the system first begins to process the SYSTEM.INI, it automatically
attaches the first job listed in the JOBS command and the terminal defined
by the first TRMDEF command. Except for this special case, however, the
system does not automatically link jobs with terminals.   (When a job is
linked to a terminal, the job and terminal are "attached." When a job is
not linked to a terminal, the job is "detached.") A detached job must have
a terminal attached to it before it can do terminal input or output. A
detached terminal, on the other hand, can be accessed through terminal
service calls or the general TRM driver. You cannot attach a job to a
detached terminal from that terminal itself; you must do it from another
terminal.

To attach jobs and terminals, you must use the ATTACH command. Once a job is attached to a terminal, it uses that terminal for input and output.

You can use the ATTACH command in several different ways:

ATTACH Terminal,Job

This command attaches the terminal and job named. If the terminal or job are already attached to other units, the ATTACH command detaches them before it attaches them to each other.

ATTACH Job

This ATTACH command attaches the user's own terminal to the job named (and detaches it from the current job).

ATTACH

This use of the ATTACH command lists the terminals that are currently attached, and the jobs to which they are attached. For more information on ATTACH, see the ATTACH reference sheet in the AMOS System Commands Reference Manual, (DWM-00100-49).

16.3 Allocating Memory

The commands that you use to allocate memory depend on the type of system you are using. If your system does not use memory management (that is, if you only have 64K of memory on your system), you will use the MEMORY command to allocate memory. If your system performs memory management via bank switching, you will use the JOBMEM command. If your system performs memory management by use of the Memory Partition Controller, you will use the JOBSIZ command.

16.3.1 Bank Switching System (JOBMEM)

Systems that bank switch memory require that you use the JOBMEM command to allocate memory to a user memory partition. Do not use JOBMEM to allocate memory to the job the system is coming up under. (For complete information on JOBMEM, refer to the JOBMEM reference sheet in the AMOS System Commands Reference Manual, (DWM-00100-49), and the document Memory Management Option in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.)

The JOBMEM command in the SYSTEM.INI takes the form:

JOBMEM Jobname Bank-#:StartAddress-EndAddress

where Jobname specifies the job you are allocating memory to. Bank-#: selects the memory bank the user partition will reside in, and StartAddress and EndAddress select the beginning and ending addresses of the memory block in that bank you want to allocate.

After the system is up and running, JOBMEM becomes a user command. You can use JOBMEM to tell you what areas of memory are allocated to your own job and to other jobs on the system.

There are some important restrictions on the use of JOBMEM; refer to the JOBMEM documentation mentioned above for complete instructions on its use. Also see the Alpha Micro Integrated Systems User's Guide, (DWM-00101-00), for information on bank switching.

NOTE ON USING THE MEMORY COMMAND ON A BANK SWITCHING SYSTEM: There are several special occasions when you use the MEMORY command on a bank switching system. Use the MEMORY command for all jobs whose memory is located in Bank Zero or sharable memory. For example, when we talk about setting up the line printer spooler, you will see that we use a MEMORY 4K command to allocate memory to the job running the spooler. That's because our sample SYSTEM.INI is setting up the spooler to run in sharable memory.

Also, the last command in the SYSTEM.INI for a bank switching system is MEMORY 0. The major purpose of this command is to deallocate the temporary 8K user partition that was automatically allocated at the top of Bank Zero for processing the SYSTEM.INI. The MEMORY 0 command also restores all sharable memory not used by the monitor within the first 64K of memory to the job the system is coming up under. For example, if your MEMDEF commands have set up a sharable area of 32K and switchable banks of 32K, if the monitor actually only takes up 29K, MEMORY 0 adds the remaining 3K (32K-29K) in the sharable portion to the job that has the first block of memory in Bank Zero.

16.3.2 Non-Memory Management System (MEMORY)

NOTE: If your system does not bank switch memory or is not an MPC system (that is, if your system does not use memory management), do not use the JOBMEM or JOBSIZ commands. If you are going to allocate memory within the SYSTEM.INI, use the FORCE and MEMORY commands. For example:

        FORCE JOB2 MEMORY 32K

16.3.3 Memory Partition Controller System (JOBSIZ)

On a system that uses the Memory Partition Controller to perform memory management, do not use the JOBMEM or MEMORY commands to allocate memory; instead, use the JOBSIZ command. For information on allocating memory on an MPC system, see the document Memory Management with the Memory Partition Controller in the "System Operator's Information" section of the AMOS Software Update Documentation Packet. Also, see the JOBSIZ reference sheet in the AMOS System Commands Reference Manual, (DWM-00100-49).

(Changed 31 October 1981)

16.4 Killing Jobs (KILL)

To properly initialize the jobs on the system, you should KILL the jobs that you have defined. A KILL command sends a Control-C to the specified job; this puts the job at the monitor level, ready to receive and send data. The KILL commands must appear after any ATTACH commands, but before any FORCE commands are used to send commands or data to the job.

Do not kill the job that the system is coming up under (i.e., the first job in the JOBS command line).

Use one KILL command for each job on the system except for the job the system is coming up under:

        KILL JOB2
        KILL JOB4

16.5 Forcing Input to a Job (FORCE)

The FORCE command gives you a way of sending input to another job. To send one line of input to another job, use the FORCE command followed by the jobname and the input. For example, our sample SYSTEM.INI contains this line:

        FORCE JOB2 LOG DSK2:2,2

The line above logs JOB2 into the system under account DSK2:[2,2]. You can also send several lines of input to a job by typing a carriage return after the jobname. Example:

        FORCE JOB2

After that point, all lines of text that follow (up to a blank line) will be sent to the specified job. (A blank line is a carriage return alone on a line.) Example:

        FORCE JOB2
        LOG DSK2:2,2
        ORDER.PRG

16.6 Mounting Disks (MOUNT)

The system automatically mounts the System Disk (DSK0) for you at the time of system start-up.   If you wish it to mount other disks as well, include one MOUNT command for each disk drive you want to mount. Example:

        MOUNT DSK1:
        MOUNT STD0:

(Changed 31 October 1981)

You MUST mount a disk before you can read or write data to it.   REMEMBER:
Never mount a disk while another user is accessing it.

16.7 Setting Head Load Time (HEDLOD)

The HEDLOD command sets the head load time for a Persci floppy disk that
runs under control of the AM-200 or AM-210 floppy disk controller boards.

The number that follows the HEDLOD command selects the number of real-time
clock ticks that the AM-200 board must keep the disk drive heads loaded
after a data transfer. The example in the SYSTEM.INI file in Section 2.0
is:

        HEDLOD 1800

This tells the floppy disk controller to keep the heads loaded for 1800
clock ticks (30 seconds when the real-time clock is operating at 60 Hz)
after any data transfer.   The HEDLOD command does not affect disk drives
that do not allow software control of head load timing. You may omit this
line if you are not running with a Persci floppy disk.

16.8 DYSTAT

DYSTAT is a system program that reports on the dynamic status of the
operating system by way of a continuous display on a video monitor.   DYSTAT
requires that your system contain a memory-mapped video board. Omit the
DYSTAT command if you do not want to run this display.

NOTE: DYSTAT runs asynchronously without regard to user job.   If you are
going to use DYSTAT, you must include it in the monitor (by using a SYSTEM
command in the SYSTEM.INI) so that it is in system memory at all times.   If
you try to execute the DYSTAT command and it is not in system memory, the
system fails the first time the job scheduler calls on DYSTAT for an update
because it sees that DYSTAT is no longer in memory. (Remember that the
system executes system commands in the partition of the user that requested
the command; after it is finished, it usually deletes the program from the
partition.)

The DYSTAT program also requires that the TODCNV routine is in system memory
so that it can display the current time in the header line of the DYSTAT
display.   Load TODCNV into system memory by using a SYSTEM command in the
SYSTEM.INI, just as you loaded DYSTAT.

For more information on DYSTAT, see the DYSTAT reference sheet in AMOS
System Commands Reference Manual, (DWM-00100-49).

16.9 Setting Up the Line Printer Spooler

You can set up the line printer spooler from AMOS command level (and doing so is a good technique for debugging the spooler if you have problems bringing it up). However, usually you will want to bring up the spooler at the time of system start-up by placing the proper instructions in the SYSTEM.INI.

The sample SYSTEM.INI in Section 2.0 shows the following 6 lines:

        FORCE SPOOL
        MEMORY 4K
        LOG 1,2
        LPTINI PRINTR.INI

        WAIT SPOOL

These six lines set up the system line printer spooler in sharable memory. (A spooler is a program that sets up a queue (or waiting line) for a particular program. When a line printer spooler is in control, requests for use of the printer are placed into a queue. As the printer becomes available, the spooler looks at the request at the top of the list, finds the file, and sends it to the printer. The spooler then removes that request from the queue. The printer prints files in the order of their requests in the queue.)

Setting up a line printer spooler is a good example of the kinds of things you can ask the SYSTEM.INI file to do at the time of system start up. For a detailed explanation of how to set up the line printer spooler, see the document Setting Up the Line Printer Spooler in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.

Index

SETTING UP THE LINE PRINTER SPOOLER

October 1981
Revision A03

This document reflects AMOS versions 4.6 and later

Table of Contents

1.0 INTRODUCTION

The line printer spooler takes care of handling requests to use the printer
by "spooling" a printer request into a queue. (As the printer becomes
available, it prints the file specified by the next request in the queue.)
This allows your job to print files at the same time as other jobs on the
system without tying up your job until the printer is free. You simply enter
your printer request (using the PRINT command) and then go on to other
things.

The spooler program has a number of special features. For example, it can:
print multiple file copies, print a banner page that identifies the listing,
delete a file after it is printed, recognize wildcard file specifications,
remember what type of form should be mounted on ˆ printer, identify a
printer by name, etc. In addition, the spooler (LPTSPL.PRG) is re-entrant,
which means that you may load it into system memory.

(NOTE: However, loading LPTSPL.PRG into system memory is not usually a good
idea.   In setting up a line printer spooler, you do not run LPTSPL.PRG
directly; instead, you run the LPTINI program which calls LPTSPL.   Since
LPTINI must be run in each memory area in which you are setting up a
spooler, and since it loads LPTSPL.PRG into memory over itself, placing
LPTSPL.PRG in system memory does not save you any memory space in the
individual area in which you are setting up the spooler and just increases
the size of your monitor area.)

You may either set up the spooler in system memory or (if your system bank
switches memory) you may place it in a bank-switched memory partition (thus
reducing the size of the monitor in system memory). You may also install
more than one line printer spooler (which you may want to do if you have
more than one printer).   Installing multiple printer spoolers gives you
multiple printer queues; one for each printer that is handled by a spooler.

For more information on the features and operation of the line printer
spooler, see the documentation on the PRINT command in the AMOS System
Commands Reference Manual, (DWM-00100-49) and the section titled "Printing
Files (PRINT)," in the AMOS User's Guide, (DWM-00100-35).

The purpose of this document is to help you set up the line printer spooler
on your own system. To that end we have included on your System Disk two
text files: SYSLPT.INI and PRINTR.INI.   SYSLPT.INI is a sample system
initialization command file much like the regular SYSTEM.INI file with which
you are already familiar, except that it includes those command lines
necessary for bringing up the spooler program. (See Section 2.0, below, for
a discussion of changing the SYSTEM.INI.) The PRINTR.INI is a sample spooler
parameter file that contains information necessary for customizing the
spooler program for your own use.

You may install multiple spoolers by setting up multiple spooler jobs. Each job can have a different spooler parameter file; this enables you to spool to several different printers on the system at the same time.

Before attempting to set up the line printer spooler, read this entire document carefully. Then modify the PRINTR.INI file that we have supplied to reflect your own needs or create your own spooler parameter file. At that point you can either use the SYSLPT.INI file to boot the system up with the new line printer spooler, or you can change your existing SYSTEM.INI file as outlined below. If you do not want the line printer spooler to be set up automatically when the system comes up under the control of the SYSTEM.INI or the SYSLPT.INI, you may set up the spooler yourself from the keyboard by following the instructions in Section 6.0, "Troubleshooting the Spooler."

2.0 MODIFYING THE SYSTEM.INI TO SET UP THE SPOOLER

The discussions below assume that you are familiar with the document titled The System Initialization Command File which appears in the "System Operator's Information" section of the AMOS Software Update documentation packet; that document contains information on the JOBS, TRMDEF, DEVTBL, ATTACH, QUEUE, MEMORY, JOBSIZ, JOBMEM and FORCE commands, and discusses how to modify a SYSTEM.INI.

Take a look at the sample system initialization command file we provide (SYSLPT.INI):

```
:T
MPCOFF
JOBS JOB1,SPOOL
TRMDEF TERM1,AM300=1,SOROC,100,100,80
TRMDEF TERM6,AM300=6:16,TELTYP,100,100,100
TRMDEF DUMMY,PSEUDO,NULL,30,30,2
DEVTBL DSK1,TRM,MEM,RES
BITMAP DSK,606,0,1
QUEUE 15
SYSTEM
CLKFRQ 60
;
ATTACH DUMMY,SPOOL
KILL SPOOL
FORCE SPOOL
MEMORY 4250
LOG 1,2
LPTINI PRINTR.INI

WAIT SPOOL
;
MOUNT DSK1:
MEMORY 0
```

(Changed 31 October 1981)

Note that this system initialization command file brings your system up as a non-memory management system. (The lack of MEMDEF statements disables bank switching, and the presence of the MPCOFF statement turns off any AM-700 Memory Partition Controller that might be in your system.) The following sections discuss this initialization command file, and we go through each step of installing a line printer spooler.

2.1 Setting Up the Spooler Job

The first step in setting up the line printer spooler is to define the job in which it is to run. The SYSLPT.INI file above creates job SPOOL in which to run the spooler program by including SPOOL in the job definition command (JOBS). You may use any unused job to control the spooler program.

2.2 Defining a Pseudo-terminal

So that we do not tie up a terminal when entering the commands to SPOOL that will set up the line printer spooler, we define a pseudo-terminal with which to run the job. A pseudo-terminal requires no actual hardware I/O device. The terminal definition (TRMDEF) command:

        TRMDEF DUMMY,PSEUDO,NULL,30,30,2

defines the pseudo-terminal named DUMMY. Be careful that the buffer sizes you specify (in this case, 30,30,2) are large enough to contain the commands that you are going to be forcing to the job via the FORCE command. (See the document The System Initialization Command File for information on the PSEUDO interface driver, the NULL terminal driver, and the TRMDEF command.)

2.3 Allocating Additional Queue Blocks

The line printer spooler uses the system queue blocks to store requests for printing. When the system is initially brought up, 20 of these queue blocks are allocated by the system. Two of the blocks are used for every printer defined on the system, and three for every file request. Because queue blocks are required by other portions of the operating system, the spooler will not allow the number of free queue blocks to go below six. Therefore, with the standard allocation of 20 queue blocks, you may queue a total of four printer requests at any one time. If you want to be able to queue a larger number of requests, use the QUEUE command in the system initialization command file to allocate additional blocks. In the example above, we have used the QUEUE command to allocate an additional 15 blocks for a total of 35.

NOTE: If you are setting up more than one line printer spooler, make sure to allocate enough extra queue blocks. For example, if you are setting up two line printer spoolers, you will need twice as many queue blocks as if you were setting up one spooler.

2.4 Attaching a Terminal to the Spooler Job

No jobs and terminals are automatically assigned to one another except for
the job under whose control the system comes up (the first job on the JOBS
line) and the terminal defined by the first TRMDEF command. We must
explicitly attach any other jobs and terminals by using the ATTACH command.
So, we attach the spooler job and the pseudo-terminal:

        ATTACH DUMMY,SPOOL

2.5 "Killing" the Spooler Job

To properly initialize a job, you must always use the KILL command on that
job after attaching a terminal to it but before forcing any commands to it:

        KILL SPOOL

2.6 Using the FORCE command to Set Up the Line Printer Spooler

Once the spooler job has been allocated, you have attached a terminal or a
pseudo-terminal to the job, and you have used the KILL command on the job,
you can start the spooler program itself. All commands issued to the
spooler job during the spooler initialization go through the FORCE command.

Now that we have attached the pseudo-terminal DUMMY to the job SPOOL, we can
send the commands to the job that will get the line printer spooler program
up and running. FORCE the following sequence to the job:

        FORCE SPOOL
        MEMORY 4250
        LOG 1,2
        LPTINI PRINTR.INI

The following sections discuss each element of this FORCE sequence:

2.6.1 "FORCE SPOOL"

The first line of this sequence tells the system that we are starting a set
of commands that we want to send to job SPOOL. (Because a carriage return
appears at the end of the job name, AMOS knows that more than one line of
input follows; a blank line signals the end of the group of commands and
data that we are forcing to the job.)

2.6.2 "MEMORY 4250"

The first thing to do is to allocate the spooler job some memory so that it
will be able to bring up the line printer spooler program. We have
determined that, on the average, 4K of memory is the usual amount of memory
required to bring up the spooler program. This amount may vary depending on
the particular device driver or terminal driver you are using to run the
printer and on the sizes of the drivers for the devices that contain the
files you want to print. For example, on one typical system, on which the
spooler uses TRM.DVR to access the printer, if TRM.DVR is not loaded into
system memory, the spooler actually uses 4250 bytes. (See Section 6.0,
"Troubleshooting the Spooler Program," for instructions on determining the
exact amount of memory required by your particular use of the line printer
spooler program.)

IMPORTANT NOTE: It is very important to make sure that you have allocated
enough memory to the spooler job. Note that if the spooler is asked to
print a file, the driver for the device containing the file must either
already be loaded into system memory (via the SYSTEM command) or there must
be enough room left in the spooler job's memory partition that the spooler
can load the driver into memory. If there is not enough memory to load the
driver into the spooler's partition, you will not see an error message, but
the file will not be printed.

On a bank-switched system, if you want to place the spooler into a
bank-switched partition rather than sharable memory, use the JOBMEM command
(not the MEMORY command) before the FORCE command. On a system that uses
the AM-700 Memory Partition Controller, use a JOBSIZ command to allocate
memory to the spooler. For example, for a bank switching system:

```
        ; Bank switching system-- put spooler into Bank Three.
        ATTACH NULL,SPOOL
        JOBMEM SPOOL 3:111530-123030
        KILL SPOOL
        FORCE SPOOL
        LOG OPR:
        LPTINI PRINTR

        WAIT SPOOL
```

Or, for a Memory Partition Controller system:

```
        ; Memory Partition Controller system.
        ATTACH NULL,SPOOL
        JOBSIZ SPOOL 4250
        KILL SPOOL
        FORCE SPOOL
        LOG OPR:
        LPTINI PRINTR

        WAIT SPOOL
```

(Changed 31 October 1981)

For information on the JOBMEM and JOBSIZ commands, see The System
Initialization Command File in the "System Operator's Information" section
of the AMOS Software Update Documentation Packet.

2.6.3 "LOG OPR:"

The spooler job is no different from any other job on the system-- it has to
be logged in to a disk account and must have memory if it is to run a
program. Here we log the spooler job into the System Operator's account,
OPR: (DSK0:1,2).

NOTE: If you do not log the spooler job into an account, it will not be able
to run the spooler program. In that event, you will not see an error
message; the system will simply not be able to bring up the spooler.

We log the spooler into account [1,2] because it must reside in the System
Operator's account in order to override the normal AMOS account protection.
This is necessary if a user wishes the PRINT command to delete files in an
account outside of his or her own project.

IMPORTANT NOTE: Although we show the system initialization command file
logging the spooler job into account [1,2], you may not wish to do so. (If
your System Operator account is password protected, you would have to
include the password in the SYSTEM.INI file, which would then make that
password accessible to ingenious but unauthorized users.) To avoid placing
the password to account [1,2] in the SYSTEM.INI file, you may log the
spooler into any disk account (most commonly [1,4]); however, note that the
LPTINI program will always transfer the spooler job to account [1,2] once
the program is running.

2.6.4 "LPTINI PRINTR.INI"

Now we bring in the LPTINI program. LPTINI reads the specified parameter
file (in this case, PRINTR.INI), and brings in the actual line printer
spooler program, LPTSPL. The parameter file, PRINTR.INI tells LPTINI how to
customize the spooler program for your own needs. (See below for information
on PRINTR.INI). The file specification that you provide to LPTINI may that
of any text file that you have created as long as it is a valid spooler
parameter file. The file specification default extension is .INI, and the
default account is [1,4].

2.6.5 "BLANK LINE"

The blank line (a carriage return alone on a line) signals the end of the
forced input to SPOOL.

(Changed 31 October 1981)

2.6.6 "WAIT SPOOL"

The next line:

        WAIT SPOOL

tells the system to wait until SPOOL is finished executing LPTINI before it
proceeds with bringing up the system; this allows the spooler program to
finish any tricky allocations and initializations.

3.0 SETTING UP MULTIPLE SPOOLERS

To set up another spooler, you may basically follow the same instructions
you followed when setting up the first spooler. Just make sure that you
define a new job for the new spooler and that you assign it its own area of
memory. (You may use the same pseudo terminal for the second spooler as you
did for the first.) The following portion of a SYSTEM.INI shows setting up
two spoolers on a bank switched system. Notice that each spooler job uses a
different .INI file-- SPOOL uses PRINTR.INI, and SPOOL2 uses PRNTR2.INI.
Each spooler parameter file specifies a different system printer.

```
        ;
        ATTACH NULL,SPOOL
        JOBMEM SPOOL 3:111530-123030
        KILL SPOOL
        FORCE SPOOL
        LOG OPR:
        LPTINI PRINTR

        WAIT SPOOL
        ;
        ATTACH NULL,SPOOL2
        JOBMEM SPOOL2 3:123032-134340
        KILL SPOOL2
        FORCE SPOOL2
        LOG OPR:
        LPTINI PRNTR2

        WAIT SPOOL2
        ;
        CLKFRQ 60
        MEMORY 0
```


4.0 THE SPOOLER PARAMETER FILE

The commands that appear in the spooler parameter file set the default
information used by the spooler program when dealing with the specific job
to which the parameter applies. You can set up several printers on the
system, all maintaining a separate printer queue, by setting aside a job for
each printer (as we did in the example above), and creating a different
parameter file for each job. By merely changing the name of the printer to

be used and the printer specification, you may initialize as many separate printer queues as you would like.

We have provided a sample parameter file for you, PRINTR.INI:

```
DEVICE=TRM:TERM6
NAME=LPT0
DEFAULT=TRUE
OPERATOR=JOB1
FORMFEED=TRUE
FORMS=NORMAL
BANNER=TRUE
HEADER=FALSE
LPP=56
WIDTH=132
```

You can create your own file (named with any valid filename) using the system text editors, VUE or EDIT. The format used by the elements of the file is:

```
command=argument
```

where argument is a value or attribute to be assigned to the command. Only one command may appear on each line of the file. Some commands take a boolean argument (a true or false value). LPTINI understands the following boolean arguments:

```
TRUE            FALSE
T               F
ON              OFF
YES             NO
Y               N
1               0
```

Here is a list of the parameter file commands:

4.1 DEVICE Command

The format of the command is:

```
DEVICE=devspec
```

where devspec is the specification of the device that is going to be used as the printer.   If you are using a terminal as a printer, your command line might look something like this:

```
DEVICE=TRM:TERM6
```

which uses the generalized terminal driver, TRM.DVR. The terminal must have been defined in a TRMDEF command line in the system initialization command file, and must have a .TDV program in account DSK0:[1,6].

If you are using a Centronics device driver, the command line might look like this:

        DEVICE=CEN:

which selects a device defined in the DEVTBL command of your system initialization command file. Whatever device specification you use for this second format, the device specified must be in your DEVTBL command line in the SYSTEM.INI, and must have a .DVR program in account DSK0:[1,6].

4.2 NAME Command

The format of this command is:

        NAME=string

where string is a one- to six-character name that you want to assign to the device specified by the DEVICE command. Since more than one line printer spooler can be set up on a system, each handling a different printer, giving a name allows you to specify a particular printer to the PRINT command.

4.3 DEFAULT Command

This command takes the format:

        DEFAULT=boolean

If the argument for this command evaluates to TRUE, the spooler program defines the printer defined by the DEVICE command as the default printer to be used when no printer is specified to the PRINT command. This command is an optional one; use it only when more than one printer is being defined on the system. If you omit DEFAULT from all spooler parameter files on the system or if all DEFAULT commands are set to false, the default printer is that printer with the least number of blocks waiting to be printed.

To specify a non-default printer when you use the PRINT command, enter the name of the printer (as specified in the NAME command above) followed by an equal sign. Then enter the specification selecting the files you want to print. For example:

        .PRINT DIABLO=*.LST (RET)

4.4 OPERATOR Command

The format of this command is:

        OPERATOR=jobname

where jobname specifies the job to which the spooler will send error messages and requests for forms changing. If you omit the OPERATOR command,

(Changed 31 October 1981)

the spooler will use the first job on the JOBS line of the SYSTEM.INI as the
operator job.

4.5   FORMFEED Command

This command takes the form:

        FORMFEED=boolean

The command sets the form feed switch default to /FF if the argument
evaluates to TRUE or sets it to /NOFF if the argument evaluates to FALSE.
(See above for a list of legal boolean arguments.) The /FF switch tells the
spooler to perform special form feed handling; this ensures that the printer
is always at top-of-form when the spooler begins to print a new listing.
For some applications (such as check printing), it is not desirable to have
a final form feed output at the end of each listing; the /NOFF switch
disables this final form feed.   If you omit the FORMFEED command, the
spooler program sets the default to /FF.

4.6   FORMS Command

The format of this command is:

        FORMS=formname

where formname is one to six characters that you choose to identify a type
of form (e.g., CHECKS, 2PART, etc.). If you omit the FORMS command, the
spooler uses the default formname of NORMAL.

The purpose of this command is to allow you specify the kind of forms that
should be mounted on the particular printer defined by this parameter file.
The PRINT command then checks printer requests against this to see if the
forms should be changed; if the printer request specifies a different form
than the one mounted on the printer, an error message occurs informing the
user that he must change the forms.

4.7   BANNER Command

The format of this command is:

        BANNER=boolean

This command sets /BANNER as the default switch if the command argument
evaluates to TRUE or sets it to /NOBANNER if the argument evaluates to
FALSE.   (See the table above for legal boolean arguments.) If BANNER is set
to true, a banner page will be printed at the front of each listing.   The
banner identifies the file printed, the printer on which the listing was
made, the date the listing was printed, etc.   If you omit the BANNER
command, the spooler sets the default switch to /BANNER.

4.8 HEADER Command

This command takes the form:

        HEADER=boolean

The HEADER command sets /HEADER as the default switch if the command argument evaluates to TRUE or sets it to /NOHEADER if the argument evaluates to FALSE. The /HEADER switch tells the spooler program to print page headers.   Page headers are titles printed at the top of every page which give the name of the file, the date on which it was printed, and the current page number. The /NOHEADER switch disables the printing of page headers. If you omit the HEADER command, the default switch is /NOHEADER.

4.9 LPP Command

This command takes the form:

        LPP=number

where number specifies the default number of lines per page. The spooler program uses this value in determining where to print page headers. If you omit the LPP command, the spooler uses the value of 56. You may override the value set by the LPP command by using the /LPP switch when you use the monitor PRINT command.

4.10 WIDTH Command

This command takes the form:

        WIDTH=number

where number specifies the default width (in characters) of the printed line. The spooler program uses this value in determining the width of page headers. If you omit the WIDTH command, the spooler uses the value of 132. You may override the value set by the WIDTH command by using the /WIDTH switch when you use the monitor PRINT command.

NOTE: The values you give to WIDTH must range between 80 and 132. If you specify a number less than 80, the spooler uses 80; if the number is greater than 132, the spooler uses 132.

5.0 SPECIFYING FORMS

You can specify the type of form to be used on the printer by including the FORMS command in the spooler parameter file at the time of spooler initialization. The PRINT command compares the forms specified in print requests against this default form to make sure that they match.

After the spooler program is up, you can change that forms default by using
the SET command:

        .SET FORMS printername formname

where printername specifies the specific printer on which the form must be
mounted, and formname gives the form type. For example:

        .SET FORMS TI810 2PART

Once a form has been set using the SET command, the PRINT command checks all
print requests sent to that printer to make sure that the proper type of
forms is being specified.

If the form you specify in the PRINT command is not mounted on the printer
when the file is selected for printing, the spooler sends this message to
the terminal attached to the operator job:

        ;LPTSPL - Please mount form formname on printername

This message repeats once a minute until the spooler is notified that the
form has been mounted (via the SET command). Since both LPTINI and PRINT
use the default formname NORMAL, you can omit the FORMS command in the
parameter file if you only use one kind of paper in your printers.

6.0 TROUBLESHOOTING THE SPOOLER PROGRAM

In the event that you have been unable to get the spooler up and running
properly, you may find this section helpful.   Two of the most probable
reasons for the spooler not coming up correctly are: 1) the printer
initialization file contains incorrect parameters (e.g., you changed the
name of the printer in the TRMDEF statement but did not change the DEVICE
statement in the printer initialization file); or 2) you did not allocate
enough memory to the spooler job, and it couldn't load in the driver for the
device you wanted to print from.

Once the system has been booted, run the SYSTAT program to see what the
spooler job is doing; if all has gone well you should see that the spooler
job is running LPTSPL in an EW state. If the job is not running LPTSPL, or
if it is in a ^C state, something has gone wrong.

You can run the LPTINI program on your own terminal to see what is
happening. Be warned that this procedure will lock up your terminal; but it
should give you some idea of what went wrong. You will have to reboot the
system to gain control of your terminal after you are finished.

To see why the spooler did not initialize properly, log into [1,2] and
allocate yourself the same amount of memory as the spooler job. Now run the
LPTINI program, specifying the same parameter file that you used earlier in
attempting to bring up the spooler. If you see nothing on the terminal
after about 20 seconds or so, everything is probably allright. If you have
another terminal connected to the system, run SYSTAT to see if the spooler

(Changed 31 October 1981)

job is now in EW state. If it is, something is wrong with the way you set
up the spooler in the SYSTEM.INI. Check over your SYSTEM.INI file for some
error in setting up the job, defining the pseudo-terminal, allocating
memory, or defining the output device. If you still see no reason why the
initialization process should have failed, try attaching the spooler to a
real terminal (not the pseudo-terminal). By doing this, you will be able to
see any error messages generated during the initialization process. If you
are using a serial output device, and a TRMDEF command exists for your
printer device (that is, it is defined as a terminal), you may attach the
spooler job directly to the printer.

Chances are, you won't get that far. Most errors are caught by the LPTINI
program which displays an appropriate error message. These are the messages
displayed by LPTINI:

MEMORY ALLOCATION FAILED
     You did not allocate enough memory to the job. LPTINI wasn't even
able to load itself and its impure areas. Allocate more memory.

?Invalid command X
     You specified command or argument "X" in your spooler parameter file
which LPTINI was not able to recognize. Check the spelling in the file
for errors. Make sure that your parameter command arguments are in the
proper form (e.g., numeric or boolean).

?Bad DEVICE specification
     The device you specified as the printer in your DEVICE command is
not a valid device specification. Make sure that the device you
specified is in the device table (the DEVTBL line of the SYSTEM.INI),
and the the device has a driver program in DSK0:[1,6].

?Nonexistent job name specified for OPERATOR
     The job you specified in your OPERATOR command does not exist.
Check your spelling, and examine a SYSTAT display to see a list of the
valid jobs on the system.

?Insufficient memory to run spooler, expand memory by n bytes
     You did not specify enough memory to run the line printer spooler
program or (if the driver is not already in system memory) to load in
in the driver for the device containing the file you want to print.
Increase your memory allocation by the amount specified. The amount of
memory required by the spooler program depends on the particular device
to which you are printing, and so may exceed the recommended amount of
4K.

Index

MEMORY MANAGEMENT OPTION

This document describes how to set up and use more than 64K of memory in
your Alpha Micro computer system by bank switching memory. It will not tell
you how to address your memory boards; in order to implement memory
management you will have to refer to documentation supplied by the board
manufacturer(s). For more information, refer to the Alpha Micro Integrated
Systems User's Guide, (DWM-00100-00).

Become familiar with this documentation before trying to implement memory
management on your system. Also, try to become familiar with the addressing
of the various memory boards you will be using in your system. If you have
problems implementing memory management, refer to Section 3.0 to see if the
problem is explained there.

For a full explanation of memory management, see Section V, "Memory
Management," of the Integrated Systems User's Guide, (DWM-00101-00).

1.0 SETTING UP A MEMORY BANK

Unless the system initialization command file (SYSTEM.INI) contains MEMDEF
commands, the operating system (AMOS) will not recognize any memory beyond
the first 64K (1K = 1024; 64K = 65536). There is one MEMDEF command for
each bank of memory to be defined. The MEMDEF commands must be before the
final SYSTEM command in the SYSTEM.INI.

Due to the way AMOS works, there must be a totally sharable portion of
memory starting at location 000000. This sharable portion must be large
enough to hold AMOS, which requires a minimum of 11K (this size increases as
more jobs, terminals, devices, and system programs are defined). You may
also place the line printer spooler in sharable memory, which requires
between 4 and 5K. Typically, the sharable portion will be 16K with no
system programs, 32K if RUN.PRG is in system memory, or 48K if BASIC.PRG and
RUN.PRG are in system memory. If sharable memory takes up some amount $\underline{x}$K,
then the largest bank-switched portion of memory can take up $(64 - \underline{x})$K of
memory. Bank-switched memory refers to that portion of memory past sharable
memory. There may be multiple banks of various sizes, as long as the
largest bank is no larger than $(64 - \underline{x})$K, and there is no sharable memory
past $\underline{x}$K.

Bank-switched memory must be physically addressed to the memory address that
it will occupy when that bank of memory is active. (Refer to the
manufacturer's specifications for the appropriate switch or jumper
settings.) For instance, assuming that sharable memory is from 0K to 31K and
bank-switched memory starts at 32K, no bank-switched portion of memory can
be addressed to a physical address below 32K. In this situation, assuming
16K memory boards, there would be one sharable memory board addressed for

OK-15K, one sharable memory board addressed for 16K-31K, and then the
bank-switched memory boards. Two banks of 32K would have two boards in one
bank addressed for 32K-47K and for 48K-63K respectively, and in the second
bank a board addressed for 32K-47K and a board for 48K-63K.

In addition to the physical memory address, all bank-switched memory boards
have an I/O port address associated with them. This I/O port is what is
used to turn the various memory boards on and off. There may be more than
one I/O port possible per board, depending on the manufacturer of the memory
board.

This scheme allows memory boards to be turned on or off by sending a value
to the I/O port of the memory board. More than one bank may use the same
I/O port, as long as the value for turning on a particular bank is different
than the value for turning on a different bank.

1.1 Defining a Memory Bank (MEMDEF)

There is one MEMDEF command for each memory bank. Memory bank numbering
starts at 0 and is incremented by one for each MEMDEF statement. Therefore,
three MEMDEF instructions in the initialization file define Banks 0, 1, and
2.

The form of the MEMDEF command is as follows:

    MEMDEF board0-adr,on-cnst,off-cnst/board1.../boardN-adr,on-cnst,off-cnst

    1. All operands are in octal, unless the SET HEX command has been
        executed.

    2. "boardX-adr" is the I/O port address that the board is set up on.

    3. "on-cnst" is the on constant-- the value that is used to turn the
        piece of memory at this I/O port on.

    4. "off-cnst" is the off constant-- the value that is used to turn the
        piece of memory at this I/O port off. This will usually turn off
        the entire board except for any memory that is sharable or set up
        for a different I/O port.

    5. The slash (/) allows multiple memory addresses to share the same
        bank of memory. This is especially convenient when using a large
        memory board (32K - 64K) that resides all at one I/O port but has
        different "on" and "off" values for smaller pieces of memory (8K to
        16K).

Assuming a 16K sharable portion residing at the front of a 64K memory board,
the remaining 48K of memory is switchable, and so are two other 64K memory
boards with only 48K active on them; the first board is set up at I/O port
100, the second board is set up at 101, the third board is set up at 102;
the "on" constant for memory on all boards in Bank Zero is 1, for memory in

(Changed 30 April 1981)

Bank 1 is 2, for memory in Bank 2 is 4; the "OFF" constant for all memory is
0; the MEMDEF commands would then be as follows:

```
        MEMDEF 100,1,0/101,1,0/102,1,0        ; Bank Zero
        MEMDEF 100,2,0/101,2,0/102,2,0        ; Bank 1
        MEMDEF 100,4,0/101,4,0/102,4,0        ; Bank 2
```

Note that memory from 0K to 15K is not set up to be switchable. Also note
that a piece of all memory boards is in all banks. The boards would be set
up in the following manner:

```
   Board 1: I/O port 100    16K-31K  On constant = 1  Off constant = 0
                            32K-47K  On constant = 2     (all boards)
                            48K-63K  On constant = 4
   Board 2: I/O port 101    16K-31K  On constant = 2
                            32K-47K  On constant = 4
                            48K-63K  On constant = 1
   Board 3: I/O port 102    16K-31K  On constant = 4
                            32K-47K  On constant = 1
                            48K-63K  On constant = 2
```

```
                 Board 1        Board 2        Board 3
                +------+       +------+       +------+
   0-15K        |shared|       |unused|       |unused|
                +------+       +------+       +------+
   16-31K       |Bank 0|       |Bank 1|       |Bank 2|
                +------+       +------+       +------+
   32-47K       |Bank 1|       |Bank 2|       |Bank 0|
                +------+       +------+       +------+
   48-63K       |Bank 2|       |Bank 0|       |Bank 1|
                +------+       +------+       +------+
```

Note that this is but one possible memory configuration and is set up in a
strange way for demonstration purposes. A more efficient use of 64K memory
boards is possible if different portions of the board are set up at
different I/O addresses or (if the entire board is at one I/O address) if
different portions have different "on" constants.

This scheme allows the use of different manufacturer's boards in the same
system, as long as they can be set up at various I/O addresses and the
largest piece turned on by a particular constant is small enough so that it
does not overflow into sharable memory.

After the system is booted, the MEMDEF command takes on a new function. The
MEMDEF command now allows you to see how memory is configured. The MEMDEF
command will display a map of all active banks as a string of characters,
with each character representing 1K of memory. Assuming a 13K monitor and
the above example memory definition, here is what MEMDEF would display after
you have booted the system:

(Changed 30 April 1981)

.MEMDEF (RET)

```
            0K      8K      16K     24K     32K     40K     48K     56K     64K
            |       |       |       |       |       |       |       |       |
            ┌───┬───┬───┬───┬───┬───┬───┬───┬───┬───┬───┬───┬───┬───┬───┬───┐
 BANK 0  MMMMMMMMMMMMMMMSSSBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
 BANK 1  MMMMMMMMMMMMMMMSSSBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
 BANK 2  MMMMMMMMMMMMMMMSSSBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
```

.

"M" stands for "monitor" memory (memory used by the monitor), "S" stands for
"sharable" memory (memory that is not monitor memory and is not bank
switched), and "B" stands for "bank-switched" memory. Since monitor and
sharable memory can be accessed by all banks, this area of memory is treated
as a part of all banks. It is always active. The area of memory used by
the monitor and sharable programs is also called "system memory."

Programs can be loaded into system memory using the SYSTEM command in the
SYSTEM.INI file. The spooler can reside in system memory or in
bank-switched memory. Any remaining sharable memory after the system is
booted up that is not used by the monitor or programs in system memory is
added to the beginning of Bank Zero and is available to the job the system
comes up under. (The last command in the SYSTEM.INI file, MEMORY 0,
allocates all available memory in Bank Zero to the job the system comes up
under.)

There is a good reason why there is no memory appearing at "64K." As you
can see, the counting of memory starts at 0K, not 1K. If there were memory
at "64K," it would really be "65K," which is 66560 bytes of memory, since 1K
= 1024.

The MEMDEF display is also a handy way to verify that you have jumpered the
boards properly. Memory which the system cannot find will be missing from
the display. Bad memory boards can sometimes be located in this fashion,
too. Of course, if the bad board is in sharable memory or Bank Zero, the
system may not come up at all.

2.0 ALLOCATING MEMORY TO JOBS

With this memory management system, the MEMORY command becomes obsolete
except for allocating memory in Bank Zero. Do not use the MEMORY command to
allocate memory for any job other than the first job in Bank Zero, the job
the system comes up under. Instead, use the JOBMEM command to allocate
bank-switched memory.

(However, remember to use the MEMORY command if you are allocating the line
printer spooler in system memory. See Setting Up the Line Printer Spooler
in the "System Operator's Information" section of the AMOS Software Update
documentation packet for information on allocating the spooler in either
system or bank-switched memory.)

(Changed 30 April 1981)

Use the JOBMEM command as follows:

    JOBMEM jobnam Bank-number:StartAddress-EndAddress

1. "jobnam" is the name of the job for which memory is to be allocated. The user's current job is used if "jobnam" is not specified.

2. "Bank-number" is the number of the memory bank that the job is to reside in.

3. "StartAddress" is the starting memory location for the job.

4. "EndAddress" is the ending memory location for the job.

If you specify no arguments, JOBMEM tells you how much memory is allocated to your job. If you specify a job name with no memory argument, JOBMEM tells you how much memory is allocated to the specified job.

Here are some sample JOBMEM statements:

```
.JOBMEM JOB2 1:40000-177376 (RET)
.JOBMEM JOB3 2:40000-77776 (RET)
.JOBMEM JOB4 2:100000-137776 (RET)
.JOBMEM JOB5 2:140000-177376 (RET)
.JOBMEM JOB5 (RET)
CURRENT MEMORY ALLOCATION IS 2:140000-177376
.JOBMEM (RET)
CURRENT MEMORY ALLOCATION IS 0:32370-177376
```

The first example gives the job called JOB2 all of the memory in Bank 1. The second through the fourth examples allocate the 48K of memory in Bank 2 in 16K partitions to three different jobs. The fourth job (JOB5) really gets 16K minus 256 bytes of memory (16128 bytes total) due to the fact that the top 256 bytes of memory are not accessible. (The top 256 bytes of memory are reserved for the I/O ports.) The fifth example lists the memory allocation for the job called JOB5. The final example lists the memory allocation for the current job.

2.1 JOBMEM Error Messages

?Memory allocation format error
    JOBMEM is confused by the format of the information you've given it; it may be that you've made a spelling error.

?Non-existent bank number
    If you've specified a bank number greater than the number of MEMDEF commands in the SYSTEM.INI, you'll see this message.

?Non-existent job
     You've specified a job that does not exist.   Check   the   JOBS   command
line in your SYSTEM.INI or run SYSTAT to see a list of the jobs on the
system.

?Allocation overlaps monitor or system memory
     Monitor and system memory is sharable between users; that   means   that
you   must   not   allocate   any part of it to an individual user. Reduce your
user allocations, or remove some of the programs that you've previously
added to system memory by deleting the appropriate SYSTEM commands in your
SYSTEM.INI.

?Illegal memory range (end is below base)
     You probably entered your ending memory address before you entered the
starting address.   The   starting   address of the memory block you want to
allocate must be less than the ending address.

?Allocation is not within requested bank's defined memory
     You've asked for memory in a valid bank, but the memory addresses that
the   bank   is   set   for   do   not   include the addresses that you've tried to
allocate. Check your MEMDEF commands, and the addressing of your memory
boards.

?Requested allocation would overlap job Jobname
     You've tried to allocate memory to one job that belongs to another.
The job you've overlapped is named Jobname.   (e.g., ?Requested allocation
would overlap job JOBA.)

3.0 HINTS, RESTRICTIONS, AND WARNINGS

The next few paragraphs discuss some of the things that you should avoid
doing when setting up your system for bank-switched memory. We also discuss
some of the miscellaneous pieces of information that you should know to set
up your system in this way.

When power is initially supplied to the system, or whenever the reset button
is pressed, Bank Zero must turn itself on and all other banks must turn
themselves off. Any other configuration may cause system failure either
immediately after AMOS starts executing the SYSTEM.INI file, or when the
MEMDEF statements in the SYSTEM.INI file are executed.

All memory in a particular bank must be contiguous.   That is, you cannot
have a piece of switchable memory starting at one location and ending at
another, a "hole" where there is no memory, and then another piece of
switchable memory. This is also true of sharable memory, and all memory in
a system that does not use bank switching. You can, however, have sharable
memory up to some point, a "hole," and then bank-switched memory. This is
not the case with Bank Zero, which must always be contiguous with sharable
memory.

If you are allocating a line printer spooler and there is not enough sharable memory for the spooler, your spooler will not work. Make sure that the spooler's "end adr" is lower than the top of sharable memory if you are placing the spooler in system memory.

If too many modules are loaded into the monitor with the SYSTEM command, monitor memory will extend into the bank-switched area, and the system will not work properly. You can use the MEMDEF command after you boot the system to check this out.

The MEMDEF commands must come before the final SYSTEM command in the SYSTEM.INI. It is recommended that the MEMDEF commands be placed right after the TRMDEF commands.

MEMDEF does not check to see if an I/O address with the same on/off constants was previously allocated. It is up to the System Operator to make sure that SYSTEM.INI is set up properly and that a single piece of memory is not allocated to two different banks.

When the MEMDEF display shows the last column of bank-switched memory at "62K" (instead of at "63K" where it belongs), you may have a problem with your disk controller board(s). The 8131 driver IC for the bootstrap PROM must be removed from any board controlling a disk device that is not the system device. PHANTOM must be removed from these boards also. Only the board which controls the system device can have the bootstrap PROM and PHANTOM enabled.

All memory addresses must be on word boundaries. In octal, memory addresses on word boundaries end with 0, 2, 4, or 6. In hex, they end with 0, 2, 4, 6, 8, A, C, or E.

With this memory management system, there is no way for a user to have a portion of memory in one bank, and another portion in another bank, unless he is running a specially written assembly language program and uses the BNKSWP monitor call. The user must be in sharable memory when he issues this call. (There is more information on the BNKSWP call in the AMOS Monitor Calls Manual, (DWM-00100-42).)

A bank of memory may be composed of part or all of several different boards, even if the boards are made by different manufacturers.

MEMORY MANAGEMENT WITH THE MEMORY PARTITION CONTROLLER (MPC)

October 1981

Revision A01

This document reflects AMOS versions 4.6 and later

'Alpha Micro', 'AMOS', 'AlphaBASIC', 'AM-100',
'AlphaPASCAL', 'AlphaLISP', 'AlphaSERV',
'AlphaVUE', and AlphaACCOUNTING

are trademarks of

ALPHA MICROSYSTEMS
Irvine, CA   92714

ALPHA MICROSYSTEMS
17881 Sky Park North
Irvine, CA   92714

Table of Contents

(Changed 31 October 1981)

MEMORY MANAGEMENT WITH THE MEMORY PARTITION CONTROLLER

1.0 INTRODUCTION

The AM-700, or Memory Partition Controller (MPC), is a controller board designed by Alpha Micro to provide a new form of memory management.

On a system without memory management, the entire amount of memory on the system is limited to a maximum of 64K. To overcome this limitation, Alpha Micro implemented its first form of memory management, called "bank switching," which allowed the System Operator to define multiple sets of memory (each limited to 64K minus the size of the resident monitor) that the monitor could switch between. Only one bank is active at any one time. Once a bank is defined, multiple user partitions can be assigned to that bank. Although each user is still restricted to the 64K limitation, the system itself is no longer restricted to the 64K maximum. (For more information on bank switching, see the document Memory Management Option in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.)

With the AM-700, Alpha Micro introduces a new form of memory management. Rather than defining multiple banks of memory, the MPC treats all memory in the system as one contiguous block. An individual user is still limited to a partition no larger than 64K minus the size of the resident monitor (also, the top 256 bytes of the last memory partition are reserved for I/O port addresses). But memory definition and allocation become much easier. The MPC method of memory management offers several advantages over the bank-switching method.

First, the MPC is much more efficient in its use of memory. Our usual system requires that memory be assigned to the monitor and to users in increments of 16K bytes on a system using Piiceon boards and 32K bytes on one using AM-710 boards. That means that on a Piiceon system, if the monitor needs 16K plus 200 bytes, you must allocate 32K bytes. The MPC system, on the other hand, assigns memory in increments of 256 bytes; therefore, regardless of the size of the monitor, almost all of the balance of the memory can be accessed for user partitions. The size of the user partitions can then be defined as needed, instead of being limited by increments of 16 or 32 kilobytes.

Aside from efficient use of memory, the MPC system provides the following advantages:

1. Simplified system set-up

2. Easy reconfiguration

3. Better use of DMA capability

(Changed 31 October 1981)

2.0 SYSTEM SET-UP

A memory management system without the MPC board requires that you know
details concerning bank sizes, which memory board or boards each bank is
located on, port addresses, on/off bytes, and the exact addresses (in bytes)
of memory available to each job. You must then define this information
using the MEMDEF and JOBMEM statements in the system initialization command
file.

With the MPC board, however, memory management is greatly simplified. You
do not need to know bank sizes, port addresses, or on/off bytes, because
division into banks (hence, bank switching) is not used. Here are the steps
for implementation:

1. Before installing the AM-700 board, alter your SYSTEM.INI file so
    that it has the minimum number of commands necessary for booting up
    when the AM-700 is installed. This SYSTEM.INI file must contain
    MPCOFF as the first command and must have no MEMDEF commands.
    Use the following SYSTEM.INI file as a model for your own.
    Depending on your system, you may have to modify the BITMAP and
    TRMDEF statements.

        :T
        MPCOFF
        JOBS JOB1
        TRMDEF TERM1,AM300=1:16,SOROC,100,100,80
        DEVTBL DSK1,TRM,MEM,RES
        BITMAP DSK,606,0,1
        SYSTEM
        CLKFRQ 60
        MOUNT DSK1:
        MEMORY 0

2. Test the SYSTEM.INI file shown above to make sure it boots
    correctly, before installing the AM-700 hardware.

3. Now install the AM-700 and the memory boards with the appropriate
    jumpers to define addresses. The jumpering process is much less
    complicated than that required by a non-MPC system. (See section
    9.0, "MPC HARDWARE," for hardware restrictions.) For details on
    installing the hardware, refer to AM-700 Installation
    Instructions (PDI-00700-01). This document also tells you how to
    address your memory boards.

4. When the hardware is set up, boot the system using the SYSTEM.INI
    file shown in Step 1. The system is now running with the MPC
    turned off.

5. Create a test initialization file (e.g., TEST.INI) with the
    commands you need for configuring the system. Figure out how much
    memory each job needs, and use the JOBSIZ command to allocate the
    correct amount of memory (see next section for an explanation of
    JOBSIZ).

(Changed 31 October 1981)

6. Use MONTST to boot off the test initialization file. If it doesn't
    work, you can do a hardware reset to boot up with the SYSTEM.INI
    file shown in Step 1, and adjust the test initialization file until
    it works correctly. NOTE: To use MONTST, the System Disk must be
    on the first fixed disk of the System Drive.

7. Once the test initialization file is working, delete the SYSTEM.INI
    file used for the original initialization and rename the test file
    to SYSTEM.INI.

2.1 The System Initialization Command File

As mentioned earlier, bank switching requires that banks be divided into
increments of at least 16 kilobytes (32K on AM-710 systems). This division
of memory into banks requires that you use the MEMDEF command in the system
initialization command file to identify the port addresses and the switches
that turn the banks on and off. If one bank overlaps from one board to
another, the MEMDEF command becomes more complex.

In addition to the MEMDEF command, the non-MPC system requires use of a
JOBMEM command to identify the bank in which each job will run, and to
allocate an exact range of bytes within that bank. Also, to define the size
and bank of switchable system memory, you must use the SYSMEM command.

In an MPC system, the MEMDEF, JOBMEM, and SYSMEM commands are not necessary
because the memory, regardless of the number of memory boards involved, is
not divided into banks. If you are using two 128-kilobyte boards (AM-710),
the MPC system recognizes only that you have 256KB; it makes no distinction
concerning banks or bank switching. For this reason, in the system
initialization command file you need only allocate the appropriate amount of
memory to each user by using the JOBSIZ command. When using JOBSIZ, you can
allocate memory in two different ways:

1. You can allocate a specific number of bytes, as long as you do not
    exceed 65535 (64K) for each user. For example, suppose you want to
    allocate 35KB of memory to each of two users and 40KB to one user
    on a system that includes 128KB of memory, and on which the monitor
    occupies 13K. The JOBSIZ commands to enter in the system
    initialization command file would be:

        JOBSIZ JOB2,35K
        JOBSIZ JOB3,35K
        JOBSIZ JOB4,40K

    Instead of allocating memory in kilobytes, you can also specify a
    number of bytes. For example, you can enter JOBSIZ JOB2,10450.
    The actual memory allocated--whether you specify kilobytes or a
    specific number of bytes such as 10450--will be rounded up to the
    next multiple of 256; in the case of 10450, the amount of memory

allocated would be 10496. If the rounding causes the memory allocation to exceed the memory available, you see an error message.

And that's all you need to do. Just use the JOBSIZ command; do not use MEMDEF or JOBMEM commands, or you get an error message.

2. The second way to use the JOBSIZ command is to let the system allocate automatically the maximum amount of memory available for users. If you have 128KB of memory, as in the example above, you could enter the following in the system initialization command file:

        JOBSIZ JOB2,MAX
        JOBSIZ JOB3,MAX
        JOBSIZ JOB4,MAX

    Jobs 2 and 3 are each automatically given 51K of memory--64K minus 13K (the size of the monitor). Job 4 gets the 13K that's left--[128-(2 X 51)]-13.

If you do not want the AM-700 board activated, you can enter MPCOFF as the first statement in the system initialization command file. If the board is installed and you do not enter this statement, the AM-700 is activated as soon as the system boots up.

Below and on the next page is a sample SYSTEM.INI file with JOBSIZ commands included.

```
:T
JOBS JOB1,JOB2,JOB3,JOB4
TRMDEF TRM1,AM300=1,SOROC,80,80,80
TRMDEF TRM2,AM300=2,SOROC,80,80,80
TRMDEF TRM3,AM300=6:16,TELTYP,50,50,20
TRMDEF TRM4,PSEUDO,NULL,20,20,2
TRMDEF TRM5,AM300=3,SOROC,80,80,80
DEVTBL DSK1,TRM,MEM,RES
DEVTBL /MTU0
BITMAP DSK,606,0,1/S
MOUNT DSK1:
QUEUE 20
SYSTEM MTSTAT.SYS
SYSTEM

JOBSIZ JOB2,MAX
ATTACH TRM2,JOB2
KILL JOB2

JOBSIZ JOB4,4K
ATTACH TRM4,JOB4
KILL JOB4
```

```
; Bring up the printer/spooler
FORCE JOB4
LOG 1,2
LPTINI PRINTR

WAIT JOB4
;
JOBSIZ JOB3,MAX
ATTACH TRM5,JOB3
KILL JOB3
JOBSIZ JOB1,MAX
```

3.0 MEMORY ALLOCATION DURING INITIALIZATION

Memory allocation during initialization of an MPC system differs from that
of a bank-switching system. When the MPC system starts booting up, MPC
determines the total memory available, the monitor is automatically assigned
the amount of memory needed to begin initialization, and JOB1--the boot-up
job--is assigned the 8 kilobytes from 56K to 64K. As the system encounters
BITMAP statements (containing the /S switch), supplemental system memory for
the bitmaps is assigned, beginning at the end of physical memory and
allocating downward. The supplemental system memory performs the same type
of function as switchable system memory on a bank-switching system. If /S
is not used with the BITMAP statement, bitmaps are placed in the monitor
area. An MPC system will not boot up with less than 64K of memory.

The diagram below illustrates the possible memory allocation for a 128KB MPC
system before any SYSTEM commands are executed. (For simplicity, assume
that the two bitmaps together occupy 1KB of memory.)

```
OK        13K                     56K  64K                            128K
 _____
|       |                      |     |                          |B |B |
|  M    |                      |  J  |                          |I |I |
|  O    |                      |  O  |                          |T |T |
|  N    |                      |  B  |                          |M |M |
|  I    |                      |  1  |                          |A |A |
|  T    |                      |     |                          |P |P |
|  O    |                      |     |                          |1 |0 |
|  R    |                      |     |                          |  |  |
|_____|_____|_____|_____|__|__|
```

Allocation for the monitor then grows as the commands in the system
initialization command file are processed. When the system encounters a
SYSTEM command without an argument, the size and boundary of the monitor are
finalized. At this point, JOB1 is moved to the 8KB segment contiguous with
the monitor area, as shown in the diagram on the next page.

OK   16K   24K   64K   128K

M O N I T O R	J O B 1			B I T M A P 1	B I T M A P 0

Notice the monitor in the above diagram now occupies 16KB of memory. As the system now processes the JOBSIZ statements of the initialization file, memory is allocated for user partitions beginning at the upper boundary of JOB1 and progressing toward high memory. Suppose JOB2 is given 32K and JOB3 is given 40K. After the JOBSIZ commands are executed, the memory allocation looks like this:

OK   16K   24K   56K   96K   128K

M O N I T O R	J O B 1	J O B 2	J O B 3		B I T M A P 1	B I T M A P 0

Now, the last command in the system initialization file should be:

    JOBSIZ JOB1,MAX

followed by a blank line.

This takes the place of the MEMORY 0 command entered at the end of the non-MPC system initialization file. It reallocates memory to JOB1, the job on which the system has been booting up. This reallocation causes the warning message shown on page 7 to appear. You need a blank line after the last command of the system initialization file so the system can execute an automatic carriage return after this warning message.

Since, as shown in the above diagrams, JOB1 is contiguous with the monitor area, assigning it the maximum amount of memory left over increases its size (in this case, to 39KB) and redefines the addresses of the other jobs on the system (see diagram on the next page).

(Changed 31 October 1981)

0K	16K	55K	87K	128K

M O N I T O R	J O B 1	J O B 2	J O B 3	B I T M A P 1	B I T M A P 0

4.0 CHANGING MEMORY ALLOCATIONS

Without the MPC, alteration of user partitions and sharable memory can require extensive changes in the MEMDEF and JOBMEM statements and time-consuming adjustments to the hardware. With the MPC, though, reconfiguration is simply a matter of entering new JOBSIZ commands in the system initialization command file or using JOBSIZ at AMOS command level. You just have to make sure your allocation, added to the size of the monitor, does not exceed the 64KB allowed, and that the total amount of memory allocated to all jobs does not exceed the total memory capacity. To find out the number of kilobytes occupied by the monitor, use the SYSTAT or SYSTEM command at AMOS command level.

To execute JOBSIZ at AMOS command level, use the same format as the entry in the system initialization command file:

.JOBSIZ Jobname,Jobsize (RET)

If other jobs are running on the system, they will be delayed for at least three seconds during the reallocation. To increase the size of one job, you may need to decrease the size of another. Use the JOBSIZ command to define a smaller partition for one user, then execute another JOBSIZ command to increase the size of the partition you want to be larger. Keep in mind, though, that any job to be affected by your JOBSIZ commands must not be running.

When you enter JOBSIZ at AMOS command level, the following message appears:

% WARNING - When using the JOBSIZ command to reallocate a job's memory, you must be aware of the following restrictions:

1) All jobs are suspended during reallocation.

2) You must not change the size of a job that is running.

4) No jobs can be using direct memory access (DMA).

Type RETURN when ready or CTRLC to exit:

5.0   USE OF DIRECT MEMORY ACCESS (DMA)

On a bank-switching system, as soon as a user executes a job that employs a magnetic tape drive, all other banks are turned off. Although other users are disabled for only a split second at a time, DMA slows system throughput.

On the MPC system, other users can continue running while one implements DMA, because the MPC translates CPU and DMA addresses separately, allowing concurrent activity.

6.0   DISPLAYING THE MEMORY MAP (MEMPLT)

MEMPLT (Memory plot) is an improved version of the AMOS command level command, MEMDEF. It displays a memory map for systems using the MPC, bank switching, or no memory management. When used on an MPC system, MEMPLT displays a memory map broken down into increments of 256 bytes instead of 1K bytes. One of the most convenient advantages of MEMPLT over MEMDEF is that MEMPLT can show what portion of memory each job occupies.

Activation of MEMPLT requires no entry in the system initialization file. You simply enter MEMPLT and press RETURN at AMOS command level. On an MPC system, the resulting display looks like the one shown below.

```
LEGEND:
A JOB001    B JOB002    C JOB003    D JOB004    E JOB005    F JOB006    G
H           I           J           K           L           M           N
O           P           Q           R           S           T           U
V           W           X           $ Supplemental system  * Monitor   + Unused


         +0K +1K +2K +3K +4K +5K +6K +7K +8K +9K +10K    +12K    +14K    +16K
          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
 0K      ****************************************************************AAAAAAAAAAAA
          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
 16K     AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABBBBBBBBBBBB
          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
 32K     BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
 48K     BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB
          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
 64K     CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
 80K     CCCCCCCCCCCCCCCCCCCCCCCDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDD
          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
 96K     DDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDDEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEEE
          |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
 112K     FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF+++++++++++++$$$$$$$$
Hit RETURN to continue
```

(Our usual practice is to show computer output by underscoring each line. However, because of the graphic nature of the MEMPLT display, we have made and exception to avoid visual confusion.)

One 24-line screen can display a map for 128KB of memory. If your system has more than 128K, press RETURN at the bottom of the first display to bring up the rest of the map.

The MEMPLT display above is the memory map for a system with:

1. A 13KB monitor.

2. One 16KB job (A); a 35KB job (B); a 22KB job (C); an 18KB job (D); an 8KB job (E); an 11KB job (F).

3. 3KB of unused memory (+).

4. 2KB of supplemental system memory ($).

The top portion of the display is a legend showing the code for each job, the monitor area, unused memory, and supplemental system memory (containing the bitmaps). The rest of the display uses these codes to show you how memory is currently allocated. Each symbol in the lower part of the display (A, B, C, *, +, etc.) represents 256 bytes occupied by the job or type of memory corresponding to that symbol.

7.0 MPC AND EXISTING PROGRAMS

Certain existing AMOS commands operate differently or do not operate at all on the MPC system. The rest of this section discusses these commands.

7.1 MEMDEF

In the MPC system, the MEMDEF command is unnecessary because banks and bank switching are not used. Use of this command when implementing the AM-700 board results in an error message.

7.2 MEMORY

The MEMORY command is not used in the MPC system. At the end of the system initialization command file, entry of JOBSIZ JOB1,MAX when using MPC replaces the entry of MEMORY 0 in a non-MPC configuration. If you do use MEMORY on an MPC system, you see an error message.

7.3 JOBMEM

JOBMEM is replaced on the MPC system by JOBSIZ. Use of JOBMEM when the MPC
board is activated causes an error message to be displayed.

7.4 BITMAP

The syntax of the BITMAP statement is the same in the MPC system as it is in
a non-MPC configuration. When you use MPC, the system allocates memory for
the bitmaps as soon as it encounters the BITMAP statements in the
initialization file. On the MPC system we call this memory "supplemental,"
distinguishable from "switchable" system memory which implies a
bank-switching system.   Supplemental system memory on an MPC system does,
however, serve the same purpose as switchable system memory on a
bank-switching system.

If you use the /S feature with the BITMAP statement, specifying supplemental
system memory, the MPC system automatically allocates memory for the bitmap,
beginning at the high memory boundary and moving downward.

For example, suppose your system has 128KB of memory, and your
initialization command file contains these statements:

        BITMAP DSK,606,0/S
        BITMAP DSK,606,1/S

Remember that in the BITMAP statements, the allocation you specify is in
words. The 606 words in each of the above statements equals 1212 bytes.

Upon initialization, the MPC system allocates the last 1280 bytes (the next
multiple of 256 above 1212) in memory for the first bitmap and the
second-to-last 1280 bytes for the second bitmap. So the bitmaps would
occupy bytes 125440 to 128K.

When you use the BITMAP command at AMOS command level, the display shows, in
octal or hexidecimal notation, the memory range occupied by the bitmaps:

        .BITMAP (RET)
        DSK0:   01175400 - 01177722
        DSK1:   01175400 - 01177722

7.5 SYSTAT

Using the SYSTAT command on an MPC system gives you the same information as
a SYSTAT on a non-MPC system, with a few variations. Examine the following
sample SYSTAT display:

(Changed 31 October 1981)

Status of AMOS Version 4.5A (MPC)

User logical base address is : 00036000

JOB1	TERM1	DSK0:44,20	023512	RN	SYSTAT	49920 bytes from 00036000
JOB2	TERM2	DSK1:20,2	024156	RN	TYPE	49920 bytes from 00177400
JOB3	TERM3	DSK0:45,1	024622	^C	VUE	20480 bytes from 00356000
JOB4	TERM4	DSK0:1,2	025266	EW	LPTSPL	6656 bytes from 00341000
JOB5	TERM5	DSK0:100,10	025732	TI	BASIC	49920 bytes from 00426000
JOB6	TERM6	DSK1:30,3	026376	RN	VUE	20480 bytes from 00567400
JOB7	TERM7	DSK1:55,5	027042	^C	BASIC	49920 bytes from 00637400

7 jobs on system

DSK0   3425 Blocks free   DSK1   5646 Blocks free

8 devices on system

Notice that the first line of the display shows MPC after the version
number. The second line is unique to the MPC system; it identifies, in
octal (or hexidecimal) notation, the monitor/user partition boundary. The
rest of the display is different from that of a non-MPC system in its
definition of job locations. Since banks are not used on the MPC system,
you see no bank numbers. Instead, you see only the size of the user
partition, in decimal (e.g., 49928 bytes), and the octal (or hex)
specification of its beginning physical extended address (e.g., 00036000).

7.6 SYSMEM

Use of SYSMEM on an MPC system is not necessary and results in an error
message. A /S option with the BITMAP statement causes the monitor on an MPC
system to automatically allocate supplementary system memory for the bitmap
at the top of memory.

8.0 ERROR MESSAGES FOR THE MPC SYSTEM

The following error messages are included as part of the MPC system:

?Insufficient free memory for request

> The job size you specified is larger than the amount of available
> memory. Use MEMPLT to display the memory map; and, if you wish, use
> JOBSIZ to decrease the size of another job before allocating memory to
> the new one.

?Insufficient supplemental memory - monitor memory will be used

In an MPC system that has been switched off, bitmaps can only be placed in the monitor area. This message lets you know that the bitmaps have been placed in the monitor area, not in supplemental system memory.

?Invalid command

The argument in the JOBSIZ command is faulty (e.g., used hex instead of decimal notation).

?Invalid job name

The job name you specified matches none of the job names listed in the JOBS statement of the system initialization file.

?JOBMEM invalid on an MPC system - use the JOBSIZ command

You used JOBMEM on an MPC system. The JOBSIZ command is all you need.

?MEMDEF invalid on an MPC system - use the JOBSIZ command

You used MEMDEF on an MPC system. The JOBSIZ command is all you need.

?MEMORY invalid on an MPC system - use the JOBSIZ command

You used the MEMORY command on an MPC system. The JOBSIZ command is all you need.

?MPCOFF cannot be run after system initialization

You tried to execute MPCOFF at AMOS command level. This command can only be used in the system initialization file.

?Not an MPC system. Use the JOBMEM command

You tried to use the command JOBSIZ on a non-MPC system. Use JOBMEM instead.

?Request overlaps sharable area

The job size requested, when added to the monitor size, exceeds 64KB.

?SYSMEM invalid on an MPC system

You used the SYSMEM command on and MPC system. SYSMEM is not necessary because the BITMAP statement on an MPC system automatically allocates supplemental system memory for the bitmaps.

9.0 MPC HARDWARE

Implementing the Memory Partition Controller board involves the following hardware restrictions:

1. You may use the MPC with the AM-100/T only, not the AM-100.

2. You may not use the 64K-byte Piiceon memory board with an MPC system, but you can use the AM-710 board. You can also use the 32K-word Piiceon board as long as you are not also using the AM-600 magnetic tape controller board.

3. The AM-600 magnetic tape controller board requires modification to be used on an MPC system.

4. If the AM-600 board is used on an MPC system, the AM-100/T must be modified.

For instructions on necessary hardware installation and the modifications mentioned in the above list, refer to AM-700 Installation Instructions (PDI-00700-01) and the AM-700 Hardware Manual (DWM-00700-00).

Index

SETTING UP MULTIPLE PIICEON 64K MEMORY BOARDS

1.0 INTRODUCTION

The Piiceon SupeRam 2 memory board provides a fairly straightforward and economical way of expanding the AMOS system beyond 64K bytes of memory. Total system memory size is expanded in 64K increments, and user banks may be set up to 16K, 32K or 48K (the SupeRam 2 board has no provisions for setting up banks in smaller increments).

This document assumes that you have read the documents "Memory Management Option" and "The System Initialization Command File" in the "System Operator's Information" section of the AM-100 documentation packet. In particular, you should know how to use the MEMDEF statement in the SYSTEM.INI file to set up user memory banks, and you should know how to use the JOBMEM command to allocate user memory. You should also be familiar with the documentation provided by Piiceon, Inc. that accompanies your SupeRam 2 memory board. You should know how to set a particular row to a specific memory address, and how to set the memory board to an I/O address.

The SupeRam 2 board assumes that the main disk controller board has the PHANTOM option enabled. Also, you may have to modify the AM-100 CPU board in order to run the SupeRam 2 board properly, since the board uses dynamic memory. (NOTE: the modified CPU will work only with SupeRam 2 boards or with fully static memory boards; Alpha Micro makes no claim that it will work with other dynamic memory boards.)

After you enter the appropriate MEMDEF commands in the SYSTEM.INI file, you may allocate memory to jobs by inserting JOBMEM commands into the SYSTEM.INI. There is no restriction to the number of jobs that you may allocate in a single bank. Keep in mind, however, that it is not possible to allocate more memory to a job than remains in a particular bank.

1.1 SETTING UP 48K BANKS

Setting up 48K banks in a 128K system with a 16K sharable area allows you to set up two banks of 48K and an extra bank of 16K. The boards should be set up as follows:

Board	Port	Row	Mem-address	Value to turn row on/off
1	100	A	0K-15K	1
		B	16K-31K	2
		C	32K-47K	4
		D	48K-63K	10
2	101	A	16K-31K	1
		B	32K-47K	2
		C	48K-63K	4
		D	16K-31K	10

The MEMDEF commands in the SYSTEM.INI file read as follows:

```
MEMDEF 100,0,16          ; 48K
MEMDEF 101,7,0           ; 48K
MEMDEF 101,10,0          ; 16K
```

Each additional board added to this configuration is set up at an incrementing port address (i.e., 102, 103,...), and you must set up the rows so that there is no memory overlap. For example, if two more boards were added to this 48K bank system, they would be set up as follows:

Board	Port	Row	Mem-address	Value to turn row on/off
3	102	A	32K-47K	1
		B	48K-63K	2
		C	16K-31K	4
		D	32K-47K	10
4	103	A	48K-63K	1
		B	16K-31K	2
		C	32K-47K	4
		D	48K-63K	10

The MEMDEF commands in the SYSTEM.INI file would look like this:

```
MEMDEF 100,0,16              ; 48K
MEMDEF 101,7,0               ; 48K
MEMDEF 101,10,0/102,3,0      ; 48K
MEMDEF 102,14,0/103,1,0      ; 48K
MEMDEF 103,16,0              ; 48K
```

Note that two of the MEMDEF commands contain specifications for multiple boards in a single bank. This is because the 48K bank size requires allocation of three rows of memory, and each board has four rows of 16K per row.

1.2 SETTING UP 32K BANKS

Setting up 32K banks in a 128K bank system with a 32K sharable area allows you to set up three banks of 32K. The boards should be set up as follows:

Board	Port	Row	Mem-address	Value to turn row on/off
1	100	A	0K-15K	1
		B	16K-31K	2
		C	32K-47K	4
		D	48K-63K	10
2	101	A	32K-47K	1
		B	48K-63K	2
		C	32K-47K	4
		D	48K-63K	10

The MEMDEF command lines in the SYSTEM.INI file appear as follows:

```
MEMDEF 100,0,14          ; 32K
MEMDEF 101,3,0           ; 32K
MEMDEF 101,14,0          ; 32K
```

Each additional board that you add is set up at an incrementing port address (i.e., 102, 103,...), and is configured in the same way as board number two in the example above. The MEMDEF commands in this case are the same as the final two MEMDEF commands in the example above, except that the port address must reflect the address of the new board.

1.3 SETTING UP 16K BANKS

Setting up 16K banks in a 128K system with a 48K sharable area allows you to set up five banks of 16K. Set up the boards as follows:

Board	Port	Row	Mem-address	Value to turn row on/off
1	100	A	0K-15K	1
		B	16K-31K	2
		C	32K-47K	4
		D	48K-63K	10
2	101	A	48K-63K	1
		B	48K-63K	2
		C	48K-63K	4
		D	48K-63K	10

The MEMDEF commands in the SYSTEM.INI file look like this:

```
MEMDEF 100,0,10        ; 16K
MEMDEF 101,1,0         ; 16K
MEMDEF 101,2,0         ; 16K
MEMDEF 101,4,0         ; 16K
MEMDEF 101,10,0        ; 16K
```

Each additional board added to this configuration is set up at an incrementing port address (i.e., 102, 103,...), and is configured in the same way as board number two in the example above. The new MEMDEF commands are the same as the final four MEMDEF commands in the example above, except that the port address must reflect the address of the new board.

1.4 FINAL NOTES

This section explains certain confusing aspects of the SupeRam 2 64K memory board, and talks about what to do when something goes wrong with a memory board in your system.

1.4.1 Power-up/Reset Requirements - All rows in bank number zero must be configured for power-up/reset "on." This means that upon power-up or reset, the rows that are in bank number zero are active; the rows in other banks are to be inactive (power-up/reset "off"). To set a specific row "on," set that row's switch at location 5C on the board to "on." To set a row "off," set its switch to "off."

1.4.2 Explanation of "Value to turn row on/off" - This number is the value
used to turn the specified row "off" for those rows in bank number zero
(generally the first board in the system), or it is the value used to turn
the specified row "on" for those rows in other banks. If multiple rows are
assigned to a specific bank, and all the rows are on the same port, the
values to turn the rows in the bank "on" are all added together.

Note that all of the first memory board (bank number zero) is turned "on"
with a value of 0. In general, the values to turn memory "on" and "off" on
bank number zero are the logical complements of those values used to turn
"on" and "off" the rest of the banks in the system.

1.4.3 The PHANTOM Option - The SupeRam 2 board has provisions for turning
"off" a row of memory when the PHANTOM signal on the S-100 bus is active.
Only row D of the particular memory board is turned "off"; this means that
row D of the board containing bank number zero must be set up for 48K-63K.
PHANTOM has almost no effect on row D on other boards, except that if
PHANTOM is still active when another board's row D is turned on, row D will
remain "PHANTOMed out."

The PHANTOM feature is not jumper-selectable; the SupeRam 2 is always
looking at the PHANTOM signal.

1.4.4 Other Configurations - Since each row is turned "on" by a different
value, and each row can be set up to a particular 16K area of memory, it is
possible to configure the individual memory boards so that the values to
turn the boards "on" (or "off") are different than the examples shown above.
Since the number of these potential configurations is almost endless, it is
up to you, the user, to determine the many possible ways to set up the
SupeRam 2 board.

1.4.5 Memory Errors - It is inevitable that some day, some how, an unlucky
user will have one row of a memory board go bad.   The other rows on the
board may still be functional. If this happens, you may "define out" of the
MEMDEF commands lines the row in question; simply subtract (or, add if bank
number zero) the value to turn that row "on" from the "on" constant in the
particular bank's MEMDEF command line.

Since the memory management system requires that all memory within a bank of
memory be contiguous, you will have to reconfigure the memory addresses of
the specific rows within the bank that contains the bad memory if a row
other than row D is damaged. This may also require changing the MEMDEF
instructions for the bank with the damaged memory. In addition, the JOBMEM
commands for any jobs in the affected bank may have to change.

DEFINING SWITCHABLE SYSTEM MEMORY

The Alpha Micro Operating System (AMOS) allows you to change the configuration of your system by changing the system initialization command file (SYSTEM.INI) that controls the system start-up process. When you add a new disk device or terminal, you add information about that device to the SYSTEM.INI.

As more device controllers and drivers become available, you will probably have a greater variety of devices on your system. Each time you add a new kind of device, however, your monitor becomes larger because the SYSTEM.INI instructs the monitor to incorporate within itself the software necessary to handle the new device. The larger your monitor becomes, the less memory is available to user memory partitions. The area of memory used by the monitor is called system memory.

If your system does not use memory management (that is, if you do not bank switch memory or use the AM-700 Memory Partition Controller), there is nothing you can do about reducing the size of your monitor except remove programs that have been loaded into system memory via the SYSTEM command in the SYSTEM.INI or change the configuration of your system.

Beginning with AMOS Version 4.3, if your system uses bank switching, you may reduce the size of your monitor by placing bitmaps into bank-switched memory. To do this, you must first set aside an area of switchable memory as system memory. (For information on bank-switched systems, refer to Memory Management Option, in the "System Operator's Information" section of the AMOS Software Update Documentation Packet and see the Alpha Micro Integrated Systems User's Guide, (DWM-00101-00).)

Beginning with AMOS Version 4.5A, if your system employs the AM-700 Memory Partition Controller (MPC), you can place bitmaps in what we call "supplemental" system memory, which is similar to switchable memory on a bank-switched system. The procedure for placing bitmaps in supplemental memory on an MPC system is different because you do not use the SYSMEM and MEMDEF commands that are necessary when bank switching. (For information on the MPC, refer to Memory Management With the Memory Partition Controller in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.)

At this time, bitmaps are the ONLY portions of the monitor that you can place in switchable or supplemental system memory.

1.0 THE SYSMEM COMMAND

The first step in defining switchable system memory is to make sure that
your SYSTEM.INI contains valid MEMDEF commands. MEMDEF commands set up
switchable memory banks that allow the system to access more than 64K of
memory. (Individual users are still restricted to a maximum memory
partition of 64K.) The switchable portions of these banks are usually
allocated to user jobs, but you will be setting aside part of this area for
your switchable system memory. (Remember, the SYSMEM and MEMDEF commands
are not used on a system implementing the AM-700 Memory Partition
Controller.)

Now that your system is bank switching memory, you can use the SYSMEM
command in the SYSTEM.INI to tell the system which part of a memory bank
outside of the sharable monitor area you want to assign to the system
bitmaps.

NOTE: When you use SYSMEM in the SYSTEM.INI, you may <u>only</u> use it to define
switchable system memory.

The SYSMEM command takes this form:

        <u>.</u>SYSMEM Bank#:StartAddress-EndAddress (RET)

where Bank# identifies the specific memory bank in which you want to
allocate switchable system memory, and StartAddress and EndAddress select
the beginning and ending addresses of the block of memory you want to
allocate to switchable system memory.

Switchable system memory may not overlap sharable memory.   That is, you
cannot allocate to switchable system memory any memory locations in the
non-switched area of memory. (The non-switched area of memory is used by
the monitor.)

You may not allocate to switchable system memory the last 256 bytes of the
64K address space. (These bytes are reserved for the I/O ports.) The
highest memory address you can allocate, then, is 177376.

Once the system is up and running, you can use SYSMEM at AMOS command level
to find out what area of memory is set aside as switchable system memory.
For example:

        <u>.</u>SYSMEM (RET)
        <u>System memory allocations are:</u>
              3:100000-166774

2.0 DEFINING SWITCHABLE SYSTEM MEMORY

To define switchable system memory on a bank-switching system, edit the SYSTEM.INI and make these changes:

1. First, you must define the area of switchable memory you want to set aside for system memory. Enter the SYSMEM command for this purpose. SYSMEM commands <u>must</u> appear after any MEMDEF commands, but before any BITMAP commands.

    NOTE: You can use more than one SYSMEM command in your SYSTEM.INI if you want to allocate more than one area of switchable system memory.

    SYSMEM uses the same format as the JOBMEM command. For example:

        SYSMEM 2:100000-116150

    The command above allocates 7272 bytes to Bank Two (the third memory bank on the system) for bitmaps. (When you use SYSMEM, remember that the BITMAP commands refer to the number of decimal <u>words</u> you need to reserve for bitmaps, NOT the number of bytes. One word is two bytes, so double the bitmap size to find the number of bytes you need to allocate with the SYSMEM command.)

2. Next, identify those bitmaps you want to place in switchable system memory by including a /S at the end of the appropriate BITMAP command lines in your SYSTEM.INI. For example:

        BITMAP DSK,1818,0,1,2,3,4,5/S
        BITMAP AMS,39,0,1
        BITMAP HWK,606,0,1/S

    The BITMAP commands above tell AMOS to reserve 39 words in the monitor area in Memory Bank Zero for the bitmaps for devices AMS0: and AMS1:. The other two BITMAP commands select the switchable option by including a /S at the end of the command line. AMOS therefore knows that the bitmaps for DSK0:, DSK1:, DSK2:, DSK3:, DSK4:, DSK5:, HWK0:, and HWK1: are to be placed in switchable system memory.

3. At the time of system start-up, AMOS automatically places the bitmaps you have previously designated (via the /S option on the BITMAP command line) into the switchable system area you have defined.

3.0 ERROR MESSAGES

If you make an error in defining switchable system memory, you see one of the following error messages:

?System memory not allocated - monitor memory will be used

You tried to place a bitmap in switchable system memory (via the /S option in the BITMAP command line), but AMOS couldn't find any switchable system memory. (You can also see this message if you did not allocate enough switchable system memory to hold the designated bitmaps.) AMOS therefore places the bitmap in the area of sharable memory used by the monitor. Check to see that SYSMEM commands are present in your SYSTEM.INI.

?Memory allocation format error

SYSMEM didn't understand the format of your SYSMEM command line. For example, did you leave out the colon after the bank number?

?Nonexistent bank number

You've given SYSMEM a bank number larger than the total number of MEMDEF commands in your SYSTEM.INI. (That is, you've referred to a bank number that does not exist.)

?Allocation overlaps monitor memory

You must not allocate to switchable system memory any of the sharable memory area. (Sharable memory is memory that contains the monitor and that all users can access.) Type MEMDEF followed by a RETURN to see the memory bank configuraton for your system. This display tells you which areas of memory you can allocate to switchable system memory.

?Illegal memory range (end is below base)

Ending address of the block of memory you allocate to switchable system memory must be greater than the starting address.

?Allocation is not within requested bank's defined memory

You've specified a valid bank number to the SYSMEM command, but that bank is not addressed for the memory locations you've requested. Check the addressing of your memory boards and check the MEMDEF statements in the SYSTEM.INI.

?Allocation overlaps memory previously allocated to a job

You've already tried to allocate to switchable system memory an area that has already been allocated (via the JOBMEM command) to a user job. Check your bank number and memory addresses in the SYSMEM command. If they're all right, check the memory allocations for the jobs on your system.

?SYSMEM invalid on an MPC system

You tried to use SYSMEM on a system that employs the AM-700 Memory Partition Controller (MPC). On an MPC system, BITMAP allocates switchable regions automatically when a /S is used.

CONFIGURING FLOPPY DISK DRIVERS

1.0 INTRODUCTION

The AMOS system supports several different kinds of disk devices. Because each type of device has its own characteristics and requirements, a separate device driver program must exist for each kind of disk device you use on your system.

The disk driver program links AMOS's generalized disk service routines with the physical disk device. One disk driver exists for each kind of hard disk that you can run under control of the AM-500, AM-400, and AM-410 hard disk controllers.

In the past, several different floppy disk drivers existed: one for each combination of drive type and disk format. A much larger number of floppy disk formats are now available. For this reason, rather than providing a separate disk driver for each possible combination of device type and disk format, we have given you the ability to configure your own floppy disk drivers.

1.1 New Alpha Micro Disk Formats

Before Release 4.2, only three floppy disk formats were available: IBM-compatible (STD), Alpha Micro format (AMS), and Image format (IMG). With the introduction of a new floppy disk controller (the AM-210) that can handle disks in double-sided, double-density formats, a new range of floppy disk formats is now possible. The FIXDVR program allows you to configure your own floppy disk drivers based on the particular disk type, floppy disk controller, and disk format you want to use. You must use FIXDVR to configure a driver for each different combination of device, controller, and format that occurs on your system.

2.0 USING FIXDVR

To run the FIXDVR program, log into the Device Driver Library account, DSK0:[1,6]. Then type FIXDVR followed by a RETURN:

        .LOG DSK0:[1,6] (RET)
        .FIXDVR (RET)

FIXDVR now begins to ask you a series of questions, so that it can determine how to configure a floppy disk driver that matches your particular combination of disk type, disk controller, and disk format:

1. Controller Type (A) AM-200, (B) AM-210 or (C) Icom:

    Enter the letter A, B, or C to select the type of floppy disk controller you are using for your floppy disk drive. If you select the Icom controller, FIXDVR skips down to question #4 (see below).

2. Drive type (A) Persci, (B) Wangco, or (C) CDC:

    Enter the letter that selects the type of disk drive you are using.

3. Double-density?

    FIXDVR asks this question only if you have already specified the AM-210 as your disk controller. Enter a Y for Yes, or an N for No, depending on whether you plan to use the driver on double-density disks.

4. Format (A) STD, (B) AMS, or (C) IMG:

    Enter the letter that selects the disk format you want the driver to use. You may not specify the AMS format if the driver is to use single-density format on a drive running under control of the AM-210 disk controller. If you have previously selected the Icom controller, FIXDVR now skips down to question #6. (If you are using the Icom controller, you may not specify the AMS format.)

5. Double-sided?

    FIXDVR asks this question only if you have already specified the AM-210 as your disk controller. Enter a Y for Yes, or an N for No, depending on whether you plan to use the driver on double-sided disks.

6. Enter new driver name:

    Enter the name that you want to give to the driver program. The standard names that you might want to use are listed below in Section 4.0.

FIXDVR now displays this message:

    New driver is now in memory, bitmap size is nn

The driver that you created is now in memory. Use the SAVE command to save it on the disk. For example:

    .SAVE DDS.DVR (RET)

(Changed 1 May 1980)

The command above saves the driver onto the disk (in the account you are logged into, DSKO:[1,6]) as the file DDS.DVR. If you do not specify an extension, the SAVE command saves the file under the extension .DVR (which indicates a device driver program).

2.1 FIXDVR Error Messages

You may see the following error messages when using FIXDVR:

1. ?Could not find xxxxxx.DVR

     FIXDVR could not find the necessary file. If you are configuring a driver for the AM-200, FIXDVR requires that 200DVR.DVR be in DSKO:[1,6]; a driver for the AM-210 requires 210DVR.DVR in DSKO:[1,6]. A driver for the Icom controller requires that ICMDVR.DVR be in DSKO:[1,6].

2. Please enter Y or N

     Several of the questions that FIXDVR asks require that you answer with a Y or N for Yes or No.

3. ?Invalid reponse

     Several of the questions that FIXDVR asks require that you enter a letter to select an option (e.g., an A to select STD format).

4. ?Invalid device

     You have a bad version of 200DVR.DVR, 210DVR.DVR, or ICMDVR.DVR in DSKO:[1,6].

5. ?Icom does not support AMS format

     You tried to format an Icom floppy diskette in AMS format. (Icom floppy drives only support STD and IMG format.)

6. ?AM-210 does not support single-density AMS format

     You may not use single-density AMS format on a device that runs under the control of the AM-210 floppy disk controller.

7. ?AM-200 does not support CDC floppy disks

     You may only run CDC floppy disks under the control of the AM-210 floppy disk controller.

3.0 MODIFYING THE SYSTEM INITIALIZATION COMMAND FILE

You must modify the system initialization command file (SYSTEM.INI) to:

1. Add the device for which you have just defined a device driver to
    the DEVTBL command line. NOTE: If you have too many devices on the
    DEVTBL command line to fit all on one line, you may use several
    DEVTBL command lines. For example:

        DEVTBL DSK1,AMS0,AMS1,AMS2,AMS3,DDS0,DDS1
        DEVTBL TRM,MEM,RES,/MTM

2. Add BITMAP commands to the SYSTEM.INI to define bitmaps for the new
    devices you are adding to the system. Use the bitmap size given in
    the final FIXDVR message. For example, if you are defining a
    driver for a two-drive device, and the format requires a bitmap of
    39, you might have seen this message:

        New driver is now in memory, bitmap size is 39

    Now you must add the appropriate BITMAP command to the SYSTEM.INI.
    For example:

        BITMAP AMS,39,0,1

4.0 STANDARD DISK DRIVER NAMES

Below is a list of some of the standard names you can assign to the device
driver defined by FIXDVR:

Driver Name	Characteristics
STD	Single-density, Single-sided, STD format
SDS	Single-density, Double-sided, STD format
DSS	Double-density, Single-sided, STD format
DDS	Double-density, Double-sided, STD format
AMS	Single-density, Single-sided, AMS format
SDA	Single-density, Double-sided, AMS format
DSA	Double-density, Single-sided, AMS format
DDA	Double-density, Double-sided, AMS format
SSI	Single-density, Single-sided, IMG format
SDI	Single-density, Double-sided, IMG format
DSI	Double-density, Single-sided, IMG format
DDI	Double-density, Double-sided, IMG format

(Changed 1 May 1980)

5.0 FORMATTING DISKS

When you format a disk, you will use either the FMT200 or the FMT210
formatting programs (depending on whether you are running the device
containing that disk under control of the AM-200 or the AM-210 Floppy Disk
Controller).

NOTE: If you are using an AM-200 controller, make sure that the AM-200
format-enable switch is turned ON before you format a disk. The Icom
controller is not capable of formatting diskettes; if you are using an Icom
controller, you must buy preformatted disks.

The device specification you give to the formatting program will identify
the driver (and thus the format) used for that device. For example, if you
want to format a disk in Drive One of a device that uses the DSS driver
(double-density, single-sided, STD format), enter:

         .FMT210 DSS1:[RET]

FMT210 then formats the disk in the proper format, and you see:

         BEGIN FORMATTING
         EXIT
         .

DISK LABELING PROCEDURES

Although removable disk cartridges provide an extremely convenient and portable method of backing up and exchanging data, problems sometimes occur in identifying just what information is on a cartridge if that disk pack does not contain a label on the outside of the pack, or if it is improperly labeled.

To solve this problem, Alpha Micro has developed a set of software which allows you to establish and verify a permanent, identifying label on each disk. The LABEL program writes data to a disk that serves as identification for that disk. Other programs (e.g., MOUNT, XMOUNT.SBR, or your own assembly language programs) read the disk identification information on the disk and tell you which disk is currently mounted.

Being able to permanently label a disk results in several advantages:

1. You may easily determine what disk is mounted, even if you are using the system from a remote site.

2. Your programs can verify that the correct disk is mounted before they begin to change data on that disk.

At the current time, the disk labeling system consists of these new and modified programs:

LABEL   Labels a disk by writing identifying data onto Block Zero of that disk. Also displays a disk's label.

MOUNT   Displays the name of the disk that was just mounted or a list of all mounted disks.

XMOUNT   A BASIC subroutine that your BASIC programs can use. It returns the contents of the disk identification field so that your BASIC programs can verify that the correct disk is mounted.

In future releases, additional programs will make use of the disk label.

This document contains information on the LABEL and MOUNT commands. Also, see the MOUNT and LABEL reference sheets in the AMOS System Commands Reference Manual, (DWM-00100-49). For information on XMOUNT, refer to XMOUNT - Basic Subroutine to Mount a Disk, in the "BASIC Programmer's Information" section of the AMOS Software Update documentation packet.

The last section of this document describes the exact format of the disk label, so that your assembly language programs can access the label information.

1.0 LABEL

The LABEL program gives you a way to label a disk with descriptive information and to display that information. Disk labels are stored in Block 0 of the disk, and are used to allow both you and your programs to verify that the correct disk has been mounted.

There are three situations in which you will want to use LABEL:

1. A disk has never been labeled and you want to give it a label.

2. A disk has a label and you want to display that information.

3. A disk has a label and you want to change it.

1.1 Labeling a Disk

To give a disk a label, enter LABEL followed by the specification of the logical device that holds the disk; then type a RETURN. For example, if you want to label the disk in logical device DSK5:, enter:

        .LABEL DSK5: (RET)

Now LABEL asks you for the following information:

Volume Name:
        Enter up to 40 characters that describe the disk. The MOUNT program displays this field when it gives the list of the mounted disks on the system. The purpose of this information is primarily to give you a way to identify the contents of the mounted disk.

Volume ID:
        Enter up to ten characters as the Volume ID. This field is used by programs to determine if the proper disk has been mounted. The XMOUNT subroutine (an assembly language routine callable by your BASIC programs) returns this field. The MOUNT program displays this field (as well as the Volume Name) when it lists the disks mounted on the system.

Installation:
        Enter the name of your installation or company. This field, which may contain up to 30 characters, identifies the site where the disk was created. This information will be particularly useful if you exchange disks among different installations.

<u>System:</u>
Enter the name of the computer system this disk was created on. This field may be up to 30 characters in length, and is especially useful when an installation has more than one computer system.

<u>Creator:</u>
The name of the person who created the disk. This field may contain up to 30 characters.

When you finish entering the requested information, LABEL returns you to AMOS command level. The disk is now labeled with the information you specified. The label also contains the date on which you labeled the disk.

1.2 Displaying a Disk Label

Once a disk contains a label, you may use the LABEL command to display that label. Just enter LABEL followed by the specification of the disk whose label you want to see. Now type a RETURN. For example:

    .LABEL DSK3: (RET)
    Currently labeled as:

    Documentation Archives (ARCHIVE01)
    Created on 1-Jan-80 at Alpha Microsystems on System B by Jack Smith
    Last access: 5-Apr-80

    Volume Name: ^C

The creation date is the date the disk was labeled; the date of last access is the date the disk was last mounted.

After LABEL displays the label information, type a Control-C to return your terminal to AMOS command level.

1.3 Changing a Disk Label

To change a disk label, enter LABEL followed by the specification of the disk whose label you want to change. Now type a RETURN. For example:

    .LABEL HWK1: (RET) ·
    Currently labeled as:

    Test Data (TESTDATA01)
    Created on 17-Apr-80 at Computer Products, Inc. on System 1 by Anne B.
    Last access: 21-Apr-80

    Volume Name:

Now, LABEL asks you for the new labeling information. Answer each question with the new information you want to place in the label. (See Section 1.1, "Labeling a Disk," for information on these fields.) When you enter the last question, LABEL returns your terminal to AMOS command level, and the disk label contains the new information.

1.4 LABEL Error Messages

If you specify a nonexistent device, you see:

> ?Cannot INIT Devn: - device does not exist

where Devn: is the disk specification you gave. Check your spelling and try again.   If you still see this message, use the DEVTBL command to see a list of the valid devices on the system.

If the LABEL command line is not in proper format (for example, if you type LABEL followed by a RETURN), you see:

> ?File specification error

Check the format of your command line and enter it again. Make sure that you include the colon after the device specification.

2.0 MOUNTING A LABELED DISK

Whenever you change a floppy disk or a hard disk cartridge (whether labeled or not), you must always use the MOUNT command to inform AMOS that the bitmap in memory for that device is no longer valid. If you do not mount a disk when you change it, AMOS has no way of knowing that it may be using the wrong bitmap when it writes data to that device; severe damage to the data on the disk could result. NOTE: Never mount or unmount a disk when other users are accessing that disk; you will damage the file structure on the disk by doing so.

You may use MOUNT to mount a disk, unmount a disk, or to display a list of all mounted disks on the system. NOTE: As of AMOS Release 4.4, you may use the /W switch to tell MOUNT to wait until the specified device is ready before mounting the disk. For more information on using MOUNT to mount or unmount a disk, see the MOUNT reference sheet in the AMOS System Commands Reference Manual, (DWM-00100-49).

When MOUNT sucessfully mounts a disk, if that disk has a disk label stored in Block 0, you see:

> x (y) mounted

where "x" is the Volume Name, and "y" is the Volume ID specified when the
disk was labeled.   (For more information on these two fields of the disk
label, see Section 1.1, "Labeling a Disk," above.)
For example:

>   .MOUNT DSK1: (RET)
>   System Disk (SYS001) mounted

If you want to see a list of all disks that are mounted on the system, type
MOUNT followed by a RETURN. For example:

>   .MOUNT (RET)
>   Disks mounted:
>
>   DSK0:   System Disk (SYS001)
>   DSK1:   Payroll Data (PRD001)
>   DSK2:   Payroll Data (PRD002)
>   DSK3:   Development Disk (DEV001)
>   DSK4:
>   DSK5:   Backup Disk #3 (BCK003)
>   AMS0:   Transfer Disk (TRN001)

Each line of the display gives the following information: device
specification, Volume Name, and Volume ID. (For example, the Volume Name of
the disk in device DSK0: is "System Disk"; its Volume ID is "(SYS001)".)

3.0 CONTENTS OF THE DISK LABEL

Systems programmers may be interested in the exact format of the disk label.
Note that not all of the defined fields are presently used by the disk
identification software; these unused fields will be used by future releases
of Alpha Micro software.

Field	Size	Contents
Header	2 words	125252 : 052525
Volume Name	40 bytes	ASCII text
Volume ID	10 bytes	ASCII text
Creator	30 bytes	ASCII text
Installation	30 bytes	ASCII text
System Name	30 bytes	ASCII text
Creation Date	4 bytes	System date format
Access Date	4 bytes	System date format
Backup Date #1	4 bytes	System date format
Backup Vol. ID #1	10 bytes	ASCII text
Backup Date #2	4 bytes	System date format
Backup Vol ID #2	10 bytes	ASCII text

The fields are used as follows:

Header
: Used to flag that this disk is labeled. If the flag words are not correct, certain programs will ignore the label.

Volume Name
: An ASCII string that describes the disk. This field is designed to be a description of the contents of the disk (e.g., Archives Disk - Jan/March 1980).

Volume ID
: A short ASCII string that describes the disk. This field is used by programs checking for the proper disk being mounted (e.g., ARCHV1).

Creator
: An ASCII string that describes who created this disk (e.g., John Hoolihan).

Installation
: An ASCII string that names the site where the disk was created (e.g., Acme Computers, Inc.). This is useful when installations interchange disks.

System Name
: An ASCII string that gives the name of the particular computer system, within an installation, on which the disk was made (e.g., Purchasing).

Creation Date
: The original date on which the disk was labeled.

Access Date
: The date the disk was last MOUNTed.

Backup Date #1
: The date of the last backup (the "father" backup). This field gives the date of the most recent backup. We do not use this field at this time, but it is reserved for future use.

Backup Volume ID #1
: The volume ID field of the disk on which the most recent backup (the "father" backup) exists. This field is not used at this time, but it is reserved for future use.

Backup Date #2
: The date of the "backup before last" or "grandfather" backup. This field is not used at this time, but is reserved for future use.

Backup Volume ID #2
: The volume ID field of the disk on which the "grandfather" backup exists. This field is not used at this time, but it is reserved for future use.

DISK MAINTENANCE PROCEDURES FOR THE SYSTEM OPERATOR

October 1981
Revision A05

This document reflects AMOS versions 4.6 and later

Table of Contents

1.0 NOTE TO THE SYSTEM OPERATOR

This document discusses some of the maintenance procedures that you, as the System Operator, must perform on the disks used on the system. We've aimed this discussion at the System Operator (the person in charge of the details of system administration and maintenance) because we assume that most general users of your system will not be concerned with disk maintenance, and because many of the commands we talk about below can be fatal to your data if used carelessly or without full understanding of command operation.

The next few pages tell you how to: 1) format disks; 2) initialize disks; 3) label and identify disks; 4) allocate and change user accounts on a disk; 5) pack a disk; 6) perform disk diagnostic tests; 7) recover from disk errors; and, 8) perform disk backup.

The System Operator should establish a regular schedule for disk backup and disk diagnostic tests. For example, an hour every Tuesday and Friday morning might be set aside for disk diagnostic tests and disk backup. If you are changing or creating a lot of data or if you are particularly concerned about the security of your data, you might want to back up your disks once or twice every day.

Because many of these procedures must be done when only the System Operator's job is accessing the disks, it is wise to schedule them during a time when other users are off the system (for example, early morning or weekends).

<u>NOTE:</u>

Previous versions of this document used the term "file record." However, the use of the word "record" can cause some confusion because it is sometimes used in other documents to mean several different things. To help make our terminology more consistent, we have adopted the following conventions:

A <u>physical record</u> is the sector on the disk. This is the actual, physical grouping of data on the disk. The hard disk devices currently supported by Alpha Micro use a physical record size of 512 bytes. Physical record sizes for floppy disk devices vary depending on the device, and may range from 128 bytes to 512 bytes.

A <u>disk block</u> is the logical grouping of data on the disk that AMOS uses when reading from and writing to the disk. AMOS always transfers data one disk block at a time. A disk block may be made up of one or more physical records. Each disk block has a number associated with it that AMOS uses to reference that block. Disk blocks (except in the special case of devices that use the IMG device driver) are always 512 bytes long.

(Changed 31 October 1981)

A logical record is the logical grouping of data on the disk as structured by your programs, and has little to do with the physical records or disk blocks on the disk, except that a logical record may not be larger than a disk block. (For example, a BASIC program might set up a data file in which every logical record is just large enough to contain customer addresses and names-- 60 bytes, for instance.)

1.1 Important Note to Phoenix Hard Disk Drive Users

As of AMOS Version 4.4, bad block handling has changed for devices that run under control of the AM-410 (e.g., the Phoenix hard disk drive). Make sure that you read the document AMOS Version 4.4 Method of Handling Bad Blocks in the "System Operator's Information" section of the AMOS Software Update Documentation Packet before you use DSKCPY, CRT410, or DSKANA on a Phoenix disk, if that disk was certified by a pre-4.4 version of CRT410.

1.2 Maintaining a Disk

There are four stages in the life of a disk:

1. When it is brand new, you set a disk up for initial use. First you format the disk; then you initialize it. IMPORTANT NOTE: Do not format or initialize disks that run under control of the AM-410 Hard Disk Controller. You cannot format disks that run under the control of the AM-420 either (but later, you must use the SYSACT I command to initialize all logical units on a single Winchester physical device except the first logical device). Instead, use the CRT410 command or the CRT420 command, respectively, to certify those disks. See the documents Software Notice for AM-410 Users or Software Notice for AM-420 Users in the "System Operator's Information" section of the AMOS Software Update Documentation Packet for more information about certifying either of those disks.

    Next you may decide to use the LABEL program to write an identifying label to the disk. (Of course, even if you don't use LABEL, you will always put a physical label on the outside of a hard disk cartridge or a floppy diskette to identify that disk.)

    Finally, you will need to establish the account structure on the disk by allocating user accounts.

2. After you have set the disk up, you use it. That is, you transfer files to the disk as well as create new files on it. During this time you may allocate additional user accounts and change existing ones. You may also use LABEL to re-label the disk as the need arises.

3. One way to maintain and protect the data on the disk is by performing frequent backups. You should be especially sure to back up your system disk onto a medium that will enable you to boot the system if your system disk cannot perform the function for some reason. Obviously, the sooner you do that particular backup the better. The medium you backup on may simply be a removable hard disk cartridge or a floppy disk, or it may be on a Video Cassette Recorder (VCR) or a 1/4" streaming tape drive. If you have a disk that runs under the control of the AM-420, you may use the WRMGEN program to create a warm boot monitor tape. For more information on warm boot monitors see the WRMGEN reference sheet in the AMOS System Commands Reference Manual (DWM-00100-49). For more information on the VCR and the 1/4" streaming tape drive see the manuals Software Information for the AM-610 Video Cassette Recorder Interface (DSS-10000-24), and Software Information for the AM-620/621 1/4" Streaming Tape Drive (DSS-10000-23).

    Beside backing up to maintain and protect your disk's data, remember to run diagnostic tests that look for device and media problems. If diagnostic tests indicate problems, you may have to reconstruct the data on a damaged disk. The sooner you catch damaged data, the better your chances of reconstructing it.

4. After a disk has been in use for some time, you may want to recycle the disk by initializing it (or re-certifying it in the case of those disks that run under control of either the AM-410 or the AM-420); this clears all data from the disk. Before re-using a disk, however, run some disk diagnostic tests on it to be sure that the disk media is healthy. After initializing or re-certifying the disk, the old account structure is gone, so allocate new user accounts using the SYSACT program.

2.0 FORMATTING DISKS

A new floppy or hard disk must be formatted and initialized (or certified) before you can write data to that disk. Formatting a disk sets the disk up so that it is organized into a specific pattern; it is then ready to receive data written in that same pattern. Whenever the system performs an action on a file, it knows the format it must use to read or write that file because of the device specification you include in the specification of the file. For example:

        .COPY DSK1:=AMSO:NEWFIL.TXT (RET)

tells the system to read the file NEWFIL.TXT in AMS format (since you specified device AMS0:), and copy it to Drive One of the System Device (in whatever format is used by DSK1:). Disks to be used in device AMS0: must be formatted in AMS format; disks to be used in device DSK1: must be formatted in whatever format is used by the System Device.

You only need to format disks that have never been formatted before or whose format you wish to change (e.g., to change a floppy disk from STD to AMS format). New floppy disks usually come preformatted in STD format. You should NOT format these disks if you plan to use them in STD format. (When you format a disk in a disk drive, that disk may take on certain characteristics of that drive. So, when you buy a disk already formatted, you should refrain from formatting the disk if you can; this helps to ensure that you can read that disk on different disk drives.)

ICOM floppy disk users may not format disks because the ICOM disk controller does not support disk formatting. ICOM floppy disk users must therefore buy their disks preformatted.

NOTE: Formatting a disk destroys all data on that disk.

2.1 Disk Certification

Before we continue with our discussion of disk formatting, it is important to mention again this important warning: if a physical disk runs under control of either the AM-410 (e.g., a Phoenix disk) or the AM-420 (e.g., a Winchester technology disk), you must not format that physical disk. Instead, you must certify it via the CRT410 or CRT420 commands. For more information about certifying these types of disks, see the documents Software Notice for AM-410 Users or Software Notice for AM-420 Users, both in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.

2.2 Disk Formats

Alpha Micro supports several different disk formats. For information on the disk formats available, and for information on converting floppy System Devices from one format to another, refer to the documents titled Disk Drivers and Formats, Configuring Floppy-Disk Drivers, and Defining Non-System Disk Devices, in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.

(Changed 31 October 1981)

2.3 The Disk Formatting Programs

After you have decided on the format you want your disks to use, you must use the proper formatting program to format those disks.   The major disk formatting programs used by Alpha Micro are:

FMT500   The formatting program for disk devices used under the control of the AM-500 Hard Disk Controller (e.g., the CDC Hawk hard disk).

FMT200   The formatting program for disk devices used under the control of the AM-200 Floppy Disk Controller (e.g., Persci and Wangco floppy disks).   Formats single sided, single density floppy disks.

FMT210   The formatting program for disk devices used under the control of the AM-210 Double Density Floppy Disk Controller (e.g., Wangco disks). Formats single- and double-sided, and single- and double-density floppy diskettes.

FMT400   The formatting program for disk devices used under the control of the AM-400 Hard Disk Interface (e.g., Century Data Trident disks).   Formats each logical unit as a separate device.

To use one of the formatting programs listed above, enter the name of the appropriate formatting program followed by the specification of the device holding the disk you want to format. Your device specification tells the formatting program which disk driver program you are going to be using on the disks in that device, and therefore which disk format to use. (For information on configuring a floppy disk driver for your particular combination of disk controller, disk drive and disk format, see Configuring Floppy Disk Drivers in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.) Suppose that you are using a Persci drive to read and write AMS-format floppy disks. To format a disk in Drive Zero in AMS format, enter:

     .FMT200 AMS0: (RET)

The formatting program now says:

     BEGIN FORMATTING

When the program has finished, you see:

     EXIT
     .

The formatting programs FMT200, FMT210, FMT400, and FMT500 all work in the same way as our example above. FMT500 requires that you mount the disk before you format it (use the MOUNT command); the other formatting programs do not require that the disk be mounted.

(Changed 31 October 1981)

IMPORTANT NOTE: No other job may run on the system while you format a disk, so make sure all users are off the system before you use one of the formatting programs listed above.

3.0 LABELING AND IDENTIFYING A DISK

Of course you will always place a physical label on the outside of a disk pack or floppy diskette to identify that disk. You may also write a label to the disk itself. The LABEL program allows you both to establish and display a disk label.

The disk label helps you to identify the disk. Another method of checking the contents of the disk is to create a hash total for that disk. (A hash total is a number that uniquely identifies a group of data.) The HASHER program allows you to generate a hash total for a specific disk. The hash totals for two disks will only be the same if the contents of those disks are identical.

3.1 Labeling the Disk

Labeling a disk results in two advantages: you can determine what disk is mounted, even if you are at a remote site, and your programs can check to be sure that the proper disk is mounted before changing data on that disk.

To label a disk, enter LABEL followed by the specification of the device that contains the disk you want to label. For example:

        .LABEL DSK3: RET

Now LABEL asks you for the following information:

Volume Name:
    Enter up to 40 characters. This field describes the disk (e.g., Payroll Data from Jan-March 1980). The MOUNT program displays this information when it mounts the labeled disk.

Volume ID:
    Enter up to ten characters as the Volume I.D. This field is used by programs to determine if the proper disk has been mounted. (MOUNT displays this field, too, when it mounts the disk.)

Installation:
    Enter the name of your installation or company. This field may contain up to 30 characters, and identifies the site where the disk was created (e.g., CompuWord, Inc.).

System:
     Enter the name of the computer system this disk was created on.
This field may be up to 30 characters in length.

Creator:
     Enter the name of the person who created the disk. This field may
contain up to 30 characters.

When you finish entering the required information, LABEL returns you to AMOS
command level.

Once a disk contains a label, you can use LABEL to display that information.
Type LABEL followed by the specification of the device that contains the disk
whose label you want to see; then type a RETURN. For example:

     .LABEL HWK1: (RET)

Since the disk already has a label, LABEL displays that information for you.
For example:

     Currently labeled as:

     Documentation Archives (ARCHIVE01)
     Created on 12-APR-80 at BeachTogs, Inc. on System C by J.K.Milne
     Last access: 8-JUN-80

     Volume Name:
     ^C

Notice that LABEL incorporates into the disk label the date the disk was
labeled (creation date) and the date the disk was last mounted ("last access").
Now LABEL begins to ask you for the various label fields. If you do not want
to change any of the information in the label, type a Control-C. Otherwise,
you may answer the questions as in the paragraphs above.

For more information on LABEL, refer to Disk Labeling Procedures in the
"System Operator's Information" section of the AMOS Software Update
Documentation Packet.

3.2 Identifying a Disk

The HASHER program allows you to generate a hash total for a specified disk.
This gives you a number that you can compare against the hash total of another
disk to see if the disks contain identical data. You will probably find HASHER
to be particularly helpful when you are making multiple copies of a disk via
DSKCPY. If the hash total of the master disk does not match the hash totals of
the disks copied to, the copies are not perfect.

To use HASHER, enter HASHER followed by a RETURN:

     .HASHER (RET)

When the system asks for the input drive, enter the name of the disk for which you want hash totals:

Input drive: DSK2:

Now you see:

[Hashing nnnn blocks]

where nnnn is the number of blocks on the specified disk.

When HASHER is finished, you see:

Hash is: nnn

where nnn is the hash total for the specified disk.

NOTE TO HAWK HARD DISK USERS: When you use DSKCPY on a Hawk disk, DSKCPY allows you to use the special fast copy mode or the slower /O mode. (If you use the DSKCPY fast copy mode for Hawk devices, no other user may run on the system while the copy is taking place; the /O mode, while slower, allows users to continue running on the system.) DSKCPY optionally generates a hash total for a disk, but it generates the hash total differently depending on whether you are using the fast copy mode or the /O mode. That means that the hash total differs for the same disk, depending on the mode in which the disk copy was made.

To allow you to use HASHER to generate a disk hash total that is compatible with one generated by DSKCPY, HASHER also has a Hawk fast copy and a /O mode. Just as with DSKCPY, HASHER does not allow you to run other users on the system if you use the default Hawk fast copy mode. You see:

%All other users will be suspended while HAWK hash is running.
Hit return to continue or control-C to abort:

If no other users are on the system, you may type a RETURN. If you want HASHER to generate a hash total in the same way that DSKCPY does when it uses the /O mode, use the HASHER /O switch. For example:

.HASHER /O (RET)

You may use HASHER with the /O switch while other users are running on the system. For more information on HASHER, see the HASHER reference sheet in the AMOS System Commands Reference Manual, (DWM-00100-49).

4.0 USING THE SYSACT COMMAND

The SYSACT command allows you to perform several disk maintenance functions. You may use SYSACT to: 1. initialize (that is, clear) a disk; 2. allocate user accounts on a disk; 3. change account passwords; 4. delete user accounts; and, 5. display a list of accounts (and passwords) on a disk.

(Changed 31 October 1981)

4.1 Initializing a Disk

After you have formatted a disk, you must initialize it if you are going to build an account structure on that disk. If you are simply going to transfer a literal image of another disk onto the newly formatted disk (via DSKCPY) you do not need to initialize it. (Remember that if you are using a disk that runs under control of the AM-410, you must not format or initialize it. Remember also that you must not originally use SYSACT to format or initialize a disk that runs under control of the AM-420, but should use the CRT420 program to certify and format the disk; however, once it is formatted and you have defined the logical devices you want to use in the device table (see the DEVTBL reference sheet), you <u>must</u> initialize all but the first logical device by using SYSACT.)

Initializing a disk writes zeros in the Master File Directory and the bitmap. If you wish to clear the data on a disk, you do not need to reformat it (unless you want to change the format that it uses); just re-initialize it.

NOTE: Once you have initialized a disk, there is no way to access the data on that disk-- you have, in effect, erased the disk.

To initialize a disk:

1. If the disk drive you are using allows it, write-protect all devices that you are not going to initialize.

2. Log into the System Operator's account:

    .LOG [1,2] (RET)

3. Mount the disk you want to initialize. Type MOUNT followed by the specification of the device holding the disk, then type a RETURN:

    .MOUNT AMS1: (RET)

4. Type SYSACT followed by the specification of the device that holds the disk you want to initialize:

    .SYSACT AMS1: (RET)

    You now see the SYSACT prompt, *.

5. Be very sure that you have specified the correct device. To make sure the disk is empty, use the SYSACT L command; type an L and a RETURN:

    *L

    If the disk has already been initialized and is in use, you see a list of PPNs and passwords.

    If the device you specified does not hold the disk you want to initialize, type an E and a RETURN to exit SYSACT and return to AMOS command level.

(Changed 31 October 1981)

6. If you have specified the correct device, type an I followed by a
    RETURN:

            *I [RET]

7. SYSACT now tells you:

        Initializing the disk clears all files - enter Y to confirm:

    SYSACT initializes the disk only if you enter a Y followed by a
    RETURN. When it has finished, SYSACT displays its prompt symbol.

8. To exit SYSACT and return to AMOS command level, type an E followed by
    a RETURN:

            *E [RET]
            .

4.2 Building the Disk Account Structure

After you have initialized a disk, that disk is now ready to receive data. You
can now begin to build accounts and files on the disk. Use SYSACT to allocate
user accounts.

All files on a disk are associated with an account on that disk. Usually, you
may not write any files to the disk until a user account exists to hold those
files.   (However, you can copy files onto an initialized disk that has no
account structure if you copy from the System Operator's account, [1,2].   In
this case, the COPY command allocates the proper accounts for you on the new
disk as it copies over the files.)

You can only access a file if you: 1. log into the account which contains the
file; or, 2. specify the account the file belongs to within that file's
specification.

Each account on the disk has a directory associated with it (called a User File
Directory or UFD) that lists the files in that account. Every disk has one
Master File Directory (called the MFD) that maintains a list of all UFDs on
that disk. When you first initialize a disk, SYSACT creates the MFD on the
disk, but no UFDs exist. As you use SYSACT to allocate user accounts, SYSACT
creates the UFDs for those accounts.

4.2.1 Allocating User Accounts

1. LOG into the System Operator's account:

            .LOG [1,2] [RET]

2. Mount the disk you want to use the SYSACT command on.

> .MOUNT DSK1: (RET)

3. Type SYSACT followed by the specification of the device holding the disk you want to allocate user accounts on. Then type a RETURN.

> .SYSACT DSK1: (RET)

You now see the SYSACT prompt, *.

4. To see a list of the SYSACT commands you can use, type an H followed by a RETURN:

> *H (RET)

5. To allocate an account, type an A followed by a space and the project number and programmer number of the account you want to allocate. (Separate the numbers with a comma.) These numbers must be octal, and must not be greater than 377. For example, suppose you want to allocate the user account DSK1:[100,1]:

> *A 100,1 (RET)

Now SYSACT asks you if you want to assign a password to that account:

> PASSWORD:

You may enter a password of six characters or less. If you enter just a RETURN, SYSACT assigns no password to the account.

6. If you want to allocate another account, use the A command again. Follow the procedure above until you have added all accounts that you wish to allocate.

7. To exit from SYSACT and return to the monitor, enter an E after the prompt; then type a RETURN:

> *E (RET)
> .

You may not create more than 63 user accounts on a single disk; that is the maximum number of entries in the Master File Directory.

4.2.2 Changing and Deleting User Accounts

Once you have created user accounts on a disk, you can begin to create and transfer files on that disk. At any time you may change or delete user accounts by using the SYSACT command.

NOTE: Always erase any files in an account before deleting it. SYSACT won't let you delete an account if there are files in that account.

1. Log into the System Operator's account:

        .LOG [1,2] (RET)

2. Mount the disk on which you want to use SYSACT:

        .MOUNT AMS1: (RET)

3. Type SYSACT followed by the specification of the device that contains
    the disk whose accounts you want to change or delete. Type a RETURN:

        .SYSACT AMS1: (RET)

    Now you see the SYSACT prompt, *.

4. If you want to see a list of all accounts already allocated on that
    disk, type an L followed by a RETURN.

5. To change the password of an account:

    a. Enter a C followed by a space and the project-programmer number of
        the account whose password you want to change. For example, to
        change the password of account [110,6]:

            *C 110,6 (RET)
            PASSWORD: MILO (RET)

        The new password for account [110,6] is now MILO.

    b. To remove a password entirely from the account, enter just a
        RETURN:

            *C 110,6 (RET)
            PASSWORD: (RET)

6. To delete a user account, enter a D followed by a space and the
    project-programmer number of the account you want to delete:

        *D 110,6 (RET)

7. To exit from SYSACT, type an E followed by a RETURN:

        *E (RET)
        .

5.0 DISK DIAGNOSTIC TESTS

There are a number of diagnostic tests you can run to check your disk media,
the disk controller and the physical device itself. (NOTE: As of AMOS Version
4.4, you may use any of the disk diagnostic programs below on a disk that
runs under control of the AM-410.)

Before running any of the tests that we discuss below, it is a good idea to use
the SET DSKERR command:

        .SET DSKERR (RET)

If you do not use SET DSKERR, the system reports only hard errors. Once you
have SET DSKERR, the system reports any soft errors that occur, and tells you
at what disk location the error occurred. (NOTE: The system makes an exception
for hard disks used with the AM-500 Hard Disk Controller. Even if you use SET
DSKERR, the system reports only hard errors for such devices. However, using
SET DSKERR does tell the system to report the disk location at which the hard
error occurred.)

REMEMBER: SET DSKERR only affects error reporting for the job that used the
SET command.

A soft error is a read-error. When a soft error occurs, the system has to
retry reading the data in a specific disk location. The system does not report
soft errors unless you use SET DSKERR. When a set number of soft errors have
occurred at the same disk location (usually eight), the system reports a hard
error. An occasional soft error is not in itself an indication of serious
problems, but frequent soft errors may indicate maladjustments in the physical
device or disk controller, or problems with the disk media itself. The
particular message the system uses to report a soft error depends upon the type
of device; check with Appendix A of the AMOS User's Guide, (DWM-00100-35),
for a list of soft error messages. For example, if the system had to retry
reading a disk block on a floppy disk drive five times, you might see something
like this:

            CRC Error — AMS1: Block 145
            CRC Error — AMS1: Block 145
            CRC Error — AMS1: Block 145
            CRC Error — AMS1: Block 145
            CRC Error — AMS1: Block 145

A hard error occurs when the system has repeatedly tried to read the same disk
location, but has failed to do so. A hard error is a serious matter, since it
indicates that the system has given up trying to read the disk block affected.

If any of the disk diagnostic tests are not able to complete an analysis
because of a hard error, they tell you so:

        ?Cannot READ Filespec — device error

If you have used the SET DSKERR command, the system tells you where on the disk
the hard error occurred. For example:

AM500 ERROR CODE 4 FOR DRIVE 1 BLOCK 12 (CYLINDER 0 HEAD 0 SECTOR 12)

To see what the error codes for a specific disk drive mean, consult the hardware documentation that accompanied that drive. (For example, the error message above occurred on a CDC Hawk hard disk running under control of the AM-500 Hard Disk Controller. For that specific disk drive, an error 4 is a CRC Error-- a Cyclic Redundancy Check error.)

For information on recovering from disk errors reported by the diagnostic tests, refer to Section 6.0, "Recovering From Disk Errors."

5.1   REDALL and RNDRED

Both REDALL and RNDRED perform read tests on a specified hard or floppy disk. REDALL reads all disk blocks (or the number you specify) beginning with the first block on the disk. RNDRED performs random-read tests.

Neither REDALL nor RNDRED alter the data on your disk; they merely read the data and report any read errors that occur.

To use REDALL:

1.   Enter REDALL followed by the specification of the device you want to read. Type a RETURN.

         .REDALL DSK1:   (RET)

     The command above tells REDALL to read all blocks on the specified disk.

2.   If you don't want REDALL to read all blocks on the disk, follow the disk specification with the number of blocks you want read:

         .REDALL DSK2:100   (RET)

     (Do not put a space between the device specification and the number of blocks.) The command above tells REDALL to read the first 100 blocks on DSK2:.

3.   REDALL now tells you the number of blocks it is reading:

         .REDALL DSK3:   (RET)
         Reading 9696 blocks
         EXIT
         .

4.   REDALL exits when it finishes reading the blocks. If any errors occurred, REDALL tells you so by displaying the appropriate error message on the screen.

To use RNDRED:

1. Enter RNDRED followed by the specification of the device you want to
    test. Type a RETURN:

         .RNDRED AMS1: [RET]

    RNDRED now randomly selects a disk track and performs a seek and read
    operation on a random block of that track. RNDRED continues on,
    selecting and reading disk locations at random. You see nothing on
    the screen unless RNDRED finds an error.

2. When you wish to exit RNDRED, type a Control-C; otherwise, RNDRED
    continues until you reset the system.

RNDRED and REDALL can display the usual system error messages that result from
an invalid device specification. For example:

         .RNDRED DSK2: [RET]
         ?Cannot READ DSK2: - device not mounted

         .REDALL ASM2:200 [RET]
         ?Cannot READ ASM2: - device does not exist

5.2 DIAG2

DIAG2 tests floppy disks by performing read/write tests. It does not verify
write operations and does not destroy the data on your diskette.

To use DIAG2:

1. Enter DIAG2 followed by the specification of the device you want to
    test. Type a RETURN:

         .DIAG2 STD0: [RET]

2. DIAG2 waits until you are ready. Then it proceeds with the test:

                     Hit return when ready: [RET]
                     Test 1 - track 0 read/write
                     Test 2 - track 76 read/write
                     Test 3 - random seek-verify 500 times
                     Test 4 - speed seek tracks 0 and 76 10 times
                     EXIT
                     .

If any errors occur, DIAG2 displays the appropriate error messages. If you
give DIAG2 an invalid device specification, you can also see standard system
error messages (e.g., - device does not exist).

(Changed 31 October 1981)

5.3 DSKANA

Use of the DSKANA command is a very important part of your disk maintenance routine. DSKANA analyzes the data on a specified disk and rewrites the bitmap. DSKANA also reports lost and mislinked disk blocks, inconsistent block counts, and other file errors. Use DSKANA frequently on every disk on the system. (You might make it a practice to use DSKANA on every disk just before you back it up.)

NOTE:

NEVER use DSKANA while other users are accessing the specified disk; to do so may damage the bitmap and the files on the disk. Make sure that the disk you are analyzing is write-enabled; DSKANA must be able to rewrite the bitmap out to the disk. Before you use DSKANA, you must log into account [1,2].

5.3.1 Displaying the DSKANA Option Summary

DSKANA operates in several different modes. For example, the default mode tells DSKANA to display a list of PPNs as DSKANA analyzes the disk accounts and then to give the final disk analysis messages that report the results of the analysis. If you would like more information (such as a display of all disk blocks and files on the disk, including the blocks and files in which file errors (if any) occurred), use the List (/L) option. If you want just a list of the disk blocks and files in which file errors (if any) occurred, use the Errors only (/E) option. If you want to send the DSKANA display to a file, specify an output file.

If you become confused about the DSKANA options, enter DSKANA followed by a RETURN:

        _DSKANA   (RET)

DSKANA displays a summary of the modes and options available to you.

5.3.2 The DSKANA Default Mode

Log into [1,2]. Now, enter DSKANA followed by the specification of the device that contains the disk you want to analyze. Then type a RETURN. For example:

        _DSKANA DSK1:   (RET)

You now see:

        [Begin analysis of Devn]

where Devn is the device you specified. If you are not using DSKANA on a device that runs under the control of the AM-410 hard disk controller, you see nothing more for some minutes, except for a list of PPNs as DSKANA

(Changed 31 October 1981)

proceeds through the accounts on the disk. Then you see some messages that
tell you the results of the analysis. For example:

```
.DSKANA HWK1: (RET)
[Begin analysis of HWK1]

[1,2]
[1,4]
[10,6]
   .
   .
   .
[110,1]
[300,20]
[The following blocks were marked in use but not in a file]
        1767     1772     2562     3456     6265     10270    11555     11567
        11661    12272    12303

[The following blocks were in a file but not marked in use]

[Rewriting BITMAP]

no file errors
 .
```

Below we discuss the messages that you see at the end of the disk analysis:

1. [The following blocks were marked in use but not in a file]

    It is quite likely that numbers will appear beneath this
    message; if they do, it is nothing to worry about. These are the
    addresses of disk blocks that the system has previously marked as
    being in use during intermediate operations. In fact, one reason
    to run DSKANA frequently is that it reclaims these temporarily
    allocated blocks so that they can be used by files.

2. [The following blocks were in a file but not marked in use]

    A list of numbers under this message is an indication of
    problems in the disk file structure. Somehow the block linking
    structure of the disk has gone astray. If you are using a disk
    that runs under control of the AM-410 and that was certified via a
    pre-4.4 version of CRT410, this message can also indicate that a
    block marked as bad in the BADBLK.SYS file was not marked in use in
    the bitmap. (The pre-4.4 version of CRT410 marks all bad blocks in
    use in the bitmap.) You must take immediate steps to restore the
    integrity of the data on your disk. (See Section 6.0, "Recovering
    From Disk Errors.")

3. [Rewriting BITMAP]

    After performing its disk analysis, DSKANA always rewrites the
    bitmap so that it reflects the true block allocation on the disk.

Before it reconstructs the bitmap, DSKANA compares the bitmap hash total with the total stored in the bitmap itself. If the two do not agree, DKSANA tells you so:

[BITMAP on disk had a bad hash total]

(A hash total is a computed value used to check the integrity of a group of data.) This message can be an indication of read/write errors, but is not necessarily anything to worry about. It would be a good idea to use the SET DSKERR command so that you will be made aware of any soft disk errors that occur in the future. If this message occurs more than once in a great while as you use DSKANA, you may have a hardware problem.

4. no file errors

    If you see the message above, you know that DSKANA has completed the disk analysis, and that the file structure on the disk is intact.

    If the message instead says something like:

    5 file errors

    you have a serious problem. The file structure on the disk is in error, and you are going to have problems in recovering the data on the disk. For information on coping with this problem, see Section 6.2, "DSKANA File Errors." The first thing you will have to do is to run DSKANA again with the /L or /E switch, so that you can see where on the disk the file errors occurred. You must use the /L or /E options if you are to see the file error messages that indicate exactly what is wrong with the file structure on the disk.

NOTE TO PHOENIX DRIVE USERS:

If you are using a device that runs under the control of the AM-410, before DSKANA begins its analysis, it looks for the file BADBLK.SYS[1,2] on the disk. This file contains a list of any bad blocks or tracks on that disk. If the disk was certifed by an AMOS 4.4 or later version of CRT410, DSKANA ignores the BADBLK.SYS file and analyzes the disk in the same way that it analyzes any other disk.

If the disk was certified by a pre-4.4 version of CRT410, DSKANA must take into account the information in BADBLK.SYS. First, CRT410 checks the hash total of BADBLK.SYS. If the hash total is bad, you see:

[BADBLK.SYS contains a bad hash code]

You then know that the original certification was not allowed to finish or that the file was damaged in some way. If you want to see if anything else is wrong with the disk, let the disk analysis continue. Or, you may exit DSKANA by typing a Control-C. In either case, you should use COPY to copy all files off the disk, since DSKANA is using information in BADBLK.SYS that is of doubtful integrity. Then re-certify the disk.

(Changed 31 October 1981)

If the BADBLK.SYS hash total was OK, and if you are using the /L switch (see below), DSKANA now prints the numbers of any bad blocks. For example:

```
[bad disk blocks]
     334    335    2035
```

Now DSKANA continues on with the disk analysis, displaying a list of PPNs as it analyzes the disk accounts.

5.3.3 Using the DSKANA List Option

If you want to see more information on how DSKANA is proceeding with its analysis of the disk besides just the PPNs of the accounts it is analyzing, select the List option by including the /L switch at the end of the command line. For example:

```
.DSKANA AMS1:/L  (RET)
```

The analysis proceeds as in the example above, but now you also see: 1. the disk address of the account, and 2. a list of all files in the account along with the disk addresses used by the blocks in those files.

For example, you might see the following information for a small account:

```
[200,1]
Directory       6627
INDEX   HLP    11430
GLOSRY  TXT     3444    3445    3446    3447    3450    3451    3452
                3607    3610    3611    3614    3615
HEADER  TXT             4130    4632
```

The account [200,1] has three files: INDEX.HLP, GLOSRY.TXT, and HEADER.TXT. The directory for the account appears at disk address 6627 and takes up only one disk block.   INDEX.HLP takes up one disk block (11430). GLOSRY.TXT takes up twelve disk blocks. HEADER.TXT takes up two disk blocks (4130 and 4632).

If a file error is on the disk, you see a file error message in the appropriate spot in the DSKANA display. For example:

```
[Begin analysis of DSK3:]

[1,4]
Directory       143    354    712    1126    4010
ME      BAS     144    145    146     147     150
Block    0 - block reserved for system use only in DSK3:ME.BAS[1,4]
```

(For a list of the file error message, see Section 6.2.1, "DSKANA File Error Messages.")

5.3.4 Using the DSKANA Errors Only Option

If you want to see a list of only the disk blocks and files in which errors occurred, use the /E option. For example:

.DSKANA HWK1:/E (RET)

The display you see looks just like the one you would see if you were using no switches at all, except that if DSKANA finds a file error, you see the disk block and file in which the error occurred. For example:

.DSKANA DSK3:/E (RET)
[Begin analysis of DSK3]

[20,1]
[30,4]
Block   731 - Block used in previous file in DSK3:GLIDX.RUN[30,4]
[30,5]
.
.
.

5.3.5 Specifying an Output File

You may tell DSKANA to send its display to a disk file by specifying an output file on the DSKANA command line. For example:

.DSKANA DSK0:ERRORS.LST=DSK1:/L (RET)

If the specified file already exists, DSKANA deletes it before beginning the analysis. When DSKANA is finished, you may use the TYPE command to look at the file or you may use the PRINT command to print it. DSKANA always writes the current date into the file as the first line of that file. For example:

Disk analysis list file on 7/12/80

6.0 RECOVERING FROM DISK ERRORS

Disk errors can come to your attention in one of two ways: 1) A program on the system (e.g., one of the disk diagnostic programs) can report a soft or hard error on the disk; or 2) DSKANA can report file errors (indicating that the blocks on the disk are incorrectly linked).

In either case, you must immediately do what you can to restore the integrity of the data on the disk. Remember: the procedures below are aimed at effecting a partial recovery of your data. Once the linking structure of your disk or the disk media itself goes wrong, retrieving the data on that disk is difficult. The most effective measure is a preventive one: that is, run DSKANA regularly, so that if trouble does occur, you catch it before it has done major damage to your data. Make frequent backups so that you can easily restore damaged data.

6.1 Handling Hard and Soft Disk Errors

If the system has trouble reading a disk location (a soft error), it retries that read operation eight times before it gives up and declares that disk block to be unreadable. If you do not have the SET DSKERR option in effect, the system does not report these retries; instead, after eight soft errors occur, the system reports a hard error.

If you begin to see soft errors when reading a particular disk, it is a good idea to follow the disk cleanup procedures below before the soft errors can develop into hard errors.

6.1.1 Cleaning Up the Disk

If you have made no changes to the files on the disk since your last backup was made, all you have to do at this point is clear the disk by formatting it (using the appropriate formatting program-- see Section 2.3, "The Disk Formatting Programs") and initializing it (see Section 4.1, "Initializing a Disk"). Then use the REDALL program to read the freshly initialized disk to make sure that the disk is OK. If the disk media seems to be healthy, you can copy your backup onto the empty disk. Remember, however, that in the case of disks which run under the control of the AM-410 or AM-420, you may NOT format or initialize the disk. Instead, you must re-certify the disk with the CRT410 or the CRT-420 programs, respectively.

If you have changed some files since your last backup, you must attempt to save the data on the damaged disk. If the system can still read the disk blocks (that is, if you have soft rather than hard errors), use the COPY command to copy all files over to a good disk. (Use the REDALL program to make sure that the output disk is good before copying over to it.)

If you are dealing with a hard error, the system is not able to read at least one of the disk blocks. Use the SET DSKERR command. Then run REDALL to try to read the disk. REDALL then tells you which disk blocks are unreadable. Write those bad block numbers down.

Now there are a couple of things you can try. You may be able to fix the hard error by simply recomputing the block CRC (Cyclic Redundancy Check). Use the DSKDDT program to do so:

1. Enter DSKDDT followed by the specification of the questionable logical device. Then type the number of the disk block you want to check. Hit RETURN. For example, to check block #20 on DSK1:, enter:

        .DSKDDT DSK1:20 (RET)

2. DSKDDT loads the specified block into memory. Now type an E followed by a RETURN (the DSKDDT Exit command):

        E   (RET)

DSKDDT now writes the block back out to the disk, recomputing the CRC in the process.

3. Do this for all bad disk blocks on the disk.

4. Now, use REDALL on the disk again. With any luck, the disk is now healthy again. NOTE: This procedure does not ensure that your data is intact. (Fixing the CRC error may actually cause some data in the block to be lost; however, correcting the CRC does allow the system to read the block.) You may want to dump the restored block with the DUMP BLOCK command to see if you need to modify the data in it.

6.1.2 Getting Rid of Bad Disk Blocks

If using DSKDDT did not fix the hard errors on the disk, you must take more stringent measures to get rid of the bad blocks. To find out what files the bad blocks belong to, use DSKANA with the /L or /E options. DSKANA exits when it finds the first bad block, but at least it tells you which file contains that first bad block. For example, if block 12 is bad in the file VARSET.BAS, the DSKANA display for that file might look something like this:

```
VARSET BAS  5      6      7      10      12 AM500 ERROR CODE 4
FOR DRIVE 1 BLOCK 12 (CYLINDER 0 HEAD 0 SECTOR 12)
```

Keep track of the files in which the bad blocks appear. You can use the DSKFIL command to check an individual file for a hard error.

Using DSKCPY to copy the damaged disk to another disk won't solve the problem-- DSKCPY exits and returns you to AMOS command level when it encounters a hard error.

Simply erasing the bad files will not help either. If you erase a file that contains a bad block and then run DSKANA again, DSKANA frees up the bad block so that the system can allocate it to another file.

The only solution is to erase the bad files and then to use the COPY command to copy the good files over to another disk.

1. Make sure that the disk you are copying to has no hard errors. (Use the REDALL program or the DSKANA program.)

2. Erase the bad files by using the ERASE command. (Be very careful not to run DSKANA again on the bad disk after this point, or the bad blocks will be freed again.)

3. Log into the System Operator's account and copy the good files on the disk over to the backup disk:

```
.LOG [1,2] (RET)
.COPY AMS2:=AMS1:[] (RET)
```

(Changed 31 October 1981)

The COPY command above copies all accounts on AMS1: over to AMS2:.
If two files exist on AMS1: and AMS2: with the same name,
extension, and account number, the command above deletes the file
on AMS2: and replaces it with a copy of the corresponding AMS1:
file. If you don't want COPY to replace duplicate files, use the
/NODELETE switch.   (For example: .COPY AMS2:=AMS1:[]/NODELETE)
Because you are logged into the System Operator's account, the
command above copies all files in all accounts on AMS1: over to the
corresponding accounts on AMS2:. If the corresponding account does
not exist on AMS2:, the COPY command creates it, transferring to it
any password associated with the source account.

4. Now is the time to restore the files that you erased. Copy them
   over to the good disk from your most recent backup disk.

5. Back up the good disk.

6. Now that you have copied all of the good files over to the new
   disk, get rid of the bad blocks on the original disk by formatting
   the disk. For example:

        .FMT500 DSK1: (RET)
        BEGIN FORMATTING
        EXIT
        .

7. After you have formatted the disk, use the REDALL program to see if
   the original hard errors resulted from actual physical damage to
   the disk.   If everything is OK, you can initialize the disk and
   begin to use it again.

If hard and soft errors are frequent occurrences on your system, you ought
to take a look at the disks themselves. (For example, if you are using
floppy disks, are you storing them correctly? Are they scratched or dusty?)
You might also check your disk controller board and the physical device
itself for maladjustments.

6.2 DSKANA File Errors

For the purposes of this discussion, we assume that you have already handled
any soft or hard errors on the disk, and that the only problem with the disk
at this point is in the linking of the disk blocks.

If DSKANA reports file errors (e.g., 5 file errors) or lists disk blocks
under the message: [The following blocks were in a file but not marked in
use], the block linking structure of the disk is in error. You must take
immediate steps to recover the data on your disk.

Run DSKANA again, but this time use the /L switch so that you can see what
disk blocks were assigned to each file or use the /E switch if you just want
a list of the blocks and files where the errors occurred.

(Changed 31 October 1981)

Look for the file error messages in the display (see Section 6.2.1, "DSKANA
File Error Messages"). Their location in the display indicates which disk
blocks are incorrectly linked, and this tells you which files are bad. For
example, if part of the DSKANA display looks something like:

```
.DSKANA DSK3:/E  (RET)
[Begin analysis of DSK3]

[20,1]
[30,4]
Block    731 - Block used in previous file in DSK3:COPYR.LST[30,4]
[30,5]
              .
              .
              .
```

then you know that a disk block used by the file COPYR.LST was also used by
another file. Two files cannot share the same disk block.

If you are fortunate enough to have a very recent backup disk, and the files
you have changed since that backup are okay, you can simply copy those
changed files from the damaged disk over to the backup disk, which now
becomes your original. (Of course, this assumes that the backup disk is all
right. You might make it a practice always to run DSKANA before backing up
a disk, so that you know that your backup disk is always good.)   Now
initialize the damaged disk; this clears the disk.

If you are not so lucky as to have a recent backup disk, you must do what
you can to salvage the data on your original disk:

1. Make a disk backup so that you have a copy with which to work.

2. Now that you have found out what files are bad, use the ERASE
    command to erase those files from the disk. If you do not do so,
    the errors in the linking structure will propagate, and you will
    lose even more data.

3. After clearing the bad files from the disk, run DSKANA again to
    make sure that all problems have been cleaned up.

4. Once no more file errors show up, you must set about restoring the
    files you have erased. If you have old backup disks that contain
    good copies of the files that you have just erased, restore those
    files on your disk by copying them over from the backup disks.

Make a final backup copy of your newly restored disk.

The discussion above assumes that damage has been done only to your files;
if the linking structure of the disk directories themselves is bad, you may
have to copy off of the disk whatever files you are able to, and then
initialize the disk and start over with a new account structure.

A systems expert may be able to reconstruct disk directories by using DSKDDT to actually change the binary data on the disk. (Directory entries are not stored in straight ASCII, but in a special packed format called RAD50; you will have to make the conversion yourself.) Using DSKDDT to reconstruct directories is very dangerous; do not try it except as a last resort, and be sure to make a backup first!

6.2.1 DSKANA File Error Messages

Below is a list of the file error messages that DSKANA can display when you use the List (/L) or Errors only (/E) options. The location of the messages in the display tells you the disk location of the bad block link.

1. Block used in previous file

    The last disk block listed in the file where this message appears also exists in another file. Since two files cannot share a disk block, this message means that the system made an error in allocating disk space to the two files.

2. Block marked as bad

    A block marked as bad in the BADBLK.SYS file has been mistakenly allocated to a file.

3. This file has a bum block count

    The actual disk block count for the file where this message appears does not match the block count assumed by the file itself. The system made an error in allocating blocks to this file.

4. Illegal block link

    A link in the file where this message appears points to an invalid disk address (e.g., to a disk block that does not exist).

5. Block reserved for system use only

    A link in the file where this message appears points to a disk block that is reserved for system use. The system has mistakenly allocated a disk block that should not be allocated to a file.

6. Block creates endless loop in file

    The linking structure of this file is such that eventually the disk blocks point back to themselves. That is, block-A points to block-B which points back to block-A.

7. Device error on Devn:

    This block contains a hard error that the system could not recover from.

8. [unable to locate BITMAP for rewrite]

    DSKANA couldn't find the bitmap area in memory for the device
    being analyzed. This means that the bitmap in memory may be
    invalid.

9. BITMAP rewrite error code XXXXX

    The bitmap could not be written back out to the disk. The
    number you see is the error code that indicates what the problem
    was.   For a list of these error codes, see Chapter 6, "The File
    Service System," in the AMOS Monitor Calls Manual, (DWM-00100-42).

7.0 PACKING THE DISK

Besides explaining the concept of "packing the disk," this section explains
why packing a disk is necessary, and gives information on the two commands
you can use to pack a disk: DSKPAK and COPY (with the /PACK option).

NOTE: Before packing a disk, make sure the linking structure of the disk is
intact. We recommend that you use DSKANA before you pack a disk.

The next few pages refer frequently to sequential and random files. If you
are not familiar with these terms, see Chapter 5, "Introduction to Files,"
in Introduction to AMOS, (DWM-00100-65).

7.1 When to Pack a Disk

When AMOS writes file blocks out to the disk, it follows this allocation
scheme:

1. If a block belongs to a sequential file, AMOS searches for the
    first free disk block on the disk beginning with the front of the
    disk, and writes the block there. Sequential file blocks thus tend
    to be located toward the front of the disk.

2. If AMOS is trying to write out a random file, it searches for the
    LAST area on the disk in which the file will fit, and writes the
    file there. Random files thus tend to be located toward the end of
    the disk.

This scheme leaves an area in the middle of the disk for new file blocks.
When you delete a file from the disk, the disk blocks that made up that file
are now free for use by other files. "Packing" the disk consolidates these
free areas on the disk by sliding the random files down toward the end of
the disk and sliding the sequential file blocks up toward the front of the
disk.   This allows the system to make efficient use of the free space on
that disk.

You especially need to reduce fragmentation of free space on the disk if you make use of a number of random files. If you only use sequential files, you will not need to pack the disk very often because the system, as it allocates disk blocks for sequential files, fills in any "holes" left by deleted sequential files.

On the other hand, if you use large random files, you will probably want to pack the disk quite often, perhaps before every disk backup. It is particularly important to consolidate free space if you use random files because when it comes time to allocate space for a random file, the total number of free blocks on a disk doesn't matter-- it's the number of free blocks that appear in a contiguous group that counts. For example, it is quite possible to get a device full error when allocating a random file of 50 blocks, even though you have 200 blocks free on that disk-- the system has to find 50 contiguous disk blocks.

7.1.1 Displaying the Bitmap

If you want to have some idea of how much in need of packing your disk is, you can take a look at the bitmap of that disk. A bitmap is a map of your disk. That is, it tells the system what blocks on the disk are available and which are in use by a file. If you look at the bitmap, you see a matrix of 1s and 0s which represents the free and used blocks on that disk. Each block on the disk is represented by a one if that block is in use and a zero if it is free.

If all of the 1s are clustered together in large groups at the beginning and end of the disk, with only occasional 0s scattered among the groups, the data on the disk is efficiently allocated. If, however, the 1s and 0s seem to be randomly mixed on the disk, you should pack the disk. To see a display of your bitmap, enter DUMP BITMAP followed by the specification of the device whose bitmap you want to see. For example:

        .DUMP BITMAP DSK0: [RET]

You now see on the screen the bitmap of the disk in device DSK0:. To freeze the display, type a Control-S; to resume the display, type a Control-Q. To interrupt the display, type a Control-C. At the end of the display, the DUMP command tells you how many blocks are available on the disk. For example:

        3411 free blocks

7.2 DSKPAK

DSKPAK packs all random files on the disk. That is, it moves all random files toward the end of the disk, consolidating the free area in the middle of the disk for new random files. It does this by sliding random files down to occupy the area left by deleted files. If there are no sequential file blocks in the random file area, DSKPAK causes all of the random files to form one contiguous area at the end of the disk.

(Changed 31 October 1981)

Do NOT run DSKPAK while other users are accessing the specified disk. To run DSKPAK, enter DSKPAK and the specification of the device you want to pack:

        .DSKPAK DSK1: (RET)

When DSKPAK is finished, you see the AMOS prompt.

7.3 COPY (the /PACK Option)

You can use the COPY command to pack the data on your disk if you use the /PACK option.

NOTE: It is VERY important that no other job be allowed to access the disk while you are packing it. If fact, it is good practice never to pack a disk while other jobs are running on the system.

To pack all files on the disk, log into the System Operator's account, [1,2]. Then type: COPY/PACK, the specification of the device you want to pack, an equal sign, the specification of the same device, and a RETURN. For example:

        .LOG [1,2] (RET)
        .COPY/PACK DSK1:=DSK1:[] (RET)

The command above copies every file on DSK1: over to itself. This causes the system to reallocate blocks for all files on the disk, writing over areas left by deleted files. (This command does NOT pack the User File Directories or the Master File Directory of the disk.) To completely pack the disk, you might have to perform this command several times.

NOTE: To achieve maximum packing of the disk, use the COPY command from the System Operator's account to copy all files from one disk to another freshly initialized disk. (In this case, you don't need the /PACK option since you are not copying files over to themselves.)

8.0 DISK BACKUP

The most important procedure in your disk maintenance routine is disk backup. You will have noticed that the discussions of the disk diagnostic tests assume that you have recent backups of all of your disks. Backups are your only assurance that you can at least partially recover from disaster.

Make disk backup a regular procedure on your system. You should encourage all users to back up the files in their own accounts to either their own data disk or to a communal data disk. You should do regular backups of the System Disk and any important data disks.

8.1 The COPY Command

Before using COPY to back up files, you will probably want to use the SET command to set DSKERR:

    .SET DSKERR (RET)

Setting this option tells the system to report any soft disk errors to your job.

Individual users may back up the files in their own accounts by using the COPY command. For example, to transfer copies of all files in account DSK0:[300,1] to the same account on DSK1:, enter:

    .LOG DSK0:[300,1] (RET)
    .COPY DSK1:= (RET)

Individual users may also back up accounts that are within the same project. For example, to back up all accounts in Project 100 from DSK0: to DSK1:, log into an account in Project 100 and specifiy wildcard PPNs in the COPY command:

    .LOG DSK0:[100,0] (RET)
    .COPY DSK1:[]=[100,*] (RET)

To back up all accounts from one device to another (for example, DSK0: to DSK1:), log into the System Operator's account:

    .LOG [1,2] (RET)
    .COPY DSK1:=[] (RET)

Because you are logged into the System Operator's account, the command above acts very differently than the usual COPY command:

1. The command above copies all files in all accounts on DSK0: over to DSK1: regardless of whether those accounts are in Project 1.)

2. Even though there is no wildcard PPN symbol on the left side of the equal sign, the command above copies all files in all accounts on DSK0: over to their corresponding accounts on DSK1: (not into the single account you are logged into, [1,2]).

3. If the destination account does not exist on DSK1:, the command above creates it. The command above transfers over to a new account any password associated with the corresponding source account that files are being copied from.

To back up specific accounts, enter the correct account specifications. For example:

    .LOG [1,2] (RET)
    .COPY DSK1:=[200,*],[100,*] (RET)

(Changed 31 October 1981)

If you use a disk that runs under the AM-420 controller (that is, a Winchester technology disk), and use a floppy disk as a backup disk, you may use COPY to transfer files between the Winchester disk and the floppy disk; however, for large backups, you will probably want to use the programs WINFLP, FLPWIN and FLPDIR. (For further information on these programs see the AMOS System Commands Reference Manual (DWM-00100-49) and the document Winchester/Floppy Backup System in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.) NOTE: Other backup systems for Winchester technology devices are the Video Cassette Recorder and the 1/4" Streaming tape drive. You cannot use COPY with these devices. Instead, you must use the programs VCRSAV, VCRRES and VCRDIR for the VCR and STRSAV, STRRES and STRVCR for the 1/4" streaming tape drive to transfer files to these devices. For further information about these programs see the manuals Software Information for the AM-610 Video Cassette Recorder Interface (DSS-10000-24), and Software Information for the AM-620/621 1/4" Streaming Tape Drive (DSS-10000-23).

8.2 The DSKCPY Command

IMPORTANT NOTE:

As of AMOS Version 4.4, you may use DSKCPY on devices that run under control of the AM-410 hard disk controller, but only if the disk you are copying from and the disk you are copying to have both been certified with an AMOS 4.4 or later version of CRT410. If either disk has been certified with a pre-4.4 version of CRT410, you must use COPY to back up the disk.

Use DSKCPY to make a literal image of one disk onto another. You may use DSKCPY on any type of disk; however, you may not copy between disks of different types (for instance, you cannot do a DSKCPY between a Winchester technology disk and a floppy disk). (We say that two devices are of the same type if they use the same device driver program. All devices that use the same device driver appear in the DEVTBL line of the system initialization command file with the same three-character device code. For example, DSK1, DSK2, and DSK3 use the same device driver; AMS0, AMS1, AMS2, and AMS3 use the same device driver.) So, you can use DSKCPY to copy between any two devices that share the same device code (e.g., from AMS1: to AMS2; STD3: to STD1:, DSK3: to DSK2:, etc.)

Besides copying and verifying data from one disk to another, DSKCPY also optionally generates a hash total for the disk copied to. To generate a hash total, use the /H switch. (For example: .DSKCPY/H.) (For full information on DSKCPY, see the DSKCPY reference sheet in the AMOS System Commands Reference Manual, (DWM-00100-49).

IMPORTANT NOTES:

1. Never run DSKCPY while other jobs are accessing the disks you are
    copying between.

2. A common mistake in using DSKCPY is to accidentally reverse the
    input and output device specifications. This has the effect of
    copying your empty disk onto your original disk. To avoid this
    situation, if your disk device allows it, always write-protect the
    disk from which you are planning to copy.

3. Before you copy to a disk, make sure that the disk is empty or that
    it does not contain any data that you need. Make sure that the
    disk is in the proper format. If the disk already contains data
    (that is, it is not brand new), you do not need to reformat it.
    (However, remember that using DSKCPY writes over any data already
    on a disk.)

4. If you use DSKCPY to copy a disk that has soft errors, DSKCPY will
    display the proper error message when it encounters the soft error,
    and then will attempt to copy the bad block over to the output
    disk, recomputing the CRC in the process. Since it had trouble
    reading the bad block, the copy that DSKCPY makes of the block may
    contain garbled data.

    If you use DSKCPY to copy a disk that has one or more hard errors,
    DSKCPY aborts the copy when it encounters the first hard error, and
    returns you to AMOS command level.

    If you suspect that your disk has hard or soft errors, follow the
    procedures outlined in Section 6.1, "Handling Hard and Soft Disk
    Errors."

5. You can use DSKCPY to copy both System Disks and data disks.   (A
    System Disk is a disk the system can boot from-- that is, it
    contains elements of the operating system necessary for system
    operation. A data disk is any disk that is not a System Disk.)

We have divided the information below into instructions for backing up hard
disk devices and floppy disk devices.

We have additionally divided those instructions into information on using
DSKCPY on: 1. Multiple-unit devices; and 2. Two-unit System Devices.

If you are using DSKCPY on a device that has more than two logical units
(for example, besides DSK0: and DSK1:, your system also has DSK2: and
DSK3:), backup procedures are the same regardless of whether that device is
a System Device as long as you avoid copying to the drive containing the
System Disk.

However, if you are using DSKCPY on a System Device that contains only two
logical units (DSK0: and DSK1:), using DSKCPY becomes more complicated
because you MUST use DSK0: (the drive reserved for the System Disk) during

the backup procedure. This can be tricky because intermediate steps in the backup may require that you write over the System Disk or replace that disk with a backup disk.

8.2.1 Important Note for Hawk Hard Disk Drive Users

DSKCPY uses a special fast copy mode for Hawk devices only. This mode gives a disk copy in approximately one-third the time of the traditional Hawk disk copy (about 6 minutes versus 18 minutes.) However, it does require that no other user be using the AM-500 controller while the disk copy is taking place. When you use DSKCPY on a Hawk device, you see:

%ALL other users will be suspended while HAWK copy is running
Hit return to continue or a control-C to abort:

If no other users are accessing disks that run under the AM-500, you may type a RETURN to continue; otherwise, type a Control-C to exit DSKCPY.

If it is not convenient for all users to stop running on the system, you may use the /O switch to tell DSKCPY to use the slower copy mode for the Hawk device. If you use the /O switch, you do not see the message above; other users may run on the system, but the disk copy will take about 13 minutes.

If you use the /H switch to generate a hash total for the copied to Hawk disk, the hash total will differ for the same disk, depending on whether you used the Hawk fast copy mode or the /O mode. The /O switch has no effect when DSKCPY is being used on a non-Hawk device.

As you read the instructions in the sections below, remember that if you are using DSKCPY on a Hawk device, you must decide whether to use the fast copy mode or the /O mode.

8.2.2 The Hard Disk Multiple-Unit Device

The instructions below assume that you are working with a disk in which at least one logical unit is an unremovable, permanently fixed disk. We assume that you are copying between disk cartridges or from the fixed disk to a cartridge. The instructions below also apply to using DSKCPY on a device that contains only two logical units (e.g., HWK0: and HWK1:) if that device is not a System Device, since you do not have to worry about writing over the System Disk. (Remember, the instructions do not apply to the Winchester technology disks.)

If you are using a hard disk device that contains more than two logical units within one physical device (e.g., the Century Data Trident which can contain DSK0: through DSK18:), you will treat that single device as a multiple-unit system. That is, you can back up each logical unit separately by copying one logical unit to another. (If you have two of these kinds of devices you will probably want to back up by copying from one drive to another. Be advised, however, that backing up several hundred megabytes of data is going to take awhile!)

(Changed 31 October 1981)

To copy either System Disks or data disks:

1. Write-protect the disk you are copying. (For example, if you are going to make a backup copy of the disk in device DSK2:, write-protect DSK2:.) If possible (that is, if you can arrange to perform the DSKCPY when no other users are going to be accessing the disks), write-protect ALL disks except the one you are copying to.

2. If you need to change disk cartridges, do so. This may require cycling-down the disk drive, but you do NOT need to turn off the drives or the computer.

3. Enter:

   _.DSKCPY (RET)

4. DSKCPY now asks you which devices to copy between:

   Input Drive:
   Output Drive:

   For input drive, enter the specification of the device you are copying from. For output drive, enter the specification of the device you are copying to. For example, if you are copying from AMS1: to AMS3:, enter:

   Input Drive: AMS1: (RET)
   Output Drive: AMS3: (RET)
   [Copying 616 blocks]
   [Duplication and verification completed]

   _.

   DSKCPY tells you how many disk blocks it is copying, and gives you progress reports on the status of the disk copy. When you see the AMOS prompt, DSKCPY is done.

8.2.3 The Hard Disk Two-unit System Device

If you are using DSKCPY on a System Device, and that device has only two logical units (DSK0: and DSK1:), using DSKCPY becomes a little more complicated because you cannot avoid using DSK0:. (If you are copying from a fixed disk to a disk cartridge, follow the procedures above for multiple-drive systems.)

To copy a data disk, you must first back up the System Disk (the fixed disk), copy the data disk from a cartridge onto the System Disk, and then copy that back to another cartridge. Then you must restore the System Disk from the backup. This procedure requires that you change cartridges several times, so to make sure that no mix-ups occur, carefully label all cartridges so you can be sure which cartridge is the System Disk backup, which is the original data disk, and which is the data disk backup.

(Changed 31 October 1981)

8.2.3.1 Backing Up the System Disk

1. Write-protect DSK0: and DSK1:.

2. Insert a backup disk cartridge to which you will copy the System
    Disk.   (This may require cycling-down the disk drive, but do not
    turn it or the system off.)

3. Write-enable DSK1:.

4. Use the DSKCPY command to copy from the System Disk (DSK0:) to the
    cartridge:

```
        .DSKCPY  (RET)
        Input Drive: DSK0:
        Output Drive: DSK1:
        [Copying 9696 blocks]
        [Duplication and verification completed]

        .
```

5. Remove the cartridge. Label and date it (e.g., SYSTEM DISK BACKUP
    1/23/80).

8.2.3.2 Backing Up the Data Disk

A. PART I:

    1. Write-protect DSK1: and DSK0:.

    2. Place the data disk you want to copy into DSK1:.

    3. Load DSKCPY into memory:

        .LOAD DSK0:DSKCPY[1,4] (RET)

    4. Before proceeding, make sure that you have backed up your
        System Disk.   The next step writes over everything on DSK0:,
        and without a valid system backup you will not be able to get
        your system up and running again.

    5. Write-enable DSK0:.

(Changed 31 October 1981)

6. Copy the data disk onto DSK0:

        .DSKCPY (RET)

        Input Drive: DSK1: (RET)
        Output Drive: DSK0: (RET)
        [Copying 9696 blocks]
        [Duplication and verification completed]
        .

7. DO NOT TURN OFF THE SYSTEM OR THE DISK DRIVES.

B. PART II:

    1. Write-protect DSK0: and DSK1:.

    2. Place the data disk backup cartridge into DSK1:, but do NOT
       TURN OFF THE SYSTEM OR THE DRIVES. (Remember, the System Disk
       on DSK0: is temporarily gone.)

    3. Write-enable DSK1:.

    4. Copy DSK0: to DSK1:.

            .DSKCPY (RET)
            Input Drive: DSK0: (RET)
            Output Drive: DSK1: (RET)
            [Copying 9696 blocks]
            [Duplication and verification completed]
            .

    5. Remove the backup cartridge. Date and label it (e.g., DATA
       BACKUP 4/20/81). DO NOT TURN OFF THE DRIVES OR THE SYSTEM.

8.2.3.3 Restoring the System Disk

After copying a data disk, you MUST restore the data originally on DSK0:.

    1. Write-protect DSK0: and DSK1:.

    2. Place the System Disk backup cartridge into DSK1:. (Do NOT turn
       off the drive or the system.)

    3. Write-enable DSK0:.

4. Copy DSK1: to DSK0:.

         .DSKCPY [RET]
         Input Drive: DSK1: [RET]
         Output Drive: DSK0: [RET]
         [Copying 9696 blocks]
         [Duplication and verification completed]
         .

5. Remove the System Disk backup disk and store it in a safe place. The system is now up and ready for normal use.

6. Delete the DSKCPY program from memory:

         .DEL DSKCPY [RET]

8.2.4 The Floppy Disk Multiple-Unit System

These instructions apply to all situations where you can avoid writing to DSK0: of the System Device—that is, you are using: 1. a floppy disk device that contains more than two logical units (e.g., you have AMS2: and AMS3: as well as AMS0: and AMS1:); 2. a floppy disk device that has only two logical units, but that device is not a System Device (since you don't have to worry about writing over the System Disk); and 3. a two-unit System Device to copy System Disks.

To copy either System Disks or data disks:

1. Do not turn off the computer or the disk drives. Do not remove the System Disk.

2. Write-protect the drives you will not be copying to.

3. Insert the original and backup disks in the drives. (The original disk is the disk you want to copy; the backup disk is the empty disk you will be writing to.)

4. Type DSKCPY followed by a RETURN.

5. DSKCPY now asks which disks you want to copy between:

            Input Drive:
            Output Drive:

   For input drive, enter the specification of the device you are copying from. For output drive, enter the specification of the device you are copying to. For example, if you are copying from AMS1: to AMS2:, enter:

              Input Drive: AMS1: (RET)
              Output Drive: AMS0: (RET)
              [Copying 616 blocks]
              [Duplication and verification completed]
              •

    6. DSKCPY is now done. Write-enable the disks.

8.2.5 The Floppy Disk Two-unit System Device

The instructions below apply to a situation where you MUST use DSK0: of a
two-unit System Device. (That is, you must copy a non-System Disk and you
only have two drives with which to do it.)

To copy a System Disk, follow the instructions above for multiple-unit
systems; leave the System Disk in DSK0:, and copy it to DSK1:.

To copy a data disk:

    1. Do not turn off the system or the disk drives. Leave the System
        Disk in DSK0:.

    2. Write-protect DSK1:.

    3. Insert the input disk into DSK1:.

    4. Use the LOAD command to load the DSKCPY program:

              .LOAD DSK0:DSKCPY[1,4] (RET)

    5. Remove the System Disk. Insert the backup disk into DSK0:.

    6. Enter DSKCPY followed by a RETURN:

              .DSKCPY (RET)

    7. For input drive, enter the specification of the device you are
        copying from.   For output drive, enter the specification of the
        device you are copying to. For example:

              Input Drive: DSK1: (RET)
              Output Drive: DSK0: (RET)
              [Copying 616 blocks]
              [Duplication and verification completed]
              •

8. When DSKCPY has finished, remove the backup disk and reinsert the
    System Disk. Mount the System Disk.

9. Delete DSKCPY from your memory partition:

        .DEL DSKCPY (RET)

Index

AMOS VERSION 4.4 METHOD OF HANDLING BAD DISK BLOCKS

1.0 INTRODUCTION

AMOS Version 4.4 supports a new method of handling bad areas on devices such as the Phoenix disk drive. This class of high-performance disk drives, because of the high density of data on such devices, may contain flaws which prevent contiguous allocation of the disk area. As disk drives incorporating the newer technologies become available, the need for error tolerance becomes increasingly important. The Phoenix disk drive is the first disk supported on the Alpha Micro system where the media is not guaranteed to be 100% good.

AMOS Version 4.3 supported an interim solution to this problem by simply marking all bad blocks as "in-use," thereby preventing their use within a file. Although this method solved most of the problems posed by the media flaws, it introduced others, such as the possible fragmentation of storage on a surface, preventing the use of that surface as a single random file. Because of the actual distribution of bad blocks on the Phoenix drive, this did not turn out to be a major problem. However, recognizing that this was not the optimal solution, we reserved fifteen spare or "alternate" tracks on each Phoenix drive for use as alternate storage. To better support future disk drives, and to eliminate the possibility of surface storage fragmentation, we have implemented a new method of handling media flaws which uses these spare tracks.

2.0 THE NEW METHOD OF HANDLING BAD BLOCKS

The new technique of handling bad blocks is to flag bad <u>tracks</u> and translate these tracks to the spare tracks at the time of disk access. This method results in several advantages:

1. The entire disk surface is available for use. That is, no "bad" blocks can exist in the middle of a surface, preventing complete use of the surface as a random file.

2. DSKCPY-type utilities can be used on the surface, speeding up backup time.

3. Disk maintenance utilities, such as DSKANA and DSPAK, do not need to treat the disk as a special case; it can be accessed the same way as any other device.

4. The actual choice of bad tracks versus bad blocks is device dependent, making the new method easily expandable for future devices.

The technique used is as follows:

1. The certification program for the disk creates the file BADBLK.SYS
    which contains a list of bad tracks or blocks, depending on the
    particular device. Currently, the Phoenix uses bad tracks.

2. The DEVTBL command reserves space for an <u>Alternate Track Table</u> in
    system memory.   DEVTBL does this automatically, with no change in
    the command format. The size of the Alternate Track Table is
    device dependent; currently, the Phoenix causes 30 bytes to be
    allocated for each surface.

3. The MOUNT command reads the BADBLK.SYS file into the Alternate
    Track Table whenever a disk is mounted.

4. Whenever a disk access is requested, the disk driver scans the
    Alternate Track Table to see if the requested block is in a track
    which is marked as bad.   If so, a translation is performed to
    access the requested block within the alternate track assigned to
    that bad track.

    The actual allocation of alternate tracks is device dependent;
    currently, the Phoenix allocates tracks 808-822 as alternate
    tracks. The first bad track on a Phoenix is assigned to track 808;
    the second is assigned to track 809, and so on.

5. When the system is being booted, the bootstrap routine must read in
    BADBLK.SYS to handle the case where SYSTEM.MON is allocated on an
    alternate track.

At the current time, the Phoenix is the only device that uses the BADBLK.SYS
technique of checking for bad areas.

3.0 CONVERTING TO THE NEW METHOD

Because of the advantages of the new method of handling media flaws, most
Phoenix users will want to convert their disks to use the new format. The
old format is still usable under AMOS Version 4.4; however, the old format
may <u>not</u> be supported in future releases.

                 YOU MUST CONVERT ALL DISKS TO THE NEW FORMAT!

(Of course, if you are a first-time Phoenix user as of AMOS Release 4.4,
your disks are already using the proper system.)

To convert a disk surface to the new format, follow this four-step
procedure:

1. First, be certain that you are running under Version 4.4 of AMOS.
    (Use the SYSTAT command to see which version of AMOS you are

  using.) The following procedures will have no effect if you are running under AMOS Version 4.3 or earlier.

2. Create a backup copy of the surface to be converted on a certified, but otherwise blank cartridge. (In the case of converting a cartridge disk, you will need to clear off a fixed surface so that you can place your backup there.)

3. Certify the surface to be converted, using the CRT410 program.

4. Copy the files from the backup disk to the newly certified disk via COPY.

Before AMOS Version 4.4, Phoenix users were restricted to using COPY when backing up a Phoenix surface. Now, <u>after</u> you have converted a disk, you may use DSKCPY from that time on to copy that converted disk to a certified, but blank, cartridge or surface when you want to do a backup, as long as the backup disk has also been certified via the 4.4 or later CRT410.

(Remember that DSKCPY makes a literal image; any data on the backup disk will be destroyed during the copy.)

DEFINING NON-SYSTEM DISK DEVICES

There are many times when you need to configure your system to access a disk device other than the System Disk. Examples of this would be adding an AM-500 and Hawk disk drive to an AM-400 Trident based system, or adding an AMS-format floppy to an STD-format system. In both of these cases, the procedure is essentially the same. In addition to defining the device on your system, it is often necessary to transfer data from one type of device to another, or to convert a System Disk from one device or format to another. This document will cover both the definition of new disk devices and the transfer of data.

1.0 DEFINING NEW DISK DEVICES

To add a new disk device to your system, follow the procedure outlined below. It is a good idea to create a copy of your SYSTEM.INI file and modify that copy, rather than to modify the SYSTEM.INI file itself. The modified copy can be tested using the MONTST program, without risking creating a version of your SYSTEM.INI file which will not properly boot the system. Then, when you are sure that your new system initialization command file works, rename it to SYSTEM.INI. Regardless of these precautions, you should have a backup copy of your System Disk, just in case your changes don't work.

1. Add the device driver program to those in the [1,6] area on DSK0:. Each device to be used on the system must have a device driver program in this area. The device driver program must have a unique three-character name and an extension of .DVR. After locating the driver program for the device you wish to use, it should be renamed to a three-character name. For instance, the driver for the AM-500 is distributed as HWK500.DVR on standard System Disks; before use it should be renamed to a three-character name such as HWK.DVR. If you are defining a floppy disk, you must use the FIXDVR program to create a new driver for your particular configuration of drive, format and controller. See the document titled Configuring Floppy Disk Drivers in the "System Operator's Information" section of the AMOS Software Update documentation packet for more information.

2. Using one of the system text editors, VUE or EDIT, make these changes to your SYSTEM.INI file:

    a. Add the new device name to your device table; that is, add the three-character device name defined in step #1 to the DEVTBL command line. This defines the new device name, and allows I/O to take place through the new device driver. Disk devices are

always sharable devices, and should therefore be defined before the / in the DEVTBL command line. If you cannot fit the new device names all on one line, you may use more than one DEVTBL command.

You must include both the device name and all valid unit numbers in the DEVTBL command line. Thus, if you are adding an AMS-format floppy driver to your system and wish to be able to reference both drives 0 and 1, you must add the device codes AMS0 and AMS1 to the DEVTBL command line.

b.   To allow the system to write information on the new device, you must define a bitmap for the new device. A bitmap is the method by which the system allocates space on the device. To add a new bitmap, place a BITMAP command for the new device immediately after any other BITMAP commands. The BITMAP command requires that you specify the device code, the bitmap size, and the units (drives) to which this bitmap will be applied.

The device code is the three-letter code you defined in step #1.

The bitmap size is dependent on the particular device you are defining. If you are working with a floppy disk drive, the FIXDVR program will tell you the bitmap size. For other devices, see the documentation accompanying the particular device.

Each BITMAP command actually defines a buffer in system memory which is used for reading and writing the bitmap to and from the device. This bitmap buffer may be shared by multiple units of the same device, resulting in a saving of system memory. In exchange for this saving, however, I/O must be done more frequently to the device to update the buffer than if the bitmap buffers were not shared.

For example, to define an AMS device (which has a bitmap size of 39 words) and to share the buffer between units 0 and 1, add the command:

        BITMAP AMS,39,0,1

If you do not wish to share the bitmap buffer, but wish to have separate buffers for each unit, add two lines which read:

        BITMAP AMS,39,0
        BITMAP AMS,39,1

c.   You may wish to include the driver program in your system memory area for faster response while working with the new device. To do this, insert the command:

(Changed 30 April 1981)

SYSTEM xxx.DVR[1,6]

into the SYSTEM.INI file just above the final SYSTEM command
line.   In almost all cases, this is not really necessary, as
the system will fetch the driver from the disk each time it is
needed if it is not already in memory, and will load it into
the memory partition of the user who is requesting access to
that device. However, including it in system memory will speed
up the response time of the system. NOTE: Never include the
driver for your System Device in system memory-- since that
driver was MONGENed into your monitor, it is always in system
memory anyway, and adding it via the SYSTEM command just takes
up uneccessary space.

IMPORTANT NOTE: There are several situations where you must put
a non-System Device driver into system memory:

i.   If you are using the /S option of the BITMAP command to
     place bitmaps for non-System Devices in switchable memory,
     you must place the drivers for those devices in system
     memory.

ii. Some programs (notably BASIC, RUN, COMPIL, and AlphaVUE),
     do not follow standard memory module conventions, and
     therefore require that the device driver of any non-System
     Device be in system or user memory if you are going to
     access that device while using those programs.   (You may
     place the driver into system memory using the SYSTEM
     command, or the individual user may load it into his own
     memory partition via the LOAD command before invoking one
     of the programs listed above.)

3.   After performing the preceding steps, boot up the system using the
     MONTST command:

     .MONTST SYSTEM,TEST.INI (RET)

     (Remember that to use the MONTST command your System Disk (DSK0:)
     must be on the fixed disk of the System Device, and your job must
     be operating in the first memory partition on the system (Bank Zero
     for bank switching systems).)

You should now be able to access the new devices for both reading and
writing. Once you are satisfied with the operation of the device and your
changes to your SYSTEM.INI file, rename yur temporary initialization command
file to SYSTEM.INI and reset the system.   You are now up and running.

2.0 TRANSFERRING DATA TO AND FROM THE DEVICE

Once a device has been defined on the system using the procedures described above, any of the standard system utilities may access it. If you are starting with a new device with no data on it, use the correct formatting program to write formatting information onto the device. You may then use the SYSACT program to initialize the device and to create accounts on that device. You may then use the COPY command to to ˘ copy data onto the new device.

If you wish to read existing data off of the new device, merely mount the device and use COPY to transfer the data.

3.0 BUILDING A SYSTEM ON A NEW DEVICE

Once you have defined the new device on your system, you may wish to create a System Disk for the device so that you can boot your system from the new device. To create a new System Disk, follow the procedure outlined below:

1. Format and initialize the device on which you want to build a system.

2. Copy your current System Disk over to the new device.

3. Log into the new device on account [1,4] and create a new monitor. Use the MONGEN program to do this. (See Generating System Monitors in the "System Operator's Information" section of the AMOS Software Update documentation packet.)

   When MONGEN asks you for the driver name, specify the driver for the device for which you are building the System Disk. After you run MONGEN, save the new monitor onto the new device, but don't reboot the system at this point.

4. Edit the SYSTEM.INI command file on the new System Disk and change the BITMAP commands for all DSK devices to match the bitmap size and configuration of the new device. Remove the old specifications for the device from the DEVTBL command. If you wish, you may add the current system device as a peripheral to the new System Disk by using a method similar to the one you used to create your current System Disk.

5. You may now use the MONTST program to test your new System Disk. If it successfully boots, you may wish to change your hardware configuration to boot off the new device; you may find information regarding this in the installation instructions accompanying the device.

DISK DRIVERS AND FORMATS

1.0 INTRODUCTION

This document discusses the disk drivers and formats supported by Alpha Micro software and how to use them. One mark of the AMOS system's flexibility is that it allows you to customize your system software for your particular hardware configuration by selecting the disk driver programs and disk formats that you want to use.

2.0 DISK FORMATS

The format of a disk refers to the way that the data on a disk is structured. This discussion concerns itself with two main areas: the physical and the logical disk format.

The physical format means the number of bytes in each sector on the disk (which in turn dictates the number of sectors on each track). The logical format refers to the way in which AMOS reads the physical format, including blocking factor and file structuring.

We call the sector a physical record, because it reflects the physical organization of data on the disk. AMOS imposes a logical organization on the disk (regardless of the physical attributes of the device); we call these logical units of data disk blocks. Except in the special case of devices that use the IMG device driver, disk blocks are always 512 bytes long. For hard disk devices currently supported by Alpha Micro, the size of the disk block is the same as the physical record size (512 bytes); floppy disk devices sometimes use a physical record that is smaller than a disk block.

The following paragraphs discuss the disk formats that Alpha Micro currently supports:

1. Standard (STD) Floppy Disk Format -- This format comes in both single- and double-density, and single- and double-sided versions. All are designed to use IBM-compatible diskette formatting. Single-density diskettes are formatted with 128 bytes per sector and 26 sectors per track. Double-density diskettes are formatted with 256 bytes per sector and 26 sectors per track. Double-sided diskettes for both densities merely have twice the number of tracks per diskette (154 instead of 77). The format for both densities is 512 bytes per block, with single-density blocked at 4 sectors per disk block, and double-density blocked at 2. Both densities have sectors mapped out to read every 5th sector to improve rotational

latency time, thereby allowing one 512-byte disk block to be read within one revolution of the disk. The AMOS file system is fully supported on this format. Single-density, single-sided diskettes are supported on both the AM-200 and AM-210 floppy disk controllers; double-density and double-sided diskettes are supported on the AM-210 only.

2. AMS Floppy Disk Format -- This format also comes in both single- and double-density, and single- and double-sided versions. All are formatted in 512-byte sectors. Single-density has 8 sectors per track; double-density has 16. The physical records are interleaved on a 3-to-1 basis to improve latency for sequential accesses. Double-sided diskettes merely have twice the number of tracks per diskette (154 instead of 77). The logical format is the same as the physical format, with each disk block being one 512-byte physical record. The AMOS file system is fully supported on this format. Single-density, single-sided AMS format diskettes are supported on the AM-200 controller; only double-density diskettes (both single- and double-sided) are supported on the AM-210.

3. Image (IMG) Floppy Disk Format -- The physical format is the same as standard format, with 128- or 256-byte sectors and 26 sectors per track. The logical format is the same as the physical format. Image format is supplied to allow the programmer to read a non-AMOS diskette (such as one created on an IBM computer) by physical record number, and to retrieve the data under the programmer's own file structure. The AMOS file system will operate under this format, but is extremely inefficient due to the small disk block size (128 bytes) and the fact that one full revolution must occur between sequential reads. Image format for single-density, single-sided diskettes is supported by the AM-200 controller; both single- and double-density, single- and double-sided diskettes are supported by the AM-210.

4. Trident Hard Disk Format -- The Century Data Trident drives have their own special format. All drives use a 512-byte physical record which is also their disk block size. The T-25, T-50, and T-200 drives all have 22 sectors per track; the T-80 and T-300 have 32. The T-25 has 408 tracks per surface; all other Trident hard disks have 815.

5. Hawk Hard Disk Format -- The CDC Hawk drive has its own special format which gives a physical record size of 512 bytes, 812 tracks, and 12 records per track.

6. AM-410 Hard Disk Format -- The CDC Phoenix drive has its own special format, which gives a physical record size of 512 bytes. Each surface has 808 tracks, with 36 records per track. (Actually, each surface has 822 tracks; fifteen of those tracks are spare or "alternate" tracks that are used in case bad tracks are detected on the surface.)

(Changed 31 October 1981)

7. <u>AM-420 Hard Disk Format</u> -- The Priam Winchester Hard Disks have their own special formats. They each give a physical record size of 512 bytes. The Priam 1070 drive has 4 surfaces of 185 tracks (plus 5 spare or "alternate" tracks) per surface, with 22 sectors per track. The Priam 3450 drive has 5 surfaces of 515 tracks (plus 10 alternate tracks) per surface, with 24 sectors per track. The Priam 6650 drive has 3 surfaces of 1101 tracks (plus 23 alternate tracks) per surface, with 36 sectors per track.

3.0 HARDWARE DISK DEVICES

The AMOS system supports several different disk devices. Each device has its own peculiar traits, and must be understood for proper system configuration. The following paragraphs describe the various hardware devices currently supported:

1. <u>Persci Floppy Disk</u> -- Runs under the control of the AM-200 floppy disk controller. The Persci supports all single-sided, single-density floppy disk formats described above. The Persci is capable of formatting diskettes.

2. <u>Wangco Floppy Disk</u> -- Runs under the control of the AM-200 or AM-210 floppy disk controller. With the AM-200, the Wangco supports all single-sided, single-density floppy formats described above. With the AM-210, it supports all single-sided, single- and double-density formats, except for single-density AMS format. The Wangco is capable of formatting diskettes.

3. <u>CDC Floppy Disk</u> -- Runs under the control of the AM-210 floppy disk controller. It supports all single- and double-sided, single- and double-density formats, except for single-density AMS format. The CDC is capable of formatting diskettes.

4. <u>Icom Floppy Disk</u> -- This is the original system offered with the AM-100 computer, but it is no longer in extensive use because of the limitations it imposes. The Icom disk runs under control of its own controller board, and does not have the capability of formatting diskettes. The Icom supports only single-density, single-sided STD and IMG formats.

5. <u>Trident Hard Disk Subsystem</u> -- The Century Data Trident is a special, large disk-pack subsystem that runs under control of the AM-400 interface board and the Century Data 1150-A formatter unit. The only format supported on this device is the Trident format. The unit will reformat disk packs (one surface at a time). The T-25, T-500, and T-80 Tridents have 5 surfaces; the T-200 and T-300 have 19 surfaces each. (The T-200 and T-300 are formatted as nineteen individual disks; the T-25, T-50, and T-80 are formatted as five individual disks.)

6. <u>Hawk Hard Disk Subsystem</u> -- The CDC Hawk is a special disk-pack subsystem that runs under control of the AM-500 Hard disk Controller board. The only format supported is Hawk format. The AM-500 is capable of formatting disks.

7. <u>AM-410/Phoenix Hard Disk Subsystem</u> -- The CDC Phoenix is a special disk subsystem that runs under control of the AM-410 Hard disk Controller board. The only format supported is Phoenix format. The AM-410 is capable of formatting disks.

8. <u>AM-420/Priam Hard Disk System</u> -- The several Priam disks are special Winchester technology disk systems that run under control of the AM-420 Hard Disk Controller board. The only format supported is AM-420 format. The AM-420 is capable of formatting disks.

3.1 DISK DRIVER PROGRAMS

A driver is a program that links the generalized disk service routines of the monitor with the physical disk device. You can generate your own version of the monitor (via the MONGEN program) that incorporates a specific disk driver into the monitor for use as the System Device (device DSK:). The disk drivers reside in account DSK0:[1,6], and can be called in as needed if you have more than one type of device on your system, or if you wish to operate with several different types of formats on one device. In other words, your System Device may be a Persci in single-density standard format, but you may use a separate driver (called AMS, for instance), to read an AMS-format diskette. In this instance you would generate a single-density standard format driver to use with MONGEN when you create the monitor. The AMS driver exists in account [1,6] of the system disk, and is available for use when you want to read AMS-format diskettes.

Drivers are supplied to you on your System Disk in account [1,6]; they have a six-character name that represents the device type they handle. Below are listed the disk drivers that are currently available, along with the devices they support:

```
200DVR - AM-200 with Persci or Wangco
210DVR - AM-210 with Wangco or CDC
ICMDVR - Icom floppy disk
TRIT25 - AM-400 with Trident T-25
TRIT50 - AM-400 with Trident T-50
TRIT80 - AM-400 with Trident T-80
TRI300 - AM-400 with Trident T-300
HWK500 - AM-500 with CDC Hawk
SMD410 - AM-410 with CDC Phoenix
420DVR - AM-420 with Priam 1070, Priam 3450, or Priam 6650
```

The drivers above all have the .DVR extension. Since all devices must be
identified by a three-character device name, you will have to rename or copy
over the particular drivers to some chosen three-character device code; this
code will also have to be added to your SYSTEM.INI file in both the DEVTBL
and BITMAP command lines. You may call these drivers any three-character
name you choose, except for the System Disk defined when you use MONGEN to
generate a new monitor, which will always have the device code of DSK.

The floppy drivers must be configured by the FIXDVR program prior to use.
See the document Configuring Floppy Disk Drivers in the "System Operator's
Information" section of the AMOS Software Update Documentation Packet for
information on configuring floppy drivers.

3.2 FORMATTING PROGRAMS

Formatting programs are provided for those devices that support formatting.
The following is a list of the programs provided:

        FMT200 — All formats supported by the AM-200
        FMT210 — All formats supported by the AM-210
        FMT400 — All formats supported by the AM-400
        FMT500 — Hawk format for CDC Hawk disks

(NOTE: Use the CRT410 program to format and initialize disks that run under
control of the AM-410, and CRT420 to format and initialize disks that run
under control of the AM-420.)

3.3 MONGEN

The MONGEN program is used to generate a new system monitor by overlaying
the disk driver area in an existing monitor with a different disk driver
from the list above. (See Generating System Monitors in the AMOS Software
Update Documentation Packet.) You do not need to rename the disk driver
used in the MONGEN procedure to a three-character device code; you may
reference it directly by its six-character name (e.g. TRIT80 for a Trident
T-80). The MONGEN procedure automatically renames this device so that it
can be referenced as DSK in the new system. Note, however, that you must
configure floppy drivers with FIXDVR, prior to the MONGEN procedure.

GENERATING SYSTEM MONITORS

1.0 INTRODUCTION

The MONGEN program allows you to generate system monitors for any disk hardware by inserting the necessary disk driver into an existing monitor.

To build a new monitor, you need an existing monitor (AMOS version 3.1 or later), and the disk driver for the specific device that you are going to use as the System Disk. The monitor that you will normally use is the file SYSTEM.MON located in account [1,4] of the System Disk. The disk driver to be used will be one of the drivers in DSK0:[1,6]. MONGEN will insert the driver into the monitor (overlaying the old driver), and leave the new monitor in memory. You may then test the new monitor directly from memory (via the MONTST command), or you may save it onto a disk (via the SAVE command).

2.0 USING THE MONGEN PROGRAM

Type MONGEN followed by a RETURN:

    .MONGEN (RET)

The MONGEN program responds with:

    INPUT MONITOR NAME:

enter the file specification of the monitor program you are going to use as the foundation of your· new monitor. You may specify the default system monitor DSK0:SYSTEM.MON[1,4] by entering just a carriage return. If you enter a file specification, the default device and account is DSK0:[1,4]; if you want to use a monitor not in that account, include device and account specifications.

MONGEN locates the specified monitor and loads it into your memory partition. Be sure that you have enough room to accomodate the monitor and disk drivers, as well as the MONGEN program itself, and also enough room for a couple of disk buffers. Typically this requires about 16K bytes of user memory.

Now MONGEN asks for the specification of the disk driver you want to insert into the monitor:

    NEW DISK DRIVER NAME:

Enter the file specification of the disk driver program you want to use. You may NOT enter just a carriage return. The default device is DSK0:, the default account is [1,6], and the default extension is .DVR.

MONGEN locates the disk driver program and loads it into memory above the previously loaded monitor. MONGEN now inserts the driver into the proper area of the monitor (thus overlaying the original disk driver). The new monitor is now complete in your memory partition.

Now MONGEN asks for a name to be given to the new monitor:

> ### NEW MONITOR NAME:

Enter a one- to six-character name (the default extension is .MON); this is now the name of the new monitor in memory. MONGEN now exits, leaving the new monitor as a module in your memory partition. You can test the new monitor by using the MONTST program, or you can save the monitor as a disk file by using the SAVE command.

NOTE: MONGEN does not affect the currently running monitor either in memory or on the System Disk. Nor does MONGEN test the new monitor; it merely builds a new monitor as a module in your memory partition.

3.0 DISK DRIVER PROGRAMS

The hard disk drivers and floppy disk drivers currently available are listed in the document called Disk Drivers and Formats, which you can find in the "System Operator's Information" section of the AMOS Software Update Documentation Packet. These driver programs are on your System Disk in account [1,6].

BUILDING A TERMINAL DRIVER (THE FIXTRM PROGRAM)

IMPORTANT NOTICE: For Release 4.5, the program to build a terminal driver was named NEWTRM. However, to more closely correspond to Alpha Micro's conventions for program naming, the name of that program has been changed as of Release 4.6 to FIXTRM.

Because of the device independent nature of the Alpha Micro system, adding a new kind of terminal is a simple matter-- you need only modify your system initialization command file to contain a terminal definition statement for the terminal, and add a terminal driver program for the new terminal to your System Disk.   The terminal driver program defines the terminal to the system, giving information about the particular characteristics of the device.   Because terminals differ so widely, certain commands may be useful only for certain types of terminals. The two main types of terminals are hard copy terminals, which output data to a permanent display via a printing mechanism, and CRT terminals (also called video display terminals, or VDT), which output data to a video screen. The terminal drivers for CRT terminals are the most difficult to write, since the programmer must worry about cursor positioning and other screen-oriented commands.

Alpha Micro supports a wide range of terminals by providing a large number of terminal driver programs (in both source and assembled form). In the past, if you wanted to use a terminal for which Alpha Micro does not supply a terminal driver, you had to try to modify an existing terminal driver to perform the commands of the new terminal, or had to laboriously write your own driver program.   Either case required that you understand quite a bit about AMOS assembly language and terminal drivers.

Beginning with Release 4.5, Alpha Micro is providing a tool to help our users develop their own terminal driver programs for any kind of terminal they would like to add to the system.   The tool is a terminal driver building program called FIXTRM.   This program provides you a skeleton terminal driver, and allows you to customize the skeleton terminal driver according to the needs of your particular terminal. You customize the terminal driver by answering a number of questions with yes or no, or values, or parameters which describe your terminal.

1.0 THE FIXTRM PROGRAM

Using FIXTRM to build a new terminal driver will take you about half an hour and will require that you are familiar with the characteristics of the terminal you need the driver for. It is a good idea to have the user's manual that accompanied your terminal at hand while you use FIXTRM so that you can look up the information that FIXTRM requests.

FIXTRM uses the following files. Do not erase or modify them.

       DSK0:ECHO.MAC[1,6]   DSK0:TDV1.MAC[1,6]   DSK0:FIXTRM.CMD[2,2]
       DSK0:TABDEF.MAC[1,6]   DSK0:TDV2.MAC[1,6]   DSK0:FIXTRM.RUN[7,6]
       DSK0:VARDEF.MAC[1,6]   DSK0:FIXTRM.SBR[7,6]

To use FIXTRM, log into the Driver Library account, DSK0:[1,6]:

        .LOG DVR: (RET)

Now, run the terminal driver building program FIXTRM by entering:

        .FIXTRM (RET)

FIXTRM now gives a set of instructions:

       Terminate all input lines with carriage returns.
       All numeric input is in decimal.
       Separate numeric answers on the same line with spaces.
       All NUMERIC answers default to 0.
       If you have any problems, refer to the FIXTRM documentation.

Next, FIXTRM asks a series of questions. Your answers to these questions
will tell FIXTRM how to customize the skeleton driver to form a terminal
driver for your particular terminal. The result will be a .MAC file that
you can assemble to produce an assembled terminal driver program.

You may enter your answers in either uppercase or lowercase letters, or a
combination of both. Also, any response of yes or no you enter may be
abbreviated to the first letter of the word. For example, to a yes-no
question you wish to answer in the affirmative, you may enter any of these:

        YES
        Yes
        Y
        y

Toward the end of this document there is a condensed Worksheet you can use
to gather the information you'll need before you begin FIXTRM.

We discuss each of the questions asked by FIXTRM below:

1.1   What is the name of the driver?

FIXTRM first asks for the name of the new terminal driver. Choose a one- to
six-character name for the driver. (Be sure no file exists in DSK0:[1,6],
the account you are logged into, that has the name you select and an
extension of .MAC; FIXTRM will abort itself and return you to AMOS command
level if the filename you select is not unique in that account.) You should
use a name that readily identifies the type of terminal supported by the
driver (for example, "HAZEL" for a Hazeltine terminal). Do not supply a
file name extension (that is, if you want to use the name SOROC, enter just
"SOROC", not "SOROC.TDV"). For example:

What is the name of the driver? SOROC

You will later add this terminal driver name to the system initialization command file TRMDEF statement that will define your terminal. (For information on modifying the system initialization command file to add a new terminal to your system, see The System Initialization Command File in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.)

1.2 Enter the number of nulls required after a line feed:

Some terminals require that one or more null characters be sent to them following a linefeed. These null characters serve as a delay so that the terminal does not print the first few characters following the linefeed on that same line, before actually performing the linefeed. Usually only hard copy terminals require this delay. For example, your terminal might require that 6 nulls be sent after a linefeed to ensure that the beginning of the next line actually follows the linefeed. Enter 0 if no delay is required.

1.3 Enter the number of nulls required after a form feed:

Some terminals, again usually just hard copy terminals, require that one or more null characters be sent to them following a form feed. These null characters serve as a delay so that the terminal does not print the first few characters following a form feed on the bottom of that same page, before actually performing the form feed. Your terminal, for instance, might require that 4 nulls be sent to it after a form feed in order to process the form feed before printing further characters. Enter 0 if no delay is required.

1.4 Does your terminal have a keyboard?

Some terminals are output-only devices, and have no keyboard. Enter yes or no. If your answer is no, FIXTRM will ask you no further questions. Instead, it will say "Please wait...", then finish building the terminal driver. For information on what to do next, once you are returned to AMOS command level, see Section 2, "Finishing Up."

1.5 Is RUBOUT a shift-underline?

Answer yes if the RUBOUT key (sometimes labeled RUB, DEL, or DELETE) is sent by pressing the Shift and Underscore keys simultaneously. This is the only way that a shifted rubout is generated. Otherwise, you must have a separate key for RUBOUT that you don't have to shift for; in this case, answer no.

1.6 Is your terminal a CRT?

Answer yes if your terminal has a video display screen. (CRT means cathode-ray tube, a common reference to the actual device which electronically displays characters on the screen. Often known as a VDT, or Video Display terminal.) Answer no if your terminal is a "hard copy" terminal; that is, if it prints a permanent copy of its interaction with you and the system on paper. If you answer no, FIXTRM will ask you no further questions. Instead, it will say "Please wait...", then finish building the terminal driver. For information on what to do next, once you are returned to AMOS command level, see Section 2, "Finishing Up."

1.7 Enter the number of rows on the screen:

Enter the number of horizontal rows that may be displayed on the terminal at one time. This number is usually 16, 24 or 25 rows.

1.8 Now enter the number of columns:

Enter the number of vertical columns that may be displayed at one time. This number is usually 32, 40, 64, 72, 80 or 132 columns.

1.9 For the cursor positioning command, is the row sent first?

Following the cursor positioning command from the terminal driver, if your terminal must receive the ROW byte from the terminal driver before it receives the COLUMN byte, answer yes. If it must receive the COLUMN byte first, then the ROW byte, answer no.

1.10 Enter the positional offset from 1,1:

FIXTRM assumes that the home position on the video dsplay terminal is 1,1 (row 1, column 1.) However, many terminals make the home position 0,0, and some use 32,32, so FIXTRM accepts an offset value relative to 1,1. This offset value is added to 1,1. If the home position on your terminal is 0,0, enter -1. If the home position on your terminal is 32,32, enter 31.

1.11 Enter the decimal ASCII value of the command leadin code:

Some terminals use command codes, while others use control codes, to perform operations such as clear screen or position cursor. If your terminal accepts command codes, it requires a command leadin code to tell it to recognize the following input as a command code. FIXTRM requires from you the decimal equivalent of the ASCII value for the leadin code. A typical leadin code, for example, is an ESCAPE. You would enter the ASCII value for

an ESCAPE, which is 27 (base 10). If your terminal uses only control codes (e.g., Control-K (^K) for cursor up), enter a 0.

1.12 Enter the delay (in clock ticks) required after commands:

Some terminals, whether using command codes or control codes, require a time delay to complete execution after receiving commands to CLEAR (clear screen), EOS (erase to end of screen), EOL (erase to end of line), or even to position the cursor. Some terminals may also require a time delay after any insertion or deletion of lines or characters. If your terminal requires a time delay after any of these operations, enter the number of clock ticks required to delay until the command taking the longest to execute is complete. You can calculate this number of clock ticks by:

1. Knowing the current baud rate setting of your system. (Say 19,200 baud.)

2. Counting the number of characters the terminal needs to generate to complete that longest operation. (Say it takes 1920 characters to do the longest operation on your terminal.)

3. Dividing the number of characters that need to be sent by the baud rate to obtain the decimal part of a second. (1920/19200 = 0.1 seconds.)

4. Knowing the cycles per second (Hertz, or Hz) of A.C. power used by your system as a reference frequency. In North America and a few other areas, 60 Hz is used, but most of the rest of the world uses 50 Hz. (Say 60 Hz.)

5. Finding the time of one cycle (one clock tick), in seconds. For 60 Hz, one cycle is approximately 0.01667 seconds. For 50 Hz, it is 0.02 seconds. (Say 0.01667 seconds per clock tick.)

6. Dividing the decimal part of a second, previously obtained, by the length of one clock tick to find the number of clock ticks required. (0.1/0.01667 = 6 clock ticks.)

The result, rounded up to the nearest whole clock tick (in our case, still 6 clock ticks), is the number of clock ticks required to delay before sending further information to the terminal while the terminal is performing a time-consuming command. FIXTRM causes the terminal driver to accomplish this delay by using the SLEEP monitor call and simply doing nothing for the duration of that operation.

If your terminal does not require a delay after receiving any command, enter a 0.

1.13   Does the terminal have insert and delete line commands?

Your terminal will have both or neither, but not just one or the other. Enter yes if it has both, or no if it has neither.

1.14   Does the terminal have erase to end of screen?

Enter yes if the terminal has the ERASE to END OF SCREEN (EOS) feature, or no if it does not.

1.15   Does the terminal have erase to end of line?

Enter yes if your terminal has the ERASE to END OF LINE (EOL) feature, or no if it does not.

1.16   Do you want function keys translated?

Some terminals have special keys called function keys. They might be labeled F1, F2, F3 and so on. When pressed, they send unique codes to the terminal driver to request the performance of equally unique commands. Those codes transmitted by these function keys can be "translated" to their equivalent codes in both the Alpha Micro text editor, AlphaVUE, and the Alpha Micro screen debugger, AlphaFIX. By answering yes to this question, you cause FIXTRM to set certain flags in the terminal driver so that VUE and FIX respond directly to the various function keys. If you do not have function keys, or do not want them translated, answer no. The next question will not be asked.

1.17   Enter the decimal ASCII value of the function key leadin code:

You are asked this question only if you answer yes to the previous one. To recognize the unique code sent from the terminal which indicates that the key pressed is a function key, the terminal driver must first receive the command leadin code. This command leadin code is typically an ESCAPE, having a decimal ASCII value of 27. If your terminal has function keys, enter the decimal ASCII value of the command leadin code it transmits to the driver. If your terminal does not have any function keys, but rather incorporates terminal commands into control codes, enter a 0.

1.18   Enter the delay required between function key characters:

When you press a function key, the terminal returns to the terminal driver a sequence of characters, one at a time, to accomplish the specified command. Each separate character is sent at the repeat rate of the terminal (the rate at which characters are repeated when, on the various terminals, you either

(Changed 31 October 1981)

hold down a key or hold down both a key and the REPEAT key). The terminal driver determines whether or not the character sequence is that of a function key by first testing for the function key leadin code character. If it finds that character, it then performs a counting loop internally, then looks at a buffer. Function key characters alone are found in the buffer at the end of that loop; if the terminal driver finds a character in the buffer, it knows a function key is being transmitted. If the length of the loop time is too short, the terminal driver won't find the character in the buffer whether or not it is a function key character. Much too long a loop time simply delays the system. The length of time the terminal driver does its timing loop is determined by your answer to this question. Determine the delay count required by:

1. Finding which is lower: 1) the "effective repeat baud rate" (the maximum repeat rate (in characters per second) the terminal can accomplish, multiplied by 10); or 2) the actual baud rate the terminal will be operating at. (Say the "effective repeat baud rate" is 600 (for 60 characters per second). Say the actual baud rate of the terminal is 19,200. The lower value is 600.)

2. Dividing 1 by that lower value. (1/600 or 0.001666....)

3. Dividing the result by a constant value depending on which processor your system uses. For the AM-100, that constant is .0000075, and for the AM-100/T, that constant is .000004333.... (Say it is an AM-100. 0.001666.../0.0000075 = 222.222 (approximately).)

4. Taking the final result and rounding it up to the nearest integer. (223.)

Enter that value (in our example, 223) in response to the question.

NOTE: Sometimes, if the value as determined above is slightly too low (due to overhead or other factors), and when the function key is pressed several times rapidly in succession (or the REPEAT key is also used), the function key translation routine cannot translate the function key properly. Changing the value determined above to a slightly higher value should solve the problem.

1.19   Are there terminating characters sent by function keys?

If you answered yes to the question, "Do you want function keys translated?", FIXTRM asks you this question (and the following question if you answer yes to this one).

Some terminal function keys send an "end-of-key" code that tells the system the function key transmission is over. This code may be a single character, or a combination of two or more characters. Typically, the code is <EOT>, <ETX>, <CR>, <CR><LF>, or something similar. The terminal driver needs to know if this terminating code is sent (but doesn't care what characters the code consists of); if the code is sent, the driver discards all of its characters. Answer yes to this question if the function keys send a

terminating code of one or more characters. Answer no if no terminating
code characters are sent.

NOTE: Some terminals allow you to set the character or characters of the
terminating code with a manual switch built into the terminal itself.   You
may set this switch into any position available; in any case, the terminal
driver discards it. However, we recommend that you do not select the <EOT>
(i.e., the <End of Transmission>, or Control-D) switch position if your
terminal operates over a telephone line. This is because <EOT> will cause
some modems, which are between the terminal and the system running the
terminal driver program, to disconnect the phone line.

1.20   Enter the number of terminating characters:

If you answered no to the question, "Are there terminating characters sent
by function keys?", this question will not appear. If you answered yes, now
enter the number of characters in the function key terminating code. For
example, if the terminating code is <ETX>, enter a 1; if the terminating
code is <CR><LF>, enter a 2. The terminal driver will ignore that many
characters following the character unique to the function key. The terminal
driver does not care what those characters are, but must know exactly how
many of them are in the terminating code.

1.21   Are there values sent by function keys that should be discarded?

Some terminal function keys send more than the command leadin code and the
character unique to the given key. Any characters coming before these two
should be discarded.   For example, the BEEHIVE DMxx series returns four
characters, <STX><ESC><ASCII character><ETX>. For that terminal, <ESC> is
the leadin character, the ASCII character is unique to a given key, and
<STX> must be discarded. <ETX> is the function key terminator, also to be
discarded, but which is handled in a different way, based on your answers to
the previous two questions.   Answer yes to the current question if your
terminal transmits characters other than the leadin character and the unique
character. If you answer no, FIXTRM skips to the second question below.

1.22   Enter the decimal ASCII value of the character to discard:

If you answer yes to the question "Are there values sent by function keys
that should be discarded?", FIXTRM asks you this question. Answer by
entering the decimal equivalent of the ASCII value of the character to be
discarded, then type a RETURN. Do not enter the value of any terminating
characters. FIXTRM will repeat the question until you identify all the
characters (other than terminating characters) to be discarded. After you
have finished, type just a RETURN when the question is asked again; FIXTRM
will go on to the next question.

(Changed 31 October 1981)

1.23 Is the XXX command implemented on the terminal?

Where XXX is a command name.   FIXTRM begins a loop of up to four questions,
which it repeats for 27 of the 28 commands your terminal may implement (that
is, electronically support based on its construction). (One of the 28
commands is standard on all terminals and FIXTRM automatically provides it
rather than asking you for it.) As an example of the four-question loop, you
first see:

        Is the CLEAR SCREEN command implemented on the terminal?

If you answer yes, and if information you gave previously to FIXTRM makes
the question appropriate, FIXTRM may ask you:

        Does the CLEAR SCREEN command require a delay?

If you answered yes to the first question in this loop, and if other
information you gave FIXTRM earlier makes the following question
appropriate, FIXTRM may ask you:

        Does the CLEAR SCREEN command use the standard leadin code?

And finally (if you answered yes to the first question in this loop), FIXTRM
requests the command code or codes your terminal transmits for the CLEAR
SCREEN command:

        Enter the decimal ASCII value(s) of the command code(s):

Now FIXTRM repeats the sequence for the second command:

        Is the CURSOR HOME command implemented on the terminal?

If you answer no this time, FIXTRM skips the three questions dependent on a
yes and inquires about the next command. And goes on in this manner until
it asks you about all but one of the 28 commands listed below, basing the
secondary questions upon your response to Is the XXX command implemented on
the terminal?. Here are the commands FIXTRM asks you about:

Command Name	TCRT Command
CLEAR SCREEN	0
CURSOR HOME	1
CURSOR RETURN (is not asked; it is an ASCII standard)	2
CURSOR UP	3
CURSOR DOWN	4
CURSOR LEFT	5
CURSOR RIGHT	6
LOCK KEYBOARD	7
UNLOCK KEYBOARD	8
ERASE TO END OF LINE	9
ERASE TO END OF SCREEN	10
PROTECT FIELD (reduced intensity)	11
PROTECT FIELD (normal intensity)	12
ENABLE PROTECTION OF FIELDS	13
DISABLE PROTECTION OF FIELDS	14
DELETE LINE	15
INSERT LINE	16
DELETE CHARACTER	17
INSERT CHARACTER	18
READ CURSOR ADDRESS	19
READ CHARACTER AT CURRENT CURSOR ADDRESS	20
START BLINKING FIELD	21
END BLINKING FIELD	22
START LINE DRAWING OR ALTERNATE CHARACTER SET	23
END LINE DRAWING OR ALTERNATE CHARACTER SET	24
SET HORIZONTAL POSITION	25
SET VERTICAL POSITION	26
SET TERMINAL ATTRIBUTES	27
CURSOR POSITIONING	none

The commands in the above list are the only ones that FIXTRM asks you about (except for cursor return, which is standard for all ASCII terminals, and which therefore is automatically provided by FIXTRM). They are the TCRT commands supported by Alpha Micro. The numbers to the right of the commands in the list above are the TCRT commands for each command. (NOTE: Cursor positioning, the last on the list, has no TCRT command.) If you wish to implement additional terminal commands, the source file which FIXTRM creates for you must be modified "by hand" to contain those additional commands.

NOTE: Some terminals (such as the ADM-31) implement the end blinking field and/or the end line drawing field as an end attribute command. If this is the case for your terminal, enter the single end attribute command sequence when either the end blinking field or the end line drawing field sequences are requested by FIXTRM.

1.23.1 Does the XXX command require a delay? - If you did not enter a 0 following the question "Enter the delay (in clock ticks) required after commands:" appropriate, FIXTRM may ask you this question. Answer with a yes or no.

1.23.2 Does the XXX command use the standard leadin code? - This second question is not asked unless you have indicated, in response to an earlier question FIXTRM asked you, that there is a leadin code. Answer yes if the command in question is implemented as a multiple character sequence. Answer no if the command is implemented as a control character.

1.23.3 Enter the decimal ASCII value(s) of the command code(s): - Enter the decimal equivalent of the ASCII value(s) of the command code(s). Do not enter the leadin code. FIXTRM automatically places it if the command requires a leadin code. Enter the values by separating them with a space, then type a RETURN. You can enter a maximum of 10 values per command.

This is the last of the secondary questions. Now FIXTRM asks you about the next command in the list, until it has inquired about 27 commands. When the loop is completed, FIXTRM continues with more questions.

1.24 Enter a code returned by a function key:

This question and the next one are only asked if you have indicated earlier that your terminal has one or more function keys that are to be translated into AlphaVUE commands. These two questions alternately repeat until you enter just a RETURN in response to this question.

Enter the ASCII value (in decimal) returned by a specific function key that uniquely identifies that key. Then type RETURN. For example, assume that the leadin code for your terminal's function keys is <ESC>. Say that one of the terminal's cursor positioning keys sends <ESC><A> when pressed. In response to this question you would enter 65, the decimal ASCII value for A. (See Section 5.0 for a complete decimal eouivalent ASCII chart.)

1.25 Now type a VUE command that corresponds to the function key:

This question follows the previous question for each of the function keys. It correlates the terminal's function keys to specific operations within VUE. In response to this question, press the actual key (plus the Control-key) that VUE uses for the specific operation you want to correlate to the function key. Do NOT type a RETURN. (FIXTRM expects you to type a VUE command, and is waiting for it. If you type a RETURN, FIXTRM does not process the RETURN until it again asks you, "Enter a code returned by a function key:", at which time FIXTRM thinks you are telling it there are no more command codes to enter.) For example, an IBM terminal sends <ESC><A> from a function key to indicate MOVE CURSOR UP. To make VUE correspond to that function key of your IBM terminal, enter a Control-K in response to this question, because ^K is the VUE command for MOVE CURSOR UP.

2.0 FINISHING UP

FIXTRM has asked its last question. You will see "Please wait...", then:

> The driver is complete. You may assemble and test it now.
>
> .

FIXTRM has created the source file for your terminal driver. Now you must
assemble the program by using the MACRO command. Enter:

>   .MACRO SOROC (RET)

where "SOROC" is the name you used in telling FIXTRM what terminal driver to
build. When MACRO finishes and returns you to AMOS command level, rename
the .PRG file it produced to the terminal driver extension, .TDV:

>   .RENAME/D *.TDV=SOROC.PRG (RET)

You now have a finished terminal driver program, customized for your
terminal.

3.0 ERROR MESSAGES

FIXTRM has two error messages. You may also see standard AlphaBASIC error
messages.   Typing a Control-C (^C) closes the partially completed output
file and aborts FIXTRM.

> ?Bad answer - try again
>   You typed an answer that did not start with a Y, y, N or n.
>   FIXTRM repeats the question.

> ?Sorry - Filename.MAC already exists
>   You specified a driver name which already exists as a macro file
>   (that is, having a .MAC extension) in DSK0:[1,6]. Following this
>   error message, FIXTRM has aborted itself. Start over by entering
>   FIXTRM and a RETURN following the AMOS prompt. Be sure to choose
>   a unique driver name.

4.0 WORKSHEET FOR YOUR TERMINAL

Use this section as a worksheet for your specific terminal. You will find
the information you need in your terminal user's manual. If FIXTRM asks
about a feature your user's manual doesn't mention, the terminal probably
does not implement that feature. You may find commands using different
terminology in the user's manual. Comparing this document's text to the
manual's, you will be able to find the information FIXTRM is asking for.

        FIXTRM WORKSHEET FOR THE _____ TERMINAL

FIXTRM uses the following files. Do not erase or modify them.

        DSK0:ECHO.MAC[1,6]   DSK0:TDV1.MAC[1,6]   DSK0:FIXTRM.CMD[2,2]
        DSK0:TABDEF.MAC[1,6]   DSK0:TDV2.MAC[1,6]   DSK0:FIXTRM.RUN[7,6]
        DSK0:VARDEF.MAC[1,6]   DSK0:FIXTRM.SBR[7,6]

Log into the Driver Library account, DSK0:[1,6]:. Run FIXTRM by entering:

        .FIXTRM (RET)

Remember, these are the rules:

        Terminate all input lines with carriage returns.
        All numeric input is in decimal.
        Separate numeric answers on the same line with spaces.
        ALL NUMERIC answers default to 0.
        If you have any problems, refer to the FIXTRM documentation.

Also, remember that you may use any of these combinations for yes or no
responses: YES, Yes, yes, Y, y or NO, No, no, N or n.

What is the name of the driver (6 char. or less; no extension)? _FIDO_
Enter the number of nulls required after a line feed: __
Enter the number of nulls required after a form feed: __
Does your terminal have a keyboard?   (Y) N
Is RUBOUT a shift-underline?   Y   N
Is your terminal a CRT?   (Y) N
Enter the number of rows on the screen: 24
Now enter the number of columns: 80
For the CURSOR POSITIONING command, is the row sent first?   Y   N
Enter the positional offset from 1,1:
Enter the decimal ASCII value of the command leadin code: __
Enter the delay (in clock ticks) required after commands: __
Does the terminal have INSERT and DELETE LINE commands?   Y   N
Does the terminal have ERASE TO END OF SCREEN?   Y   N
Does the terminal have ERASE TO END OF LINE?   Y   N
Do you want function keys translated?   Y   N
Enter the decimal ASCII value of the function key leadin code: ___
Enter delay required between function key characters: __
Are there terminating characters sent by function keys?   Y   N
Enter the number of terminating characters: --
Are values sent by function keys that should be discarded?   Y   N
Enter the decimal ASCII value of the character to discard: ___

Now summarize your terminal's commands by placing checkmarks and codes in the appropriate columns for each command below.	IS THE COMMAND IMPLE- MENTED? Y : N	DOES THE COMMAND REQUIRE A DELAY? Y : N	USE THE STANDARD LEADIN CODE? Y : N	PLEASE ENTER UP TO 10 COMMAND CODES IN DECIMAL ASCII VALUES (SPACED APART):
CLEAR SCREEN	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
CURSOR HOME	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
CURSOR UP	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
CURSOR DOWN	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
CURSOR LEFT	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
CURSOR RIGHT	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
LOCK KEYBOARD	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
UNLOCK KEYBOARD	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
ERASE TO END OF LINE	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
ERASE TO END OF SCREEN	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
PROTECT FIELD (red. intensity)	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
PROTECT FIELD (norm. intensity)	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
ENABLE PROTECTION OF FIELDS	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
DISABLE PROTECTION OF FIELDS	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
DELETE LINE	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
INSERT LINE	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
DELETE CHARACTER	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
INSERT CHARACTER	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
READ CURSOR ADDRESS	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
READ CHAR. AT CURSOR ADDRESS	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
START BLINKING FIELD	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
END BLINKING FIELD	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
START LINE DRAW OR ALT.CHAR.SET	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
END LINE DRAW OR ALT.CHAR.SET	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
SET HORIZONTAL POSITION	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
SET VERTICAL POSITION	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
SET TERMINAL ATTRIBUTES	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __
CURSOR POSITIONING	✓ : __	__ : __	__ : __	__ __ __ __ __ __ __ __ __ __

Enter a code returned by a function key: __
Now press a VUE command that corresponds to the function key: __
Enter a code returned by a function key: __
Now press a VUE command that corresponds to the function key: __
Enter a code returned by a function key: __
Now press a VUE command that corresponds to the function key: __
Enter a code returned by a function key: __
Now press a VUE command that corresponds to the function key: __
Continue to repeat for more function keys to correspond.

Assemble and test the driver now:

        .MACRO drivername (RET)

Rename the .PRG file produced to the terminal driver extension, .TDV:

        .RENAME/D *.TDV=drivername.PRG (RET)

(Changed 31 October 1981)

5.0 DECIMAL EQUIVALENT ASCII CHART

VALUE	CHARACTER	VALUE	CHARACTER	VALUE	CHARACTER
0	NULL	48	0	96	`
1	SOH	49	1	97	a
2	STX	50	2	98	b
3	ETX	51	3	99	c
4	ECT	52	4	100	d
5	ENQ	53	5	101	e
6	ACK	54	6	102	f
7	BEL	55	7	103	g
8	BS	56	8	104	h
9	HT	57	9	105	i
10	LF	58	:	106	j
11	VT	59	;	107	k
12	FF	60	<	108	l
13	CR	61	=	109	m
14	SO	62	>	110	n
15	SI	63	?	111	o
16	DLE	64	@	112	p
17	DC1	65	A	113	q
18	DC2	66	B	114	r
19	DC3	67	C	115	s
20	DC4	68	D	116	t
21	NAK	69	E	117	u
22	SYN	70	F	118	v
23	ETB	71	G	119	w
24	CAN	72	H	120	x
25	EM	73	I	121	y
26	SS	74	J	122	z
27	ESC	75	K	123	{
28	FS	76	L	124	|
29	GS	77	M	125	}
30	RS	78	N	126	~
31	US	79	O	127	DEL
32	SP	80	P		
33	!	81	Q		
34	"	82	R		
35	#	83	S		
36	$	84	T		
37	%	85	U		
38	&	86	V		
39	'	87	W		
40	(	88	X		
41	)	89	Y		
42	*	90	Z		
43	+	91	[		
44	,	92	\		
45	-	93	]		
46	.	94	^		
47	/	95	_		

SYSTEM PROGRAMMER'S INFORMATION

This section contains the following documents:

I/O Programming on the Alpha Micro Computer, Revision A01

Terminal Service System

Software Information for the User of the AM-960 Status Display

AMOS VERSION 4.6 SOFTWARE UPDATE DOCUMENTATION PACKET

'Alpha Micro', 'AMOS', 'AlphaBASIC'
'AlphaPASCAL', 'AlphaLISP', 'AlphaSERV',
'AlphaVUE', and 'AlphaACCOUNTING'

are trademarks of

ALPHA MICROSYSTEMS
Irvine, CA   92714

I/O PROGRAMMING ON THE ALPHA MICRO COMPUTER

Whenever a program must communicate with an I/O device, whether within an interface driver, a device driver, or a hard-coded program, it is said to be using I/O programming. The Alpha Micro central processors (the AM-100 and the AM-100/T) both use a memory-mapped I/O technique. Except for a few minor differences, these two processors are completely compatible; those minor differences occur in I/O programming. This document outlines the various methods and techniques used in I/O programming, and explains the differences between the two Alpha Micro processors.

The Alpha Micro processors both contain 256 separate I/O locations, each eight bits wide. These I/O locations are mapped into memory locations FF00 (hex) through FFFF (hex). These locations may be accessed as bytes (AM-100 and AM-100/T) or as words (AM-100/T only). When byte operations are performed on the I/O locations, the two processors behave identically; it is only when word operations are performed that differences appear.

For word operations, the AM-100 uses only the low-order half of the word during a write operation. During a read, the I/O byte is replicated in both halves of the resulting word.

The AM-100/T, on the other hand, will properly read and write words to the I/O locations, providing that the I/O controller handles word operations. Currently, no Alpha Micro interface boards handle 16-bit transfers. However, to leave this feature available for future use, a slight change in the way that I/O word operations are performed was necessary. On the AM-100/T processor, a read-word operation to a controller handling byte-reads only, returns the resultant byte in only the upper half of the word. This differs from the AM-100, which returns the byte in both halves on a read-word operation. Likewise, word-mode I/O writes use only the data in the upper byte when using byte-oriented controllers.

Therefore, when programming byte-oriented devices (all currently available devices are byte-oriented), you should avoid word-mode instructions whenever possible to avoid any confusion. If word-mode instructions must be used, then only the high-order byte should be used. If word-mode I/O to byte-oriented controllers is restricted to the high-order byte, the AM-100 and AM-100/T are completely compatible.

Instructions that generate a read-modify-write sequence (format 7 and 10 op codes using destination modes 1 through 7) ignore the read portion of the sequence and do only the write. This was engineered into the design of the processors to avoid confusion with I/O devices which use the same I/O location for two different functions based on whether a read or a write operation is being performed. Therefore, if you wish to read the contents of an I/O location with a format 10 op code, it may only be done in the source field of the instruction. When writing to an I/O location, any format 7 op code or format 10 op code with destination modes 1-7 may be used if such an instruction makes sense without the read operation.

Examples:

1. CLRB   @#^HOFF04   Clears I/O location 4.

2. CLR   @#^HOFF04   AM-100: Clears I/O location 4.
                              AM-100/T: Clears I/O location 4 on
                              a byte-oriented controller. Clears
                              locations 4 and 5 on word-oriented
                              controllers.

3. TSTB   @#^HOFF05   Won't work without read-- writes junk
                              to I/O location 5.

4. BIT   #2,@#^HOFF06   Won't work without read-- writes junk
                              to I/O location 6.

5. BIT   @#^HOFF06,#2   AM-100: Will work fine. AM-100/T:
                              Will not work-- contents of I/O
                              location 6 are read into upper byte
                              only.

6. BIT   @#^HOFF06,#2_8.   Will work fine on all processors.
                              The AM-100 brings data into both
                              halves; the AM-100/T into the upper
                              half.

7. CMP   @#^HOFF07,R0   Will not work unless controller is
                              word-oriented and the controller is
                              on the AM-100/T.

8. CMPB   @#^HOFF08,R0   Will work fine on both processors.

9. BISB   #2,@#^HOFF09   Will not work. Requires a read-
                              modify-write sequence for proper
                              operation. Writes junk into I/O
                              location 9.

Example #6, above, is the only change likely to be required in converting an existing driver to run on the AM-100/T. (The underscore, _, used in example #6 is a MACRO expression operator which tells MACRO to shift the expression (2) the specified number of bits. Hence, a shift of eight (decimal) bits puts the 2 into the upper half of the word.)

TERMINAL SERVICE SYSTEM

1.0 INTRODUCTION

The terminal service system incorporated in the AMOS monitor is a flexible and efficient set of routines and drivers for interfacing a variety of different terminals with different interface boards. You may write your own drivers for terminals and interfaces not supported by Alpha Micro. This document describes the general structure and function of the terminal service system, but does not go into details on how to write user-defined drivers. We have made available the sources to TRMSER, FILSER and several terminal drivers for those individuals who want to write their own terminal driver programs. Details on the monitor calls used within terminal and interface drivers are in the AMOS Montitor Calls Manual, (DWM-00100-42). For a general overview of the Alpha Micro Operating System, see Part III of Introduction to AMOS, (DWM-00100-65).

2.0 GENERAL STRUCTURE

The monitor contains a general terminal processing routine called TRMSER whose function is to link user programs and monitor processes to the outside world of interactive terminals; this is done purely on a data basis, without regard to terminal or interface hardware. TRMSER processes data on a character-by-character basis. Monitor calls are available to your programs for passing characters and full buffers of data between the terminals and the system. Think of TRMSER as a telephone operator who switches calls back and forth between sources and destinations without regard to the type of telephone in use or the name of the person using that telephone. TRMSER also provides the synchronous link to the asynchronous world of the terminal hardware.

TRMSER is a monitor routine that is embedded in the operating system skeleton monitor, SYSTEM.MON. In addition to the general TRMSER routine in the monitor, there must exist one or more routines called drivers that take the data from TRMSER and translate it as necessary into the specific codes required by the hardware and then route it to the terminal through the interface board. These drivers reside in account [1,6] of the System Disk, and are automatically loaded into system memory in response to the terminal definition (TRMDEF) command lines in the systm initialization command file (SYSTEM.INI file) at the time of system startup. Driver programs MUST be reentrant; only one copy of a driver is loaded into memory regardless of the number of terminals or interface boards of that type defined on the system.

The terminal service system uses two general types of drivers: interface drivers and terminal drivers. Interface drivers contain the routines necessary to get data characters to and from the interface boards that plug into the S-100 bus. Terminal drivers contain routines that process each

character that goes to or from the terminal. Terminal drivers handle code
character conversions, echoing functions, line-feed null characters, cursor
control, and special functions as required by the type of terminal in use.

3.0 INTERFACE DRIVERS

Interface drivers link the TRMSER routines and the actual hardware
responsible for getting characters to and from the terminal device. The
interface drivers are assembly language programs with an .IDV extension.
The filename of the interface driver appears in the TRMDEF command line of
the SYSTEM.INI file, and tells the system what kind of interface is being
used by the terminal defined by that command line. Typical drivers are for
the AM-300 board, the Processor Technology 3P+S board and the IMSAI SIO-2
board. The interface driver handles all initialization sequences for the
board if required, and also sets up interrupt processes if the board
supports it. Those boards which are not interrupt driven get put into the
clock scanner queue for asynchronous access every clock tick. A special
interface driver exists on the system called the PSEUDO interface driver; it
controls no hardware at all, but instead represents a software-controllable
interface for inter-job communication and control.

4.0 TERMINAL DRIVERS

Terminal drivers customize the handling of character input and output based
on the type of terminal being used. They are assembly language programs
that have the .TDV extension. The filename of the driver is the name by
which the terminal type is referenced in the TRMDEF statements in the
SYSTEM.INI file. Typical terminal drivers are for the ADM3, the Soroc, the
Teletype, the Multiterm and the Silent 700.

The terminal driver processes all input and output chracters, and determines
if these characters need special handling because of the type of terminal
being used. The terminal driver handles echo control and different methods
of character deletion. For example, most CRT terminals have the ability to
back up and erase the character being deleted, while hard copy terminals
(such as the Teletype) must explicitly echo the character, usually in a
format that distinguishes the characters from those accepted as input.

Terminal drivers may also be written for software-controlled ports, and two
such drivers are built into the monitor already. The PSEUDO and NULL
terminal drivers are used in conjunction with the PSEUDO interface driver,
and provide a means for passing characters straight through to the
controlling job or discarding output characters that are unimportant.
Terminal drivers are usually unconcerned with the type of interface used to
physically tie the terminal to the computer.

5.0 INTERSYSTEM DRIVER LINKS

The relationship between the different elements of the terminal service system can seem confusing at first; nevertheless, efficient systems-level programming requires a thorough understanding of the links that exist between these items.   The following units are referenced in further discussions:

1. JOB - A job is the unit that controls the operation of one task or a series of tasks running on the system. A job is independent of any other jobs running on the system unless it is tied to them by special user software. Every job on the system has a unique name one to six characters long.

2. TERMINAL - A terminal is the hardware device used to physically transfer data into the system, and get data from the system to the user on a character-by-character basis.   Terminals do not themselves have names.   Typical terminals might be a Teletype, ADM3, Soroc, etc.

3. TERMINAL DEFINITION - A terminal definition unit is a block of memory in the system area set up by a TRMDEF statement. It is the basic unit by which a terminal in the system is referenced when attaching that terminal to a specific job, or when using the terminal as an I/O device under control of the TRM device driver. The terminal definition unit has a unique name one to six characters long.

4. INTERFACE DRIVER - An interface driver is the program that transfers characters back and forth between the terminal and the hardware interface board to which the terminal is physically connected.   The interface driver has a name one to six characters long that is referenced by the TRMDEF statements in the SYSTEM.INI file. Interface drivers reside in account DSKO:[1,6], and have the extension .IDV.

5. TERMINAL DRIVER - A terminal driver is the program that performs the character code conversions required by the terminal in use. The program has a name (one to six characters long) that is referenced only in the TRMDEF statement of the SYSTEM.INI. Terminal drivers reside in account DSKO:[1,6] and have an extension of .TDV.

6. DEVICE DRIVER - A device driver is a program that allows the system to communicate with any I/O device connected to the system. Device drivers are written for disks, tape units, printers and terminals. The handling of terminals as devices for use by the generalized file service system is done through the TRM device driver, and not through the terminal drivers themselves. Device drivers have a one to three character name that is referenced in the device table statement (DEVTBL) in the SYSTEM.INI, and in user file specifications (e.g., AMS1:FILNAM.TXT). Device drivers reside in account DSKO:[1,6] and have the extension of .DVR.

The terminal definition unit contains the links to the defined interface driver and to the defined terminal driver; it thus is the basic unit by which terminals are referenced on the system. When a terminal is attached to a job, the JCB (Job Control Block) and the terminal definition unit become linked to each other. A job is considered to be detached if it is not linked up to a terminal definition unit, and a terminal is considered to be detached if it is not linked to a JCB. A job may only be linked to one controlling terminal, and vice versa.

A job performs I/O operations through the particular device driver referenced by the device specified in the file specification. A job performs terminal operations through the linked terminal definition unit for the terminal that is controlling that job. A detached job is placed into terminal wait state if it attempts to perform a terminal input or output operation. Since I/O operations differ in structure and usage from terminal operations, performing I/O operations to a terminal must be done through some mechanism other than directly into the terminal definition unit. From a system standpoint, the terminal definition unit performs differently than a device driver. To allow this, a general device driver has been written called TRM which allows terminals to be accessed as devices, as opposed to being accessed only as job controlling terminals. This operation will be described later.

5.1 Terminal Input Characters

Terminal input characters are processed through a complex chain of events. When a terminal keyboard character is struck by the operator, it is transferred to the hardware interface which then passes it to the interface driver routine. The interface driver routine reads in the character and then passes it to the TRMSER processor. TRMSER puts the character into the input buffer to wait for pickup by the program or monitor. As an asynchronous event, if echoing is not suppressed or is local to the terminal, TRMSER passes the character back to the terminal driver (when it is about to be echoed) to again allow the terminal routine to perform special functions. An example of this is the special echoing of Control-U characters for line deletion or rubouts for character deletion. The terminal routine then passes the character (or the converted character) back to TRMSER to be sent to the output processor.

5.2 Terminal Output Characters

Terminal output characters can come from two main sources: 1. characters to be echoed from the input processor; and, 2. characters to be output (generated by the monitor or user program) as messages or data to the user. Both are handled differently from a buffering standpoint, but eventually are presented to a common output routine in TRMSER to be sent to the terminal. Each character for output goes from TRMSER to the terminal driver for possible output code conversion or character translation. An example of this would be the null sequence sent after every line-feed for timing

purposes to the Silent 700 terminal driver. The terminal driver processes the character and then sends it back to TRMSER for position processing. TRMSER then passes the output character (or converted character) from the terminal driver to the interface driver where it is physically output to the terminal.

6.0 USING TERMINALS AS I/O DEVICES

Most programs (including the print spooler) perform input and output operations to I/O devices rather than to the controlling terminal. In some instances it is desirable to perform these operations on a terminal rather than a specific I/O device defined by its own device driver. One example would be the printing of data on a Multiterm or Teletype, or the use of these terminals as the output device of the printer spooler. Any terminal may be accessed as a device through the general device driver called TRM. The TRM device driver acts as a software link between the format required by the FILSER file service system and the TRMSER terminal service system. Any terminal can be considered a device by using the device code TRM and using the name given the terminal definition unit as the filename (the extension and PPN are ignored in the file specification).

For example, suppose you have a Teletype connected to an AM-300 board on port number two. Your TRMDEF command in the SYSTEM.INI file might look like this:

```
TRMDEF TELLY,AM-300=2:2,TELTYP,80,100,20
```

This Teletype may then be accessed as an output device by the file specification for any I/O operations requiring a specific device:

```
TRM:TELLY
```

To output directly to this device from BASIC, you would first open the device:

```
OPEN #file,"TRM:TELLY",OUTPUT
```

and then display data as follows:

```
PRINT #file,variable-list
```

The variable list may contain text and variables that you want to print as well as PRINT USING masks.

If you are planning on using terminals as I/O devices or to spool to them,
it might save some space to include the TRM.DVR program in system memory
during system startup. You would use the command:

        SYSTEM TRM.DVR[1,6]

before the last SYSTEM command in the SYSTEM.INI file. It is not necessary
to do this, however, since FILSER automatically loads the TRM driver into
your memory partition when it is needed if it is not in system memory.

NEW METHOD OF ALLOCATING CONTIGUOUS FILES

1.0 INTRODUCTION

A new method of allocating space on the disk for contiguous files has been implemented in AMOS versions 4.0 and later. The new algorithm will help to avoid excessive fragmentation of available disk space.

Previous versions of the operating system allocated space for contiguous files by finding the first available free space on the disk that was big enough to hold the file. Because of this, contiguous files were randomly distributed on the disk surface, and free blocks were not located in one contiguous group. The result of this problem was that the operating system often failed to find space for a contiguous file even though SYSTAT would indicate that enough total free blocks were available; there simply weren't enough free blocks in one contiguous group to hold the contiguous file. A way had to be found to reduce the fragmentation of available disk space.

1.1 THE NEW METHOD OF DISK ALLOCATION

AMOS versions 4.0 and later use a new disk allocation algorithm. Now contiguous files are allocated at the end of the LAST group of free blocks on the disk big enough to hold the file. Contiguous files will tend to be located at the end of the disk. Because AMOS still allocates linked files from the beginning to the end of the disk, linked and contiguous files will tend to separate; linked files at the beginning of the disk, and contiguous files at the end. This leaves a free area in the middle of the disk for allocation of new contiguous files.

1.2 PACKING THE DISK

When you delete a contiguous file from the disk, there will be a "hole" in the middle of the contiguous area. To remove this hole (i.e., to consolidate or "pack" the contiguous file area), use the DSKPAK program. DSKPAK packs the contiguous file area of the disk by moving the contiguous files towards the end of the disk. It does this by moving contiguous files down to occupy the area left by deleted files. If there are no linked files in the contiguous file area, DSKPAK causes all of the contiguous files to be packed into one continuous block at the end of the disk. To use DSKPAK, type:

    .DSKPAK devn:

where devn: is the specification of the disk you want to pack.

NOTE: Do not use DSKPAK on disks that have not been converted to the new format. (Using DSKPAK on unconverted disks will not damage the contents of the disk. On the other hand, it won't work either.)

1.3 CONVERTING TO THE NEW SYSTEM

To convert your disks to the new format, use the COPY command to copy your old format disks to an empty disk while the system is running under AMOS version 4.0 or later. The copy of the disk will have all linked files at the beginning of the disk, and all contiguous files at the end of the disk. The free space will be in one big chunk in the middle of the disk.

USING THE MAGNETIC TAPE UTILITY PROGRAMS

October 1981
Revision A01

This document reflects AMOS versions 4.6 and later

'Alpha Micro', 'AMOS', 'AlphaBASIC'
'AlphaPASCAL', 'AlphaLISP', 'AlphaSERV',
'AlphaVUE', and 'AlphaACCOUNTING'

are trademarks of

ALPHA MICROSYSTEMS
Irvine, CA   92714

ALPHA MICROSYSTEMS
17881 Sky Park North
Irvine, CA   92714

Table of Contents

.

1.0 INTRODUCTION

Alpha Micro has developed several magtape utility programs that support the AM-600 Magnetic Tape Transport Formatter Interface. These programs-- TAPE, REWIND, SET, and SKIP-- are all programs that you can invoke at AMOS command level by entering the proper command. For information on accessing the magnetic tape driver (MTU.DVR) directly from within your own assembly language programs, see the AMOS Monitor Calls Manual, (DWM-00100-42). For brief summaries of the tape utility programs, refer to the AMOS System Commands Reference Manual, (DWM-00100-49).

You will be most likely to use your magnetic tape transport to transfer files to the AMOS system from other computer systems. Magnetic tape also provides inexpensive backup media. The TAPE program works only with sequential disk files that contain fixed-length records. It does not support any random-access file format on the tape. TAPE can convert your data from ASCII to EBCDIC form and vice versa. TAPE can also optionally append a carriage return/linefeed character pair to the end of every record it transfers from tape to an AMOS disk. (That is a feature installed in TAPE to meet requirements of programs like AlphaBASIC.) See Section 5.0, "Transferring Data Between Disk Files and Tape," below, for more information on the kind of data you can transfer.

2.0 SETTING YOUR SYSTEM UP TO USE MAGNETIC TAPE UNITS

Make sure that you have defined the magnetic tape units as system devices in the DEVTBL command line of the system initialization command file, SYSTEM.INI. Magnetic tape units are not sharable, so place them after a slash on the DEVTBL command line. For example:

        DEVTBL DSK1,STD0,STD1,MEM,RES,/MTU0,MTU1,MTU2,MTU3,PTP

In the example above, the devices MTU0, MTU1, MTU2, and MTU3 are the magnetic tape units. (The example above includes a paper tape punch, PTP, which also is not sharable.) The AM-600 supports up to eight tape units, numbered 0 through 7.

The program MTSTAT.SYS must be in system memory. (You must place it in system memory by entering a SYSTEM MTSTAT.SYS command line in the SYSTEM.INI.) The magnetic tape driver program uses MTSTAT.SYS to determine tape density and tape speed. If MTSTAT.SYS is not in system memory, when you try to perform a magnetic tape unit operation, you see:

        MTSTAT.SYS NOT FOUND

Make sure that the magnetic tape driver, MTU.DVR, is in account DSK0:[1,6]. Once you have modified your SYSTEM.INI to include information on the magnetic tape units, you are ready to use the magnetic tape utility programs. Reboot the system with your new SYSTEM.INI.

(Changed 31 October 1981)

2.1 Using System Commands to Access the Magnetic Tape Unit

Besides the magnetic tape utility programs discussed in this document,
MTU.DVR allows you to use several AMOS commands to access the magnetic tape
units. To look at files on magnetic tape, use the ASCDMP command. Enter:

        .ASCDMP Devn: (RET)

where Devn: selects the magnetic tape unit you want to access. ASCDMP
displays (in ASCII form) the tape block currently in position at the tape
transport read heads. To display the next block, use ASCDMP again.   NOTE:
If you try to read past the end of the file, ASCDMP blows up; type a
Control-C to exit ASCDMP.

You may use the TYPE command to display a magnetic tape file. Enter:

        .TYPE Devn: (RET)

where Devn: selects the magnetic tape unit you want to access. TYPE
displays the file currently in position at the transport read heads. To
display the next file, use TYPE again.

You may use COPY to copy from a disk file to the magnetic tape. Enter:

        .COPY Devn:=Filespec (RET)

where Devn: selects the magnetic tape unit you want to access and Filespec
selects the disk file you want to copy to the tape. Make sure that Filespec
is a valid AMOS file specification. You may not use COPY to copy data
from the magnetic tape to a disk file.

2.2 Concepts

Before we describe the tape utility programs, it might be a good idea to
define a few of the terms that appear in the following discussion:

    1. Magnetic Tape Transport - A transport is the hardware that reads
        and writes data on a magnetic tape (also known as a tape unit).
        (The mechanisms within the transport that actually read and write
        the data are called heads.)

    2. Magnetic Tape - The magnetic tape is the media on which the
        transport acts. Tapes are either 7-track or 9-track; the tape
        utility programs work correctly on either type. The utility
        programs assume that your tape is capable of recording at least 800
        bits per inch. The standard magnetic tape length is 2,400 feet,
        but the tape utility programs function on any length of tape
        supported by your tape transport.

    3. BPI - The BPI (bits-per-inch) value specifies the density of the
        data on the magnetic tape (either 1600 or 800 bits per inch). When

(Changed 31 October 1981)

the system starts up, the magnetic tape driver assumes that the BPI
you will be using is 1600. If this value is incorrect, you must
use the SET BPI command to enter the correct BPI. You may record
at either 800 or 1600 BPI, but when you read a tape you must read
it at the BPI it was recorded at.

4. Density - When we talk about tape density, we are talking about
   the BPI a tape was recorded at. This figure gives the density of
   the data on the disk.

5. Load Point - When the magnetic tape is positioned so that the
   metallic film at the start of the tape is at the transport read
   heads, that tape is said to be at load point. Until the tape is at
   load point, you cannot read or write data on the tape.

6. Mounted - The magnetic tape is mounted when it is correctly
   threaded on the tape drive. See the documentation accompanying
   your tape transport for information on mounting tapes. (NOTE:
   Mounting a tape is not the same as mounting a disk. When you mount
   a disk, you must actually use the MOUNT command to mount it. If
   you use MOUNT on a magnetic tape unit, it does no harm, but
   accomplishes no purpose, either.)

3.0 SETTING TAPE DENSITY

When you use the tape utility programs to read or write data on a magnetic
tape, the system must know the density of the data you want to read or
write. When the system starts up, the default data density is 1600 BPI. If
you use the TAPE command to transfer data from the tape to disk files, make
sure that the current BPI value is the same as the density value under which
the data was originally recorded. If you use the TAPE command to write data
from the disk to the magnetic tape, make sure that the current BPI value is
set to the density you want to have the data recorded at.

To set the tape density value, use the SET BPI command. At AMOS command
level, enter:

        .SET BPI Devn:NNNN (RET)

where Devn: selects the magnetic tape unit you want to affect. NNNN is the
BPI value you want to use. For example:

        .SET BPI MTU0:800 (RET)

(SET BPI sets a flag in the magnetic tape module, MTSTAT.SYS. The tape
density does not get changed to the specified value until you actually read
from or write to the magnetic tape unit. At that time, MTU.DVR reads the
BPI value from MTSTAT.SYS and sends it to the transport.)

3.1 SET BPI Error Messages

If you give an invalid device specification or incorrect BPI value to the
SET BPI command, you see:

> The format for the command is: SET BPI MTU*:XXXX
> Where * = a tape drive in the range 0 thru 7 and
> XXXX is either 800 or 1600.

Check your device specification. (You can see a list of the valid system
devices, including the magnetic tape units, by typing DEVTBL followed by a
RETURN.) The AM-600 supports up to eight tape drives (unit numbers 0-7).
Make sure that you have specified a BPI value of either 800 or 1600.

4.0 REWINDING A TAPE

To rewind a magnetic tape, use the REWIND command. Enter:

> .REWIND Devn: (RET)

where Devn: selects a magnetic tape unit. Make sure that no other user is
accessing the specified tape unit.

This program returns a tape to load point. If the tape is already rewound,
using REWIND does no harm.

Make sure that the tape you want to rewind is correctly mounted. Type
REWIND followed by the specification of the unit whose tape you want to
rewind. Type a RETURN. For example:

> .REWIND MTU1: (RET)

You now see one of the following messages:

1. Tape is rewinding now
        This message is the most common and indicates that the tape is
    in the process of rewinding.   The rewind will be finished in a
    moment.

2. Tape is already rewinding
        Someone else has already begun to rewind the tape. The rewind
    will be finished in a moment.

3. Tape cannot be rewound - it is at load point.
        The tape has already been rewound (or has never left load
    point).   Check to see that you are specifying the correct tape
    unit.

REWIND returns you to AMOS command level while the tape is still rewinding.

(Changed 31 October 1981)

4.1 REWIND Error Messages

If you use REWIND incorrectly, you see one of the following error messages:

> ?File specification error
>> The system didn't understand the format of your command line. For example, you typed REWIND followed by a RETURN, without a device number specified. Re-type the line.

> ?Cannot OPEN Devn: - device does not exist
>> The system didn't recognize your specification of the magnetic tape unit. Check your spelling. Type DEVTBL followed by a RETURN for a list of the devices (including magnetic tape units) defined on the system.

> Tape cannot be rewound - it is at load point
>> REWIND cannot rewind the tape; it is already at load point. Check to see that you have specified the correct device.

5.0 TRANSFERRING DATA BETWEEN DISK FILES AND TAPE

The TAPE utility program copies data from disk files to magnetic tape and vice versa.

If you are reading from a magnetic tape created on another system, chances are that the data is incompatible with the AMOS system. For example, AlphaBASIC requires that all data records it handles end with a carriage return/linefeed pair; if these are not present, BASIC may not be able to read the tape data you have written into a file. As another example, the screen-oriented text editor, VUE, requires that each line end with a carriage return/linefeed pair, and that the line be 510 or less characters. For these reasons, TAPE will optionally append a carriage return/linefeed character pair to the end of every record as it transfers the records from tape to an AMOS disk.

Beyond optionally adding that carriage return/linefeed pair, and converting data from ASCII to EBCDIC and vice versa, any data conversion and manipulation required to make your data compatible with the AMOS system is left up to you.

5.1 How TAPE Writes Data to a Tape

Before we discuss the actual operation of the TAPE program, it is important
to understand how TAPE actually copies data to the tape from a disk file and
vice versa.   Records on tape are organized sequentially; that is, records
are written one after the other and are read back the same way. (TAPE does
not support a random-access file format.) Records are grouped on the tape
into blocks. (The exact number of records per tape block is set by you when
you originally write the data to the tape; this is called the <u>blocking</u>
<u>factor.</u>)

5.1.1 <u>Organization of Data on the Tape</u> - When it writes data to a tape,
MTU.DVR magnetizes a short amount of tape at the end of every tape block.
This magnetized interval is called an Inter-record Gap (IRG), and when the
tape drive reads the data back again, the tape drive automatically
recognizes an IRG as marking the end of a tape block. The hardware itself
begins reading after the IRG and stops when it reaches the next one.

Because an IRG takes up room on the tape that might otherwise be used to
store data, it is a good idea to make the tape blocks relatively long in
comparision to the IRG. In this way, you can make more efficient use of the
tape by minimizing the amount of wasted space on it.

When TAPE finishes copying the data from your disk file to the magnetic
tape, it writes a special symbol to the tape called an <u>EOF,</u> or end-of-file
marker. This marker designates the end of the file so that when you use
TAPE to copy the data back to a disk file, TAPE knows when the end of the
file has been reached.

The tape transport detects the end of the tape. When you use TAPE to write
data to the tape, TAPE rewinds the tape one block when it reaches the end of
the tape; then it writes an EOF to indicate the end of the file.

5.1.2 <u>Reblocking Data</u> - When you write data to a tape from disk, several
of the <u>disk file records</u> form one tape block. The number of records per
block is called the blocking factor. For example, if you write 25 records
per tape block, the blocking factor is 25. Of course, the actual amount of
data per tape block also depends upon the number of characters in each disk
file record.

When you write data to a tape, the TAPE program asks you if you want to do
reblocking.   What this means is that you can choose the number of data
records to write to a tape block. (This number does not have to be the
number of data records per disk block that appear in the disk file.) TAPE
is able to calculate the number of characters per tape block after it asks
you the length of your data records and the blocking factor you want to use.

(Changed 31 October 1981)

The ability to change the blocking factor of the data records when transferring data from disk to tape is an important factor in making efficient use of the magnetic tape.

If you do not want to do reblocking, TAPE assumes that your data records are 512 bytes long, and that you want to write one record per tape block. When you write data to a tape, make sure that you make note of the length of your data records and the blocking factor you use. Also remember to write down the number of tape blocks you are transferring. When you transfer the data back to a data file you will have to know all of this information.

5.2 Transferring Data from a Disk File to Tape

When you write data from a disk file to a tape, you may translate the data from ASCII to EBCDIC or vice versa; you may also append a carriage return/linefeed pair to the end of each record.

TAPE requires that the file be a sequential file with fixed-length records. Make sure that the current system BPI value for the unit you are using is the data density you want to use to record the data on the tape.

To write data to a tape from a disk file, enter:

     .TAPE (RET)

Now TAPE begins to ask you a series of questions. You can exit TAPE and return to AMOS command level at any time by typing a Control-C in response to any of the questions below.

        1. The screen clears and you see:

                    This is the magnetic tape program.
                    It can copy files from disk to tape or it can copy files
                    from tape to disk.
                    *
                    *
                    *
                    Which do you want to do?
                    1- Copy a file from disk to tape.
                    2- Copy a file from tape to disk.

                    Type the number of the option you wish.
                    Answer?

            Enter a 1, followed by a RETURN. (If you enter anything but a 1 or a 2, TAPE asks the question "Which do you want to do?" and then displays the example above again.)

            (If you change your mind about wanting to write to tape, type a Control-C to interrupt TAPE and return you to AMOS command level.)

2. Now the screen clears and TAPE displays the message:

> Tape can do character code conversion.
> What type of conversion do you want to do?
> 1 - None.
> 2 - Convert the ASCII file to an EBCDIC file.
> 3 - Convert the EBCDIC file to an ASCII file.
>
>
> Type the number of the option you wish.
> Answer?

TAPE is able to translate your date from ASCII to EBCDIC form, and vice versa. Because you are transferring data from a disk file to the tape, this question is asking if you want to translate the data in the file before you write it to tape.

Enter the number that selects the conversion (if any) that you want to perform. You must enter a 1, 2, or 3 (or a Control-C to exit to AMOS command level); any other answer causes TAPE to re-display the example above.

3. The screen clears and you see:

> Type in the name of the disk file you want copied to tape.
> Answer?

Enter a valid AMOS file specification. You may only enter the specification of a sequential file.

4. Now the screen clears and TAPE displays this message:

> Which tape drive are you using?
> (For example - MTU0: MTU1: MTU2: etc.)
> Answer?

Give the specification of the magnetic tape unit you want to use (e.g., MTU5:). If you decide at this point that you do not want to continue with the TAPE operation, type a Control-C to exit to AMOS command level.

5. The screen clears and you see:

> Tape can do reblocking of the disk file.
> This means you can specify the length of the record
> and the number of records in a tape block.
>
>
> Do you want to do reblocking?
> Answer yes or no -

Answer YES or NO, depending on whether or not you want to do reblocking. (See Section 4.1., "Reblocking," for a discussion on reblocking.) If you want to do reblocking, the screen clears and TAPE displays this message:

What is the size of the record?
(For example: 10, 20, 512 or any size you want)
Answer?
*
*
*
What is the number of records in a block?
Answer?

TAPE assumes that the sequential file you want to copy to tape has fixed-length records. Give the length of the record (in bytes, including any record delimiters such as carriage returns and linefeeds.)

If you enter just a RETURN, TAPE assumes a record length of 512.

Now give the number of these records you want to fit into one tape block. This is called the blocking factor.

6. TAPE now clears the screen and tells you how many characters you are writing in a tape block. For example:

    You are writing 2500 characters in a tape block.

If this number does not reflect what you really want to do, you may type a Control-C to exit from TAPE to AMOS command level. For example:

    You are writing 0 characters in a tape block.

The message above indicates that you entered invalid data to the reblocking questions. Exit TAPE via a Control-C, and try again.

7. Next you see:

    *
    *
    *
    Is the tape loaded?
    Type return if it is

When you are sure that the tape is physically mounted on the tape drive, and that the tape is at load point, type a RETURN. (If you do not want to start writing at the beginning of the tape, the tape does not have to be at load point.)

(Changed 31 October 1981)

8. Now TAPE asks:

        <u>Is the tape drive on-line?</u>
        <u>Type return if it is</u>

    Make sure that the tape drive is on-line. (Check with the
    documentation that accompanied your tape transport to see what
    button or switch to push to put the drive on-line.)

    Now the TAPE program begins its data transfer. When the copy is
    finished, TAPE clears the screen and you see a message that tells
    you how many tape blocks were written and how many errors occurred.
    For example:

        <u>There were 100 tape blocks written or read.</u>
        <u>There were 0 errors.</u>

    Your tape drive detects any parity or CRC errors and reports them
    back to the AM-600 controller.

9. Remember to make note of the size of the data records, the number
    of data records per tape block, and the total number of tape blocks
    you wrote. Also remember the BPI value you recorded the tape at.
    When you transfer the data back to a disk file, you will need this
    information.

10. If you do not rewind the tape, you can write another file beginning
    at the next position on the tape. (You can keep using TAPE without
    rewinding the tape until the tape is full.)

5.3 Transferring Data from Tape to Disk Files

When you copy data from tape to a disk file, TAPE requires that the file be
a sequential file with fixed-length records. Make sure that the current BPI
value for the drive you are using is the value originally used to record the
data.

To write data to a disk file from a magnetic tape, enter:

    <u>.TAPE</u> (RET)

Now TAPE begins to ask you a series of questions:

1. The screen clears and you see:

>   This is the magnetic tape program.
>   It can copy files from disk to tape or it can copy files
>   from tape to disk.
>   *
>   *
>   *
>   Which do you want to do?
>   1- Copy a file from disk to tape.
>   2- Copy a file from tape to disk.
>
>   Type the number of the option you wish.
>   Answer?

    Enter a 2.   (If you enter anything but a 1 or a 2, TAPE asks the
    question "Which do you want to do?" and then displays the example
    above again.)

    (If you change your mind about wanting to write to a disk file,
    type a Control-C to interrupt TAPE and return to AMOS command
    level.)

2. Now the screen clears and TAPE displays the message:

>   Tape can do character code conversion.
>   What type of conversion do you want to do?
>   1 - None.
>   2 - Convert the ASCII file to an EBCDIC file.
>   3 - Convert the EBCDIC file to an ASCII file.
>
>
>   Type the number of the option you wish.
>   Answer?

    TAPE is able to convert your tape data from EBCDIC to ASCII form,
    and vice versa. Because you are transferring data from tape to a
    disk file, this question is asking if you want to translate the
    data on the tape before you write it to disk. Enter the number
    that selects the conversion (if any) that you want to perform.
    (You must enter a 1, 2, or 3; any other answer causes TAPE to
    re-display the example above.) To interrupt TAPE and return to
    AMOS command level, type a Control-C.

3. The screen clears and you see:

>   Type in the name of the disk file the tape is to be
>   copied into.
>   Answer?

Enter   a   standard   AMOS   file   specification   (e.g.,
DSK0:MAIL.DAT[110,2]).   (The   default   extension   is   .DAT.   TAPE
assumes the account and device you are currently logged into.)

4. The screen clears and you see:

> Which tape drive are you using?
> (For example - MTU0: MTU1: MTU2: etc.)
> Answer?

Give the specification of the magnetic tape unit you   want   to   use
(e.g., MTU5:).   If you decide at this point that you do not want to
continue   with the TAPE operation, type a Control-C to exit to AMOS
command level.

5. TAPE now displays this message:

> You must supply blocking information on the tape file.
> This means you must say how many characters are in a
> record and how many records are in a tape block.

> What is the size of the record?
> (For example: 10, 20, 512, or any size you want)
> Answer?

Give the number of bytes in each data record. Include any record
delimiter bytes, such as carriage returns and linefeeds.

NOTE: When writing data from a disk file to tape, blocking
information is optional. However, when you write data from tape to
a disk file, you MUST supply blocking information.   TAPE has no
way of knowing the blocking factor and record length you supplied
when you originally wrote the data onto the tape.   You must now
supply that information to successfully retrieve your data from the
tape.

After you answer the question above, TAPE asks:

> *
> *
> *
> What is the number of records in a block?
> Answer?

Give the blocking factor you used when you originally wrote the
data to the tape. (For example, if you wrote 100 data records   per
tape block, enter the number 100.)

You MUST enter an answer for the questions above on blocking
information. If you enter just a RETURN, TAPE re-displays the
questions.

6. Now you see a message that tells you how many characters TAPE has calculated to be in a tape block. For example:

You are reading 10000 characters in a tape block.

If this number is not correct (for example, it is 0), there was something wrong with your answers to the questions on blocking above. Type a Control-C to exit to AMOS command level and try again.

Now TAPE asks:

```
*
*
*
Do you wish to insert carriage return/line-feed pairs
after each record?
Answer?
```

Answer with a Y for yes or an N for no, followed by a RETURN. You may wish to have TAPE append carriage return/linefeed characters at the end of each record if you will be reading the data file via an AlphaBASIC program or AlphaVUE.

After this message, TAPE asks you for the number of blocks you want to read from the tape. For example:

```
You are reading 2500 characters in a tape block.
*
*
*
How many tape blocks do you want to read?
Answer?
```

Enter the number of tape blocks you want to transfer to a disk file. TAPE will read no more than the number of blocks you specify, but under certain circumstances will read less. For example, if the end of the tape is reached before TAPE has read the specified number of blocks, it will stop. If an end-of-file marker is reached, TAPE stops reading.

NOTE: Some tapes contain a special header block at the front of the tape as the first file. This header file may be only one block long. If you use TAPE on such a tape, it reads this file first and stops, reporting that it read only one block. Use TAPE again without rewinding the tape, and you will read the first of your files on the tape.

(NOTE: See Section 6.0., "The SKIP Command," for information on using SKIP to skip over files and header files.)

7. Now TAPE asks:

    ```
    *
    *
    *
    Is the tape loaded?
    Type return if it is
    ```

    Type a RETURN when you are sure that the tape is physically mounted on the tape drive reel. Make sure that it is at or past load point. Type a RETURN. (If you do not want to start reading at the front of the tape, it does not have to be at load point.)

8. Now you see:

    ```
    Is the drive on-line?
    Type return if it is
    ```

    Type a RETURN when you are sure that the tape drive is on-line.

9. TAPE now transfers over the data from the tape to the disk file. It clears the screen and tells you how many tape blocks it read from the tape and how many errors occurred. For example:

    ```
    There were 200 tape blocks written or read.
    There were 0 errors
    ```

    Your tape drive detects any parity or CRC errors and reports them back to the AM-600 controller.

10. If you do not rewind the tape, you can use TAPE again to copy the next file on the tape. In this way, you can read a tape that contains many files, even if those files consist of records of different sizes grouped with different blocking factors.

5.4 TAPE Error Messages

If you give TAPE invalid device specifications you see the standard system error messages. For example:

    ```
    ?Cannot READ Devn: - device does not exist
          You gave an invalid specification to TAPE's question: Which tape
    drive are you using?. Make sure that the magnetic tape units are
    defined as devices in the system device table.
    ```

?Memory allocation failed
        When you specified the number of characters per tape block, you
gave too large a number. (For example, you specified a record size of
100,000.) Check the validity of the blocking information you gave to
TAPE.

You are writing 0 characters in a tape block.
        This message indicates that you made an error when answering
TAPE's questions on data blocking. Type a Control-C to exit to AMOS
command level and try using TAPE again.

?Cannot OPEN Filespec - file already exists
        You tried to write data from the tape to a file, but a file with
the same specification already exists. TAPE returns you to AMOS
command level. Check your file specification to make sure that you
entered the correct device and account specifications. Try using TAPE
again, but enter a different filespec.

?Cannot READ Filespec - file not found
        You tried to write data from a disk file to the tape, but the
system cannot find the specified file. TAPE returns you to AMOS
command level. Check your file specification to make sure that you
entered the correct account and device specifications. Try using TAPE
again, but enter a different filespec.

?Cannot READ Devn: - file not found
        You are trying to write data from the magnetic tape to a disk
file, but TAPE cannot read the tape. You usually see this message when
you try to read from a blank tape or when the tape is not positioned at
the start of a file.

?Cannot WRITE Devn: - device not ready
        The tape drive is not on-line. Push the on-line switch or button
on the tape transport.

?Cannot READ Devn: - device error
        You see this message if the system BPI value is different than
that at which the tape was recorded. (Use the SET BPI command if you
need to change the system BPI value.) You can also see this message if
there is a bad spot on the tape.

6.0 THE SKIP COMMAND

The SKIP command tells MTU.DVR to skip to the next end-of-file marker on the
tape.   Use SKIP to skip over initial header files at the front of the tape
and to skip over files on the tape you don't want to read.

IMPORTANT NOTE: Make sure that no other user is accessing the tape drive
you are using when you use SKIP.

To skip to the beginning of the next file, enter:

        .SKIP Devn: (RET)

where Devn: selects the magnetic tape unit you want to use. For example:

        .SKIP MTU4: (RET)

If the tape is currently positioned at the front of a file, the SKIP command
causes the tape drive to skip over the entire file to the beginning of the
next one. If the tape is currently positioned in the middle of a file, the
SKIP command causes the tape drive to move to the beginning of the next
file.

When SKIP finishes, it returns you to AMOS command level.

6.1 SKIP Error Messages

If you supply an invalid device specification, you see:

        ?File specification error

Make sure that the magnetic tape unit you specified is a valid system
device. (Type DEVTBL followed by a RETURN to see a list of all system
devices, including the magnetic tape units.)

Index

THE MAGNETIC TAPE FILE BACKUP PROGRAMS

1.0 INTRODUCTION

The Alpha Micro magnetic tape backup system consists of three programs:
FILTAP, the program to transfer files from disk to tape; TAPFIL, the program
to restore files to disk from tape; and TAPDIR, the program to list the
contents of a magnetic tape. These three programs when used with the AM-600
Magnetic Tape subsystem allow you to easily and rapidly back up and restore
both sequential and random files. The programs have full wildcarding
capability and allow you to back up multiple disk surfaces on a single tape.
The software also allows you to split a single backup across multiple reels
of tape.

NOTE: To use these programs, you must use Version 4.5 or later of the
magnetic tape driver, MTU.DVR[1,6], supplied on the disk with the programs.

These programs store and read data on the magnetic tape in a special
variable-length record format developed by Alpha Micro. This format was
optimized for the characteristics of the Alpha Micro computer and its
magnetic tape subsystem; it was not intended that this format be used for
data interchange with other, non-Alpha Micro computers. If you wish to
transfer data to other computers using the magnetic tape subsystem, you
should use the TAPE program; see Using the Magnetic Tape Utility Programs
in the AMOS Software Update Documentation Packet for more details on that
program.

IMPORTANT NOTE: It is important to remember that this set of programs was
designed as a mechanism for backing up disk files. That means that FILTAP
writes files to the tape along with their full device and account
specifications. (It also writes the date and time of backup.) Therefore,
when using TAPFIL to read a tape, if you want to access a specific file you
MUST specify the disk and account from which the file was backed up, as well
as the magnetic tape drive containing the tape you want to access. For
instance, if you stored a file on tape from DSK2:, you must specify the
device specification "DSK2:" when you restore the file from the tape. We
give specific examples in the sections below.

1.1 Wildcarding Features

All three programs have been designed to function as much as possible like
their disk-oriented counterparts (COPY and DIR) to make the magnetic tape
software as easy to use as possible.

FILTAP, TAPFIL, and TAPDIR use wildcard symbols and specification defaults
in the same way that COPY and DIR do. This is because, like COPY and DIR,
FILTAP, TAPFIL, and TAPDIR are "wildcard commands."

Wildcard commands differ from other AMOS commands in that besides accepting the standard AMOS file specification:

    Devn:filename.extension[p,pn]{/switches}

they also accept a variety of wildcard specifications. A wildcard file specification allows you to select multiple files with only one file specification. These file specifications can contain the special wildcard symbols *, ?, [], and ALL:. For example, to specify all files in all accounts of DSK0: that have a .BAS extension, wildcard commands permit the file specification:

    DSK0:*.BAS[]

(In the example above, the symbol * stands for all possible filenames.) In addition, the wildcard commands allow you to set account, device, and switch specification defaults. For example, the wildcard command line:

    .FILTAP [100,2]*.BAS,*.TXT[117,6],*.LST (RET)

sets the default account specification to [100,2]. The command line thus selects all .BAS and .LST files in account [100,2], and all .TXT files in account [117,6]. (Notice that the account specification setting the default occurs before the filename and extension rather than after it, as is the case with a standard AMOS file specification.)

1.1.1 Switches and Wildcard Commands - "Switches" are option requests. Each switch must begin with a slash, /. Remember that wildcard commands recognize two types of switches: "file" switches and "operation" switches. An operation switch applies to all of the file specifications on an entire command line no matter where it appears on that command line.

A file switch may apply to only specific file specifications, depending on where it appears on the command line. If a file switch appears at the end of a file specification, it applies only to the files selected by that specification. For example:

    .FILTAP *.BAS,*.TXT/QUERY,*.LST (RET)

in the command line above, the /QUERY switch (a file switch) applies only to those files selected by the *.TXT file specification. If a file switch appears before a file specification, it becomes the default switch, and applies to all of the files selected by the following file specifications unless it is overridden for a particular file specification by another switch, or until a new default is set. For example:

    .FILTAP *.BAS,/QUERY*.TXT,*.LST,*.MAC/NOQUERY,*.PRG (RET)

the /QUERY switch affects all files selected by the specifications *.TXT, *.LST, and *.PRG, but not the files selected by the specification *.MAC.

(Changed 1 July 1981)

Because wildcard file specifications are extremely powerful, they can have very widely ranging effects. If you are not familiar with the way in which wildcard commands work, be sure to read Chapter 9 of the AMOS User's Guide, (DWM-00100-35) before attempting to use wildcard specifications.

2.0 WRITING DISK FILES TO TAPE - THE FILTAP PROGRAM

The FILTAP program writes disk files to the magnetic tape. This program accepts a standard wildcard file specification which specifies the files to back up.

The FILTAP program operates in two modes: /APPEND and /NOAPPEND mode. In /APPEND mode, FILTAP searches for the last file on the tape and starts writing the new files immediately after any data already on the tape. In /NOAPPEND mode, FILTAP does not look to see if any data is already on the tape, but just starts writing files at the beginning of the tape. The default mode of operation is /APPEND. Of course, if the tape is blank (i.e., a new tape), you should specify the /NOAPPEND mode.

Call the FILTAP program by giving a wildcard specification for the files you wish to back up:

     .FILTAP{/switches} Filespec1{/switches}{,Filespec2{/switches}...} (RET)

The Filespec default is *.* and the account and device you are logged into. The default switches are /NOQUERY/APPEND.

The FILTAP program now asks you for the device code and unit number of the magnetic tape drive you want to write to:

     Enter tape unit number:

(FILTAP assumes a device code of MTU.)

FILTAP now writes the files to tape, listing each file as it transfers it. (NOTE: Because FILTAP also writes the date and time of the backup to the tape, you should use the system commands DATE and TIME to set the system date and time before you use FILTAP.)

2.1 Example

For example, to back up all files from disk device DSK0: to the magnetic tape drive MTU0: (starting at the beginning of the tape), use the following command:

```
.FILTAP DSK0:[]/NOAPPEND (RET)
Enter tape unit number: 0 (RET)
AMSORT.SYS[1,4] to MTU0:AMSORT.SYS[1,4]
                       .
                       .
                       .
Total of 1282 files transferred
.
```

2.2 Writing to Multiple Tapes

If all of the files you specified will not fit on one tape, FILTAP displays
the following message:

```
%Tape is full, please mount another tape then type RETURN to
%  continue, or type control-C to abort copy
```

If you wish to continue backing up files on another reel of tape, wait for
the current tape to finish rewinding, mount a new reel of tape, then type
RETURN on your terminal. The backup will continue on the new reel. If you
wish to abort the backup, type a Control-C.

2.3 FILTAP Switches

FILTAP provides the switches below:

/QUERY	or /Q	Ask user for confirmation before copying files (file switch).
/NOQUERY	or /NOQ	Don't ask for confirmation (default, file switch).
/APPEND	or /A	Write files to tape at the end of existing files (default, operation switch).
/NOAPPEND	or /NOA	Write files at beginning of tape (operation switch).

2.4 Error Messages

You may see the following error messages when using the FILTAP program:

```
?Cannot find DSK0:SCNWLD.SYS[1,4] or MEM:SCNWLD.SYS
```
> The FILTAP program needs this file to be able to process wildcard
> symbols in your file specification. This message can indicate
> that SCNWLD.SYS does not exist, or that you do not have enough
> memory to load the file into your partition.

?Cannot READ Devn - device does not exist
?Cannot READ Devn - device is not mounted
>   You tried to copy to or from a device that is not listed in the
           DEVTBL command in your SYSTEM.INI, does not have a driver in area
           [1,6] of the System Disk, is not file-structured, or is not
           mounted. ("Devn:" is the device you specified.)

%No file-oriented device corresponding to Devn: is mounted
>   You specified a device, but left off the unit number. FILTAP
           cannot find a logical unit that matches your specification. Try
           mounting the device.

%Tape is full, please mount another tape then type RETURN to
% continue, or type control-C to abort copy
>   There is no more room on the current reel of tape. Mount another
           reel and type RETURN to continue the backup process, or type a
           control-C to abort the backup procedure.

3.0 RESTORING DISK FILES FROM TAPE - THE TAPFIL PROGRAM

Use the TAPFIL program to transfer files back to the disk from tape. The
files must have been written to the tape via FILTAP. TAPFIL provides full
wildcarding, allowing easy selection of the files to be restored, as well as
automatic renaming facilities.

To use TAPFIL, enter the TAPFIL command followed by an output specification,
an input specification, and any optional switches:

   .TAPFIL{/switches} outspec = {inspec1{/switches}{,inspec2{/switches}...}} [RET]

The output specification defaults to the input specification. The input
specification defaults to *.* and the device and account you are logged
into. The default switches are /NOQUERY/DELETE.

TAPFIL now asks you for the device code and unit number specifying the tape
drive you wish to read from:

      Enter tape unit number:

TAPFIL assumes a device code of MTU.

TAPFIL then rewinds the tape and starts searching for the specified files.
As TAPFIL finds them on the tape, it transfers the files to the disk and
accounts you have specified.

The output specification you supply to TAPFIL is the specification of the
file(s) you wish to create. TAPFIL provides full wildcarding. Just as with
the COPY command, you may not copy files from one account to another unless:
1) the account you are copying from is in the same Project as the account to
which you are copying; or, 2) you are logged into the disk account into
which you are copying the files; or, 3) you are logged into the System

Operator's account, [1,2]. (As with the COPY command, logging into account [1,2] gives you certain privileges. The default account specification of the Outfilespec when you are logged into [1,2] is the wildcard account, []. Also, if you are logged into [1,2], TAPFIL will create the account you are copying to if it does not exist.)

The input specification is a list of the files you wish to copy from the tape. The input specification must give the exact specification of the file you wish to copy, including device and account of the file as it is stored on the tape.

3.1 Example

For example, assume you are logged into DSK0:[140,1] and you wish to copy a file from tape that was backed up from your own account (DSK0:[140,1]). Enter:

```
.TAPFIL = FILE.DAT (RET)
Enter tape unit number: 0 (RET)
MTU0:FILE.DAT to FILE.DAT
Total of 1 file transferred
```

Note that in the example above, the output specification defaulted to the input specification, and the input device and account defaulted to the device and account you are currently logged into.

If you want to copy a file from tape that was backed up from another device and account (DSK2:[1,4] for example), you would enter the following command:

```
.TAPFIL = DSK2:TEST.BAS[1,4] (RET)
Enter tape unit number: 0 (RET)
MTU0:DSK2:TEST.BAS[1,4] to TEST.BAS
Total of 1 file transferred
```

If you want to restore all the files stored on a tape to their original device and account, you would log into [1,2] and enter the following command:

```
.TAPFIL = ALL:[] (RET)
```

If you want to return all the files stored on a tape back to their accounts of origin in Project 110 of DSK0:, you would log into an account in Project 110 and enter the following:

```
.TAPFIL DSK0:[]=DSK3:[110,*] (RET)
```

If you include a filename and/or extension in your output specification, you can rename the copies of the files you are writing to disk. For example:

```
.TAPFIL DSK3:[]*.OLD = DSK1:[300,1?]*.MAC (RET)
```

copies all .MAC files from the tape backed up from accounts   [300,1?] on
DSK1: to the same accounts on DSK3:, and renames the file extensions from
.MAC to .OLD.

3.2 Restoring from Multiple Tapes

If you are restoring files from multiple tapes created by the FILTAP
program, you must enter separate TAPFIL commands for each tape.

3.3 TAPFIL Switches

TAPFIL provides the following switches:

/QUERY	or /Q	Ask user for confirmation before copying files (file switch).
/NOQUERY	or /NOQ	Don't ask for confirmation (default, file switch).
/DELETE	or /D	Copy over to an existing file, thereby deleting it (default, file switch).
/NODELETE	or /NOD	Don't copy over to any existing files (file switch).

3.4 Error Messages

You may see the following error messages when using the TAPFIL program:

?Cannot find DSK0:SCNWLD.SYS[1,4] or MEM:SCNWLD.SYS
    The TAPFIL program needs this file to be able to process wildcard
    symbols in your file specification. This message can indicate
    that SCNWLD.SYS does not exist, or that you do not have enough
    memory to load the file into your partition.

?Cannot READ Devn — device does not exist
?Cannot READ Devn — device is not mounted
    You tried to copy to or from a device that is not listed in the
    DEVTBL command in your SYSTEM.INI, does not have a driver in area
    [1,6] of the System Disk, is not file-structured, or is not
    mounted. ("Devn:" is the device you specified.)

%No file-oriented device corresponding to Devn: is mounted
    You specified a device, but left off the unit number. TAPFIL
    cannot find a logical unit that matches your specification. Try
    mounting the device.

?Tape is not file structured
> The tape you are trying to read was not written by the FILTAP
> program. The TAPFIL program can only read tapes written by
> FILTAP. Check to make sure you have mounted the correct reel of
> tape.

?Missing output specification
> You omitted the equal sign in your TAPFIL command line; TAPFIL
> couldn't tell which information was your input specification and
> which was your output specification.

?More than one output specification
> You may not supply more than one output specification.

?Files may not be transferred to RES:
> You may only add programs to system memory by using the SYSTEM
> command within your system initialization command file,
> SYSTEM.INI.

%Not copied - Destination file already exists
> You tried to copy to an existing file while the /NODELETE option
> was in effect.

?You are not logged in under [1,2], can't create [p,pn]
> You cannot copy from an account to a nonexistent account unless
> you are logged in under [1,2]. If you copy to a nonexistent
> account while logged under [1,2], TAPFIL will create the account.

?Output MFD is full
> The Master File Directory only has room for 64 entries. The
> transfer in progress would have created a new account, but there
> is no room in the MFD.

%Bypassing BADBLK.SYS[1,2]
%   BADBLK.SYS exists to prevent bad blocks
%   on a device from being allocated, and
%   should never be directly accessed.
> You cannot copy the BADBLK.SYS file, since this would lead to
> corruption of the file system.

?Device full
> There is no more room on the disk.

?Cannot OPEN Devn: - protection violation
> TAPFIL could not transfer the files. You are not allowed to write
> into an account you are not logged into unless the project
> number of that account is the same as the project number of the
> account you are copying from. You must either log into the System
> Operator's account, [1,2], or the account you are copying into to
> accomplish the transfer.

4.0 LISTING THE CONTENTS OF A TAPE - THE TAPDIR PROGRAM

The TAPDIR program displays a list of the files that have been stored on a
tape.   (NOTE: You may only use TAPDIR on a tape that was written via
FILTAP.) The TAPDIR program has been designed to be as similar to the DIR
program as possible.

To create a list of the files on a tape, enter the following command:

    .TAPDIR{/switch}{Listfilespec=}{inspec1{/switch}{,inspec2{/switch}}...} (RET)

Where the optional Listfilespec specifies where the output listing is to be
placed. If you specify no listfile or equal sign, the display goes to your
terminal. By specifying a listfile, you can send the display to a disk file
or printer.

The optional inspec allows you to select the files you wish to include in
the directory listing. The default is the device and account you are logged
into, and a file specification of *.*.

TAPDIR now asks for the device code and unit number of the drive containing
the tape that you want to get a directory listing of:

        Enter tape unit number:

Now TAPDIR creates the output listing, showing the relative position of each
file on the tape, the full file specification, the size of the file, whether
the file is a linked (L) or contiguous (C) file (that is, whether it is a
sequential or a random file), and the date and time that the file was
written to the tape.   At the end of the listing, TAPDIR gives the total
number of files and blocks that it has found.

4.1 Example

For example, to list all the files on a tape on your terminal, enter the
following command:

        .TAPDIR ALL:[] (RET)
        Enter tape unit number: 0
            1   DSK0: SYS   MAC   140,1   16 L   14-May-80 14:52:23
            2   DSK0: NBSORT MAC   140,1   4 L   14-May-80 14:52:25
            3   DSK0: FILTAP MAC   140,1   23 L   14-May-80 14:52:25
            4   DSK0: JANE   DAT   140,1   99 C   14-May-80 14:52:27
        Total of 4 files in 142 blocks

To create a file (DIRECT.LST) in the account and device you are logged into
that contains a list of all data (.DAT) files on the tape, enter the
following command:

        .TAPDIR = ALL:*.DAT[] (RET)

4.2 TAPDIR Switch

TAPDIR provides the following switch:

/KILL   or /K   Delete and replace existing Listfile if it has the
                        same specification as your Listfilespec. (Operation
                        switch.)

4.3 Error Messages

You may see the following error messages when using the TAPDIR program:

?Cannot find DSKO:SCNWLD.SYS[1,4] or MEM:SCNWLD.SYS
        The TAPDIR program needs this file to be able to process wildcard
        symbols in your file specification. This message can indicate
        that SCNWLD.SYS does not exist, or that you do not have enough
        memory to load the file into your partition.

?Cannot READ Devn - device does not exist
?Cannot READ Devn - device is not mounted
        You tried to copy to or from a device that is not listed in the
        DEVTBL command in your SYSTEM.INI, does not have a driver in area
        [1,6] of the System Disk, is not file-structured, or is not
        mounted. ("Devn:" is the device you specified.)

%No file-oriented device corresponding to Devn: is mounted
        You specified a device, but left off the unit number. TAPDIR
        cannot find a logical unit that matches your specification. Try
        mounting the device.

?Tape is not file structured
        The tape you are trying to read was not written by the FILTAP
        program.   The TAPDIR program can only read tapes written by
        FILTAP. Check to be sure you've mounted the correct reel of tape.

?More than one output specification
        You may not supply more than one output specification.

?Device full
        There is no more room on the disk.

%No such files
        TAPDIR was unable to find any files matching your input
        specification.

(Changed 1 July 1981)

SOFTWARE INSTALLATION INSTRUCTIONS FOR THE AM-120

1.0 INTRODUCTION

This document describes the software that accompanies the AM-120 Auxiliary I/O Controller board and gives instructions for installing that software. Some of the programs described below are new and some are programs from previous software releases that we have modified to take advantage of the features of the AM-120 board. This document also discusses the power fail detection and handling features of the AM-120 board for AM-100 and AM-100/T systems.

The AM-120 Auxiliary I/O Controller board provides the following features:

1. Two full RS-232 serial I/O ports (one of which contains a remote reset line). Connecting a terminal to the port with the remote reset line allows you to reboot the system from that terminal.

2. Three 8-bit parallel output ports and two 8-bit parallel input ports.

3. An interval timer. The length of the timer's interval is software selectable.

4. Power failure detection and handling. (See Section 7.0.)

5. A clock/calendar with battery backup. This clock/calendar maintains both the time and the date (including day of the week) even when your system is not on.

6. Memory error interrupt.

For more information on the AM-120, and for information on physically installing the board in your system, see the document Installation Instructions AM-120, (PDI-00120-XX).

2.0 THE SOFTWARE INCLUDED WITH THE AM-120

Below is a list of the software we supply with the AM-120 board:

CAL120.PRG	Calibration program for time-of-day clock oscillator on the AM-120.
DATE.PRG	Reads and sets system date from the AM-120.
TIME.PRG	Reads and sets system time from the AM-120.
AM120.IDV	Interface driver for the AM-120.
AM120.MAC	The source code for the AM120.IDV interface driver.

Note the command reference sheets in the AMOS System Commands Reference Manual, (DWM-00100-49) for the CAL120, DATE, and TIME programs. The .PRG files listed above must appear in account [1,4] of your System Disk, and the .IDV file must appear in account [1,6] of your System Disk.

3.0 THE AM120.IDV INTERFACE DRIVER

The AM120.IDV interface driver allows you to use the two serial ports on the AM-120 for terminals or printers. You must also use the driver if you want to use the parallel ports in interrupt driven mode, the memory error interrupt, or the interval timer. (The AM120.IDV interface driver handles all of these interrupts.)

We have designed the AM120.IDV driver so that installing an AM-120 board in a system that uses an AM-100 CPU allows you to use the AM-120 parallel ports, memory error interrupt and interface timer in exactly the same way as you would use those features on the AM-100/T CPU board. This means you may use software originally written for the AM-100/T on the AM-120 with no modification. NOTE: To use the parallel ports in interrupt driven mode, or the interval timer on a second AM-120 board in your AM-100 based system, you must modify the AM120.IDV program by inserting your own code that handles the function you want to perform. We have provided the source code for the interface driver to allow you to do so.

If your system uses an AM-100 CPU, the AM120.IDV driver allows up to two AM-120 boards in the same system. If your system uses an AM-100/T CPU, the interface driver allows one AM-120 board in your system; that board must be physically addressed to the "alternate second board address," and you must reference the serial ports on it as ports 2 and 3. (In this case, the AM-100/T acts as the first auxiliary I/O controller board in the system.) (See the document Installation Instructions AM-120 for information on installing the AM-120 board.)

4.0 USING THE AM-120 SERIAL I/O PORTS

To tell the system that you want to use a terminal or printer that is connected to one of the serial ports on the AM-120, you must include an appropriate TRMDEF statement in your system initialization command file that references the AM-120 interface driver. (For information on the TRMDEF statement and on modifying the SYSTEM.INI file, see the document The System Initialization Command File in the "System Operator's Information" section of the AMOS Software Update documentation Packet.) The TRMDEF statement for the AM120.IDV is very similar in form to the TRMDEF statement for the AM-310 interface board. The interface statement portion of the TRMDEF statement takes this form:

        AM120=I/O-port-address{:baud-rate-code}

(Changed 31 October 1981)

where you may optionally supply a baud rate code for the terminal. The I/O-port-address is a number from 0-3 and selects the serial port you want to use. (If you are using one AM-120 board with an AM-100, you have serial ports 0 and 1 available; if you are using two AM-120 boards, you have serial ports 0-3 available.) As an example, a TRMDEF statement for a terminal connected to the AM-120 on a system using an AM-100 CPU might look like this:

         TRMDEF TERM1,AM120=1:37316,SOROC,100,100,100

The optional baud rate codes that you can specify are:

          :30316   50   baud
          :30716   75   baud
          :31316   110   baud
          :31716   134.5 baud
          :32316   150   baud
          :32716   300   baud
          :33316   600   baud
          :33716   1200   baud
          :34316   1800   baud
          :34716   2000   baud
          :35316   2400   baud
          :35716   3600   baud
          :36316   4800   baud
          :36716   7200   baud
          :37316   9600   baud
          :37716   19200   baud

5.0 SUPPORTING UTILITY PROGRAMS

We have modified the DATE and TIME programs so that you can set the system date and time from the AM-120 clock/calendar, and you can reset the date and time maintained by the AM-120.

On a system with an AM-120, DATE also maintains the day of the week as part of the date. The DATE program now supports European date format as well as American format. (A date in European format looks like this: 26 January 1982; a date in American format looks like this: January 26, 1982.) You may optionally request that DATE automatically reset the date at midnight.

A new program, CAL120, allows you to use a high-frequency counter to calibrate the AM-120 oscillator.

For information on DATE, TIME, and CAL120, see the command reference sheets for DATE, TIME, and CAL120 in the AMOS System Commands Reference Manual (DWM-10010-49).

5.1 Important Note on DATE and TIME

Note that as of AMOS Release 4.6, if you use TIME and DATE within the system initialization command file before the final SYSTEM command, you must have an AM-120 in the system or your system may have trouble booting under that SYSTEM.INI.

6.0 INTERNAL FORMATS FOR DATE AND TIME

The formats that AMOS uses for storing the system date and time are the same whether or not your system contains an AM-120 board.

AMOS stores the time of day in two words in system memory as the number of system clock ticks since midnight. (The "system clock" is not the AM-120 clock/calendar or interval timer, but is the clock on the CPU board derived from the AC line frequency.) For information on these two words in system memory (TIME and TIME+2), see the AMOS Monitor Calls Manual, (DWM-00100-42).

AMOS stores the system date in two words in system memory. These two words contain the month, the day, the year, and some flag bits. The byte at the word DATE contains the month, the byte at the word DATE+1 contains the day, and the byte at DATE+2 contains the last two digits of the year. The byte at DATE+3 contains some flag bits, only two of which are currently used. Within the flag byte, bits 0-6 are currently unused; a one in bit 7 means that an AM-120 is present in the system; and, a one in bit 6 means that DATE is using the European date format.

NOTE: Some programmers have in the past assumed that DATE+3 will always contain zero. This is not true; programs that depend on this assumption may experience problems when used on a system that uses the AM-120 software.

7.0 DOUBLE AND SINGLE POWER FAIL DETECTION AND HANDLING

One of the most important characteristics of the AM-120 board is its optional power fail detection and handling features.

Beginning with AMOS Release 4.5, the monitor supports a two-level power fail (also called "double power fail") option for an AM-100/T CPU (Revision D and later) used with the AM-120 board. (You must also use an Alpha Micro chassis that supports our power fail signal.)

A power fail can cause serious problems because it can result in spurious disk writes or garbling of data currently being written to the disk when the power is restored. NOTE: A power fail is only dangerous when it causes the AC line to drop long enough for the DC power line to drop below the minimum required by the computer's circuits.

Below we discuss both single and double power fail detection and handling. For information on configuring your hardware to implement the power fail

options discussed below, see the hardware manual <u>AM-120 Auxiliary</u> <u>Input/Output Circuit Board</u>, (DWM-00120-00). Also see <u>the installation</u> instructions for the AM-100/T and the AM-120 contained in the <u>Alpha Micro</u> <u>Integrated Systems User's Guide</u>, (DWM-00101-00).

7.1 Single Power Fail Option

If your system contains an AM-100 CPU and an AM-120, your system can detect when DC power <u>or</u> AC power (depending on board setup) drops below a preset limit. The AM-120 generates a reset so that when power is restored, the system reboots. This helps to minimize the chance of damage to the data on your disk drives.

If your system contains just an AM-100/T, your system can detect a drop in the AC power line. If the AC power line drops below a safe level, the AM-100/T asserts the reset line. Then, when power to the AC line is restored, the system automatically reboots. This gives even better protection than checking DC power since an AC power fail is detected before the DC power fails. However, short duration AC power drops may not affect DC power but will still cause the system to reboot, resulting in unnecessary rebooting of the system.

7.2 Double Power Fail Option

If your system contains an AM-100/T (Rev D and later) and an AM-120, your system can make use of the double power fail option. The AM-100/T checks the AC power line while the AM-120 checks the DC power line. If the AC line drops below a certain level, rather than automatically rebooting, the AM-100/T pauses without initiating any new disk transfers until it detects that the faulty AC line condition has been corrected. The AM-120 generates a reset only if the DC power line has also dropped below a safe level (causing the system to reboot when both AC and DC power are restored). If the DC power has not been affected and the AC power is restored, the AM-100/T continues processing without rebooting the system.

This protects your system against unnecessary rebooting while still minimizing the chance of damage to the data on your disk from spurious and garbled writes.

SOFTWARE INSTALLATION INSTRUCTIONS FOR THE AM-710 MEMORY BOARD

The AM-710 memory board is a 128K byte memory board produced by Alpha Micro. One of the features of this board is its ability to detect parity errors. This document discusses the changes you must make to your system initialization command file when installing the AM-710 memory board. We also discuss what happens when a parity error occurs.

For technical information on the AM-710 board, see the Alpha Micro Integrated Systems User's Guide, (DWM-00101-00). That manual will also tell you how to set the AM-710 I/O port address and how to address the boards.

NOTE FOR BANK SWITCHING SYSTEMS: If your system bank switches memory, you will also want to consult the Alpha Micro Integrated Systems User's Guide for information on what on- and off-constants to specify in the MEMDEF commands that define the memory banks made up of your AM-710 memory boards.

1.0 ENABLING PARITY ERROR DETECTION FOR THE AM-710

To enable the AM-710 memory board's parity error detection feature, you must use the PARITY command. Although you may use this command at AMOS command level, you will almost always want to use it within your system initialization command file, SYSTEM.INI. (For information on changing your system initialization command file, see The System Initialization Command File, in the "System Operator's Information" section of the AMOS Software Update Documentation Packet.)

Place the PARITY command in your system initialization command file after the MEMERR command. Type PARITY followed by the I/O port addresses of the AM-710 boards in your system (separating the addresses by commas). For example:

        PARITY 100,101,102

This I/O port address is three digits for octal numbers and two digits for hexadecimal numbers. (NOTE: Do not enter hex numbers unless you have used the SET HEX command for your job.) If you have more I/O port addresses than will fit on one command line, you may enter multiple PARITY command lines. For example:

        PARITY 101,102,103,104,105
        PARITY 106,107,110,111,112

2.0 THE MEMERR COMMAND

The AM-710 memory board requires the presence of the MEMERR command in the
system initialization command file as well as the use of the PARITY command.
If your system uses only AM-710 boards, do not specify an address after the
MEMERR command. For example:

          MEMERR
          PARITY 100,101,102

If your system contains Piiceon 32K word boards as well as one or more
AM-710s, specify the MEMERR port address required by the Piiceon memory
boards. For example:

          MEMERR 250
          PARITY 100,101,102

(If your system uses Piiceon 32K word memory boards, refer to the Alpha
Micro Integrated System User's Guide for information on the error I/O
address to supply to the MEMERR command.)

3.0 PARITY COMMAND ERROR MESSAGES

Section 4.0 below discusses the error messages you can see when a parity
error actually occurs.   In addition, there are several other messages you
can see that result from the improper use of the PARITY command itself:

-    ?There is no AM-710 at port address xxx
          Where xxx is an I/O port address you specified on your PARITY
     command line. This address did not match the jumpered I/O port address
     of any of the AM-710 boards in your system. Check the PARITY command
     line to make sure that you entered the port address correctly; then
     check the memory boards to make sure that their port addresses are
     jumpered correctly.

     ?Command format error
          You did not supply any I/O port addresses, or in some other way used
     an improper format.

4.0 WHAT HAPPENS WHEN A PARITY ERROR OCCURS?

Your system can respond in a variety of ways when a parity error occurs
depending on your system's particular configuration:

AM-100 CPU and AM-710 memory boards - Memory error detection is not
supported.

If your system contains an AM-100/T CPU (Pre-Revision F) and one or more
AM-710 memory boards:

If a parity error is detected, the monitor sends a "5" to the status display port and halts the CPU.

If your system contains an AM-100/T CPU (Revision F or later) (or an AM-100 CPU and an AM-120 Auxiliary I/O Controller) and one or more AM-710 memory boards:

If a parity error occurs, the monitor will send a "9" to the status display board, and will display the message:

?AM-710 parity error for job xxxxxx

on the Operator's terminal. (The Operator's terminal is the terminal attached to the job the system came up under.) If the job in whose memory parition the parity error occurred is attached to a real terminal (as opposed to a pseudo terminal), that user sees the message:

?Parity error

The monitor then halts that job.

When a parity error is reported, the System Operator should inform all users still running that a memory error has occurred, and ask them to finish up what they were doing and logoff. Then the System Operator should open the computer and inspect each AM-710 memory board.

When a parity error occurs, you can only identify which memory board encountered the error by looking at each AM-710 memory board to see if its LED indicators are lit. Make note of the memory board on which the parity error occurred.

At this point you must decide whether to continue operating with the memory board that generated the parity error or whether to swap the board out for another one. Parity errors can reflect a transient problem that surfaces once and then is never seen again. You may therefore wish to continue running the system as is until a parity error occurs again. If the parity error occurs on the same memory board more than once or twice in a great while, you probably will want to remove the questionable memory board and return it to your dealer for inspection and/or repair.

To reset an AM-710 memory board's parity error indicator, you must perform a hardware reset on your system.

4.1 Other Memory Errors

Note that if your system uses Piiceon 32K word memory boards as well as AM-710 boards, in addition to parity errors reported by AM-710 boards, you may also encounter other types of memory errors reported by the Piiceon boards. If a Piiceon 32K word memory board (only for use with the AM-100/T or AM-100/AM-120) detects an uncorrectable memory error, the monitor sends a

"5" to the status display port and halts the system. You may see which
Piiceon board encountered the error by looking at the boards' LED indicators
as discussed in the section above.

MAGNETIC TAPE DRIVER MONITOR CALLS

1.0 INTRODUCTION

This document discusses the AMOS monitor calls you can include in your assembly language programs that allow your programs to access the magnetic tape unit driver, MTU.DVR. For information on the use of the magnetic tape utility programs, refer to the document Using the Magnetic Tape Unit, (DWM-00100-52), in the "User's Information Section" of the AM-100 documentation packet. That document also defines some of the terms we use in the discussion that follows. Before you begin to use MTU.DVR, make sure that 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).

For information on the format of an AMOS monitor call in your assembly language program, see Chapter 1, "Communicating with the AM-100 Monitor," in the AMOS Monitor Calls Manual, (DWM-00100-42).

2.0 THE MONITOR CALLS

Below are the monitor calls you can use to access MTU.DVR. In addition to the calls below, you can use the READ and WRITE monitor calls to input and output data to and from the magnetic tape unit. Use READ and WRITE for this purpose in the same way that you would use them to perform disk I/O.

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 that you want to issue a REWIND to. If an error occurs as the result of this call, you see the standard system file operation error messages. (For example: ?Cannot INIT Devn: - device does not exist.)

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 that you want to write a file-mark to. If an error occurs as a result of this call, you see the standard system file operation error messages.

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 that you want to issue a find file-mark command to. The FMARK call causes the MTU driver to read forward on the specified tape until it finds a file mark. Standard file operation error messages can occur when you use FMARK.

4. 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 status bits TAPST returns are:

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.
3	Load point	Indicates that tape is at load point.
4	File protect	Indicates that write-enable ring is not on tape reel, and that you cannot write to the tape.
5	Rewinding	Indicates that tape unit is in process of rewinding.
6	On-line	Indicates that tape unit is on-line.
7	Ready	Indicates that tape unit is ready for reading or writing.

SOFTWARE INFORMATION FOR THE USER OF THE AM-960 STATUS DISPLAY

The AM-960 Status Display is currently available with the AM-100/T Winchester-based disk systems offered by Alpha Micro. The AM-960 provides a two-digit hexadecimal display which is visible through the front panel. The monitor uses this display to report current system status. This provides a way for you to diagnose various problems that can occur on the system. In addition, your software may also use the AM-960 to report program and system status by sending user-defined status codes to the display.

1.0 REQUIREMENTS

The hardware requirements for the AM-960 are:

1. A new-style CPU chassis front panel that provides bezel mounting for the display.

2. An AM-100/T CPU revision F00 or later. No other Alpha Micro CPU supports the AM-960.

(Refer to the document Installation Instructions Status Display, (PDI-00960-00), for hardware information on the AM-960.)

The only software requirement for the AM-960 is that your system must be running AMOS Version 4.5 or later.

2.0 USING THE AM-960

If the system is operating normally, the AM-960 shows a blank display. (This is code 00.)

Various system functions (such as boot up) can cause other codes to be sent to the display if a problem occurs. (See the section below for the codes currently used by Alpha Micro software.)

In addition, your software can also use the AM-960 to report activities or problems. To use the AM-960, issue a write command to I/O ports 0-3. A software write to any of these four I/O ports will produce the same result-- the number you specify will be sent to the AM-960 display.

The status codes available to your programs are the hexadecimal numbers 80 to FF.

3.0 STATUS DISPLAY CODES DEFINED BY ALPHA MICRO

Alpha Micro has reserved the status codes 0 through 7F (hexadecimal) for reporting system status. Of those numbers, the following codes currently have been allocated and indicate the following conditions:

Display Codes	Condition
00	System functioning normally (display blanked)
01	No DSK0:[1,4] found during bootup
02	No DSK0:SYSTEM.MON[1,4] found during bootup
03	Device error during bootup
04	System ran out of queue blocks
05	Memory error or AM-710 parity error (with AM-100/T, Rev C)
06	Bootstrap loader checksum error (bad loader PROM or bad memory)
07	Bootstrap device did not pass self-test
08	AC power below threshold
09	AM-710 parity error (with AM-100/T, Rev F or later)
0A	Second device table search failed

4.0 DISPLAY DIAGNOSTIC (LEDIAG)

To make sure that the AM-960 properly displays all possible status codes, use the LEDIAG program:

    .LEDIAG (RET)

This program sends all possible status codes to the AM-960 display for your visual inspection.

To give you an idea of how to use the AM-960 through your own software, we provide the sample listing below of the LEDIAG program. This program is written in AlphaBASIC and uses the IO function to access the I/O ports. (Refer to the AlphaBASIC User's Manual, (DWM-00100-01), for information on the IO function.)

(1 July 1981)

```
! LEDIAG program to test the AM-960 Status Display
!
!   This program wwrites 0 to FF to the display port, thus testing the
!   display and port decoding.  A 1/4 second pause occurs between each
!   display.
!
START:  PRINT : PRINT "* * AM-960 display board test * *" : PRINT
        INPUT "Type <CR> to start test: ";C$
        !
        ! Loop for I = 0 to FF and output each value to display
        FOR DISPLAY'VALUE = 0 TO 255
           !
           ! Output value to the display
           IO(0) = DISPLAY'VALUE
           !
           ! Tell user what we did.
           PRINT "Testing "; STR$(DISPLAY'VALUE); CHR$(128+13);
           !
           ! Pause for 1/4 second or so.
           FOR TIMER = 1 TO 100 : NEXT TIMER
           !
        NEXT DISPLAY'VALUE
        !
        ! We're done.
        PRINT : PRINT : PRINT "* * End of test * *"
        END
```

SOFTWARE NOTICE FOR AM-420 USERS

October 1981
Revision A01

This document reflects AMOS versions 4.6 and later

Table of Contents

1.0 INTRODUCTION

This document contains software information for users of those Alpha Micro
systems employing disks that run under the control of the AM-420 Hard Disk
Controller.   Disks that run under the control of the AM-420 Hard Disk
Controller differ somewhat from the other hard disks supported by Alpha
Micro.   They therefore require different techniques for data backup, disk
formatting, and disk initialization; the AM-420 Controller performs the
specialized control these disks require.

The reason the AM-420 controlled disks differ is because they are built
using what is known as "Winchester technology."   A Winchester technology
disk drive is sealed from the environment and permits no dust or other
pollutants to reach the media surface. This allows the data read/write
heads of the drive to be much closer to the media surface than are the
heads of conventional disks, which causes an increased overall data density
and therefore a more cost-effective system.

Alpha Micro provides special techniques and equipment for backup of the
Winchester technology disk systems.   Winchester technology disk drives
cannot have removable disk cartridges, because the platters themselves are
completely sealed within the drives. Backing up a Winchester technology
device requires another kind of device for backup, which of course must also
be as cost-effective as possible. Alpha Micro offers three backup systems
for most of its Winchester technology devices, or two backup systems for the
very high capacity disks. One of the backup devices is a built-in floppy
disk drive. Another is based around a common Video Cassette Recorder (VCR).
A third backup system employs the inexpensive 1/4" streaming tape drive.
The VCR and 1/4" streaming tape drive backup systems are very effective for
handling large storage backup. You can find more information about the VCR
and 1/4" streaming tape drive backup systems in the manuals Software
Information for the AM-610 Video Cassette Recorder Interface
(DSS-10000-24), and Software Information for the AM-620/621 1/4" Streaming
Tape Drive (DSS-10000-23).

1.1 Summary of System Types

The term "system," when referring in this document to a data-processing
device using Winchester technology storage, means a complete Alpha Micro
computer, including CPU, chassis, mother board, backup device hardware, all
memory, control and other necessary cards installed, connecting cables,
power supplies, operable software, and so on. "Subsystem" is defined in
this document as a disk drive unit configured in such a way that it can be
(indeed, must be) attached to an Alpha Micro computer having all the above
components, where the Winchester disk drive unit may act as a dependent
storage device.

The following table summarizes the types of Winchester technology-based systems and subsystems available from Alpha Micro.

Alpha Micro Desig- nation	Yield (in Mega- bytes)	Sur- faces	Priam Drive Model	Type (Sys) or (Sub)	Backup System Type	CPU Type Required	Drive Dia- meter
AM-1020F	8.5	4	1070	Sys	Floppy	AM-100	8"
AM-1021F	8.5	4	1070	Sys	Floppy	AM-100T	8"
AM-1020V	8.5	4	1070	Sys	VCR	AM-100	8"
AM-1021V	8.5	4	1070	Sys	VCR	AM-100T	8"
AM-1020S	8.5	4	1070	Sys	1/4" Stream.	AM-100	8"
AM-1021S	8.5	4	1070	Sys	1/4" Stream.	AM-100T	8"
AM-1041F	32	5	3450	Sys	Floppy	AM-100T	8"
AM-1041V	32	5	3450	Sys	VCR	AM-100T	8"
AM-1041S	32	5	3450	Sys	1/4" Stream.	AM-100T	8"
AM-1061V	60	3	6650	Sys	VCR	AM-100T	14"
AM-1061S	60	3	6650	Sys	1/4" Stream.	AM-100T	14"
AM-421	8.5	4	1070	Sub	Optional	Either	8"
AM-422	32	5	3450	Sub	Optional	Either	8"
AM-423	60	3	6650	Sub	Optional	Either	14"

VCR = Video Cassette Recorder. 1/4" Stream. = 1/4" Streamer Tape Drive

1.2 Concepts

In the following paragraphs, we define some words and concepts you need to be familiar with while reading this document.

1. Physical address - The AM-420 controller has four "physical addresses," or ports, to which Winchester technology disk devices may be physically attached by cables. When the system boots, AMOS tells the controller which peripheral disk drive is associated with address 0, which is associated with address 1, and so on. Then the devices are given names, and similar devices are given numbers.

2. Physical device - A physical disk device includes one or more iron-oxide coated, magnetically sensitive disks, and read/write heads to record and read back data that is recorded in tracks on the disk(s).

3. Logical device or logical unit - A logical device is an abstraction. The computer cannot correctly place or access groups of data on a disk unless it has a name for the disk, and master and file directories of the contents of the disk available to it. So while the physical disk may be a Winchester technology AM-1020 hard disk to you, it is PLD0:, PLD1: and PLD2: or the like to the computer. It is important to remember that the logical devices that a disk system is configured to have are only incidentally related to the physical disk drive. A single logical unit may

(Changed 31 October 1981)

contain far fewer records than one physical device can store, or can contain as many records as entire physical device can store. Exactly how the logical devices are configured is up to you, within certain physical limitations, as you will see in the pages below.

4. DSKO: - The first logical device of a system (even if the whole drive is configured as one logical device) is known as the System Disk, or DSKO:.

1.3 Important Notes

There are several things you must keep in mind before you begin to use an AM-420 controlled disk:

1. Alpha Micro's Winchester technology disk systems and subsystems may be configured in a number of ways. The 102X series and the 104X series systems and the AM-421 and AM-422 subsystems may each be configured as a single logical device, or as up to 8 logical devices. The 106X series systems and the AM-423 subsystem may be configured as from 2 to 8 logical devices.

2. You may use the AMOS disk diagnostic tests REDALL, RNDRED, and DSKANA on an AM-420 controlled disk system. You may also use DSKCPY to copy from one logical unit to another.

3. Before a Winchester technology disk drive can be trusted to perfectly store data, it must be certified. Any flaws in the disk's iron oxide coating or other physical attributes must be discovered and compensated for by the computer before the drive is put into normal service.

    You will use the CRT420 command to certify all disks tht run under the control of the AM-420. Note that CRT420 certifies the entire physical device, even though you may have multiple logical units on that device. Do not use the SYSACT initialize command to initialize the first logical unit on a Winchester technology device-- the CRT420 program initializes the first logical unit for you. However, you will use the SYSACT Initialize command to initialize every logical unit on the physical device except the first. You will also use SYSACT to add user accounts to any Winchester technology logical unit.

    CRT420 makes a very thorough and reliable search for bad disk areas. Because the amount of storage on a Winchester technology disk is so great, CRT420's certification process takes longer to accomplish than do those of other AMOS certification programs.

2.0 THE PRIAM DISK CONFIGURATION

There are currently 16 configurations of systems and three configurations of subsystems available that run under control of the AM-420.   The Winchester technology disk drives that are used are manufactured by Priam. They are the Priam 1070 (in the AM-102X series systems and making up the AM-421 subsystem), the Priam 3450 (in the AM-104X series systems and making up the AM-422 subsystem), and the Priam 6650 (in the AM-106X series systems and making up the AM-423 subsystem).

The Priam 1070 disk contains 4 surfaces of 2.125 megabytes each. These surfaces may make up one logical unit and be accessed as such, or can be configured to be up to 8 separate logical units. Each Priam 1070 surface contains 185 (#0-184) tracks, with 22 sectors per track, to give a total of 16280 (decimal) sectors (or disk blocks) on the drive. The Priam 1070 also contains 5 spare (or "alternate") tracks per head. Each disk block contains 512 bytes. The Priam 1070 requires a total bitmap size of 1018 words.

The Priam 3450 disk contains 5 surfaces of 6.4 megabytes each. These surfaces may also make up one logical unit and be accessed as such, or can be configured to be up to 8 separate logical units. Each Priam 3450 surface contains 515 (#0-514) tracks, with 24 sectors per track, to give a total of 61800 (decimal) sectors (or disk blocks) on the drive. The Priam 3450 contains 10 spare (or "alternate") tracks per head. Each disk block contains 512 bytes.   The Priam 3450 requires a total bitmap size of 3863 words.

The Priam 6650 disk contains 3 surfaces of 20 megabytes each. These surfaces must make up a minimum of two logical units and be accessed as such, or can be configured to be up to 8 separate logical units. Each Priam 6650 surface contains 1098 (#0-1100) tracks, with 36 sectors per track, to give a total of 118584 (decimal) sectors (or disk blocks) on the drive. The Priam 6650 contains 23 spare (or "alternate") tracks per head. Each disk block contains 512 bytes. The Priam 6650 requires a bitmap size of 3706 words. (Note that the bitmap size of the Priam 6650 is smaller than that of the Priam 3450. Since the larger capacity Priam 6650 must be split into at least two logical devices, the largest bitmap size you need to allocate for a logical disk is 3706 words.)

If you place a device's bitmap in switchable system memory (on a bank switching system) or in supplemental system memory (on a system that uses the AM-700 Memory Partition Controller), you must remember to load the driver for that device into system memory via the SYSTEM command within the system initialization command file. (See The System Initialization Command File in the "System Operator's Information" section of the AMOS Software Update Documentation Packet for information on placing bitmaps in switchable or supplemental system memory.)

The device driver program for the Priam disk is 420DVR.DVR[1,6] on the System Disk. You will use the FIX420 program to configure the 420DVR.DVR driver for your particular drive.   The bootstrap loader program is PRILOD.PRG[1,4] on the System Disk.

(Changed 31 October 1981)

3.0 USING AM-420 CONTROLLED DEVICES AS SYSTEM AND NON-SYSTEM DEVICES

The paragraphs below discuss adding an AM-420 controlled disk drive to your system either as the System Device or as a non-System Device.

IMPORTANT NOTE: Hardware settings on the disk controllers on your system define the device that the system attempts to boot from. However, on a Alpha Micro 102XF series or 104XF system with the built-in floppy disk drive, the system first tries to boot from the floppy disk if a disk is installed in the floppy disk drive and the door is closed. This is so you can keep a backup System Disk and can boot off it if required. Therefore, you must not leave any floppy disk in the floppy disk drive with the door closed if you do not intend to boot from the floppy disk.

Similarly, if the enabling boot PROMs are installed on the system, every AM-10XXV and AM-10XXS system will try to boot from its backup device if a boot cassette or boot cartridge is installed on the backup device. For instance, if you have a boot cassette in your AM-10XXV's VCR backup device (that is, you have a warm boot monitor generated by WRMGEN on the cassette currently in the Video Cassette Recorder), and the VCR is set up to operate, the system will boot from the VCR. (For further information on booting from these backup devices, see the manuals Software Information for the AM-610 Video Cassette Recorder Interface (DSS-10000-24), and Software Information for the AM-620/621 1/4" Streaming Tape Drive (DSS-10000-23).)

3.1 Using an AM-420 Controlled Device as the System Device

The AMOS system software is already on your AM-420 controlled System Device when it is installed. To boot from that device:

1. Turn on the computer, holding down the reset button. Turning on CPU power causes the system disk drive to sequence up. Winchester technology subsystems on your system have their own power supplies, so you must turn them on individually. When everything you want on is powered up, press the reset button again. If your system initialization command does not automatically MOUNT the subsystems upon booting, you may use the MOUNT program at AMOS command level to sequence up the subsystems. (When you sequence down your subsystems, you must sequence down all disks (the System Disk, DSK0: last) before turning off your CPU. Use MOUNT and the /S switch to sequence down each drive. If you do not follow this procedure, it is quite possible that the data on your disks will be damaged.) For more information on MOUNT and the switch /S, see Section 8.0, "Using MOUNT."

    NOTE: We recommend that you always leave your system powered up, even when you are not using it. Although the disk surfaces are sealed off from the atmosphere, the filter system protects the entire device from pollutants. Also, an electronic device can usually be left on for several days with less wear and tear than if it is powered down and then powered back up later.

(Changed 31 October 1981)

2.   The drive may be a single logical unit, 2 logical units, or up to 8
     logical units. But the logical disk unit the system boots off of
     is ALWAYS known as DSK0:. (If, in the case of the 10XXF series
     systems, the system boots off the system-backup floppy disk, the
     built-in floppy disk drive is known as DSK0:.) The System Disk
     contains the programs SYSTEM.MON[1,4] and   SYSTEM.INI[1,4];
     therefore the system recognizes it as a System Disk and tries to
     boot the system from it.

     Remember that the system can do a warm boot off a VCR or a 1/4"
     Streaming tape drive. However, these special devices do not become
     any sort of System Device. The warm boot monitor is incomplete,
     designed for error recovery and the like. For more information on
     warm boots, see the manuals Software Information for the AM-610
     Video Cassette Recorder Interface (DSS-10000-24), and Software
     Information for the AM-620/621 1/4" Streaming Tape Drive
     (DSS-10000-23).)

3.2 Using an AM-420 Controlled Device as a Non-System Device

If you boot the system from some other type of disk device, the AM-420
controlled device is called a non-System device. You will need to follow
some simple steps before you can access the AM-420 controlled disk:

1.   Generate a driver program by logging into DSK0:[1,6] and running
     FIX420.   See the discussion in this document on FIX420, and also
     the FIX420 reference sheet in the AMOS System Commands Reference
     Manual (DWM-00100-49), for details on how to use FIX420.

2.   Now you need to define the Winchester technology disk as from 1 to
     8 logical devices (or from 2 to 8 if the disk is a Priam 6650
     disk.) To do this, edit your system initialization command file,
     SYSTEM.INI, and place the device name and the logical unit
     configuration on a DEVTBL command line. This adds the Winchester
     device to your System Device table. (For details on using the
     DEVTBL command to define logical devices and on using the BITMAP
     command to define disk bitmaps, see Section 7.0, "Modifying the
     System Initialization Command File," in this document. For more
     general information, see The System Initialization Command File
     in the "System Operator's Informaticn" section of the AMOS Software
     Update Documentation Packet.)

3.   Add BITMAP commands to the SYSTEM.INI to define bitmap areas for
     the AM-420 controlled device.   Remember that you can set up your
     system so that part of system memory resides in a switchable area
     of a memory bank (for a bank switching system) or in supplemental
     system memory (for an MPC system). You can thus place your bitmaps
     in switchable or supplemental system memory and so reduce the size
     of the monitor.   (If you do place your bitmaps in switchable or

supplemental system memory, you <u>must also</u> load the driver into
system memory.) Now reboot the <u>system with</u> your new SYSTEM.INI.

4. Now, if the disk contains no data, run CRT420 and then use SYSACT
    to initialize all but the first logical unit. <u>Then</u> you can
    access the AM-420 controlled device.

If you want more information on adding new disk devices to your system,
refer to the document <u>Defining Non-System Disk Devices</u> in the "System
Operator's Information" section of the AMOS Software Update Documentation
Packet. Section 3.0 of that document, "Building a System on a New Device,"
explains how to convert the Winchester technology device to be the System
Device after you have added it to your system as a non-System Device.

4.0 THE CERTIFICATION PROGRAM, CRT420

Winchester technology, which provides a much higher density of data than
more conventional drives (e.g., the CDC Hawk), also makes media flaws (i.e.,
tracks containing bad disk blocks) a much more likely possibility.
Therefore it is necessary to use the CRT420 command to certify a Winchester
technology drive before you use it. This certification process identifies
any bad tracks and heads on the disk, formats the disk, and initializes it.

The CRT420 program certifies an AM-420 controlled physical unit by writing
to, reading from and verifying every track and head of the disk. It also
formats and initializes the disk. It creates the file BADBLK.SYS[1,2] and
lists any bad tracks and heads on the media in that file. AMOS assigns
spare tracks (called <u>alternate tracks</u>) to the bad tracks, to be used in
their place whenever the system tries to access a block on the bad tracks.
If you want to see what those bad tracks are, you may use the command BADBLK
to display the contents of BADBLK.SYS[1,2], where the bad track information
is stored. (See Section 5.0 of this document, and also the BADBLK reference
sheet in the <u>AMOS System Commands Reference Manual</u> (DWM-00100-49), for
details on how <u>to use BADBLK.</u>)

4.1 Hints and Restrictions

There are several things you should keep in mind before using the
certification program:

1. Only the System Operator may run CRT420. Log into the System
    Operator's account, [1,2], before using the program.

2. You may ONLY use CRT420 on disks that run under control of the
    AM-420 Controller.

3. A physical unit does not have to be mounted before you certify it;
    CRT420 mounts the disk for you.

4. CRT420 writes data in every byte of the physical unit you are
    certifying.   That is, it overwrites all data on all the logical
    units of the physical drive. If there is data anywhere on the
    physical unit that you need to preserve, make sure that you back it
    up before using CRT420.

5. CRT420 communicates directly with the AM-420 Controller without
    going through the driver program.   Therefore, you MUST NOT run
    CRT420 at the same time as any other program that accesses devices
    that run under control of the AM-420.

4.2 Sample Use of CRT420

Below is a sample output of a typical disk certification. The next section
discusses the questions that CRT420 asks and the messages that it displays.

```
.CRT420 PLD0: (RET)
CAUTION: THIS PROGRAM WRITES TO ALL BLOCKS

Enter maximum acceptable number of bad tracks per surface: 20 (RET)
?5 bad tracks is maximum
Enter maximum acceptable number of bad tracks per surface: 5 (RET)
Display current track? (Y or N): Y (RET)
Current track is: 0
Enter serial number (10 char. max): PAYROLL1 (RET)

Begin certification of PLD0:
Current track is: 1
Current track is: 2
              .
              .
              .
?Track 7, Head 1, Sector 3 did not verify
?Track 8, Head 2, Sector 4 did not verify
              .
              .
              .
Current track is: 14
Current track is: 15
              .
              .
              .
Current track is: 184

?2 bad tracks detected
Certification complete
.
```

4.3 Using CRT420

To use CRT420, enter CRT420 followed by the specification of the logical
unit you want to certify. For example:

        .CRT420 PLD0: (RET)

You now see the following message:

                CAUTION: THIS PROGRAM WRITES TO ALL BLOCKS

(If you do not want to continue the certification, you may enter a Control-C
at this point.)

CRT420 now asks you several questions:

        1. Enter maximum acceptable number of bad tracks:
                Give CRT420 the maximum number of bad tracks that you will
            accept on the disk you are certifying. The maximum CRT420 will
            allow is 5 bad tracks for the Priam 1070, 10 bad tracks for the
            Priam 3450, and 23 bad tracks for the Priam 6650. If the number of
            bad tracks that CRT420 finds exceeds the value you specify, CRT420
            tells you so and then aborts the certification, returning you to
            AMOS command level:

                ?Device has exceeded maximum number of errors

                .

        2. Display current track? (Y or N):
                If you want CRT420 to tell you as it verifies each track,
            enter a Y; otherwise, enter an N. If you answer Y, CRT420 now
            displays this message:

                Current track is: 0

            NOTE: Asking CRT420 to display the number of the track it is
            currently verifying greatly increases the length of time it takes
            to certify a disk surface.

        3. Enter serial number (10 char. max):
                You may optionally give CRT420 a ten-character alphanumeric
            I.D. for the logical unit you are certifying. CRT420 writes this
            identifier into the BADBLK.SYS file for that logical unit.

After answering the questions above, CRT420 begins to certify the disk. You
see this message:

                Begin certification of Devn:

where Devn: is the device specification you supplied on the CRT420 command
line.

(Changed 31 October 1981)

If you asked CRT420 to tell you its track position as it certifies, you now
see a list of messages that can look something like this:

> Current track is: 1
> Current track is: 2
> Current track is: 3
>   .
>   .
>   .

and so on.   When CRT420 encounters a track that does not verify, it tells
you so. For example:

> ?Track 15, Head 1, Sector 2 did not verify
> ?Track 16, Head 3, Sector 4 did not verify
>   .
>   .
>   .

When finished certifying the disk, CRT420 tells you how many bad tracks it
found. For example:

> ?2 bad tracks detected
> Certification complete
> •
> ¯

4.4   How CRT420 Certifies a Disk

CRT420 follows these procedures when it certifies a disk:

1.   CRT420 creates a file named BADBLK.SYS in account [1,2] on the disk
     you are certifying.

2.   If you have specified a serial number, CRT420 writes that
     information to this file.

3.   CRT420 writes a 32-bit pattern in every record on the first track
     of the disk.   Then it reads each byte on that track and verifies
     the data.   (CRT420 checks for CRC errors as well as data
     verification errors.)

     If any data does not verify, CRT420 places the number of the track
     and head in the BADBLK.SYS file and tells you that the track is
     bad. For example:

     > ?Track 35, Head 1, Sector 21 did not verify

     CRT420 writes a total of four data patterns, following the
     procedure above for each data pattern.

4. CRT420 now moves on to the next disk track and performs these operations again, reporting any bad tracks that it finds (and entering their numbers into the BADBLK.SYS file). CRT420 verifies every track, including the spare tracks. When it finishes, CRT420 tells you that it is done and how many tracks are bad. For example:

> **?2 bad tracks detected**
> **Certification complete**

5. Now CRT420 computes a hash total for the BADBLK.SYS file and stores it in the file. This value provides a validity check that other programs (e.g., BADBLK) can use to make sure that the BADBLK.SYS file is complete. When AMOS goes to access a disk block, it checks system memory (where a copy of BADBLK.SYS was placed when the disk was mounted) to see if that block occurs on a bad track; if so, AMOS accesses the assigned alternate track instead.

If you interrupt the CRT420 program by typing a Control-C, you see:

> **?Certification incomplete**
> **^C**
>
> **.**

and CRT420 intentionally writes a bad hash total to the BADBLK.SYS file. (A bad hash total tells other programs that may later look at the file that the data in the file is not complete and is not to be trusted.)

4.5 CRT420 Error Messages

Below are the error messages you can encounter when using CRT420.

?You must be logged into PPN [1,2] to run CRT420
 Because it writes data into every byte on the disk, CRT420 is a dangerous program to run. You must be logged in as the System Operator to certify a disk. Log into the System Operator's account, [1,2], before trying to use CRT420.

?NN bad tracks is maximum
 where NN is 5 bad tracks for the Priam 1070, 10 tracks for the Priam 3450, and 23 tracks for the Priam 6650. You specified a number greater than NN as the maximum number of bad disk tracks you would accept on the certified disk. However, the BADBLK.SYS file cannot handle more than NN bad disk tracks; so, enter the maximum number NN or a number less than NN.

?Track 0, Head 0 did not verify. (First track and head must verify.)
 The first track and head of the disk did not verify. CRT420 cannot continue the certification if the first disk track does not

verify, so it now stops the certification and returns you to AMOS command level.

?Track n, Head n, Sector n did not verify.
     CRT420 marked track n, head n in the BADBLK.SYS file as a bad track.

?Device has exceeded maximum number of errors
     CRT420 found more bad tracks per surface than the value you specified as the maximum number of bad tracks you would accept. CRT420 now aborts and returns you to AMOS command level.

?S100 data transfer error
     An error occurred with the AM-420 controller board. CRT420 aborts and returns you to AMOS command level. If you receive this error several times, you may have hardware problems.

?Nonexistent device
     Your device specification on the CRT420 command line is invalid; the system believes that the device does not exist. Check your spelling and try again.

?Certification incomplete
     You typed a Control-C to interrupt the disk certification. CRT420 now intentionally writes a bad hash total to the BADBLK.SYS file to let other programs know that the data in the file is incomplete and not to be trusted.

5.0 THE BADBLK PROGRAM

The BADBLK program allows you to see the contents of the BADBLK.SYS file created by the certification program, CRT420. BADBLK also verifies the BADBLK.SYS hash total.

CRT420 creates the BADBLK.SYS file to flag bad tracks. The BADBLK.SYS file is accessed when the system is booted, when a device is mounted, and when you use the command BADBLK at AMOS command level. If the track to be accessed is flagged as bad, the assigned spare track is accessed instead.

The technique used is as follows:

1. The DEVTBL command automatically reserves space for an Alternate Track Table in system memory.

2. The MOUNT command reads the BADBLK.SYS file into the Alternate Track Table whenever a disk is mounted.

3. Whenever a disk access is requested, the disk driver scans the
    Alternate Track Table to see if the requested block is in a track
    which is marked as bad. If so, a translation is performed to
    access the requested block within the alternate track assigned to
    that bad track.

    The actual allocation of alternate tracks is device dependent;
    currently, the AM-420 controlled Priam devices actually translate
    the alternate tracks sequentially. For example if track 4 is bad,
    track 5 is used.

4. When the system is being booted, the bootstrap routine reads in
    BADBLK.SYS to handle the case where SYSTEM.MON is allocated on an
    alternate track.

5.1 Using BADBLK

To use BADBLK, type BADBLK followed by the specification of the device whose
BADBLK.SYS file you want to see. Then type a RETURN. For example:

> .BADBLK PLD0: (RET)

If BADBLK found the BADBLK.SYS file, it tells you so:

> PLD0: BADBLK.SYS[1,2]

BADBLK now tells you the serial number associated with that device and the
number of tracks marked as bad on that disk. For example:

> Serial number: INVENTORY2
>
> Number of bad tracks: 2
>
> Head 1, Track 30
>
> Head 2, Track 140

NOTE: Track and head numbers are decimal. BADBLK exits and returns you to
AMOS command level:

> EXIT
> .

5.2 BADBLK Error Messages

You can see the following BADBLK error messages:

?File not found: Devn:BADBLK.SYS
> BADBLK was not able to find the BADBLK.SYS file for the disk you specified. Make sure that the device you specified is an AM-420 controlled disk and that the disk has been certified (i.e., you've run CRT420 on that device).

CAUTION: HASH TOTAL DID NOT VERIFY
> The BADBLK.SYS file contained a bad hash total. This indicates that the data in that file is not to be trusted. Use COPY to copy all files off the logical unit containing that BADBLK.SYS file, being careful not to overwrite the BADBLK.SYS file on the new disk. Then re-certify the disk.

You may also see several system error messages if your device specification is invalid. For example:

?Cannot INIT Devn: - device does not exist
> The system did not recognize the device specification you gave. Check your spelling and try again.

?Cannot READ Filespec - disk not mounted
> The system is unable to read the device you specified on the BADBLK command line because it is not mounted. Use the MOUNT command to mount the disk and try again.

6.0 THE DRIVER GENERATION PROGRAM, FIX420

A driver is a program that links the generalized disk service routines of the monitor with the physical disk device. Many drivers are supplied to you on your System Disk in account [1,6]; their names are representative of the equipment you may use them for. However, so many disk formats have become available that, for certain disks, Alpha Micro provides you the means to configure a disk driver of your own if it is not already in account [1,6].

The disk driver configuration program for all of Alpha Micro's Winchester technology disk systems and subsystems is called FIX420, so named because the Winchester disks are controlled by the AM-420 board.

FIX420 builds upon the skeleton driver called 420DVR.DVR, also in the Device Driver Library account, DSK0:[1,6]. The result is any of dozens of potential configurations, depending on which physical disk you are configuring, and your choice of configuration options.

6.1 Using FIX420

To run the FIX420 program, log into the Device Driver Library account,
DSK0:[1,6]. Then type FIX420 followed by a RETURN:

        _.FIX420 (RET)

FIX420 now begins to ask you a series of questions, so that it can determine
how to configure a Winchester technology disk driver that matches your
particular combination of disk type and disk format. First FIXDVR lists the
available drive types:

        Drive type:

            (A) 8.5 Mb 8 in.
            (B) 32 Mb 8 in.
            (C) 60 Mb 14 in.

        Which drive ?

Enter the letter A, B or C to select the drive type you have.

If you have a Priam 1070 or Priam 3450 drive, FIX420 asks you:

        How many logical drives per physical unit (1-8) ?

If you have a Priam 6650 drive, FIX420 asks:

        How many logical drives per physical unit (2-8) ?

(NOTE: 2 is the minimum number of logical devices you can configure the 60
megabyte drive for.)

Specify how many logical devices you want the physical unit to be split
into. (NOTE: Alpha Micro hardware and software permit the disks that are
controlled by the AM-420 controller to be configured so that the logical
devices are, for the most part, independent of the physical disk.   The
physical platters within the drive do not determine the configuration in any
way.)

Now FIX420 continues:

        Enter new driver name:

Enter the name you want to assign to the driver program. Each driver must
have a unique, three-character name. The standard names that you might want
to use are:

        Driver Name   Characteristics

           PLD   8.5 Megabyte 8 inch
           PMD   32 Megabyte 8 inch
           PBD   60 Megabyte 14 inch

(Changed 31 October 1981)

FIX420 now configures the driver and displays this message:

New driver is now in memory. Bitmap size is nnnn Words.

where nnnn is the number of decimal words you must assign to the bitmap for
the new device. The driver you created is now in memory. It is not yet a
file on the disk. Save it to the disk by using the SAVE command.   For
example:

.SAVE PLD.DVR (RET)

The command above saves the driver onto the disk (in the account you are
logged into, DSK0:[1,6]), as the file PLD.DVR. If you do not specify an
extension, the SAVE command saves the file under the extension .DVR (which
indicates a device driver program).

6.2 FIX420 Error Messages

You may use the following error messages when using FIX420:

1. ?Could not find 420DVR.DVR
        FIX420 couldn't find the necessary file. FIX420 requires that
    420DVR.DVR be in DSK0:[1,6].

2. ?Invalid response
        The question that FIX420 asks requires that you enter a letter
    to select an option. Check your typing.

3. ?Invalid device
        You have a bad version of 420DVR.DVR in DSK0:[1,6].

4. ?You must be logged into PPN 1,6 to modify the driver.
        Log into DSK0:[1,6] before trying to use FIX420.

7.0 MODIFYING THE SYSTEM INITIALIZATION COMMAND FILE

You must modify the system initialization command file (SYSTEM.INI) to:

1. Add a DEVTBL command line for the device you have just defined a
    device driver for. NOTE: If the first device on a DEVTBL command
    line is a Winchester technology device, all subsequent devices on
    that line share the same alternate track table. Therefore, you
    must not combine the AM-420 controlled device on the same DEVTBL
    command line as any other kind of device; you should specify only
    the logical devices associated with a single physical drive on a
    single DEVTBL command line; and you should keep all the logical
    devices of a physical drive together on the same DEVTBL command
    line.

(Changed 31 October 1981)

For example, say you have as your System Disk drive an AM-423, 60 megabyte Winchester technology drive. Also, you have two drives of the AM-421, 8.5 megabyte variety. Finally, you have one AM-422, 32 megabyte drive. That is a total of four physical drives, so you need four DEVTBL command lines. Remember in examining the following DEVTBL command lines that DSK0 (of the AM-423) is not shown because the system already knows about the System Disk. However, DSK1 is shown because it is a logical split of that physical drive. (NOTE: If, in another configuration, you had separate physical disks and wanted to name one of them DSK1, you would still arrange the DEVTBL command lines as in the example below. The system can tell what you are doing by looking at the driver program you created previously.)

The DEVTBL example:

```
DEVTBL DSK1
DEVTBL PLD1
DEVTBL PLD2
DEVTBL PMD6,PMD7
```

The first DEVTBL command line indicates a physical disk logically split in two (DSK0, assumed, and DSK1). The second command line indicates a physical device, PLD1, which is also a single logical device (i.e., no logical split). The third, PLD2, like the second, is a single logical device (separate devices with the same 3-character name must have the same number of logical units!). The fourth is logically split into two devices (PMD6 and PMD7).

Note the drive numbers, underlined, in this copy of the above example:

```
DEVTBL DSK1
DEVTBL PLD1
DEVTBL PLD2
DEVTBL PMD6,PMD7
```

Because there may be more logical devices than physical devices for AM-420 controlled drives, the drive numbers (0 (assumed) and 1, 1, 2, 6 and 7) in the DEVTBL command lines are derived from a special allocation formula.

To implement the formula, you tell the AM-420 controller to associate its own physical addresses to physical disk drives by the order in which the DEVTBL command lines are encountered. (The physical addresses, or ports, on the controller are numbered 0, 1, 2 and 3; 0 is associated by the controller with the first DEVTBL command line encountered, 1 is associated with the second DEVTBL line encountered, etc.) Then you tell the controller the names and numbers of the logical devices on the physical device listed on each DEVTBL command line.

You should use the above mentioned allocation formula when adding
DEVTBL command lines to the SYSTEM.INI for AM-420 controlled
devices.   The formula for determining the beginning logical unit
number for a specific physical device is:

> "Physical address times number of logical units
> on that physical device equals the drive number
> of the first logical unit on that physical
> device."

Using the example above, the device on the first DEVTBL line
encountered is physical device 0. There are two logical units on
that physical device (including the assumed one not shown on the
DEVTBL command line). The drive number of the first logical unit
on that physical device is 0 x 2 or 0. The subsequent logical unit
is numbered 1, which is shown:

DEVTBL DSK1

The device on the second DEVTBL line encountered is physical device
1.   There is one logical unit on that physical device. The drive
number of the logical unit on that physical device is 1 x 1 or 1.
It is important to remember that when two separate physical devices
(such as this and the next one) have the same 3-character name,
they must both have the same number of logical units:

DEVTBL PLD1

The device on the third DEVTBL line encountered is physical device
2.   There is one logical unit on that physical device. The drive
number of the logical unit on that physical device is 2 x 1 or 2:

DEVTBL PLD2

Notice that if the second and third DEVTBL command lines both
specified two or more logical devices, the drive numbers on the
second DEVTBL command line would be 1 and 2, while on the third
DEVTBL command line they would be 2 x 2 or 4, and then 5.

The device on the fourth DEVTBL line encountered is physical device
3.   There are two logical units on that physical device. The
drive number of the first logical unit on that physical device is 3
x 2 or 6. The subsequent logical unit is numbered 7.

DEVTBL PMD6,PMD7

2. Add BITMAP commands to the SYSTEM.INI to define bitmaps for the new
    devices you are adding to the system. Use the bitmap size given in
    the final FIX420 message. For instance, above we suggested an
    example where you have four Winchester technology drives on your
    system. The BITMAP statements in the SYSTEM.INI for that situation
    should read:

```
BITMAP DSK,3706,0,1
BITMAP PLD,1018,1
BITMAP PLD,1018,2
BITMAP PMD,1932,6,7
```

The numbers 3706, 1018 and 1932 which you see in the BITMAP command lines were given to you by FIX420, the driver configuring program, as it concluded each of three drivers. These designate the number of words of system memory which are reserved when the system boots up, and which are used by each logical device's bitmap.

The numbers following the bitmap word sizes, of course, must correspond exactly to the drive numbers found in the DEVTBL command lines.

8.0 USING MOUNT

The MOUNT command has taken on a new feature that is very important to users of systems using Winchester technology disks.

8.1 One Conventional Use of MOUNT

Conventionally, when using a floppy disk and most hard disk systems, it becomes desirable to change hard disk cartridges or floppy disks to access removable storage. The system must know that you have replaced one disk with another, so it can update its bitmap information. Otherwise it may write over data on the new disk by storing it in a place that was available on the old disk. You must use the MOUNT command to inform the system of the change.

8.2 Three New Uses for MOUNT

The Winchester technology drives have no removable disks, so of course the above described use of MOUNT is not useful to these disks. And turning power on sequences up and mounts the System Disk of an AM-10XX system. However, turning on power neither sequences up nor mounts a subsystem running under the control of the AM-420 controller. Therefore, MOUNT has three applications of paramount importance to Winchester technology drive users.

1. Before the drive is sequenced up, the read/write heads actually rest on the media surfaces on "landing points," and are latched into position there in case the drive is to be moved.

    When you use MOUNT and specify a Winchester technology drive, the drive is sequenced up; that is, it is brought to its ready

condition.   The   read/write heads are unlatched, the disk is spun
up, and the heads actually "fly" over the media surfaces.   An
example of using MOUNT to sequence up a Winchester technology
device might be:

> .MOUNT PLDO: (RET)

The drive normally takes about 30 seconds to sequence up.   If the
drive has not sequenced up properly in 90 seconds, you receive back
an error message.

2.   When you wish to isolate a drive from the system, you may unmount
     that drive using MOUNT and the /U switch. For example:

> .MOUNT PLDO:/U (RET)

The drive is unmounted, but is not sequenced down.

3.   You absolutely must sequence down all Winchester   technology
     drives before shutting off power to the system.   During the
     sequencing down process, the disks stop rotating, the read/write
     heads come to rest on the "landing points" and are latched into
     position, and other vital shutdown routines are performed. You
     sequence down a Winchester technology drive using MOUNT and the /S
     switch. Since the physical drive is sequenced down, you may
     specify any logical unit on the physical drive to sequence down.
     For example:

> . MOUNT PLDO:/S (RET)

IMPORTANT NOTE: When sequencing down more than one Winchester technology
disk drive, you must not power down (and thus sequence down) the drive
containing the System Disk, DSKO:, until all other drives are sequenced
down.   .

8.3 Reference

You can obtain more information on MOUNT from the MOUNT reference sheet in
the AMOS System Commands Reference Manual (DWM-00100-49)

9.0 SUPPORT PROGRAMS

There are a number of programs used to support the AM-420 controlled disk
systems. They are:

1.   The disk certification program, CRT420, described in Section 4.0 of
     this document, which certifies a physical disk controlled by the
     AM-420 controller.

2. FIX420, described in Section 6.0 of this document, which builds on
    the AM-420 disk driver skeleton called 420DVR.DVR, resulting in a
    specific driver configuration.

3. PRILOD, the bootstrap loader for a system that uses one of the
    three Priam hard disks as the System Disk.

4. WINFLP, the program to do backup from a Winchester technology disk
    to a floppy disk.

5. FLPDIR, the program to do a directory on a floppy disk that is
    being used as a backup for a Winchester technology disk.

6. FLPWIN, the program to restore to a Winchester technology disk the
    files backed up on a floppy disk.

7. VCRSAV, the program to do backup from a Winchester technology disk
    to a Video Cassette Recorder.

8. VCRDIR, the program to do a directory on a Video Cassette Recorder
    that is being used as a backup for a Winchester technology disk.

9. VCRRES, the program to restore to a Winchester technology disk the
    files backed up on a Video Cassette Recorder.

10. STRSAV, the program to do backup from a Winchester technology disk
    to a 1/4" streaming tape drive.

11. STRDIR, the program to do a directory on a 1/4" streaming tape that
    is being used as a backup for a Winchester technology disk.

12. STRRES, the program to restore to a Winchester technology disk the
    files backed up on a 1/4" streaming tape drive.

13. MOUNT/S, the MOUNT program with the /S switch that sequences down a
    Winchester technology disk.

For the most part, the method for handling disk flaws, described in Section
4.0, "The BADBLK Program," is transparent. For that reason, commands such
as DSKANA, DSKCPY, REDALL, DSKPAK, and so on, are able to treat any disk
that was certified with a 4.5A or later version of CRT420 in the same way
that they do a device that does not contain a BADBLK.SYS[1,2] file.

For information on WINFLP, FLPWIN and FLPDIR, see the document
Winchester/Floppy Backup Programs in the "System Operator's Information"
section of the AMOS Software Update Documentation Packet.

For information on VCRSAV, VCRDIR and VCRRES, see the manual Software
Information for the AM-610 Video Cassette Recorder Interface (DSS-10000-24)

For information on STRSAV, STRDIR and STRRES, see the manual Software
Information for the AM-620/621 1/4" Streaming Tape Drive (DSS-10000-23).

Index

July 1981

THE WINCHESTER/FLOPPY BACKUP PROGRAMS

1.0 INTRODUCTION

The Alpha Micro Winchester/Floppy backup system consists of three programs:
WINFLP, the program to transfer files from a Winchester technology hard disk
to a floppy disk; FLPWIN, the program to restore files to a Winchester
technology disk from a floppy disk; and FLPDIR, the program to list the
contents of a floppy disk. These three programs, when used with disks that
run under control of the AM-420 Hard Disk Controller, allow you to easily
and rapidly back up and restore both sequential and random files.   The
programs have full wildcarding capability and allow you to split a single
backup, or even a single file, across multiple floppy disks.

These programs store and read data on the floppy disk in a special record
format developed by Alpha Micro. This format was optimized for the AM-420
Winchester technology disk-based system; it was not intended that this
format be used for data interchange with other, non-Alpha Micro computers.

IMPORTANT NOTE: It is important to remember that this set of programs was
designed as a mechanism for backing up disk files. That means that WINFLP
writes files to the floppy disk along with their full device and account
specifications. (It also writes the date and time of backup.)   Therefore,
when using FLPWIN to read a floppy disk, if you want to access a specific
file you MUST specify the disk and account from which the file was backed
up, as well as the device containing the floppy disk you want to access.
For instance, if you stored a file on a floppy disk from DSK0:, you must
specify the device specification "DSK0:" when you restore the file from the
floppy disk. We give specific examples in the sections below.

1.1 Wildcarding Features

All three programs have been designed to function as much as possible like
their disk-oriented counterparts (COPY and DIR) to make the Winchester
technology backup software as easy to use as possible.

WINFLP, FLPWIN, and FLPDIR use wildcard symbols and specification defaults
in the same way that COPY and DIR do. This is because, like COPY and DIR,
WINFLP, FLPWIN, and FLPDIR are "wildcard commands."

Wildcard commands differ from other AMOS commands in that besides accepting
the standard AMOS file specification:

     Devn:filename.extension[p,pn]{/switches}

they also accept a variety of wildcard specifications. A wildcard file
specification allows you to select multiple files with only one file

specification.   These file specifications can contain the special wildcard
symbols *, ?, [], and ALL:. For example, to specify all files in all
accounts of DSK0: that have a .BAS extension, wildcard commands permit the
file specification:

    DSK0:*.BAS[]

(In the example above, the symbol * stands for all possible filenames.) In
addition, the wildcard commands allow you to set account, device, and switch
specification defaults. For example, the wildcard command line:

    .WINFLP [100,2]*.BAS,*.TXT[117,6],*.LST (RET)

sets the default account specification to [100,2]. The command line thus
selects all .BAS and .LST files in account [100,2], and all .TXT files in
account [117,6]. (Notice that the account specification setting the default
occurs before the filename and extension rather than after it, as is the
case with a standard AMOS file specification.)

1.1.1 Switches and Wildcard Commands

"Switches" are option requests. Each switch must begin with a slash, /.
Remember that wildcard commands recognize two types of switches: "file"
switches and "operation" switches. An operation switch applies to all of
the file specifications on an entire command line no matter where it appears
on that command line.

A file switch may apply to only specific file specifications, depending on
where it appears on the command line. If a file switch appears at the end
of a file specification, it applies only to the files selected by that
specification. For example:

    .WINFLP *.BAS,*.TXT/QUERY,*.LST (RET)

in the command line above, the /QUERY switch (a file switch) applies only to
those files selected by the *.TXT file specification. If a file switch
appears before a file specification, it becomes the default switch, and
applies to all of the files selected by the following file specifications
unless it is overridden for a particular file specification by another
switch, or until a new default is set. For example:

    .WINFLP *.BAS,/QUERY*.TXT,*.LST,*.MAC/NOQUERY,*.PRG (RET)

the /QUERY switch affects all files selected by the specifications *.TXT,
*.LST, and *.PRG, but not the files selected by the specification *.MAC.

Because wildcard file specifications are extremely powerful, they can have
very widely ranging effects. If you are not familiar with the way in which
wildcard commands work, be sure to read the AMOS User's Guide,
(DWM-00100-35) before attempting to use wildcard specifications.

(1 July 1981)

2.0   WRITING WINCHESTER DISK FILES TO A FLOPPY DISK - THE WINFLP PROGRAM

The WINFLP program writes disk files to a floppy disk. This program accepts
a standard wildcard file specification which specifies the files to back up.

The WINFLP program operates in two modes: /APPEND and /NOAPPEND. In /APPEND
mode, WINFLP searches for the last file on the floppy disk and starts
writing the new files immediately after any data already on the floppy disk.
In /NOAPPEND mode, WINFLP does not look to see if any data is already on the
floppy disk, but just starts writing files at the beginning of the floppy
disk. The default mode of operation is /APPEND. Of course, if the floppy
disk is blank (i.e., a new floppy disk), you should specify the /NOAPPEND
mode.

Call the WINFLP program by giving a wildcard specification for the files you
wish to back up:

        .WINFLP{/switches} Filespec1{/switches}{,Filespec2{/switches}...} (RET)

The Filespec default is *.* and the account and device you are logged into.
The default switches are /NOQUERY/APPEND.

The WINFLP program now asks you for the device code and unit number of the
floppy disk drive you want to write to:

        Enter backup device:

(WINFLP assumes a device code of DDA0: if you just type a RETURN.)

WINFLP now writes the files to the floppy disk, listing each file as it
transfers it. (NOTE: Because WINFLP also writes the date and time of the
backup to the floppy disk, you should use the system commands DATE and TIME
to set the system date and time before you use WINFLP.)

2.1   Example

For example, to back up all files from disk device DSK0: to the floppy disk
drive DDA0: (starting at the beginning of the floppy disk), use the
following command:

        .WINFLP DSK0:[]/NOAPPEND (RET)
        Enter backup device: DDA0: (RET)
        AMSORT.SYS[1,4] to DDA0:AMSORT.SYS[1,4]
                        .
                        .
                        .
        Total of 1282 files transferred
        .

2.2 Writing to Multiple Floppy Disks

If all of the files you specified will not fit on one floppy disk, WINFLP displays the following message:

      %Device is full. Please mount another disk, then type RETURN to
      % continue, or type Control-C to abort copy

If you wish to continue backing up files on another floppy disk, insert a new floppy disk, then type RETURN on your terminal. The backup will continue on the new floppy disk. If you wish to abort the backup, type a Control-C.

2.3 WINFLP Switches

WINFLP provides the switches below:

    /QUERY   or /Q   Ask user for confirmation before copying files (file
                           switch).

    /NOQUERY   or /NOQ   Don't ask for confirmation (default, file switch).

    /APPEND   or /A   Write files to floppy disk at the end of existing
                           files (default, operation switch).

    /NOAPPEND   or /NOA   Write files at beginning of floppy disk (operation
                           switch).

2.4 Error Messages

You may see the following error messages when using the WINFLP program:

    ?Cannot find DSK0:SCNWLD.SYS[1,4] or MEM:SCNWLD.SYS
            The WINFLP program needs this file to be able to process wildcard
            symbols in your file specification. This message can indicate
            that SCNWLD.SYS does not exist, or that you do not have enough
            memory to load the file into your partition.

    ?Cannot READ Devn - device does not exist
    ?Cannot READ Devn - device is not mounted
            You tried to copy to or from a device that is not listed in the
            DEVTBL command in your SYSTEM.INI, does not have a driver in area
            [1,6] of the System Disk, is not file-structured, or is not
            mounted. ("Devn:" is the device you specified.)

    %No file-oriented device corresponding to Devn: is mounted
            You specified a device, but left off the unit number. WINFLP
            cannot find a logical unit that matches your specification. Try
            mounting the device.

    %Device is full. Please mount another disk, then type RETURN to

(1 July 1981)

<u>% continue, or type Control-C to abort copy</u>
There is no more room on the current floppy disk. Mount another
disk and type RETURN to continue the backup process, or type a
control-C to abort the backup procedure.

<u>%Did not copy Filespec</u>
The file was not copied because you typed a Control-C to abort the
backup instead of inserting a new floppy disk and continuing by
splitting a file between the two floppy disks.

<u>?DSK is not a valid backup device</u>
You cannot use these programs to write to the system device.

<u>?No room to perform append; device full.</u>
An attempt was made to append to a disk that was already full.

3.0 RESTORING DISK FILES FROM A FLOPPY DISK - THE FLPWIN PROGRAM

Use the FLPWIN program to transfer files back to the disk from a floppy
disk.   The files must have been written to the floppy disk via WINFLP.
FLPWIN provides full wildcarding, allowing easy selection of the files to be
restored, as well as automatic renaming facilities.

To use FLPWIN, enter the FLPWIN command followed by an output specification,
an input specification, and any optional switches:

  <u>.</u>FLPWIN{/switches} {outspec}={inspec1{/switches}{,inspec2{/switches}...}}

The output specification defaults to the input specification. The input
specification defaults to *.* and the device and account you are logged
into. The default switches are /NOQUERY/DELETE.

FLPWIN now asks you for the device code and unit number specifying the
floppy disk drive you wish to read from:

       <u>Enter backup device:</u>

(FLPWIN assumes a device code of DDA0: if you just type a RETURN.)

FLPWIN starts searching for the specified files at the beginning of the
floppy disk. As FLPWIN finds them on the floppy disk, it transfers the
files to the disk and accounts you have specified.

The output specification you supply to FLPWIN is the specification of the
file(s) you wish to create. FLPWIN provides full wildcarding. Just as with
the COPY command, you may not copy files from one account to another unless:
1) you are in the same project as the account you are copying files to; or
2) you are logged into the System Operator's account, [1,2]. (As with the
COPY command, logging into account [1,2] gives you certain privileges. The
default account specification of the Outfilespec when you are logged into
[1,2] is the wildcard account, []. Also, if you are logged into [1,2],
FLPWIN will create the account you are copying to if it does not exist.)

The input specification is a list of the files you wish to copy from the floppy disk. The input specification must give the exact specification of the file you wish to copy, including the device and account of the file <u>as it is stored on the floppy disk</u>.

3.1 Example

For example, assume you are logged into DSK0:[140,1] and you wish to copy a file from a floppy disk that was backed up from your own account (DSK0:[140,1]).   Enter:

        .FLPWIN = FILE.DAT (RET)
        Enter backup device: DDA0: (RET)
        DDA0:FILE.DAT to FILE.DAT
        Total of 1 file transferred

Note that in the example above, the output specification defaulted to the input specification, and the input device and account defaulted to the device and account you are currently logged into.

If you want to copy a file from a floppy disk to the account you are logged into that was backed up from another device and account (DSK2:[1,4] for example), you would enter the following command:

        .FLPWIN = DSK2:TEST.BAS[1,4] (RET)
        Enter backup device: DDA0: (RET)
        DDA0:DSK2:TEST.BAS[1,4] to TEST.BAS
        Total of 1 file transferred

If you want to restore all the files stored on a floppy disk to their original device and account, you would log into [1,2] and enter the following command:

        .FLPWIN = ALL:[] (RET)

If you want to return all the files stored on a floppy disk back to their accounts of origin in Project 110 of DSK0:, you would log into an account in Project 110 and enter the following:

        .FLPWIN DSK0:[]=DSK3:[110,*] (RET)

If you include a filename and/or extension in your output specification, you can rename the copies of the files you are writing to disk. For example:

        .FLPWIN DSK3:[]*.OLD = DSK1:[300,1?]*.MAC (RET)

copies all .MAC files from the floppy disk backed up from accounts [300,1?] on DSK1: to the same accounts on DSK3:, and renames the file extensions from .MAC to .OLD.

3.2 Restoring from Multiple Floppy Disks

If you are restoring files from multiple floppy disks created by the WINFLP
program, you must enter separate FLPWIN commands for each floppy disk,
unless the floppy disks contain files split across multiple floppy disks.

If this happens FLPWIN will prompt you with the message:

>   Remove old disk and insert new one.
>   Type RETURN when ready, Control-C to abort.

If you attempt to copy split files out of sequence, FLPWIN will give you one
of three error messages:

>   ?Attempt to copy files out of sequence.
>   ?New file is number 2, last file was number 0

You attempted to retrieve split files in a sequence other than the sequence
in which they were written onto the floppy disk. For example, one file was
split among several floppy disks, but when restoring that file you did not
enter the floppy disks in the proper order. Place the correct disk into the
drive and continue.

>   ?File unsplit mismatch error

You attempted to retrieve a split file, but the next disk you entered
contained a new file rather than the next portion of the split file.   Place
the correct disk in the drive and continue.

>   ?Device is not properly file structured

The disk currently being read was not created using WINFLP. Place the
correct disk in the drive and continue.

3.3 FLPWIN Switches

FLPWIN provides the following switches:

>   /QUERY   or /Q   Ask user for confirmation before copying files (file
>   switch).
>
>   /NOQUERY   or /NOQ Don't ask for confirmation (default, file switch).
>
>   /DELETE   or /D   Copy over to an existing file, thereby deleting it
>   (default, file switch).
>
>   /NODELETE ·or /NOD Don't copy over to any existing files (file switch).

3.4 Error Messages

You may see the following error messages when using the FLPWIN program:

?Cannot find DSK0:SCNWLD.SYS[1,4] or MEM:SCNWLD.SYS
> The FLPWIN program needs this file to be able to process wildcard symbols in your file specification. This message can indicate that SCNWLD.SYS does not exist, or that you do not have enough memory to load the file into your partition.

?Cannot READ Devn - device does not exist
?Cannot READ Devn - device is not mounted
> You tried to copy to or from a device that is not listed in the DEVTBL command in your SYSTEM.INI, does not have a driver in area [1,6] of the System Disk, is not file-structured, or is not mounted. ("Devn:" is the device you specified.)

%No file-oriented device corresponding to Devn: is mounted
> You specified a device, but left off the unit number. FLPWIN cannot find a logical unit that matches your specification. Try mounting the device.

?Device is not properly file structured
> The floppy disk you are trying to read was not written by the WINFLP program. The FLPWIN program can only read floppy disks written by WINFLP. Check to make sure you have mounted the correct floppy disk.

?Missing output specification
> You omitted the equal sign in your FLPWIN command line; FLPWIN couldn't tell which information was your input specification and which was your output specification.

?More than one output specification
> You may not supply more than one output specification.

?Files may not be transferred to RES:
> You may only add programs to system memory by using the SYSTEM command within your system initialization command file, SYSTEM.INI.

%Not copied - Destination file already exists
> You tried to copy to an existing file while the /NODELETE option was in effect.

?You are not logged in under [1,2], can't create [p,pn]
> You cannot copy from an account to a nonexistent account unless you are logged in under [1,2]. If you copy to a nonexistent account while logged under [1,2], FLPWIN will create the account.

?Output MFD is full
>   The Master File Directory only has room for 64 entries. The
     transfer in progress would have created a new account, but there
     is no room in the MFD.

%Bypassing BADBLK.SYS[1,2]
%   BADBLK.SYS exists to prevent bad blocks
%   on a device from being allocated, and
%   should never be directly accessed.
>   You cannot copy the BADBLK.SYS file, since this would lead to
     corruption of the file system.

?Device full
>   There is no more room on the disk.

?Cannot OPEN Devn: - protection violation
>   FLPWIN could not transfer the files. You are not allowed to write
     into an account you are not logged into unless the project
     number of that account is the same as the project number of the
     account you are copying from. You must either log into the System
     Operator's account, [1,2], or the account you are copying into to
     accomplish the transfer.

?Attempt to copy files out of sequence.
?New file is number 2, last file was number 0
>   You attempted to retrieve split files in a sequence other than the
     sequence in which they were written onto the floppy disk. For
     example, one file was split among several floppy disks, but when
     restoring that file you did not enter the floppy disks in the
     proper order. Place the correct disk into the drive and continue.

?File unsplit mismatch error
>   You attempted to retrieve a split file, but the next disk you
     entered contained a new file rather than the next portion of the
     split file. Place the correct disk in the drive and continue.

?Device is not properly file structured
>   The disk currently being read was not created using WINFLP. Place
     the correct disk in the drive and continue.

4.0 LISTING THE CONTENTS OF A FLOPPY DISK - THE FLPDIR PROGRAM

The FLPDIR program displays a list of the files that have been stored on a
floppy disk. (NOTE: You may only use FLPDIR on a floppy disk that was
written via WINFLP.) The FLPDIR program has been designed to be as similar
to the DIR program as possible.

To create a list of the files on a floppy disk, enter the following command:

     .FLPDIR{/switch}{Listfilespec=}{inspec1{/switch}{,inspec2{/switch}}...} (RET)

Where the optional Listfilespec specifies where the output listing is to be

placed.   If you specify no listfile or equal sign, the display goes to your
terminal. By specifying a listfile, you can send the display to a disk file
or printer.

The optional inspec allows you to select the files you wish to include in
the directory listing. The default is the device and account you are logged
into, and a file specification of *.*.

FLPDIR now asks for the device code and unit number of the drive containing
the floppy disk that you want to get a directory listing of:

        Enter backup device:

(WINFLP assumes a device code of DDA0: if you just type a RETURN.)

Now FLPDIR creates the output listing, showing the relative position of each
file on the floppy disk, the full file specification, the size of the file,
and whether the file is a linked (L) or contiguous (C) file (that is,
whether it is a sequential or a random file). Both of these letters may be
followed by a split (S) code, and the date and time that the file was
written to the floppy disk. At the end of the listing, FLPDIR gives the
total number of files and blocks that it has found.

4.1   Example

For example, to list all the files on a floppy disk on your terminal, enter
the following command:

```
        .FLPDIR ALL:[]
        Enter backup device: DDA0:
            1  DSK0:  SYS     MAC   140,1        16 L  14-May-80  14:52:23
            2  DSK0:  NBSORT  MAC   140,1         4 L  14-May-80  14:52:25
            3  DSK0:  WINFLP  MAC   140,1        23 L  14-May-80  14:52:25
            4  DSK0:  JANE    DAT   140,1        99 C  14-May-80  14:52:27
        Total of 4 files in 142 blocks
```

To create a file (DIRECT.LST) in the account and device you are logged into
that contains a list of all data (.DAT) files on the floppy disk, enter the
following command:

        .FLPDIR = ALL:*.DAT[] (RET)

4.2   FLPDIR Switch

FLPDIR provides the following switch:

    /KILL   or /K   Delete and replace existing Listfile if it has the
                          same specification as your Listfilespec. (Operation
                          switch.)

(1 July 1981)

4.3 Error Messages

You may see the following error messages when using the FLPDIR program:

?Cannot find DSK0:SCNWLD.SYS[1,4] or MEM:SCNWLD.SYS
>The FLPDIR program needs this file to be able to process wildcard symbols in your file specification. This message can indicate that SCNWLD.SYS does not exist, or that you do not have enough memory to load the file into your partition.

?Cannot READ Devn - device does not exist
?Cannot READ Devn - device is not mounted
>You tried to copy to or from a device that is not listed in the DEVTBL command in your SYSTEM.INI, does not have a driver in area [1,6] of the System Disk, is not file-structured, or is not mounted. ("Devn:" is the device you specified.)

%No file-oriented device corresponding to Devn: is mounted
>You specified a device, but left off the unit number. FLPDIR cannot find a logical unit that matches your specification. Try mounting the device.

?Device is not properly file structured
>The tape you are trying to read was not written by the WINFLP program. The FLPDIR program can only read floppy disks written by WINFLP. Check to be sure you've mounted the correct floppy disk.

?More than one output specification
>You may not supply more than one output specification.

?Device full
>There is no more room on the disk.

%No such files
>FLPDIR was unable to find any files matching your input specification.

?DSK is not a valid backup device
>You cannot use these programs to write to the system device.

SOFTWARE NOTICE FOR AM-410 USERS

April 1981
Revision A03

This document reflects AMOS versions 4.4 and later

'Alpha Micro', 'AMOS', 'AM-100',
'AlphaBASIC', 'AlphaPASCAL', and 'AlphaLISP'

are trademarks of

ALPHA MICROSYSTEMS
Irvine, CA   92714

ALPHA MICROSYSTEMS
17881 Sky Park North
Irvine, CA   92714

Table of Contents

1.0 INTRODUCTION

Disks that run under the control of the AM-410 Hard Disk Controller differ somewhat from the other hard disks supported by Alpha Micro. They therefore require different techniques for data backup, disk formatting, and disk initialization. (As of this release, the only disk that runs under control of the AM-410 is the 90-megabyte CDC Phoenix hard disk.) The purpose of this document is to acquaint you with these devices.

The technology that gives the Phoenix disk a much higher density of data than more conventional drives (e.g., the CDC Hawk), also makes media flaws (i.e., bad disk blocks) a much more likely possibility. Therefore it is necessary to use the CRT410 command to certify a Phoenix logical unit before you use it. This certification process identifies any bad tracks on the disk, formats the disk, and initializes it.

IMPORTANT NOTE:

    As of AMOS version 4.4, the system handles devices that contain media flaws differently than earlier versions of the operating system. (That is, AMOS now uses a "bad track" instead of a "bad block" method of keeping track of media flaws.) For more information, refer to AMOS Version 4.4 Method of Handling Bad Disk Blocks, in the "System Operator's Information" section of the AMOS Software Update documentation packet.

    The information in this document reflects AMOS 4.4; it assumes that if you have Phoenix devices that were certified with a pre-4.4 version of CRT410, you have converted your Phoenix disk surfaces via the instructions given in AMOS Version 4.4 Method of Handling Bad Disk Blocks. If you have not done this conversion, refer to the previous version of this document in your 4.3 AMOS Software Update documentation packet, Software Notice for AM-410 Users, for instructions on backing up and analyzing your Phoenix surfaces.

CRT410 certifies the Phoenix media by checking each disk track. It creates the file BADBLK.SYS[1,2] and lists any bad tracks on the media in that file. AMOS assigns spare Phoenix tracks (called alternate tracks) to the bad tracks, to be used in their place whenever the system tries to access a block on the bad tracks. The BADBLK command displays the contents of BADBLK.SYS[1,2].

1.1 Important Notes

There are several things you must keep in mind before you begin to use the Phoenix disk:

1. You must not format or initialize a disk that runs under control of the AM-410. That means that you must not use the SYSACT initialize command. (After you use CRT410 to certify the disk, however, you may use SYSACT to add user accounts to the disk.)

2. As of AMOS Version 4.4, you <u>may</u> use the AMOS disk diagnostic tests REDALL, RAZA, and RNDRED on the Phoenix.

3. As of AMOS version 4.4, you <u>may</u> use DSKCPY or COPY to back up a Phoenix surface.   Both DSKCPY and COPY have been changed so that they will not write over the BADBLK.SYS[1,2] file on the backup disk, and will not write into bad disk tracks.

4. We have increased the thoroughness and reliability with which CRT410 searches for bad disk areas. This means that certifying a disk now takes longer than under AMOS Version 4.3. certification.

2.0 THE PHOENIX DISK CONFIGURATION

The CDC Phoenix disk contains 5 fixed platters of 15 megabytes each and a removable 15-megabyte cartridge. Each of these platters is a separate logical unit and must be accessed as such. The cartridge always bears the highest unit number for that physical device. (For example, if you have a single Phoenix you are using as a non-System Device, then SMD0: - SMD4: are fixed disks; the cartridge is SMD5:.) The one exception to this rule is the Phoenix physical device being used as the System Device. If the system is running off the cartridge, the System Disk cartridge is DSK0: and the fixed disks are logical units DSK1:-DSK5:. (A second Phoenix on this system follows the general rule above; the cartridge is DSK15: and the fixed disks are DSK10:-DSK14:.)

Each Phoenix logical unit contains 808 (#0-807) tracks, with 36 sectors per track, to give a total of 29088 (decimal) sectors (or disk blocks) on each unit. The Phoenix also contains 15 spare (or "alternate") tracks, #808-822. Each disk block contains 512 bytes. The Phoenix requires a bitmap size of 1818 words.

Note that AMOS allows you to place bitmaps in switchable system memory. (See the document <u>Defining Switchable System Memory</u>, in the "System Operator's Information" section of the AMOS Software Update documentation packet, for information on using this technique to reduce the monitor size.)

The device driver program for the Phoenix disk is SMD410.DVR[1,6] on the System Disk. The bootstrap loader program is SMDLOD.PRG in account [1,4] on the System Disk.

(Changed 30 April 1981)

The paragraphs below discuss adding the Phoenix to your system either as the System Device or as a non-System Device. (Hardware settings on the disk controllers on your system define the device that the system attempts to boot from.)

If you have several Phoenix disks, note that the first physical device contains logical units zero through five (e.g., DSK0:-DSK5:), a second Phoenix device contains logical units ten through fifteen (e.g., DSK10:-DSK15:), a third Phoenix contains logical units twenty through twenty-five, and so on.

2.1 Using the Phoenix as a Non-System Device

If you are not using the Phoenix disk as the System Device (that is, if you boot the system off some other type of device), you will need to follow some simple steps before you can access the Phoenix disk:

1. Rename the Phoenix driver program from DSK0:SMD410.DVR[1,6] to DSK0:SMD.DVR[1,6].

2. Now you need to define the Phoenix disk as a system device. To do this, edit your system initialization command file, SYSTEM.INI, and add the devices SMD0, SMD1, SMD2, SMD3, SMD4, and SMD5 to the DEVTBL command line. This adds the Phoenix to your system device table. (For information on using the DEVTBL command to define devices and on using the BITMAP command to define disk bitmaps, see The System Initialization Command File, (DWM-00100-09, Rev A03), in the "System Operator's Information" section of the AMOS Software Update documentation packet.)

3. Add BITMAP commands to the SYSTEM.INI to define bitmap areas for the Phoenix. (The Phoenix requires a bitmap size of 1818 decimal words.) Remember that you can set up your system so that part of system memory resides in a switchable area of a memory bank. You can thus place your bitmaps in switchable system memory and so reduce the size of the monitor. Now reboot the system with your new SYSTEM.INI. You can now access the Phoenix.

If you want more information on adding new disk devices to your system, refer to the document Defining Non-System Disk Devices in the "System Operator's Information" section of the AMOS Software Update documentation packet. Section 4.0 of that document, "Building a System on a New Device," explains how to convert the Phoenix to the System Device after you have added it to your system as a non-System Device.

The configuration of your Phoenix non-System Device is as follows: The five fixed platters are SMD0:, SMD1:, SMD2, SMD3:, and SMD4:. The removable cartridge is SMD5:.

2.2 Using the Phoenix as the System Device

If the Phoenix is your System Device, you will want to copy all of the
system software from the System Disk cartridge down to one of the fixed
disks:

1. Turn on the computer, holding down the reset button. CPU power
    must be on whenever the Phoenix disk is cycled up. Therefore,
    always turn the computer on before cycling up the Phoenix; always
    cycle down the Phoenix before turning off your CPU. If you do not
    follow this procedure, it is quite possible that the data on your
    Phoenix disk will be damaged. NOTE: We recommend that you always
    leave your Phoenix drive powered up when you are not using it.
    This allows the filter system to continually guard the disk
    surfaces. (Of course, if you are not using your system, you will
    probably want to write-protect the drive, cycle it down, and turn
    off your CPU even though you leave the drive powered up.)

2. Insert the Phoenix System Disk Update cartridge according to the
    instructions accompanying the drive. Cycle it to READY status.

3. The logical unit the system boots off of is ALWAYS known as DSK0:.
    The System Disk Update cartridge contains the programs
    SYSTEM.MON[1,4] and SYSTEM.INI[1,4]; therefore the system
    recognizes it as a System Disk and tries to boot the system off the
    cartridge. Because you are running off the cartridge, the fixed
    platters are units DSK1:-DSK5: and the cartridge is DSK0:. (NOTE:
    Even after you install the system software onto a fixed platter, if
    you reset the computer with the System Disk Update cartridge
    mounted, the system tries to boot off the cartridge.)

4. To install the system software on the first fixed platter, use the
    CPY410 command. This command certifies DSK1: and copies the
    contents of the cartridge down onto that fixed disk.

5. You can now remove the cartridge and insert a data pack. The next
    time you reset or turn on the system, the system boots off the
    fixed disk on which you have installed the system software.

    NOTE: If you do not want to remove the cartridge from the drive,
    but do not want the system to boot off the cartridge, erase file
    SYSTEM.MON[1,4] from the cartridge. Now the system cannot
    recognize the cartridge as a System Disk and therefore will not
    attempt to boot off it.

    However, it is most important that you always keep an intact System
    Disk cartridge (i.e., it contains recent system software and
    SYSTEM.MON and SYSTEM.INI) around so that you can re-install the
    system should the fixed disk become damaged.

6. The fixed platters are now logical units DSK0:-DSK4: and the
    non-System Disk cartridge is DSK5:.

(Changed 30 April 1981)

3.0 THE CERTIFICATION PROGRAM, CRT410

The CRT410 program certifies a Phoenix logical unit by reading, writing, and verifying every block of the disk. It also formats and initializes the disk.

3.1 Hints and Restrictions

There are several things you should keep in mind before using the certification program:

1. Only the System Operator may run CRT410. Log into the System Operator's account, [1,2], before using the program.

2. You may ONLY use CRT410 on disks that run under control of the AM-410 Controller.

3. Run CRT410 on every logical unit of a Phoenix disk before the first use of that device. Make sure that the logical unit is not write-protected.

4. A logical unit does not have to be mounted before you certify it; CRT410 mounts the disk for you.

5. CRT410 writes data in every byte of the logical unit you are certifying. If there is data on that unit, make sure that you back it up onto another device before using CRT410.

6. CRT410 communicates directly with the AM-410 Controller without going through the Phoenix driver program. Therefore, you MUST NOT run CRT410 at the same time as any other program that accesses devices that run under control of the AM-410.

3.2 Sample Use of CRT410

Below is sample output of a typical disk certification. The next section discusses the questions that CRT410 asks and the messages that it displays.

```
.CRT410 SMD0: (RET)
CAUTION: This program writes to all blocks

Enter maximum acceptable number of bad tracks: 40 (RET)
?15 bad tracks is maximum
Enter maximum acceptable number of bad tracks: 10 (RET)
Display current track? (Y or N):Y (RET)
Current track is: 0
Enter serial number (10 char. max): PAYROLL1 (RET)

Begin certification of SMD0:
Current track is: 1
Current track is: 2
                .
                .
                .
?Track 7 did not verify
?Track 8 did not verify
                .
                .
                .
Current track is: 14
Current track is: 15
                .
                .
                .
Current track is: 807

?2 bad tracks detected
Certification complete
_
```

3.3 Using CRT410

To use CRT410, enter CRT410 followed by the specification of the logical unit you want to certify. For example:

```
.CRT410 SMD5: (RET)
```

You now see the following message:

CAUTION: This program writes to all blocks.

(If you do not want to continue the certification, you may enter a Control-C at this point.)

CRT410 now asks you several questions:

1. Enter maximum acceptable number of bad tracks:
        Give CRT410 the maximum number of bad tracks that you will
   accept on the disk you are certifying. If the number of bad tracks
   that CRT410 finds exceeds this value, CRT410 tells you so and then
   aborts the certification, returning you to AMOS command level:

              ?Device has exceeded maximum number of errors
              :

2. Display current track? (Y or N):
        If you want CRT410 to tell you as it verifies each track,
   enter a Y; otherwise, enter an N. If you answer Y, CRT410 now
   displays this message:

              Current track is: 0

   NOTE: Asking CRT410 to display the number of the track it is
   currently verifying greatly increases the length of time it takes
   to certify a disk surface.

3. Enter serial number (10 char. max):
        You may optionally give CRT410 a ten-character alphanumeric
   I.D. for the logical unit you are certifying. CRT410 writes this
   identifier into the BADBLK.SYS file for that logical unit.

After answering the questions above, CRT410 begins to certify the disk. You
see this message:

        Begin certification of Devn:

where Devn: is the device specification you supplied on the CRT410 command
line.

If you asked CRT410 to tell you its track position as it certifies, you now
see a list of messages that can look something like this:

        Current track is: 1
        Current track is: 2
        Current track is: 3
              .
              .
              .

and so on. When CRT410 encounters a track that does not verify, it tells
you so. For example:

?Track 15 did not verify
?Track 16 did not verify
                    .
                    .
                    .

When finished certifying the disk, CRT410 tells you how many bad tracks it
found. For example:

        ?4 bad tracks detected
        Certification complete
        .

3.4 How CRT410 Certifies a Disk

CRT410 follows these procedures when it certifies a disk:

1. CRT410 creates a file named BADBLK.SYS in account [1,2] on the disk
    you are certifying. This file will contain a list of all of the
    bad tracks on the disk. If you have specified a serial number,
    CRT410 writes that information to this file.

2. CRT410 writes one data pattern in every byte on the first track of
    the disk. Then it reads each byte on that track and verifies the
    data. (CRT410 checks for CRC errors as well as data verification
    errors.)

    If any data does not verify, CRT410 places the number of the track
    in the BADBLK.SYS file and tells you that the track is bad. For
    example:

            ?Track 35 did not verify

    CRT410 writes a total of four data patterns, following the
    procedure above for each data pattern.

3. CRT410 now moves on to the next disk track and performs these
    operations again, reporting any bad tracks that it finds (and
    entering their numbers into the BADBLK.SYS file). CRT410 verifies
    every track. When it finishes, CRT410 tells you that it is done
    and how many tracks are bad. For example:

            ?3 bad tracks detected
            Certification complete

4. Now CRT410 computes a hash total for the BADBLK.SYS file and stores
    it in the file. This value provides a validity check that other
    programs (e.g., BADBLK) can use to make sure that the BADBLK.SYS
    file is complete and healthy. When AMOS accesses a disk block, it
    checks to see if that block occurs on a bad track; if so, AMOS uses
    the alternate track assigned to that bad track instead.

If you interrupt the CRT410 program by typing a Control-C, you see:

?Certification incomplete
^C

.

and CRT410 intentionally writes a bad hash total to the BADBLK.SYS file. (A bad hash total tells other programs that may later look at the file that the data in the file is not complete and is not to be trusted.)

3.5   CRT410 Error Messages

Below are the error messages you can encounter when using CRT410.

?You must be logged into PPN [1,2] to run CRT410
        Because it writes data into every byte on the disk, CRT410 is a dangerous program to run. You must be logged in as the System Operator to certify a disk. Log into the System Operator's account, [1,2], before trying to use CRT410.

?15 bad tracks is maximum
        You specified a number greater than 15 as the maximum number of bad disk tracks you would accept on the certified disk. However, the BADBLK.SYS file cannot handle more than 15 bad disk tracks; so, enter a number less than 15.

?Track 0 did not verify. (First track must verify.)
        The first track of the disk did not verify. CRT410 cannot continue the certification if the first disk track does not verify, so it now stops the certification and returns you to AMOS command level.

?Track n did not verify.
        CRT410 marked track n in the BADBLK.SYS file as a bad track.

?Device has exceeded maximum number of errors
        CRT410 found more bad tracks than the value you specified as the maximum number of bad tracks that you would accept. CRT410 now aborts and returns you to AMOS command level.

?S100 data transfer error
        An error occurred with the AM-410 controller board. CRT410 aborts and returns you to AMOS command level. If you receive this error several times, you may have hardware problems.

?Nonexistent device
        Your device specification on the CRT410 command line is invalid; the system believes that the device does not exist. Check your spelling and try again.

?Certification incomplete
        You typed a Control-C to interrupt the disk certification. CRT410
now intentionally writes a bad hash total to the BADBLK.SYS file to let
other programs know that the data in the file is incomplete and not to
be trusted.

4.0 THE BADBLK PROGRAM

The BADBLK program allows you to see the contents of the BADBLK.SYS file
created by the certification program, CRT410.   BADBLK also verifies the
BADBLK.SYS hash total. In future releases, BADBLK will allow you to modify
BADBLK.SYS to rebuild a damaged disk.

BADBLK checks the specified disk to see if it was certified by a pre-4.4
version of CRT410; if it was, BADBLK displays the number of bad blocks on
the disk. If the disk was certified by a 4.4 version of CRT410, BADBLK
displays the number of bad tracks. (The messages you see will tell you if
BADBLK is displaying the number of bad blocks or tracks.)

NOTE: Although AMOS version 4.4 handles disks that have been certified via
the "bad block" or the "bad track" method, future releases may not support
the old bad block system. Therefore, you must convert all Phoenix surfaces
certified by pre-4.4 CRT410 over to the new system by certifying those
surfaces with CRT410 versions 4.4 or later.

4.1 Using BADBLK

To use BADBLK, type BADBLK followed by the specification of the device whose
BADBLK.SYS file you want to see. Then type a RETURN. For example:

        .BADBLK SMD1: (RET)

If BADBLK found the BADBLK.SYS file, it tells you so:

        SMD1: BADBLK.SYS[1,2]

BADBLK now tells you the serial number associated with that device and the
number of tracks or blocks marked as bad on that disk. For example:

        Serial number: INVENTORY2
        Number of bad blocks: 0

If there are any blocks or tracks listed in the BADBLK.SYS file, BADBLK
lists them for you. For example:

(Changed 30 April 1981)

<u>Number of bad tracks: 3</u>

<u>35   36   37</u>

BADBLK exits and returns you to AMOS command level:

<u>EXIT</u>

<u>.</u>

NOTE: Track numbers are decimal; block numbers are octal.

4.2   BADBLK Error Messages

You can see the following BADBLK error messages:

<u>?File not found: Devn:BADBLK.SYS</u>
      BADBLK was not able to find the BADBLK.SYS file for the disk you
specified.   Make sure that the device you specified is a Phoenix disk
and that the disk has been certified (i.e., you've run CRT410 on that
device).

<u>CAUTION: HASH TOTAL DID NOT VERIFY</u>
      The BADBLK.SYS file contained a bad hash total. This indicates
that the data in that file is not to be trusted. Use COPY to copy all
files off the logical unit containing that BADBLK.SYS file, being
careful not to overwrite the BADBLK.SYS file on the new disk.   Then
re-certify the disk.

You may also see several system error messages if your device specification
is invalid. For example:

<u>?Cannot INIT Devn: - device does not exist</u>
      The system did not recognize the device specification you gave.
Check your spelling and try again.

<u>?Cannot READ Filespec - disk not mounted</u>
      The system is unable to read the device you specified on the
BADBLK command line because it is not mounted. Use the MOUNT command
to mount the disk and try again.

5.0   THE DSKANA PROGRAM

For the most part, the new method of handling disk flaws is transparent.
For that reason, DSKANA is able to treat any disk that was certified with a
4.4 or later version of CRT410 in the same way that it does a device that
does not contain a BADBLK.SYS[1,2] file.

Index