

[This is a facsimile of a 1977 Bell Labs memo on Mini-UNIX that came as part of our Mini-UNIX distribution. I have another, similar, document entitled "Mini-UNIX Summary", which is not reproduced here.]

Bell Laboratories

Subject: The MINI-UNIX System

Date: January 3, 1977

From: H. Lycklama

MEMORANDUM FOR FILE

1. Introduction

The MINI-UNIX Operating System was written to run on all PDP-11 processors without the memory management unit available on the PDP-11/40, 11/45 and 11/70 processors and is therefore restricted to a 28K word address space. The operating system itself (MX) is basically a modified version of the Version 6 UNIX Operating System (1) and as such supports most of the standard UNIX system calls. The system resides in 12K words of memory and is written in the C language. An instruction emulation package is included in the system for those machines which do not support the extended instruction set (e.g. mul, div, ash, etc.). The system supports up to four users using a simple round-robin time-slicing algorithm. It supports all of the UNIX user programs unmodified. User programs which have been slightly modified are discussed in a later section. MINI-UNIX thus provides an inexpensive software development system in a UNIX time-sharing environment for those installations with a minimum amount of hardware which is insufficient to support the full Version 6 UNIX Operating System.

Other software tools are also available for easing the transition from software written under the DEC DOS Operating System. This includes a macro assembler and linker.

2. Hardware

The MINI-UNIX system runs on any PDP-11 processor with 28K words of memory. The PDP-11 computer is a 16-bit word mini-computer with a UNIBUS for interfacing DEC peripherals to the CPU. The typical configuration consists of a PDP-11/10 CPU with 28K words of memory, a console terminal and an RK05 moving-head disk controller with two removable disk cartridges for swapping and file system storage. Each RK05 disk pack has 2.5 Megabytes (8-bit byte) of storage. However, the MX system also supports the RF fixed-head disk (1 Megabyte) and the RP03 and RP04 moving-head disk controllers with 40M bytes and 80M bytes, respectively. Other peripherals supported include line-printer, Dectape, magtape and various asynchronous and synchronous interface units.

The system is normally configured to be 12K words in size. This includes an emulation package for interpreting the 10 extended instructions normally performed by the EIS hardware available as an option on some PDP-11 processors and standard on the PDP-11/45 processor. A minimum system has room for 6 or 7 system buffers. As new drivers are added to the system, the number of system buffers must be decreased if the system size is maintained at 12K words. Thus it is recommended that for some applications it may be appropriate to add the drivers for only a few peripherals on any one version of the system and thus maintain a few versions of the system, one for each set of drivers desired concurrently in the system. This keeps the system size at 12K words in order to be able to support all of the user software of Version 6 UNIX.

3. System Features

The Operating System itself is written in the high-level language, C (2) and as such bears a strong resemblance to the standard UNIX Operating System which runs on the PDP-11/40, 11/45 and 11/70 computers. Because of the memory address space limitation due to the lack of a segmentation unit, the system size is generally kept to 12K words in size. Thus its capabilities are somewhat less than those of the standard UNIX system, especially in the area of interprocess communication and interactive debugging. The number of processes (and users) which MINI-UNIX (MX) can comfortably support is in general much less than for the standard UNIX system. No more than four users are recommended for MX. Currently the number of processes allowed is thirteen. This is a 'sysgen' parameter. The address space available to a user program is 16K words. This enables almost all user programs which run under Version 6 UNIX to run unmodified under MX. The exceptions are noted in a later section. The capabilities of MX fall somewhere in between those of standard Version 6 UNIX and those of the LSX operating system (3). LSX is a single-user UNIX-compatible system for the LSI-11 microprocessor using floppy disks as secondary storage.

MX supports all of the UNIX system calls of the Version 6 UNIX operating system with the exception of: ptrace, pipe, prof, setgid and getgid. For the sake of completeness the status of implementation of the system calls are summarized in an Appendix. Thus MX supports all UNIX user programs with the exception of programs which use the above mentioned system calls. User programs are compiled and relocated to start at address 060000 and may occupy up to 16K words of memory. Up to 13 processes may exist at any one time, although only one process may be in core and running. This should be sufficient to handle up to four simultaneous users. Since no memory relocation is available, the complete user program image must be swapped out to bring in a new user program. Hence no sharing of text is possible. No software memory management is required. Scheduling is done on a simple round-robin basis with each process in the run state receiving a two second time slice.

The file system supported by MX is identical to that provided by the UNIX time-sharing system. The structure of the super-block and of the file inodes is maintained. MX supports the identical file system hierarchical structure and makes the same distinctions between ordinary files, directories and special files. Removable file systems are supported as well. Hence 'mount' and 'umount' system requests are treated identical to those in standard UNIX. Under MX, a file's size is limited to one megabyte. Large files are supported, but huge files (two level indirect block) are not. Groups are not supported in MX. Thus a file has only a given owner user ID and no group ID. Read, write and execute permission bits are available for both owner and non-owner of a file. The set-user ID bit is also supported in MINI-UNIX.

No interactive debugging is possible in MX since the 'ptrace' system call is not implemented. One may still use the C debugger 'cab' but cannot plant breakpoint traps in the running image of a child process. It may only be used for post-mortem debugging on core images. The profiling of a process to determine where it is spending its time is also not permitted. Pipes have not been implemented in the MX system for two reasons. One is that it requires too much code in the system which is already butting its head against the top. Two is that the overhead required in switching from the process writing the pipe to the process reading the pipe involves a process swap. Filters, however are simulated at the command level as discussed in the next section.

Some other features have been stripped out of the system in MX in the interest of minimizing system address space. Upon reading a file block, "read-ahead" is not invoked. This contributes only slightly to a loss in throughput. Physical I/O has also not been implemented. This precludes the ability to read large contiguous pieces of disk directly into the user's address space without system side-buffering. The system buffering scheme is much simpler than in standard UNIX but a maximum of only eight buffers may be allocated in a minimum system configuration.

4. UNIX User Program Modifications

Since the MX system uses no segmentation unit, all user programs must be compiled and relocated to start at address 060000. Under MX, the C programs are relocated automatically using the 'ld' program which has been modified as described below. For assembly-language source programs, the 'a.out' program may be relocated either using the 'reloc' program or the 'ld' program. The various user programs which have been modified from the standard Version 6 UNIX or new programs which are supported but not mentioned in the UNIX Programmer's Manual (4) are described below.

4.1. ar

The archive program supported is a version which is updated from that described in the UNIX Programmer's Manual. It is written in the C language and has some new keys added. However the documentation still applies.

4.2. bc

The 'bc' command works as in the standard UNIX system but cannot be used interactively since pipes are not implemented in the system. In generating the source for 'bc' using the 'yacc' compiler-compiler, the source must be edited slightly to make it simulate the use of pipes.

4.3. check

The check program is not described in the UNIX Programmer's Manual. However, it combines most of the individual features of dcheck, ickcheck and ncheck as described in the manual. A nag of -I will do the equivalent of ncheck with no flags. All other flags described for dcheck, ickcheck and ncheck apply.

4.4. Debugger

The debugger db has been slightly modified in order to ease debugging under the MINI-UNIX system. The default relocation address of all programs is assumed to be 60000(8). If the flag '-a' is given when the debugger is invoked, the relocation address is assumed to be zero. This is useful for debugging the system.

The default address can be changed by editing 'dbl.s' and setting the variable 'uorg' equal to the new relocation address. Then the source code should be re-assembled. If none of the changes are desired, the variable mx in dbl.s should be set to 0 and the source code reassembled.

4.5. Kdmp

This program is used to extract a dump of the complete system (28K words) from the swap area on disk. The core image produced in the file "kore" may then be debugged post-mortem. Consult system source code for the actual disk tracks used. If a system crash occurs and the core image of the system is written out to disk, the system should be booted up single-user. Immediately, one should execute "kdmp" (preferably on an uncorrupted file system) before the swap area is over-written.

4.6. Ld

The ld program has been modified to relocate the program origin to 060000 automatically. This is to facilitate the compilation of C programs under the MINI-UNIX system. Note that all assembly programs, after being assembled, must be relocated to 060000 either by using the link editor ld or the relocation program reloc. The relocation origin may be changed from 060000 by changing the value of the TOPSYS parameter in ld.c to the appropriate value.

Using the '-a' option flag with ld turns off the relocation of a program's origin. The program is then assumed to start execution at location 0. This is useful for link editing the MINI-UNIX system itself or

any other program which starts execution at location 0.

4.7. Mkpt

The mkpt program constructs a prototype file given a specification file for direct input to the mkfs program. Run "man 8 mkpt" for details.

4.8. Ps

The process status command outputs basically the same information as that of the Version 6 UNIX ps command. It has been modified to take into consideration the different process table layout and the different swapping technique used in MINI-UNIX.

4.9. Reloc

This Program is used to relocate all relocatable symbols in a program. Thus "reloc a.out 60000" will relocate all relocatable symbols in "a.out" and relocate the starting address of the program (absolute zero) to 060000. It must be used with all assembler output programs if the linkeditor is not required.

4.10. Sh

Since the MINI-UNIX system does not support pipes, the shell has been modified to simulate pipes through the use of disk files. When a command line which requires the use of a pipe is detected, a disk file is created and opened for reading and writing. The file is immediately unlinked, so that the name is available for another pipe right away. The file is called '._pf'. Hopefully this name will not conflict with any user file names. Thus, the symbol '|' in a shell command line becomes equivalent to '> ._pf; < ._pf'. The command:

```
% prog 1 | prog2
```

translates into:

```
% progl > ._pf; prog2 c ._pf
```

The process writing on the pipe writes everything into the file, and when it exits, the reader process is swapped in. It reads what the writer has written on the pipe file. The only danger in this type of implementation is due to the limitation of space on the file system being used. A very large amount of information going through a pipe could fill up the disk. Except for this, the pseudo-pipe code is transparent to the user.

The sync command has been added to the shell. Thus the system call is made directly from the shell and no other process is spawned.

4.11. Typo

This program checks for the most likely spelling errors in a document. It has been modified somewhat to enable it to fit in 16K words of memory.

4.12. Yacc

The source for the yacc compiler has been edited to change some table sizes to make it run in 16K words of memory. Further editing of symbols at the beginning of yO.c may be required to make up a yacc compiler with different table sizes for a particular application.

5. Summary

Since MX uses no segmentation unit, no protection is provided for the user program. Thus new user programs must be debugged carefully. In practice, the use of the C language limits the user's program's use of the program counter and stack pointer thus limiting damage and usually causing a bus error before anything drastic happens. The lack of a segmentation unit does have its advantages. It means the user can directly access all I/O registers on the UNIBUS and does not have to write special device drivers interfaced with the file system to control the special peripherals. Thus in cases where a real-time program is to be run, one may disable the system clock to inhibit unwanted clock interrupts and also swapping of processes. One may also catch clock interrupts during the running of user programs if the clock is to be used for user program timing control. The clock should then be restored to system control upon exit from the user program.

There is also another set of user programs available under MX which may be used to ease transition from DOS, the DEC operating system, to UNIX for those installations now using the DEC DOS operating system on a PDP-11 CPU. This package consists of a macro-assembler and a linker-loader for assembling programs written under DOS for the DEC macro-assembler. The result is a UNIX 'a.out' file.

The normal configuration for MX includes a PDP-11/10 CPU with 28K words of memory and two RK05 disk cartridges for secondary. The PDP-11/10 processor is slower than the PDP-11/40 processor and does not have the full instruction set of the PDP-11/40 processor, thus requiring the emulation of the missing instructions. A typical C compilation requires about twice the total time of that required on the equivalent PDP-11/40 configuration. However, response to the editor commands is not significantly longer than on a more powerful CPU. The cost of a minimum configuration:

- PDP-11/10 CPU
- 28K words memory
- 2 RK05 disk drives
- KL11 interface to control console
- DL11E interface to dial-up line
- 60 cycle clock

is of the order of \$20,000 at today's prices (December 1976). This provides an inexpensive tool for software development in a UNIX time-sharing environment for those configurations which have insufficient hardware to support a full Version 6 UNIX system.

H. Lycklama

MH-1352-HL [???

Atts. References Appendix

Reproduced by [Mark Riordan](#).