

I N T R O D U C T I O N T O T R I C E P

Copyright (C) 1984 by Morrow Designs, Inc
All rights reserved.

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without prior written permission of Morrow Inc.

DISCLAIMER

No representations or warranties, express or implied, are made with respect to the contents hereof, including, but not limited to, the implied warranty of merchantability or fitness for a particular purpose. Further, Morrow Inc., reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation to notify any person of such revision.

Morrow
600 McCormick St.
San Leandro, CA 94577

IMPORTANT WARRANTY INFORMATION

LIMITED WARRANTY

Morrow Designs, Inc. warrants its products to be free from defects in workmanship and materials for the periods indicated below. This warranty is limited to the repair and replacement of parts only.

This warranty is void if, in the sole opinion of Morrow Designs, Inc., the product has been subject to abuse or misuse, or has been interconnected to other manufacturer's equipment for which compatibility has not been established in writing.

Circuit boards - Parts, including the printed circuit board, purchased as factory assemblies, are warranted for a period of ninety (90) days from the original invoice/purchase date.

Electro-mechanical peripherals - Peripheral equipment such as floppy or hard disk drives, etc., not manufactured by Morrow Designs, Inc., are included in the limited warranty period of 90 days from the original invoice date when sold as part of a Morrow system.

Exception - Expendable items such as printer ribbons, software media, and printwheels are not covered by any warranty.

Software/Firmware - Morrow Designs, Inc. makes no representations or warranties whatsoever with respect to software or firmware associated with its products and specifically disclaims any implied or expressed warranty of fitness for any particular purpose or compatibility with any hardware, operating system, or software/firmware. Morrow Designs, Inc. reserves the right to alter or update any program, publication or manual without obligation to notify any person of such changes.

LIMITATION OF LIABILITY: THE FOREGOING WARRANTY IS IN LIEU OF ALL OTHER WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL MORROW DESIGNS, INC. BE LIABLE FOR CONSEQUENTIAL DAMAGES EVEN IF MORROW DESIGNS, INC. HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

WARRANTY RETURN PROCEDURE

Should a buyer experience a defect in either workmanship or materials during the warranty period, any Morrow Authorized Service Center will replace or repair the product at its expense only if the product is promptly returned to the dealer or Service Center with dated proof of purchase.

Should factory repair be necessary, the Service Center shall contact Morrow Customer Service for a Return Materials Authorization (RMA) number.

T A B L E O F C O N T E N T S

1.	FORWARD TO THE TRICEP MANUAL	1
2.	INSTALLING YOUR TRICEP SYSTEM.	3
2.1	Installation Overview.	3
2.2	STEP 1: UNPACK AND INSPECT	4
2.2.1	Checking for Hidden Damage	5
2.3	STEP 2: SELECT A LOCATION	8
2.4	STEP 3: PLUG IN THE POWER CORDS	9
2.5	STEP 4: TURN THE POWER ON, FIRST TIME	10
2.6	STEP 5: SETTING UP THE FIRST TERMINAL	11
2.6.1	Terminal Settings.	11
2.7	STEP 6: CONNECTING THE CONSOLE CABLE.	12
2.8	STEP 7: BOOTING, THE FIRST TIME	14
2.8.1	Troubleshooting Booting Problems	15
3.	FIRST TIME USE	21
3.1	CHECK YOUR FILE SYSTEM!.	21
3.1.1	Checking for Bad Blocks.	22
3.2	File System Check.	23
3.3	First Backup	24
3.3.1	Floppy Disk Backup	24
3.4	Quick Tour of the File System.	26
3.4.1	Trees: the Basic Philosophy.	26
3.4.2	Why Trees?	27
3.4.3	Moving Through the File System (Made Easy)	27
3.5	Making Room: "Unnecessary" Files.	32
4.	CONFIGURING YOUR TRICEP.	33
4.1	From Boot to Multi-user.	34
4.1.1	Inittab Controls INIT.	36
4.1.2	Inittab Example.	40
4.1.3	Modified inittab Example	41
4.1.4	Sending Signals to init.	42
4.1.5	Processes: A Few Words.	43
4.1.6	gettydefs: Sets Up Baud Rate	44
4.1.7	etc/passwd Controls login.	46
4.1.8	Controlling Shells: /etc/profile, /etc/cshrc.	47
4.2	Connecting Additional Terminals.	47
4.3	/etc/passwd: The User Database	50
4.3.1	/etc/passwd, The User Name File.	50
4.3.2	Editing /etc/passwd.	53
4.3.3	Editing /etc/group	54
4.3.4	Creating the New HOME Directory.	55
4.4	Removing Users	56
4.5.	58
4.6.	Connecting Printers.	60
4.6.1	Serial Printers.	60
4.6.2	Parallel Printers.	61
4.7.	LP, Print Spooling Facility.	62
4.7.1	LPADMIN, the lp Configuration Command.	62
4.7.2	Using lp	67

4.8.	Adding Disks to Your System.	70
4.8.1	Mounting File Systems.	71
4.8.2	Unmounting File Systems.	72
4.9.	Connecting Modems.	72
4.9.1	Modem to TRICEP Cable.	73
5.	MORROW ENHANCEMENTS.	76
5.1.	Diskette Specialists: far and dar.	76
5.1.1	far and dar: Command Line Syntax	77
5.1.2	far Differences.	78
5.1.3	dar Differences.	79
5.2.	ddt, Disk Debugging Tool	80
6.	TRICEP SYSTEM SPECIFIC VALUES.	81
6.1.	Major and Minor Device Numbers	82
6.1.1	TTY Devices (SIO4-DMA)	83
6.1.2	Mini Winchester Devices (HDC-DMA).	83
6.1.3	Floppy Disks (DJ-DMA).	85
6.2.	Every Day Procedures	88
6.2.1	Booting.	88
6.2.2	Going Multi-User	88
6.2.3	Setting the Date	89
6.2.4	Running FSCK	90
6.2.5	Checking for Free Disk Space	90
6.2.6	Going Single-User and Backing Up	92
6.2.7	Turning the System Off	94
7.	USING THE REST OF THE MANUALS.	95
7.1.	Divisions in the Documentation	96
7.1.1	TRICEP Installation and Maintenance Guide.	97
7.1.2	User Guide and Manuals	97
7.1.3	Document Processing Guide.	98
7.1.4	Programming Guide.	98
7.1.5	Support Tools Guide.	98
7.1.6	Administration Guide and Manual.	99

1. FORWARD TO THE TRICEP MANUAL

Welcome. You have found the portion of the manuals that deals directly with your TRICEP system. The instructions for installing and configuring your system are in these pages. Other information that applies to your system hardware, the TRICEP, are found here.

The bulk of the documentation that you received was originally written by Bell Labs or the University of California, Berkeley. This material was edited by UniSoft so that it reflects UniPlus+ System V running on Motorola 68000 processors. This information is as correct as is humanly possible. It is also of little help to anyone new to the UNIX system.

Besides providing TRICEP specific information, I will attempt to guide you to an understanding of the UNIX system. You will probably need more help than this manual provides if you are new to UNIX. There really is so much to UNIX that it takes several books just to explain most of the things available.

I will also try to convey to you some of the excitement I had when I started working with UNIX. Here, at last, was an operating system that could be controlled through files written in English (almost). Here was a wealth of commands, an easily organized file system and a flexibility of devices (printers and terminals). No more poking bytes into mysterious locations in memory. No more directories with 200 incomprehensible file names. I was really pleased.

So, here we go. I will try to remember what was mystifying when I first met UNIX and explain it to you. Some of these explanations are elementary, and not intended for UNIX experts. I don't want to lose anyone by too complex an explanation. I also don't want to put anyone off by being too simplistic, so please bear with me.

Finally, I realize that no one likes reading manuals. Manuals are the source of last recourse. Since no one reads entire manuals, this one is organized so that you can pick out and read what you need. Most topics are covered quickly in an overview, then beaten more or less to death in a wordy version. Diagrams are included where I felt they'd be helpful. There also is an index in the back that can help you find what you need quickly and a table of contents in the front.

Questions and comments about this manual can be mailed to me in care of the Documentation Department at Morrow.

Much Good Fortune,

Rik Farrow
August 9, 1984

2. INSTALLING YOUR TRICEP SYSTEM

The following sections of this manual explain how to install and test the basic TRICEP system. This includes the TRICEP computer and the first terminal. Installation of additional terminals and a printer are explained later.

Anyone who has ever set up a computer system knows that it is easy to make a mistake. There is, of course, a lot of room for error. Even if you know what you are doing you still will make mistakes.

We are going to make things easier for you by explaining everything two or three times. The first explanation is a bare bones description that covers all the steps involved. You can use the overview by itself, but you will probably need a little more explanation.

Many pages of in depth explanation follow the overview. Each step in the overview is explained in great detail. And, illustrations are used (the third explanation) where they are most helpful. If you don't understand something in the overview, look in the wordy-with-pictures section that explains that step. There is a section for each step in the overview.

We provide checkpoints with most steps so that you can be sure that you have succeeded in understanding and following instructions. The things that we expect the most problems with are steps 5 and 6, setting up and connecting the terminal. Step 7 has checkpoints that cover the boot process.

We'll also do our best to explain briefly and simply what you are doing to your system as you begin using the software. This should provide you with a better understanding of how UniPlus+ System V works.

2.1 Installation Overview

Before we get started, this section provides a brief description of the procedures involved in installing your TRICEP system hardware. This section may be all that you need if you have previously installed a TRICEP system. For those of you first-timers, the installation steps are explained in greater detail in the sections that follow. Here goes...

1. Unpack and check for external damage; remove the insert from the floppy drive;
2. Choose a suitable location for the computer and its peripherals;
3. Plug the TRICEP into a grounded socket; plug a terminal into another grounded socket;

4. Turn on the power (turn TRICEP keyswitch to clockwise): are the fan and keyswitch panel light on?
5. Set the terminal to 19200 baud, 8 data bits, 2 stop bits and no parity; turn it on; can you see your cursor?
6. Connect this terminal to the RS-232 connector located above the three RS-232 connectors along the lower edge of the back of the TRICEP; there are (at least) four of these connectors: this connector may be labeled "console";
7. Turn the keyswitch to "RESET", and press the RETURN key when asked to do so after the colon. This should boot the UNIX system (you will see the "#" prompt when this is done). Go on to FIRST TIME USE.

2.2 STEP 1: UNPACK AND INSPECT

The instructions for unpacking your system are essentially the same as for unpacking a piece of stereo equipment: carefully remove the cabinet from the box and save all packing materials in case you ever want to ship the system again. Besides the cabinet containing the TRICEP, you should have also found:

- o the UniSoft documentation and these instructions,
- o a single diskette,
- o a power cord, and
- o a cable for printers with centronics interface.

The diskette contains a standalone floppy UNIX system that is used for repairing damaged file systems on the hard disk.

THIS IS IMPORTANT:

There is a cardboard insert in your floppy drive to protect the heads during shipping. **You must remove it before applying power to the system.** Lift the door latch and slide it out. Leave the door open. Easy enough?

The power cord and Centronics cable are the only external parts that are packed with the TRICEP. Since you are going to connect at least one terminal to the system, you also need an RS-232 cable with two male ends for each terminal you are connecting. Hopefully, your dealer was savvy enough to realize this and sold you one with each terminal. RS-232 cables are standard supplies at any computer store.

The minimum requirements for the RS-232 cable are that it has at least pins 2, 3, 7 and 20 connected at both ends. Pin 2 gets attached to pin 2, pin 3 to pin 3, etc. The section on installing printers has more details on RS-232 cables. If you bought a Morrow terminal, a RS-232 cable will be in the box with the terminal.

Keep your carton and packing materials! If you ever need to ship your terminal or computer, you will thank yourself a thousand times for saving the containers and packing materials. Morrow requires that the TRICEP be shipped in the original container.

If you do see blatant physical damage to the TRICEP notify the carrier immediately for filling out a damage claim. If you decide to return your system, notify the people you are returning it to before reshipping.

2.2.1 Checking for Hidden Damage

By the time you have removed the TRICEP cabinet from the packing materials, you should have discovered any gross damage suffered during shipping. In spite of being well packed, it has been our experience that shippers sometimes drop boxes hard enough to dislodge even securely fastened PC boards. This means that you may have loose, rather than broken, parts that you can reattach yourself. And by "hidden", we mean inside the TRICEP cabinet. Opening the cabinet is simple and not dangerous, if due caution is exercised.

You do NOT need to check inside your TRICEP unless you want to or are directed to do so by the troubleshooting section in Step 7.

If you have already plugged in your TRICEP, please unplug it. The cover to the metal cabinet is attached by four screws, two along the bottom of each side, that you need to remove before you can slide off the cover. The cover slides off toward the front.

When you have removed the cover, you will have revealed the printed circuit boards. Each board looks like a green or blue rectangle, about 5 by 10 inches (13 x 25 cm) in size, covered with black integrated circuits on one side. Two brass rails secure the boards at the top. The boards should be neatly standing in a row, with white plastic guides holding the boards upright, and a black plastic connector fastening them along their lower edge.

By looking down between the pc boards (there is room for several more boards right in the middle) you will see another, much larger, green printed circuit board that extends from the front to the rear of the cabinet. This board is called the WUNDERBUS I/O, and holds 14 black connectors, each about 7 inches (18 cm) long. These connectors work identically to one another, that is, they provide the same data and signals to each printed circuit board that is plugged into them. To work properly, the lower edges of the smaller pc boards must be firmly attached to the black connectors. This connection, between the smaller printed circuit boards and the WUNDERBUS I/O, is sometimes a problem when shipping computers.


Checkpoint:

We are going to check for two things: that the boards are held between the white plastic guides, and that all the boards are firmly connected.

1. To see whether each pc board is still held between the two white plastic guides, look along the left and right ends of the boards. The white guides begin 3/4" (2 cm.) below the top of the boards, and its easy to see whether each board is resting in the guides, unless you keep your computer under the table or in a dark room.
2. It's pretty hard to see the connection between the boards and the WUNDERBUS, but you can do two things. You can see that the top edge of each board is about 3/8" (.9 cm.) above the metal sides, and that the tops of all the boards appear to be about the same height. And you can press down on the tops of the boards to reseat them. To reseat boards, press gently with your thumbs on both ends of the top of a board. You can use up to about 10 lbs. (4 kilos) of force pressing down, and can alternate the force from your left thumb to your right to gently rock the board into place. Even if all the boards look seated, try this with at least one board.

If you find a board that is completely unseated, well, reseat it. You will probably need to loosen the retainer rails that are held in place by two phillips head screws. If you find a loose board or cable, it should be readily apparent where it connects. If not, the drawing on the facing page shows the proper location for boards and cables.

If you can't figure out where to connect a cable, get on the horn to your dealer. These situations are very unlikely, but we want you to be as self-reliant as possible in case they occur.



2.3 STEP 2: SELECT A LOCATION

The choice of the location for your TRICEP is up to you. However, there are four requirements that must be followed. Your TRICEP should be installed where it will remain cool, level, stable and dry. The level and stable requirements refer mainly to the hard and floppy disk drives. These both contain moving parts that are designed to be operated on a level surface for maximum life expectancy.

And, if you drop or jar a hard disk, the head alignment inside the sealed unit may change, making everything on the disk unreadable. Or, the surface of the disk may be scratched, making parts of it unreadable. Hard disks are SUPPOSED to be able to withstand some abuse, but why take chances?

The definition of "dry" for electronic equipment is that the environment must be non-condensing. That is, the humidity should never be so high that water condenses on the metal parts of your computer. Obviously, rain, or spilled coffee, is too wet.

"Cool" means that the room temperature where the TRICEP is installed must remain below 85 degrees Fahrenheit (28 Celsius) for proper operation. Some components of the system will behave erratically when the temperature rises above this point. Remember to leave room for ventilation (4 inches, 10 cm.) around the top, sides and bottom of your TRICEP.

And, if it is really dry in your environment, you may have a problem with static. The static-charge that you can build up in your body walking across a rug might amount to thousands of volts, but very little current. What this means is that what feels like a little shock to you is like electro-convulsive therapy for your computer. Computer-brain damage will be the likely result. Problems with static can be vanquished by purchasing a rug that has copper wires woven into it. Strange though it may seem, rugs with metal content are not too uncommon, and you may be able to buy a small one from a rug store. Or, you can get an anti-static rug from an office supply or computer retailer.

One final suggestion about location. Like a stereo system, all of the connections for your system are located on the back of the cabinet. If you can arrange the system so that you can stand behind it, at least while installing it, everything will be a lot easier. Once the physical installation is completed, you can move the system so that the back faces a wall. Remember to leave at least 4" (10 cm.) at the back and over the top of the cabinet for ventilation.


2.4 STEP 3: PLUG IN THE POWER CORDS

... but don't just jump right in and do that. Read this section first, and be sure that the power switches to the TRICEP and its peripherals are off. The TRICEP has a keyswitch on the front panel. Turn it fully COUNTER-CLOCKWISE (to the left) to insure that your system is off.

Your TRICEP comes configured from the factory for 115 volt operation. Just to be absolutely sure that your available voltage matches what the computer needs, check the switch on the rear of the TRICEP. The switch is next to the power socket and tells which voltage level the computer expects. Never change this switch when the power to the TRICEP is on.

The power socket is a hole with three prongs in it. Plug the female end of the power cord into this socket. The socket is designed so that there is only one way the plug will fit in, so don't worry about doing it wrong. Plug the other end of the power cord into the outlet you have selected.

Your TRICEP can be plugged into any grounded outlet. Grounded outlets for 115 volts look something like this:



They will receive the three-pronged plug that comes with the TRICEP. (If you have a 230 volt system, you will also have a three pronged socket, but in a different shape.) If grounded receptacles are not available, use an adapter to plug the power cord into the wall, and plug accessories, such as your terminal and printer, into the outlets on the back of the TRICEP. When using a three pronged adapter, be sure to fasten the metal tab or green wire from the adapter to the screw that secures the socket cover to the wall. And **don't just break the ground lead off of the power plug.** This is a computer, please, not a sabre saw.

To make matters worse, not all three pronged sockets are actually grounded. In older houses and buildings, all the wiring was done with two wires - no ground wire! Later, three pronged sockets were added by someone and "grounded" to the metal box they were installed in. If you want to check your ground (the third hole, which is round, or the screw that you attach the adapter to), use a volt-ohmmeter and check the resistance between your alleged ground and a cold water pipe. The resistance should hover near zero.

If you need more outlets, get a power strip and plug everything into the power strip. Computer systems and their accessories are like stereos in that they work best if all components are using the same ground. If you are going to have terminals that are located a distance away from the TRICEP, use the full RS-232 specification cables which include a ground on pin 7. (Morrow cables supply this).

Besides a shared ground, the TRICEP and its peripherals, (terminals, printers, etc.) require a steady power supply. The minimum suggested voltage for the TRICEP is 105 volts, and the maximum is 125 volts. For 230 volt systems, the tested voltage range is 208 to 265 volts. In the United States and most of Europe, the power companies are pretty reliable at delivering electricity in this range, so most of you won't have trouble with this.

One problem that you need to try to avoid is voltage variations due to the startup of large electrical appliances like air conditioners, space heaters, and freezers. If an electric lamp connected to the outlet you want to use dims noticeably when one of these appliances turns on, you may be plagued with erratic performance from your computer. The only solutions are running a dedicated power line from your main service panel to either the computer or the appliance, or getting an uninterruptable power supply.

2.5 STEP 4: TURN THE POWER ON, FIRST TIME

Now, ready for the big step? If you have an external hard disk or other peripherals (printers, terminals, etc.), turn them on first. Apply power by turning the keyswitch one click clockwise. The power is now ON to the TRICEP.

Checkpoint 1:

A pilot lamp inside the panel next to the keyswitch should now be glowing. The fan inside of the the back of cabinet should also be on. You have a 5 and 1/4 inch hard disk, known as a mini-wini or mini-winchester drive below the floppy drive. The red indicator light on the drive will also be on. (This light is visible through the ventilation holes below the floppy drive). You can hear noises from the hard disk as it comes up to speed. The two grounded power outlets on the rear panel of the TRICEP now have power applied as well. If everything checks out here, go on to the next section, setting up the first terminal.

Troubleshooting 1:

If neither the pilot light nor the fan came on, check:

1. The connection of the power cord to the back of the TRICEP. This socket is about 3/4 inch (2 cm.) deep, and the female end of the power cord must be fully inserted for a good connection.
2. The grounded outlet. Plug a desk lamp into the outlet to make sure that you have power.
3. The circuit breaker on the back of the TRICEP. It is of the pushbutton type seen on stereo receivers and speakers. Push it all the way in and release with a snap.
4. The power switch itself. Flick the keyswitch a couple of times.

If the fan comes on, but the pilot light doesn't come on, the TRICEP may well be operational with the exception of the lamp itself. If the fan is not working, do NOT operate the TRICEP as it will overheat.

2.6 STEP 5: SETTING UP THE FIRST TERMINAL

This section covers attaching ONLY the first terminal. If you're going to have more, these procedures apply in principle, but there is further information under "Adding Terminals" in the **CONFIGURING YOUR TRICEP** division of this manual. For now, let's concentrate on getting just the first one working.

The first terminal connected to a TRICEP system is called the console. On mainframe (large) computer systems, the console is used by the system operator. TRICEP's console is used both by the system administrator and ordinary users. When you first bring up your TRICEP system, the console is also the ONLY terminal that you can use. In fact, whenever the system is in Single User mode, only the console will be operative. (Making other terminals usable is known as "going multiuser".)

2.6.1 Terminal Settings

Terminals have many switch-selectable features, of which only a few concern us at present. The first, and most important, is the baud rate, that is, the transmission rate in bits per second. The TRICEP system expects the console terminal to be set to 19200 baud. If the terminal that you plan to use will not work at 19200 baud, you have a problem. The TRICEP requires a terminal set to 19200 baud during booting.

9600 is the most common baud rate used for terminals connected directly to computers. But, you need 19200, twice as fast. Look up the baud rate selection, or option switch settings, in the operation manual for the terminal you are using, and set the terminal for 19200 baud. While you are looking at this section of your terminal's manual, also set the switches for:



- o full duplex - since the TRICEP echos the characters you type back to your screen, if you see two of everything later, you should double check this setting;
- o 8 bits - the number of bits used to send a character;
- o 2 stop bits - the pattern that marks the end of a character;
- o no parity - parity is an option that is used for checking transmission accuracy (usually not a problem at 19200 baud); and
- o on line - as opposed to local, means that the keys you type are transmitted to the computer; sometimes this switch will be on the keyboard.

Once again, the most likely cause of trouble here is the baud rate. The usual sign of using the wrong baud rate is that the sign-on messages (after booting or resetting) will be a nonsensical collection of assorted characters and punctuation. The TRICEP has software-controlled baud rates, so you must set your terminal to 19200 baud initially in order to use the software to change the baud rate after booting, if you so desire.

2.7 STEP 6: CONNECTING THE CONSOLE CABLE

Now, look at the back of your terminal. There will be one or more female RS-232 connectors there, (like the capital letter "D" lying on its rounded side), the same size and shape of the end of the RS-232 cable we told you you'd be needing. If there is only one connector on the back of the terminal, great. If there are two, use the one labeled "modem", "RS-232", "main port", or "P1", not the one labeled "printer" or "aux". The second connector is not designed to be connected to a computer, and you don't need it. Connect one end of the RS-232 cable (either end is fine) to the correct connector on the back of the terminal.

The other end of the RS-232 cable is attached to the connector that is above the row of three similar connectors on the back of the TRICEP. This connector may be labeled "syscon" or "console".



Now, plug in the terminal and turn it on. You can use one of the two sockets on the back of the TRICEP to plug in your terminal, although we suggest that you use these sockets for extra hard disks or other peripherals. Use a grounded socket, preferably the twin to the one the TRICEP is plugged into.

When the terminal is turned on, it may "beep", and you should be able to hear it warming up (sounds like a TV set, a high-pitched whine.) After warming up, (5 seconds or so), you should see a marker on the screen of your console called the cursor. The cursor marks the place where the next character will appear on the screen. (The cursor precurses the next character.) It will probably be in the upper left-hand corner of your screen.

If you don't see the cursor on the screen, you may need to adjust your terminal's brightness control. Before you leave this paragraph, find your cursor. If you can't find the cursor, something is wrong with your terminal, so you might as well just mark your place in this manual, and return when your terminal is okay. Thanks.

2.8 STEP 7: BOOTING, THE FIRST TIME

Everything ready? Grasp the keyswitch and turn it clockwise to "RESET" and release it. The following message will appear on your screen:

```
Morrow TRICEP System (prom 2.1)
Testing memory.
```

If you don't see this, but you do see something, you have terminal problems. Gibberish on the screen indicates bad baud rate. Set your baud rate to 19200 and try RESET again.

If you don't see anything on your screen after resetting, the terminal may still be the problem. Check the connections at the backs of the terminal and the TRICEP computer, as described in the previous section.

Testing memory takes between 10 and 30 seconds to complete. Then, several more messages will be sent to the console:

```
Memory OK
512000 memory
5 inch floppy not ready.
8 inch floppy not ready.
Press RETURN to continue
Standalone boot
```

```
: []
```

The memory test and the attempted floppy disk accesses are part of a program stored in the PROM on the processor board. The amount of memory found is given as the number of decimal bytes.

The floppy disk access attempts allow booting the system without using the hard disk. If a floppy disk has been inserted in the first 5 1/4 or 8 inch floppy drive and the door closed, the system will attempt to boot from the floppy. This is used for repairing file systems on the hard disk and will be explained later.

After the PROM program has checked the floppy drives, it loads the "Standalone boot" program from the hard disk and prints the message on the screen. The CPU, memory, hard disk and disk controller have checked out okay if you reach this point successfully. If you have other messages on your screen, please jump ahead to troubleshooting.

The Standalone boot program is now waiting for you to press the RETURN key. Please press it now. A new set of messages will be added to the console:

```
mw(0,0)unix
Loading at 0x400: 90268+8929+74042
Type RETURN to start at 0x400
```

This messages reports the successful loading of the UNIX kernel and the initializing of its data areas. Once again, the Standalone boot program pauses and waits for you to press a RETURN before continuing. Please press RETURN now.

The final set of messages report the copyright notice, the version creation date and amount of memory remaining after loading the kernel. The system then checks the hard disks and finally places the superuser prompt on the screen:

```
#
```

If you have this symbol (#) at the bottom of your screen, you have successfully installed your TRICEP system. The UNIX operating system is running in single user mode. You should skip over the next section, and move onto FIRST TIME USE.

2.8.1 Troubleshooting Booting Problems

This section explains what problems might occur while booting, what the symptoms of these problems are and how to cure them. This information is not only useful for solving problems the first time you boot your TRICEP, but anytime you have a problem booting.

As you begin to follow these instructions, we have assumed that you have the power on to both your computer and your terminal, you can see your cursor and you have connected your terminal to the computer. This was explained earlier in Steps 5 and 6.

We will also assume that you have RESET the computer by turning the key in the keyswitch clockwise, and releasing it. The TRICEP system also performs an automatic RESET on power on.

Nothing on screen after RESET

Most of the information that you need for diagnosing problems with your system comes from your console. This poses a problem if when you turn your system on nothing appears on the console. Is the problem in the terminal or the computer?

The solution to this is to listen to your computer. Your TRICEP system will make certain noises if it is able to follow the PROM program. If you can hear these noises, the computer is working and the terminal is not.

Immediately after resetting, the PROM program runs a memory check. This is silent, and takes from 10 to 30 seconds, depending on how much memory you have. (Don't forget to wait for this!) Next, the program tries to read a diskette in the 5 1/4" drive. This involves loading the heads, and that makes a clicking sound in the floppy drive mechanism.

Reading from the floppy disk will fail (unless you have inserted a formatted diskette and closed the drive door). The PROM program then tries to read the Standalone boot program from the innermost tracks of the hard disk. You should be able to hear the stepper whirr as it moves the heads into this position. This is not easy to hear, so you may need to hold your head near the hard disk, which is located just below the floppy drive.

Did you hear the disk drive noises? If you did, you have some problem with your terminal. Go back to Steps 5 and 6 and check out your terminal's settings and the connection between the terminal and the computer. It's easy to use the wrong connector, on either the TRICEP or the terminal.

The problem might also be in the RS-232 cable itself. The TRICEP system looks for a hardware signal that comes back from the terminal via pin 20. This signal is known as Data Terminal Ready. If your cable doesn't provide for a connection between pin 20 of the terminal and pin 20 of the TRICEP console connector, nothing will be transmitted.

If you don't hear any noises from the disk drives, you may have a hardware problem in the computer. Re-read the section on "Checking for Hidden Damage". The problem may be loose pc boards or cables. If this doesn't help, you should contact your vender for assistance.

Garbled Characters Appear on Screen

This is an easy one. You have set your terminal to the wrong baud rate. You MUST set your console terminal to 19200 baud to RESET the TRICEP. This cannot be altered, as it is part of the PROM program. See Step 5 for more details.

Insufficient Memory Message Appears


Your TRICEP systems requires 512 kilobytes of memory to boot. The number of bytes of read/write memory available is counted during memory testing. If no memory at all is found, or memory is incorrectly addressed, the message:

RAM is missing or incorrectly addressed.
Please refer to the installation manual for further instructions.

A second message is displayed if less than 512 kilobytes of RAM are found:

Warning: Insufficient system RAM available.
UNIPLUS System V requires a minimum of 1/2 megabyte of RAM.

The proper response to both of these messages is to follow the instructions given in the section called "Checking for Hidden Damage". If your memory boards are well seated, then it is possible that the settings that establish the location of the boards in memory is incorrect. On Morrow 256K DRAM memory boards the jumper settings for the first four boards look like:



These slide on jumpers are located in the lower left corner of the memory boards. To examine these jumpers, you must first remove the retainer rails and slide out the memory boards. Memory boards are customarily located in the frontmost slots of the Wunderbuss motherboard. Memory boards can be distinguished by their neat rows of chips and lack of cable connectors. (See section on "Checking for Hidden Damage" for details on removing and replacing boards.)

To have correct memory addresses requires that the addresses start at 0, and continue contiguously. If you still fail the memory test after checking the memory boards, please contact your vender for assistance.

Can't Boot Message

This message may appear after memory has been checked, and the floppy disk drives accessed. It may be preceded by the word "Time-out." What has happened is that the PROM program has been unable to access the hard disk. Occasionally, this is caused by insufficient power to the hard disk during power on, and can be cured by turning off the system. Wait 15 seconds or so, and try powering up again. This may be a reoccurring problem during power on, but it doesn't seem to affect system performance after booting.

The other possibility is that a board or cable is loose. Once again, please follow instructions in the section on "Checking for Hidden Damage".

System Doesn't Respond to Pressing RETURN

This is a manifestation of a simple problem. The correct messages have appeared on the terminal's screen, but pressing RETURN does not continue the Standalone boot. The problem is that data can get from the computer to the terminal, but not from the terminal to the computer. This is caused by one of three things. First, you may have your terminal in "Local" mode, allowing it to receive but not send. Pressing RETURN causes the cursor to move to the left margin, but not to the next line. Check your terminal's settings.

The second cause is that your terminal may have received the lock keyboard character. This has the unpleasant side effect of making a system appear to be crashed if it occurs after a successful boot. The solution to this is to turn the terminal off, wait a few seconds, then turn it back on again.

The other cause for this is an incorrectly built RS-232 cable. The pins that carry transmitted data from the terminal to received data at the computer are not connected. Please refer to Step 5 (page 11).

Loading Fails

This is much harder to diagnose because once you see the message

```
mw(0,0)unix
Loading at 0x400: 86268+...
```

most of the system has been used successfully. One possibility is that the hard disk is still not fully up to speed. This is the same cause of failure mentioned earlier in the "Can't boot" case. Simply turn the system off, wait 1 minute for disk to stop spinning, and try again. If this doesn't work after two or three attempts, contact your vendor for assistance.

System Doesn't Start after Pressing RETURN

This is the final checkpoint. The hard disk has succeeded in loading memory with the UNIX kernel. This has involved the cooperation of the CPU board, the PROM, the console terminal and memory. If the system fails at this point, the file system on the hard disk may have been corrupted or damaged, possibly through shipping or incorrect shutdown. Please follow the instructions in the section on "Repairing the Hard Disk from the Standalone Floppy".

System Stops After Checking Hard Disks

Checking the hard disks is the last task performed by the kernel before interrupts are turned on. You can tell when the kernel has finished checking the hard disks when it has displayed a message like

```
mw0:  csr  0xFF0054  1 drive
```

Once this message has been displayed, interrupts are enabled. Systems that crash at this point are having a problem with an interrupt. For example, a board that has the interrupt line stuck so that it is constantly generating interrupts. Boards that are missing interrupt jumpers, or have incorrect interrupt jumpers can also cause this problem.

3. FIRST TIME USE

This section of the installation manual is broken into two parts: testing and backup of the file system and a quick tour of the file system. Details on Configuration follow in the next chapter. The testing and backup will take several hours, but is necessary to the reliability of your system. During this time, you may, if you like, read ahead and familiarize your self with your UNIX system.

The quick tour of the file system provides a map of the file system and the commands used to move around in it. You may want to skip this section if you have worked extensively with UNIX systems in the past. However, if you are responsible for configuring this system, understanding the arrangement and organization of directories will be of immense assistance to you.

The UNIX system is largely configured through human readable files. That is, you can use an editor to change these files, and thereby reconfigure your UNIX system. But, you need to know the rules for changing these files, and what they control. This is what is explained in the Configuration section.

Much of the material presented here is covered by the UniSoft manuals, and by books on UNIX. The difference between these other sources and this manual is that here you are given advice specific to your TRICEP, rather than for some generic UNIX box. We can't and won't try to duplicate Bell Lab's documentation, but we will tell you the best way to handle TRICEP specific operations, like backing up and configuration.

3.1 CHECK YOUR FILE SYSTEM!

Here we are, me writing and you reading, while your TRICEP system is just sitting there, hot to trot. Well, a little patience, please. Just because you can see that superuser prompt, #, on your console doesn't mean everything is hunkydory. What we are concerned with now is that there may be hidden damage (oh no, no more HIDDEN damage) on your hard disk. Here's the scoop: we at Morrow have shipped you a complete 12 megabyte or so UNIX System V port by UniSoft on your hard disk. But, is it all there?

The hard disk that you have was selected partially because of its ability to survive the hazards of shipping. Now, it's time to check this. We will proceed to do this in three stages:

badblks - a surface check for damaged sectors,

fsck - a consistency check of the file system, and

tar - for the creation of your master backup copy.

As previously mentioned, this will take at least four hours. Please don't ignore this section. You'll be sorry if you do. So, let's get things started without further ado.

3.1.1 Checking for Bad Blocks

Please type the following command at the console:

```
# badblk /dev/rmw0hw0
```

Now that that's started, we can relax, and read ahead. Badblk starts out by reporting the number of badblocks previously discovered on your disk. Then, it attempts to read every sector on your hard disk. This takes about 28 minutes for a 34 megabyte hard disk. Badblk reports any bad (unreadable) sectors (blocks) it finds, assigns an alternate sector and updates the badblock map in block 0.

While badblk is running you should copy down the numbers that appear on the screen. These are damaged sectors that have appeared since your system was tested. If you have several damaged sectors that aren't being used, they will not affect your file system. Badblocks with block numbers greater than 12000 (24000 for 512 byte blocks) are probably in the empty (unused) portion of the file system.

Now, what if badblk discovers bad blocks with numbers less than 12000 (24000)? If you have many (more than 20) of these bad blocks, you have a serious problem with your hard disk and should report to your vendor the number of unreadable blocks. If there are less than 20 bad blocks, you can probably continue to use your system. The files that have the bad blocks can be discovered through normal system use. For example, if you typed the command `cal` and the error message

```
mw0a error bn=1234 sr=0x4
```

appeared on the system console screen, you would know that the file `/bin/cal` had a bad block, number 1234, in it. Then, you can add this block number to the badblock list with the command

```
mwbad
```

and replace the file `/bin/cal` with a backup from your vendor.

If you don't like this hit or miss type of approach, you can try reading the entire file system, file by file, to discover the files with bad blocks with the command

```
/etc/vchk -kc > vchk.results
```

We suggest that you run `fsck` (see next section) first. Then run `vchk -kc`. This will take about 34 minutes, and will attempt to read every file in the file system, and produce an error report for every file with an unreadable block. Carefully note the block numbers and files where errors occur. Add the bad blocks to the bad blocks list and replace the files from backups.

`Vchk`, or version checker, checks a file system and compares it to a special list, called a V-check tree. We supplied your TRICEP with a V-check tree (`/etc/vchk_tree`) that tests to see that all necessary files, directories, ownership and permissions are intact. The `-k` option causes `vchk` to do a checksum on files in the `vchk_tree`. One option of `vchk` allows it to copy (take) missing files from a host machine. We will talk more about the `vchk` utility later.

3.2 File System Check

The `badblk` command attempts to read every block in the file system. The ability to read a block says nothing about the validity of the information found there. The file system check program, `fsck`, understands the structure of file systems. `fsck` not only understands file systems, it can repair them. This makes `fsck` an indispensable tool in preventing file system problems. You should use `fsck` every day. To start `fsck` the first time, you can enter its name,

```
fsck
```

by itself. `fsck` will look in the file `/etc/checklist` for the names of the file systems to check.

During file system checking, the `fsck` program will produce reports about what it is doing on the system console. If problems are discovered, you will be prompted before any action is taken. Unless you are already experienced with UNIX file systems, you, the operator, should always answer "Y" for yes.

If the message "BOOT UNIX (NO SYNC!)" appears, reset your TRICEP. This message means that the superblock has been modified.

`fsck` takes about 4 minutes to complete on a 34 megabyte drive. If `fsck` cannot check your file system, (for example, the message "Can not stat root" appears), you have a serious problem and should contact your vendor.

Further information about running fsck is given in the section intitled "Every Day Procedures", and also in the System Administrator's Guide by UniSoft.

3.3 First Backup

Up until now, the check programs have been pretty simple. You simply started the program, and perhaps provided a yes response or two. Well, you need to provide a bit more assistance for the first backup operation.

The standard (and least expensive) form of backup medium is floppy diskettes. Unfortunately, this is also the slowest and most (human) error prone form of backup. The section that follows describes first time backup procedures for this medium.

3.3.1 Floppy Disk Backup

This procedure requires about four hours of time and 40 formatted diskettes. You, the operator, need to be standing by with your supply of formatted diskettes. Other than removing and inserting diskettes, your task will be essentially mindless, allowing you to continue reading through this manual (or whatever).

Right off, you need to have 40 formatted floppy diskettes. The TRICEP system uses

5 1/4" double sided, soft-sectored, diskettes

Please use only diskettes that are certified double-sided.

To format diskettes, insert a diskette in the floppy disk drive, close the latch and type

```
diskformat /dev/rdj0
```

for each diskette. This takes about one minute per diskette. Label these diskettes, after they are formatted, with the date and the words "MASTER TRICEP BACKUP". Number the diskettes sequentially from 1 to 40.

Now you are ready to begin the backup process. Insert the first formatted diskette, number 1, in the floppy drive, close the latch, and type

```
tar cvfB /dev/dj0a 798 /etc
```

This begins to copy the file system to floppy diskettes in tar format. After each file is copied, the filename and number of blocks (its size) will be reported. The tar command can only copy entire files to a diskette, so when a file is too large to fit in the space remaining on a diskette, it will request that you insert another.

Change disks and press return.

Put in the next diskette and press return. The amount of storage required on floppies will be greater than the number of blocks of files on the hard disk because some space will be wasted at the end of each floppy, and by tar's extravagant padding.

You should record the names of the first and last files that are stored on each floppy diskette. You can also include the names of directories on each diskette. Directory names consists of the list of letters between the first and last slash (/). For example, in the name /usr/src/sys/config.c, the directory name is /usr/src/sys, dropping the part after the last slash. This will help you find the correct diskette if you need to replace a single file. Otherwise, you will have to go through all the diskettes sequentially until you find the correct file.

After you have backed up the /etc/ directory, go on to /bin, then /lib, /usr, as in

```
tar cvfB /dev/dj0a 798 /bin
tar cvfB /dev/dj0a 798 /lib
tar cvfB /dev/dj0a 798 /usr
```

Remember to have your supply of formatted, numbered and labeled diskettes handy BEFORE you begin. If you are one diskette short, or try to use an unformatted diskette, you will have to start the tar command over from the beginning. (Ahhrrrgg!!!) Go back to the diskette that begin with that tar command and start over (with more diskettes in hand).

It takes tar about seven minutes to fill a diskette. You can use this time to copy down the first and last filenames, and the directory names, that are being copied to each numbered diskette. If you later lose a file, say, /usr/lib/lpadmin, you can insert the correct diskette, and type

```
tar xvf /dev/dj0a /usr/lib/lpadmin
```

Otherwise, you can put the diskettes in one at a time, until you have searched through all 40 to find the right one.

NOTE: The floppy device for using tar has a special designation: /dev/dj0a as opposed to /dev/dj0. Tar is not very clever, and overwrites the boot block unless the special device dj0a with a 1k offset is used. This leaves 798 512 byte blocks per diskette.

3.4 Quick Tour of the File System

Still working on that first backup? This section gives you a quick look around your new home. The UNIX file system has a well deserved reputation as an elegant solution for organizing many files on huge disks. This system is not hard to follow once you get a grasp of it, and can certainly work to your advantage if you bother to understand it. I will try to explain in the perspective that I met it (the file system): coming out of a floppy disk and CP/M environment.

If you have worked as a system administrator or programmer on a UNIX system before, this section may seem pretty old hat to you. System V's file system is laid out essentially the same as the file systems in Versions 6 and 7, and in System III. The obvious changes to UNIX are in `init`, `getty` and `login`, and many new commands.

3.4.1 Trees: the Basic Philosophy

This section provides an overview to the discussion that follows in the next section. The UNIX file system is usually compared to an upside down tree, with the root at the top and branches heading downward. The root directory, represent as / (slash) is the beginning of the file system. Nine directories branch off from the root, providing the basic subdivisions of the rest of the file system. These are:

- bin the commands
- dev the device connections
- etc the system administration commands and files
- lib the libraries (for C and other languages)
- lost+ a place used by `fsck` to hold files and directories that found have lost their names and their parents
- t a hook for mounting file systems
- tmp the place for temporary or scratch files
- users the place where user directories will be located
- usr the beginning of a large subtree, with more commands, more libraries, spooling directories, administration and accounting files

There are several files in the root directory: `unix`, (the kernel program loaded by the Booter code), `block0`, (an image of block 0 of the root file system) and the superuser's shell setup files (like `.profile`).

3.4.2 Why Trees?

So much for the overview. Now for the explanations.

This is a story of organization. The problem is simple: the capacity for tens of thousands of files. If you are used to working with floppy disks, organization is simple because only a handful of files will fit on each floppy. The limited capacity of floppy disks provides a natural division and keeps directories small. You can keep all the files you are currently working with on one or two diskettes. When a diskette gets full, you copy the files you need to a new diskette, and the old diskette becomes the backup.

The problem is that it takes a long time to fill a hard disk. Suddenly faced with as much as 100 times the capacity of a single floppy, you can go a long time before running out of room. Directories can grow impossibly large, so that it can take five minutes just scan through them. Then, suddenly, you have a serious problem. The disk is full and you don't know which files to delete.

The UNIX file system helps you in three ways. First, you create directories and use them as if they were file folders. Groups of related files fit in each directory. If a directory has too many files, you create a subdirectory and move files into it.

Second, you use easy to understand names for directories and files. The UNIX systems allows names of up to 14 characters, so you can have real names for filenames. This makes it easier to find files.

Third, the UNIX system makes entries in the inode (file descriptor) every time it accesses (reads) or modifies (writes) to a file. This can be used in picking out the straw from the gold when you are cleaning up file systems.

I suggest that you use short names for directories. You will have fewer directories than files, so their names will be easier to remember. Directory names become part of filenames when used in commands, and having several long directory names can make a filename very long very quickly. Having to retype long filenames like `"/usr/rik/textfiles/systemV/tour"` gets old very quickly. I use `"/usr/rik/Txt/SV/tour"` and save myself a dozen keystrokes. Using capital letters for directory names makes them stand out in directory listings.

3.4.3 Moving Through the File System (Made Easy)

Okay, enough of the reason for being. Now, let's learn how to move around and see what's here.

There are three basic commands for moving around the file system:

`cd` change directories moves you from one directory to another

`pwd` tells you the name of the directory (how you got there)

`ls` shows you what's there (a directory listing)

For this tour of the file system, I will be using the allegory of "the office of the future". Instead of describing the file system as it really is, I will pretend that is an office space that corresponds exactly to the file system. This will help those of you who are visually oriented get a feel for the file system.

The tour starts in the reception area, also known as the root directory. As you walk into the room, the receptionist asks you for your name (your login) and your reason for being there (your password). Then, she busies herself with setting up your visit, so you have a chance to look around.

The room you are in is almost devoid of detail (minimalist interior decorating, lots of chrome, subdued colors). The feature that catches your eye immediately is that there are nine doors leading from the room, each of them labeled: `bin`, `dev`, etc, etc. In the corner, there is an ordinary looking file cabinet labeled "unix". At that moment, the door labeled "usr" opens, and I step out, tall, serious and somewhat goodlooking (I'm giving this tour, remember).

We greet each other, and I launch into my introductory spiel:

"Welcome to the UNIX file system. As you may have already noticed, we are here in the root directory, the beginning of the file system. Later on, when you have your own HOME directory, that will be the beginning of your file system."

"There is a simple way to find out where you are in the file system at anytime by typing "pwd". Let's try it now"

Leaning over the secretary's desk, I pick out the keys `p`, `w`, `d` on her terminal.

`pwd`

/

"The slash is the symbol representing the root directory. As we move through the doors that represent other directories, the directory names will be added after the slash of root. pwd stands for "path way to directory". Obviously, the path way to a directory will get much longer the further we move away from root."

Wanting to impress you immediately, I start moving you toward the door labeled "bin". To move through the "bin" door, it is necessary to type:

```
cd bin
```

The "bin" door opens and we step into a very large, dark room. To get a look around, you need to type:

```
ls7
```

turning on the lights and displaying the names of some 230 files. The files appear to be tools, each different from the next, all with some controls on top, and places where something goes in and the result comes out.

"This is the bin room, the workplace of most of the famous UNIX tools. There are tools for creating files, editing files, altering files and making new tools. Each tool is designed so that it can be hooked together with other tools. This way, new tools can be improvised simply by temporarily connecting them together."

"You won't actually be visiting here very often. When you use a command, this room, ..er, directory, will be automatically searched for the tool. To help you remember the name, bin stands for binary, and binary is short for "binary executable files", that is, programs. Look, there's pwd. Let's try it, even though we know where we are."

```
pwd  
/bin
```

"You can see that the door we went though to get to this directory has been added to the slash of root. Okay, let's go back to the root and check out another directory."

```
cd /  
cd dev  
ls7
```

Leaving bin, we return to the root, then we walk over to the door labeled "dev" and walk in. This room has a very different character to it. There are about fourty files here, arranged around the walls of the room. Some files have a small, character sized, connectors, resembling phone jacks. Others have a large, squarish, block sized openings in the front. The files with the small connectors are attached to terminals, printers and even a

clock. The files with the big block-sized openings are hooked into quietly humming hard disks and idle floppy disks.

One of the files with block-sized openings, labeled mw0a, is connected with a transparent tube that disappears into the ceiling of the room. Even as we watch, blocks flash through the transparent tube, coming and going to some place above the ceiling.

"This is the "dev" room. dev is short for device. This is the directory where the connections between the file system and devices are made. You may have noticed that each file here represents only one physical device. You, or the system administrator, can control which connections are made to devices. Other connections, like the system console and the root file system, mw0a, are designed into the kernel and can't be changed."

"You could, for example, connect your terminal to someone else's terminal and type right onto their screen. There are, however, much more polite ways of talking to people than taking over their personal terminals. Now, let's look at the etc directory."

```
cd etc
No such file or directory
```

"Whoops, made a mistake. Since we're in the dev directory, and the door into etc is located in root, our change directory command failed. We can take a shortcut, however. We can specify to the cd command that we want to get to etc AND we need to start at the root by typing:"

```
cd /etc
ls7
```

This room is filled with a bunch of office workers and a couple of long-haired programmers. There are several filing cabinets around the walls, with names like passwd, group, motd and gettydefs. There is also an organizational chart on the wall labeled with the funny name "inittab".

"This is the etc directory. The easiest way to remember the name is, of course, to think of et cetera. And so on."

"This directory houses the files and programs that configure your UNIX system. For example, the passwd file will contain the login names and passwords of all the users you will have on your system. That chart on the wall, inittab, determines what programs will be run initially. The gettydefs file contains login speeds for terminals, and motd has any message you want every one to read after logging in."

"Those long hairs over there represent your resident system programmers, fsck and fsdb, that check your file systems and repair damage to them. You'll quickly become familiar with fsck,

because you'll use him every day. fsdb is a bit harder to communicate with, but can also be helpful."

```
cd /  
ls7
```

"We're back in the root directory. I wanted to point out to you that several of the doors currently have empty rooms behind them. The t directory is used for attaching temporary file systems on other disks. The tmp directory is for temporary files. If you make a habit of creating your temporary files in tmp, you won't have any trouble remembering which files you can safely erase because they are short term."

"The lost+found directory is used by fsck. If fsck finds any files or directories that aren't listed in any directory, fsck manufactures a name for them and leaves them in lost+found. Makes sense, doesn't it?"

"There's lots of empty space behind the users door. The /users directory is where you will be creating HOME directories for your system users. Keeping your users together makes backing up simpler."

"Now, let's check out the last door, usr."

```
cd usr  
ls7
```

Behind the door labeled usr is a long corridor with more doors. Two of the doors have familiar names on them, bin and lib. Other doors are labeled spool, admin and man.

"The usr directory is the entry point into the rest of the file system."

"The bin and lib directories have more tools and libraries within them, respectively. The man directory has all the manual entries stored in nroff form. Before we finish the tour, let's move into the guest room so that I can make a point above pathways to directories."

```
cd /users  
cd guest  
pwd  
/users/guest
```

"Okay, our pathway to directory is now /users/guest. We started in root, /, moved into users, then into guest. Additional slashes are used to separate directory names. You don't have to use entire path names to specify a file, because you can form a file name relative to your current directory, that is, the room that you are currently visiting. Now, back to the root."

cd

Using the `cd` command without specifying a pathway to directory means changed directory to the user's home. In this case, the user is the root, so we flash back to the root directory. The secretary logs you out, and you step back into the duller reality of this manual.

3.5 Making Room: "Unnecessary" Files

Owners of TRICEP systems with 16 megabyte hard disks are immediately faced with a shortage of disk space. The UNIX system and swap space use up almost the entire hard disk. The solution is to remove some files from the file system. Since you have completed your backup, these files are still available if you find you need them later.

There are two major candidates for removal: the games and the manuals. The `/usr/games` directory is the sixth largest in the file system, and is not absolutely necessary. You may wish to compromise by leaving some games and their supporting files.

The `/usr/man` directory contains the source of the documentation produced by UniSoft. Yes, you have, on line, the equivalent of many inches of paper, stacked flat. The advantage is that you can use the `man` command to review manual entries. The disadvantage is twofold: the `man` command is painfully slow and doesn't allow you to see the beginning of an entry after it has passed. It would be nice if there were a better way of viewing manual pages, but there is not. A compromise here is to remove everything but `/usr/man/u_man/man1` and `/usr/man/a_man/man1`. These two subdirectories contain the most often used manual entries.

There are several other candidates for removal. The `nroff` and `troff` macros in `/usr/lib/macros` aren't necessary unless you use `nroff` and `troff`. The `/bin/efl` program, Enhanced Fortran Language, is useless unless you have a Fortran compiler. There are also spelling dictionaries (`/usr/lib/spell` and `/usr/dict`) on line that you may wish to dispense with.

If you are interested in uncovering the largest files and directories in your file system, you can use the following commands to produce a list sorted with the largest file first:

```
du -a / | sort -rn > /tmp/filesizes
```

Please be certain that you have backed up before you remove any file. Sometimes removing a file will have unexpected side-effects, and you will need to restore the file from backups.

```
rm -r /usr/games
```

will remove the entire `/usr/games` directory.

4. CONFIGURING YOUR TRICEP

Your TRICEP system runs the UniPlus+ operating system. This is essentially Bell Lab's System V UNIX with some enhancements. The documentation provided by UniSoft was written by Bell Labs and edited by UniSoft to be somewhat more applicable to Motorola 68000 systems. UniSoft also added entries for the Berkeley UNIX programs that are included, such as the C-shell (csh), and exchanged the name UniPlus+ with the name UNIX. Unfortunately, the Bell Labs documentation is a bit thick for anyone that doesn't ALREADY know what they are doing with UNIX.

We will do our best to fill the gap between the generalized information provided by the UniSoft manuals, and the procedures that will work best on your TRICEP.

The UNIX system has a simply wonderful feature. Most of its activities can be controlled by editing HUMAN READABLE FILES. This accounts for some of the enthusiasm that people feel for UNIX. This chapter explains how the files that control your TRICEP fit in, and how to edit these files.

There is, of course, a small problem with the human readable files. Your TRICEP must be able to read them also. This entails rigid structuring for these files. You must compromise with the machine, which is neither as clever or as creative as you are. You have access to all the flexibility offered, but only if you obey the rules. We will carefully explain the rules to you, so that you and your TRICEP will get along splendidly.

The human readable files are introduced in the same order as they are used by the TRICEP. Most of these files are involved in the process of going multi-user. We will start with an overview from boot to the appearance of the user's prompt (\$ or %). Then, we will talk about each of the files used in this process, all in the /etc directory:

- inittab - Controls which terminals are allowed to login, starts background processes, like the print scheduler, and the daemons, through /etc/rc;
- gettydefs - Controls initial terminal baud rate and set up;
- passwd - Establishes login names, holds passwords, user and group id's, home directory and shell program name;
- profile - Works with the shell (not csh) to initialize every shell user's environment;
- cshrc - Works like profile for csh.

We will also fit in connecting terminals, modems, additional disks and printers. The lp commands will be explained, so that you will be able to use the print spooling facilities of UNIX.

4.1 From Boot to Multi-user

The diagram named Booting UNIX outlines the process your TRICEP goes through to become a multi-user system. This diagram has been simplified somewhat by leaving out the scheduler (created by the kernel) and the single user state (the superuser shell started by init).

The first several programs, the PROM, Booter code and kernel, are not alterable simply by editing files. These form an invariant system initialization pattern that can be modified only by editing source files and remaking the kernel, PROM or Booter.

The PROM code has three tasks: check memory, check for floppy disks and to load the Booter code. The memory tester determines whether you have enough working memory to boot. The floppy disk checker looks at the first 5 1/4" and 8" drives. The PROM code will attempt to boot from any floppy disk that is inserted in a drive with the door closed. If there is enough memory (512 K) and no floppy disk is ready, the PROM code loads the Booter code from the hard disk. The Booter code is located on 18 blocks in the alternate block area of the hard disk.

The Booter code understands how the file system works and how to load the kernel or a standalone program. When it is ready to go, it prints a colon (:) prompt and waits for either a RETURN or a file name to boot. Typing RETURN enters the file name "unix" from the root device (mw(0,0)) as the file to boot. You can enter the name of a kernel program that you are testing here.

The Booter code locates the file name in the appropriate directory, looks up the inode (file descriptor) and loads the file into memory. The Booter then waits for another RETURN before starting to execute the kernel or standalone program.

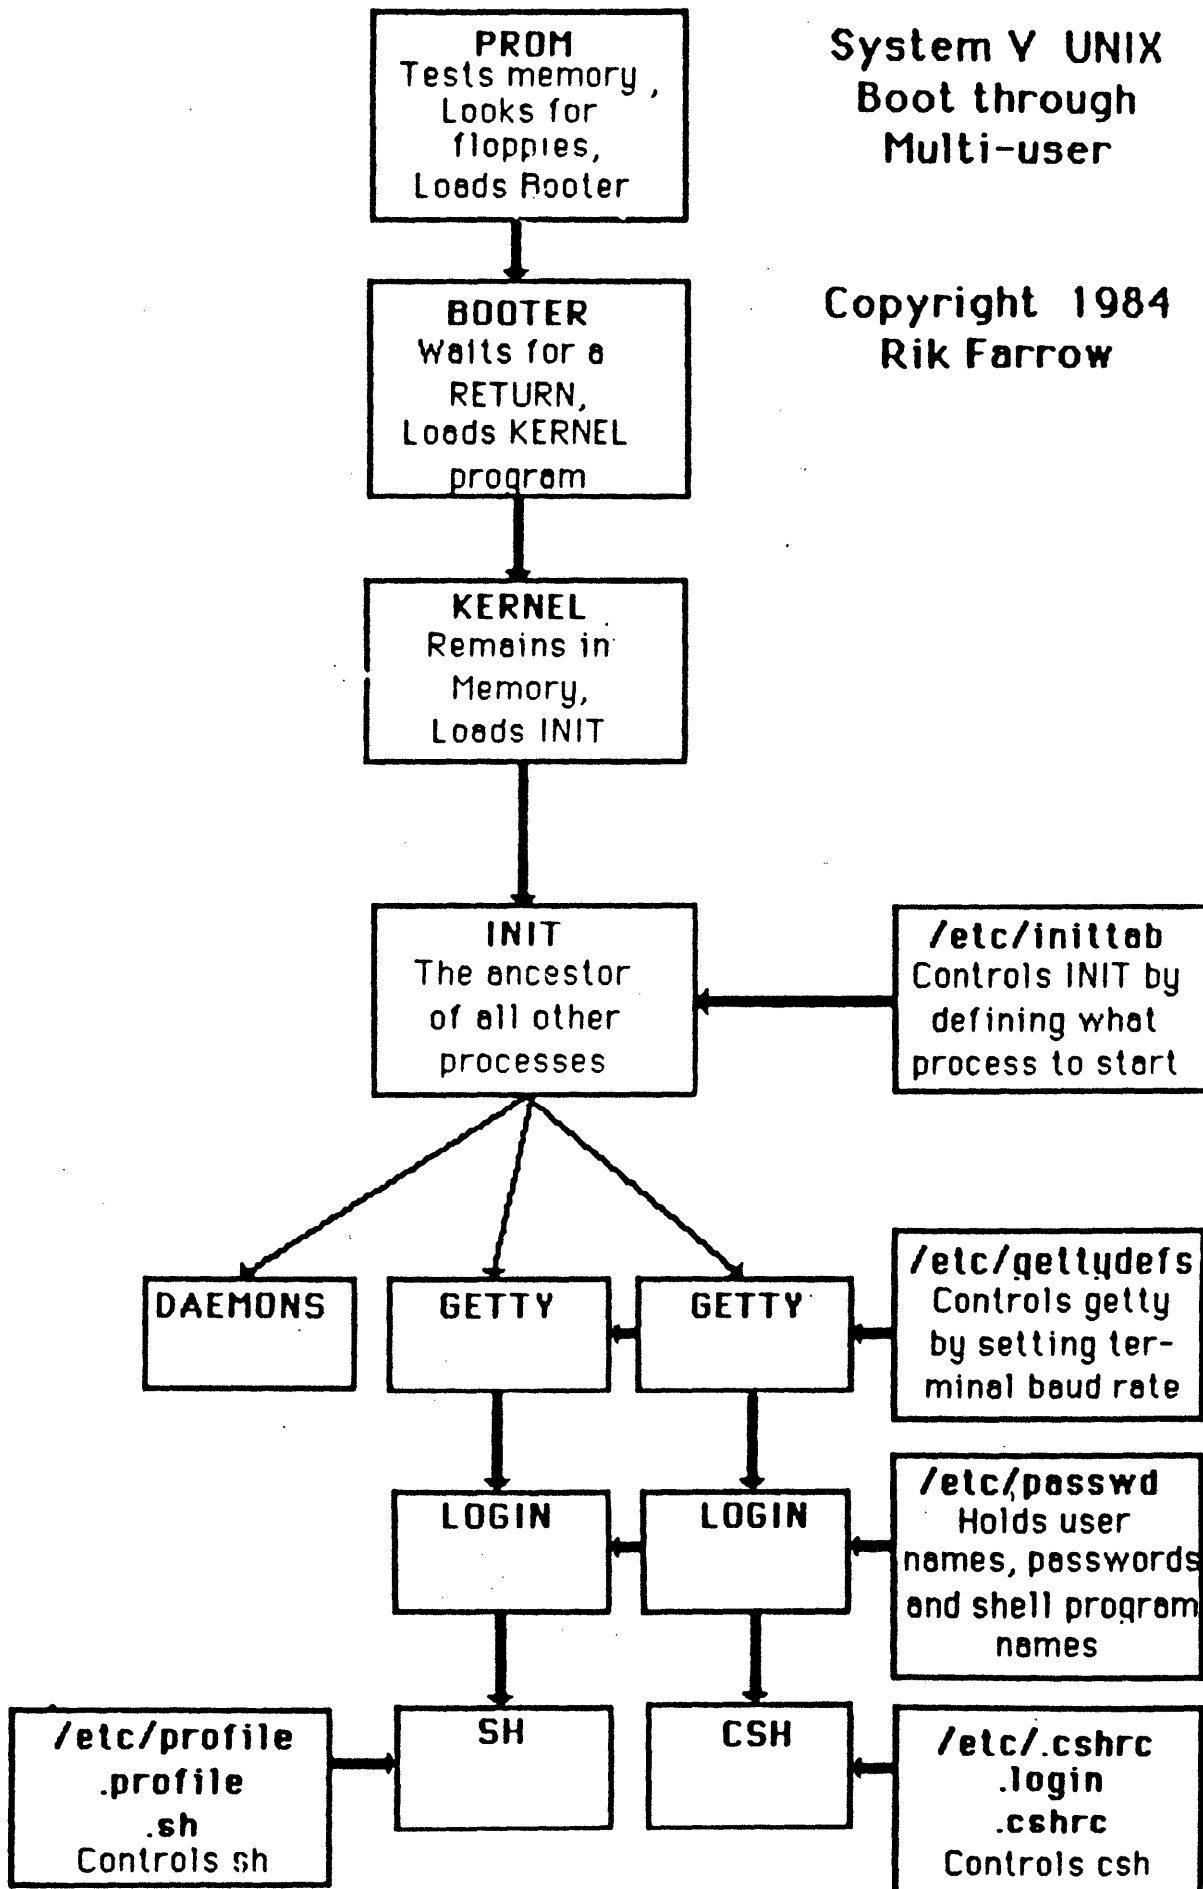
The kernel is the part of the operating system program that interfaces between user software and system hardware. User programs can access memory in their own code and data areas, but make requests (system calls) to the kernel to communicate with disks, terminals or other memory. The kernel is never swapped out (it's what performs the swapping).

The kernel initializes itself, prints copyright information, looks for the serial ports (SIO-4 boards) and looks for hard disks. The message "drive timeout" means that the kernel did not receive a response from a hard disk within a reasonable amount of time. Missing hard disks are discovered this way. The kernel then enables interrupts for the first time.

The kernel next creates the zeroeth process, the scheduler. The scheduler, as its name suggests, determines which process will be next to run or swap. The scheduler process, not shown in our drawing, runs when the system is idle, when a process goes to sleep and every clock tick. The scheduler, like the kernel, doesn't get swapped out of memory.

System V UNIX Boot through Multi-user

Copyright 1984
Rik Farrow



The next process to be created is the **init** process. The **init** process controls the birth of the next level of processes. The processes that **init** will create (through **fork** and **exec** system calls) are defined in the file **/etc/inittab**. **Init** looks through **inittab** (initialization table) every time it is awakened, and initiates or kills processes according to entries in **inittab**.

When **init** wants to initiate a login process, it begins by starting **getty**, the program that establishes a connection with a terminal (get tty, get it?). **Getty** uses the file **/etc/gettydefs** (**getty** definitions) to determine the baud rates for terminal lines and line condition information for the terminal. **Gettydefs** is set up so that **getty** will change to a different baud rate if **getty** sees a **BREAK** or framing error while trying to make contact. **Getty** also looks for lower case characters. Logging in with all upper case characters causes **getty** to consider your terminal as upper case only, and sets a flag that converts lower to upper case characters, and vice versa.

If **getty** succeeds in reading a name, it passes the name it read to the login program. The login program uses the file **/etc/passwd** to determine valid login names. The login program also requests passwords for login names with password entries. Then, the login program starts a shell program in the home directory for valid login names.

The two most common shells, the Bourne shell (**sh**) and the C-shell (**csh**), are also set up by certain files. The Bourne shell uses the **/etc/profile** file to set up conditions that are shared by all Bourne shell users. Then, the **.profile** and **.sh** files belonging to the individual user finish setting up the local environment.

The C-shell uses **/etc/cshrc** to initialize all C-shells, and the **.login** and **.cshrc** files to initialize the local environment. The files **/etc/profile**, **/etc/cshrc**, **.profile** and **.login** are executed immediately after logging in, and not during subsequent invocations of shells.

4.1.1 Inittab Controls INIT

The **/etc/inittab** file controls the **init** program through its line entries. Each line in **inittab** (short for **init** table) controls one process that may be initiated by **init**. These lines contain four types (fields) of information: the name of the entry (**id**), the run level for this entry, the type of action to be performed and the name of the process to be executed. The line that starts the console shell running in multi-user mode is shown below:

```
co:2:respawn:/etc/getty console co_19200
```

The name of this entry is **"co"**. This is the **id** that appears in the **/etc/utmp** file, processed by the **who** command.

Colons (:) separate each field of information. The second field contains a "2", meaning that this line will be active when the run level is 2. The third field is respawn, meaning that when this line is active (according to the run level), the action taken will be to create the process if it doesn't already exist. The last field is the process to be executed, getty.

To get a better idea about how inittab works, let's look at a diagram that represents the information available to init. The three boxes labeled BOOT, NORMAL and POWERFAIL, in the INIT diagram (next page) represent the three states that init can be in. Init can only be in one state at a time. Init starts out in the BOOT state when it's first loaded by the kernel. Init leaves BOOT state for the NORMAL state when it enters any of the numbered run levels. In the event of a power failure, init temporarily enters the POWERFAIL state, then returns to NORMAL after carrying out any "power" or "powerwait" actions.

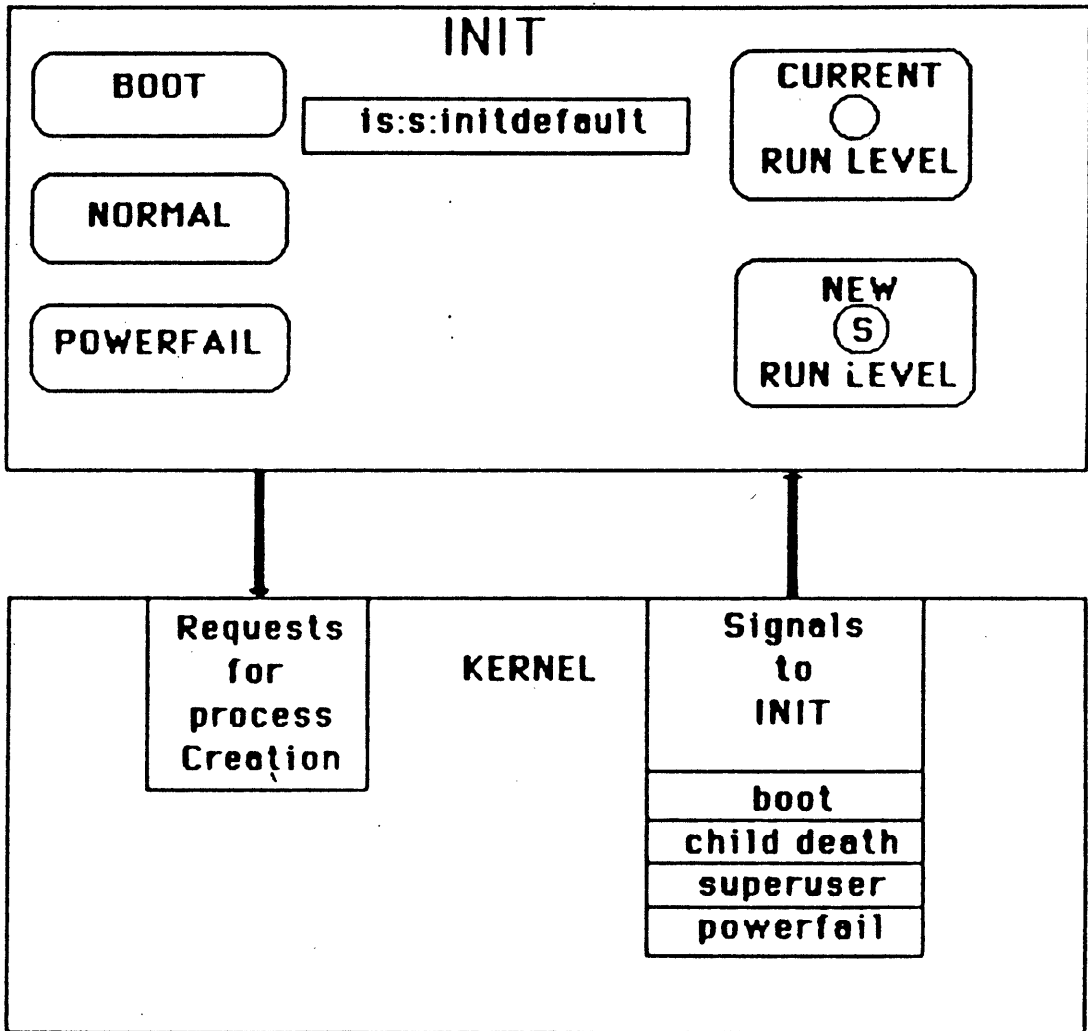
Init also keeps track of the current run level and the new run level. The run levels are the numbers from 0 to 6, and the letter s. Only the single user run level, represented by the "s", has a predefined meaning. The system administrator defines the meanings of the numbered run levels when he or she edits /etc/inittab, so their significance is a matter of organization. The system administrator can use one numbered run level to wake up one group of terminals, and use a different run level to include more terminals for login.

The first line of the file inittab appears in a window in our illustration of init. When inittab is awakened, it examines each line in inittab, always starting with the first line. Init compares the run level field with the new run level. If the new run level matches any of the run levels in the run level field, init follows the action described in the action field. If the run level field does not match the new run level, init signals any living process associated with this line that it has 20 seconds to live, and then kills any process still around at the end of 20 seconds. Then, init looks at the next line in inittab.

The actions that init will take depends on its current state. There are two actions possible while in BOOT state. These actions are triggered when init leaves the BOOT state and goes into a numbered run level the first time:

- boot Create the process named in the process field;
- bootwait Create the process named in the process field, and wait until it completes (dies) before looking at the next line inittab.

Since init can only leave the boot level once after booting, boot and bootwait actions will only be done once.



After init has entered the NORMAL state, there are four actions that it can follow. Remember, that actions are ignored if the run level for the line doesn't match:

off	Kill the process if it exists;
once	Create this process once when entering the run level;
wait	Create this process once when entering the run level, and wait until it completes before reading the rest of inittab;
respawn	Create this process, and restart it every time it dies.

The respawn action is used for initiating the login process. So, when a user logs off, his shell has died and init respawns a new getty-login-shell sequence.

There are two actions that init can perform when entering POWERFAIL state:

powerfail	Create this process once; and
powerwait	Create this process once, and wait until it dies before continuing.

Your TRICEP system cannot not detect power failure. Even if it could, the loss of power will have erased most of memory anyway. The powerfail state is designed for larger systems where some parts are more sensitive than others to loss of power.

There is one other name that can appear in the action field, called **initdefault**. By including this action in the inittab, init knows which run level to enter when it is started the first time. For example,

is:s:initdefault:

causes init to start in single user run level, s. This entry appears in the inittab file that comes with your TRICEP, and is the reason why your system comes up in single user mode. You can make your system come up multi-user instead by simply editing inittab and changing the "s" to a number from 0-6. If there is no initdefault entry, init asks at the system console for the run level to use.

4.1.2 Inittab Example

Let's take a look at the inittab file that you got with your system. We've already explained the first line, where the initdefault action sets the initial run level. The next two actions are tied to the BOOT state. The line with the id "b1" will start executing the /etc/bcheckrc shell file when the system leaves single user mode. Since this is a "bootwait" action, no other line in inittab will be processed until /etc/bcheckrc completes.

Notice that the second field, the run level, is empty (two colons together, ::). This indicates that this is line valid for any run level, from 0 to 6, and is equivalent to having the numbers :0123456: in the run level field.

```
is:s:initdefault:
bl::bootwait:/etc/bcheckrc </dev/console >/dev/console 2>&1 #bootlog
bc::bootwait:/etc/brc 1>/dev/console 2>&1 #bootrun command
sl::wait:(rm -f /dev/syscon;ln /dev/systty /dev/syscon;) 1>/dev/console 2>
rc::bootwait:/etc/rc 1>/dev/console 2>&1 #run commmands
co::respawn:/etc/getty console co_19200
tty0::respawn:/etc/getty tty0 co_19200
tty1::respawn:/etc/getty tty1 co_19200
tty2::respawn:/etc/getty tty2 co_19200
tty3::off:/etc/getty tty1 co_19200
tty4::off:/etc/getty tty2 co_19200
tty5::off:/etc/getty tty1 co_19200
tty6::off:/etc/getty tty2 co_19200
```

/etc/bcheckrc handles setting the date and checking file systems. The next line, with the id bc, passes the file /etc/brc to a shell to execute. This file, also with the action bootwait, erases the old mount table (where mounted file systems are logged).

The next two lines have the action "wait" and "bootwait" and an empty run level field, meaning that these lines are valid for all numeric run levels. The line sl removes the old /dev/syscon file and links /dev/systty to /dev/syscon. Since this is a wait action, no other processes are started before this process completes. This is the line that can result in the system console being linked to a remote terminal.

The next line starts the /etc/rc shell script. The rc script (run commands) creates a new mount table, mounts additional file systems and starts various daemons (pronounced like demons). Daemons are independent processes that perform tasks in the background. The accounting programs, the cron daemon, (used to start processes according to times in the crontable), and the lp daemon (lpsched) can be started now. You can add a command here to erase temporary files in /tmp:

```
rm -r /tmp/*
```

The next four lines spawn `getty-login` processes for the console and three other terminals. Since no run level is specified, `getty-login` processes will be initiated when any run level except single user is entered. And, because these have the action `respawn`, any time a user logs out (killing the child of `getty-login`, the shell), a new `getty-login` process will be started.

The last four lines of this file have been turned off with the action `"off"`. If you have a second SIO4 board allowing four more terminals, you can change `"off"` to `"respawn"` to permit logins on your last four terminals.

Do not turn on (with `respawn`) these lines if you do not have the capacity for eight terminals. The message that you will see if you do this will look like

```
getty: cannot open "tty3". errno: 2
```

If you accidentally change these lines, your TRICEP will thrash about trying to open the terminals attached to the non-existent second SIO4 board, and require resetting and the editing of the `/etc/inittab` file while in single-user mode.

So, this is the `inittab` file that initially controls `init`. When you boot up, you will be in single user mode (because of the `is:s:initdefault` entry). You must signal `init` with

```
telinit 2
```

to go multi-user. `init` then starts the bootwait actions, then the wait actions and finally the `respawn` actions that bring all the terminals to life.

4.1.3 Modified `inittab` Example

Now, suppose that you wanted to change `inittab` to agree with your system's configuration. As an example, let's imagine that you have two terminals, one used as the console and the other connected to port `tty1`. You also have a modem connected to `tty0`, and a serial printer connected to `tty2`. Obviously, you don't want to attempt to login the printer. You might also want to have independent control over the modem line, so that it is not always on when the other two terminals are on. And, to add a little protection to your system, you might want to come up in multi-user mode instead of single user, so that the superuser is required to enter the correct password after resetting. The edited `inittab` file that does all this follows.

```

is:2:initdefault:
bl::bootwait:/etc/bcheckrc </dev/console >/dev/console 2>&1 #bootlog
bc::bootwait:/etc/brc 1>/dev/console 2>&1 #bootrun command
sl::wait:(rm -f /dev/syscon;ln /dev/systty /dev/syscon;) 1>/dev/console
rc::bootwait:/etc/rc 1>/dev/console 2>&1 #run commands
co::respawn:/etc/getty console co_19200
tty0:34:respawn:/etc/getty tty0 co_1200
tty1::respawn:/etc/getty tty1 co_19200
tty2::off:
tty3::off:/etc/getty tty1 co_19200
tty4::off:/etc/getty tty2 co_19200
tty5::off:/etc/getty tty1 co_19200
tty6::off:/etc/getty tty2 co_19200

```

Notice that the `initdefault` run level is now set to 2. Run level 2 is the multi-user run level by convention. All of the lines with empty run level fields will match run level 2. The modem, connected to `tty0`, won't be allowed to login until the superuser changes the run level to 3 or 4. `Tty2`, with the serial printer attached, won't have a `getty`-login process started for it. The last four lines are still turned off forcing `init` to ignore these lines.

4.1.4 Sending Signals to `init`

The superuser sends signals to `init` by using `telinit` (tell `init`). `Telinit` accepts one argument, the new run level, which can be a number between 0 and 6, or the letters `s`, `a`, `b`, `c` or `Q`. The superuser could allow a login on the modem (as in our Modified `Inittab` Example) by sending `init` a message with

```
telinit 3
```

This changes the run level to three and does not terminate the processes already running at run level 2, since their run level field is empty (matching all run levels). To turn off the modem connection, the superuser can signal `init` to return to run level 2 with

```
telinit 2
```

The `telinit` command can also be used to send other signals to `init`. For example, `init` can be forced to look at the `inittab` file by sending

```
telinit Q
```

This is used to get `init` to look at an edited version of `inittab`. You can terminate a login process (that was a `respawn`) by changing the action (`respawn`) to "off" and waking up `init` with `telinit Q`.

The letters a, b, and c are used as pseudo-run levels. Normally, when the run level changes, processes that don't match the run level are killed (gracefully). You can have lines in `inittab` with the letters a, b, or c for the run level field if you want `init` to start a process without affecting the current run level or killing other processes.

4.1.5 Processes: A Few Words

Before leaving the topic of `init`, it may be useful to you to have an understanding of what a process is, how it is created and what its death means. This goes a little outside the scope of this manual, but may be helpful to you.

We have already mentioned that the kernel creates the first two processes, the scheduler and `init`. What the kernel did to create these processes was to load their program code into memory, initialize their data areas, and set up a special data area known as the upage. The upage is to processes what the inode (file descriptor) is to files. The upage holds information including the user and group number of the process owner, the process id of the parent process, the size of the code and data areas, environmental variables (like `HOME`, `PATH`, etc.) and the stack. The upage defines the process: its memory boundaries, owner, and condition.

When `init` spawns a new process, for example, `getty`, it follows a two step procedure. First, `init` makes a system call to the kernel called a `fork`. The `fork` system call makes an exact duplicate of `init` and `init`'s upage with one difference: the new copy of `init` is called the child of `init` in the new upage. Then new copy of `init` makes a second system call, `exec`. The `exec` call copies a new program, in this case `getty`, in the place where the child `init` was and modifies the upage to reflect the new code and data area. We now have the parent process, `init`, and a child process of `init`, `getty`.

`Getty` waits for someone to type in a login name. As soon as `getty` reads a possible login name, `getty` replaces itself with the login program by using the `exec` system call. In other words, the child process of `init`, `getty`, has been replaced by `login`. `Login` checks the `/etc/passwd` file and asks for a password (if one is required). If the login is successful (a match in the `/etc/passwd` file), the `login` program fixes up the upage area with the users environment, and execs the program (usually a shell) in place of itself. Now, the user's shell has replaced `login`, that relaced `getty`, that was a child of `init`. So, the user's shell is now a child of `init`.

Each time you use the shell to execute a command, the shell forks a new copy of itself, and this copy execs the command. Thus, your shell is still around waiting for the child process to finish executing the command. The shell wakes up when the command finishes (dies). If you execute a command in the background (with &), the shell process doesn't wait for the child to die, but continues independently.

When a process exits for the last time, it dies. For example, when you log out of a shell, the shell dies. Since the shell is a child (indirectly) of init, a signal, child death, is sent to init. The dead child gives up its code and data space, and its upage goes away. Then, init forks a new copy of itself, and execs a getty that waits for the next login.

4.1.6 gettydefs: Sets Up Baud Rate

Init begins the login procedure for terminals by creating getty processes. The getty program attempts to make the initial contact with the user through a terminal. The conditions required by the terminal, such as baud rate, parity, number of bits, are taken from the /etc/gettydefs file.

You may have noticed that getty gets two arguments passed to it in the "respawn" entries in inittab, the name of the terminal line to open and the name of a line in /etc/gettydefs. The line in gettydefs is the first in a chain of entries that specify which baud rate and what line conditions to use. If getty receives a NULL character from the terminal device driver (indicating that the baud rate is incorrect), it will try the next entry in the chain of entries. You can also force getty to change the baud rate by typing the BREAK key.

There are five fields in each line of the gettydefs file, with each field being separated by the "#" character. A blank line separates entries. The five fields are:

label	The name used as an argument by getty to find a line in gettydefs, like co_19200;
initial flags	The line conditions used when contact is first made, usually simply a baud rate, e.g., EXTA;
final flags	The line conditions given to the terminal before getty turns over control to login, (TABS, SANE);
login message	The message that the user sees on his screen when he is prompted for a login name; and
next label	The next gettydefs line to use if the connection fails at this baud rate.

The gettydefs file that comes with TRICEP systems follows.

```
co_19200# EXTA # EXTA SANE TAB3 #
  \r\n\nUniplus+ System V\r\n\nlogin: #co_9600
co_9600# B9600 # B9600 SANE TAB3 #
  \r\n\nUniplus+ System V\r\n\nlogin: #co_4800
co_4800# B4800 # B4800 SANE TAB3 #
  \r\n\nUniplus+ System V\r\n\nlogin: #co_2400
co_2400# B2400 # B2400 SANE TAB3 #
  \r\n\nUniplus+ System V\r\n\nlogin: #co_1200
co_1200# B1200 # B1200 SANE TAB3 #
  \r\n\nUniplus+ System V\r\n\nlogin: #co_300
co_300# B300 # B300 SANE TAB3 #
  \r\n\nUniplus+ System V\r\n\nlogin: #co_19200
```

Please notice that each entry is linked to the entry that follows by the next label field. The last entry, co_300, is linked to the first entry, co_19200, forming a closed circle. Thus, getty can attempt to make contact at each of the following baud rates:

19200, 9600, 4800, 2400, 1200, 300

If you want to change this chain, you can edit the next label field. The baud rates and line conditions honored by getty are defined in the manual entry for termio(7) in the System Administrator's Manual.

The final flag SANE has a special meaning not mentioned in termio(7). SANE is a set of "normal" line conditions: ECHO, ICANON and CRLF. These translate as echo characters, process backspace and kill characters and translate carriage returns into linefeeds (the UNIX system's default line terminator). The EXTA flag stands for 19200 baud (which is not standard UNIX).

The \r and \n are special characters that are translated into a carriage return and line feed, respectively. You can customize your TRICEP system login prompt by changing the message in the login message field. For example,

```
co_19200# EXTA # EXTA SANE TAB3 #
  \r\n\nTRICEP System V\r\n\nuser name? #co_9600
```

produces the login message

```
TRICEP System V
user name?
```

You can create gettydefs entries that are unique for each login port if you like, with special login messages. Or, you can identify a particular TRICEP if you are using several by editing the login message field.

4.1.7 etc/passwd Controls login

The login program is executed by getty when getty has collected a possible login name. The login program reads the /etc/passwd file and searches for a match to the name passed to login by getty. Login next requests a password if a matching entry has an encrypted password. Login also asks for a password if there is no matching entry for the name in the /etc/passwd file. This is designed to foil attempted system breakers who are trying random names as possible logins. Login exits (and dies) if a matching /etc/passwd entry is not found, and init starts a new getty process to wait for the next login.

If login has successfully matched a /etc/passwd entry, it uses the entry to set up the user's environment. The environment is part of the process upage. login changes the user and group numbers in the environment (and of the controlling terminal) to match the entry's user and group numbers. The HOME variable is set to the home directory field in the entry. Finally, the program that is the last field of the entry is exec'ed, replacing login. Usually this program is a shell.

The /etc/passwd file that comes with your system has only three users that you can log into: root, rootcsh, and guest. root and rootcsh are the superuser login names. You need to add a password to both these names immediately. Type the command

```
passwd root
```

while logged in as root (or rootcsh) and add a passwd if you have not already done so. Passwords should be at least 8 characters long and have some non-alphabetic characters in them. For example,

```
fog56@end      holy/fat      .to.be.30
```

are acceptable and easy to remember passwords. Your last name, nick name, wife's or boyfriend's name, are not good passwords. The root password is very important. Don't forget it! And don't share it with everyone, but give it only to people who need to know. You should also add the same password to the rootcsh (or root) login name at this point by typing

```
passwd rootcsh
```

and adding the password.

The guest entry is the example entry in the /etc/passwd file. It allows an ordinary user to log in as guest in the /users/guest home directory. This login name also does not have a password. Passwords are necessary on UNIX systems that require any security.

Editing the /etc/passwd file is explained in detail later.

4.1.8 Controlling Shells: /etc/profile, /etc/cshrc

If the login program execs a shell program, more variables are added to the environment. Each shell program reads a shell script in the /etc directory that contains commands to set up the environment. Every user's shell is thus initially set up by these files. The user has no control over this, so that system administrator can give all users the same initial environment. The users can override this initialization by explicitly resetting these variables in their own set up files.

/etc/profile Initializes the environment of all login Bourne shells (sh)

/etc/cshrc Initializes the environment of all login C-shells (csh)

\$HOME/.profile The users own initialization script for the Bourne shell (sh)

\$HOME/.cshrc The users own initialization script for the C-shell (csh)

One of the things the a system administrator can do with /etc/profile and /etc/cshrc is to establish a system wide umask. The umask is used to mask out certain permission bits when new files are created. For example, a umask of 022 prevents write permission from being automatically granted to the Group and Others when new files are created. The user can override the umask by setting his own umask in .profile or .cshrc. Examples of user configuration files are contained in the root directory. These files all begin with a dot (.) and can be listed with:

```
ls7 -a
```

4.2 Connecting Additional Terminals



This section should be easy for you to follow, since you have successfully connected the console. The TRICEP system can have three to seven additional terminals, depending upon whether you have one or two SIO-4 boards. A quick glance at your back panel is all you need to ascertain this. Systems with one SIO-4 board have four DB-25 female connectors (like the letter D lying on its rounded side). Systems with two SIO-4 boards have 8 DB-25 female connectors.

Terminals are connected in the same way as the console was. The instructions for doing this are included in steps 5 and 6, in the Installation part of this manual. There are two differences that you need to be aware of. The first is rather obvious, that is, you will be connecting each additional terminal to a DB-25 connector labeled tty0 through tty6. Only the console is connected to the connector labeled console. (See the Back Panel illustration on the next page.)

The second difference is that additional terminals may be setup to work at different baud rates. Remember that the console must be set to 19200 baud. You can use other baud rates for your additional terminals, but 19200 is recommended. If you want to use a different baud rate for a terminal, you need to edit the line in the /etc/inittab file that starts the getty-login-shell process for it. For example, to make the terminal connected to tty2 login at 9600 baud, you would change 19200 to 9600 in the line starting with tty2:

```
tty2::respawn:/etc/getty tty2 co_9600
```

Remember that co_9600 is the label of a line in gettydefs and that the BREAK key can be used to change baud rate at login.



4.3 /etc/passwd: The User Database

/etc/passwd is the user database. All the valid user names for the system are kept here. Programs that use user names, like `ls`, `chown` and `find`, search /etc/passwd to cross reference user ids (a number) and user names. Editing the /etc/passwd file is an essential part of system configuration.

The next two sections go into the details of adding and removing users. There is more to this than just editing an entry in /etc/passwd. It is also possible to create special users by having passwd entries that execute a restricted shell (`rsh` or `ushell`) or some other program than a shell.

Adding New User Names

The superuser is responsible for adding new users to a system. The superuser decides where the new user's home directory will be and selects user and group numbers. Then, the superuser:

- o Edits the /etc/passwd file and adds an entry;
- o Edits the /etc/group file and adds the user name to the groups it is a member of (optional);
- o Makes the home directory, changes its ownership and adds a login profile;
- o Tests the new user login.

These steps are all that is necessary to create new user accounts.

4.3.1 /etc/passwd, The User Name File

The /etc/passwd file can be read by anyone, but it can be written to only by the superuser. This arrangement allows people to search the password files for their own and other user's entries. Certain commands, like `ls` with the `l` option, or `find`, use the passwd file and substitute user names for user numbers.

Editing the passwd file is restricted to the superuser because anyone who can edit the passwd file can give themselves superuser powers. The superuser can access, change or erase every file in the file system. It is necessary to have the superuser so that the file system can be maintained. It is also necessary to restrict access to the superuser to avoid accidental (or malicious) damage to the files system.

Every entry in the `/etc/passwd` file has seven fields for information. They are:

- o The user name, the name typed in during login;
- o The password; when an entry is added, the password field is left blank and the `passwd` command is used to add the password;
- o The user id number, a unique number between 1 and 65235 assigned to the user name;
- o The group id number, a number (1-256) shared with others in the same group; members of the same group have special access permission for other members files;
- o A comment field; often the user's name is put here;
- o The pathname of the HOME directory; the directory assigned to the user for login;
- o The pathname of a program to execute; usually, the shell is executed, and the Bourne shell is the default.

The information in password entries is separated by colons (":"). Two blank colons appearing together indicate a null field. Here is an example entry from a password file:

```
rik:A5hjp0vnh9up:34:2:rik farrow:/users/rik:/bin/csh
```

User names can be up to eight letters long. The name must be in lower case letters, so that `getty` doesn't consider the terminal to be upper case only. Usually, people choose their own user names, and base them on their real names or nicknames.

Passwords appear after the user name. This field is left empty (two colons together, `::`) when the entry in the `passwd` file is created. The `passwd` command is used to add a password in an encrypted form.

The superuser or the user can assign a password using the `passwd` command after editing the `/etc/passwd` file. In systems that require any security at all, passwords are mandatory except for restricted users, like `who`, or `uucp`.

The two numbers that follow the password are the user and group id. These numbers are used to check the user's file access permissions whenever files are accessed. Every one sharing the same group id number, 2 in this example, has the same group access permission. In the same way, if two user names have the same user id number, they will have the same user permissions. This is not a good way to share permissions. User numbers should be unique. The group id number is provided for sharing file access permission.

The information following the user and group numbers is the comment field, in this case, a full name. Another use for this space is for comments, such as the user's location and phone number, or the purpose for the password entry.

The next field is the pathname of the home directory. When the login process is complete, the user finds himself in this directory. This information is also assigned to the environment variable HOME during login. Anytime the `cd` command is used without arguments, the current directory is changed to the HOME directory.

The last field is the program to execute when the user logs in. The program executed is, by default, the Bourne shell. In the example, "rik" logs in with the C shell instead. In the following example of your `passwd` file, several more programs to execute are shown.

```
# cat /etc/passwd
root:v2xDWX6hxB5J6:0:0:::/bin/sh
rootcsh:CHBjPcMVBi36Q:0:0:::/bin/csh
daemon:xxxxxxxxxxxx:1:1::/
bin:xxxxxxxxxxxx:2:2::/bin:
sys:xxxxxxxxxxxx:3:3::/bin:
adm:xxxxxxxxxxxx:4:4::/usr/adm:
uucp::5:5::/usr/spool/uucppublic:/usr/lib/uucp/uushell
check:xxxxxxxxxxxx:6:6::/
lp:xxxxxxxxxxxx:7:7:/usr/spool/lp:
who::22:1:who command:/bin:/bin/who
guest::100:100:guest account:/users/guest:
```

The uucp user name is used by the UNIX to UNIX copy program. Logging in on this account executes the `uushell` program, a restricted shell. This command expects to receive input from a similar shell on a remote UNIX system. Access to the local system is controlled by forcing the remote system to log as a restricted user.

The restricted user "who" displays the output from the `who` command and then logs out. The guest account allows password free access to the system and pose potential security hazards. Guest, with the empty command field, will get the default Bourne shell.

4.3.2 Editing /etc/passwd

The superuser adds new users by editing /etc/passwd. It is a good idea to work on a backup copy of special files, such as /etc/passwd, instead of editing them directly. For example,

```
cp passwd passwd.bak
```

provides an immediately available copy of the passwd file for editing. The following example illustrates the addition of a user to /etc/passwd.bak:

```
$ su
Password:
# cd /etc
# cp passwd passwd.bak
# ex passwd.bak
"passwd.bak" 23 lines 567 characters
a
ralph::223:100:Ralph Minkler:/users/ralph:/bin/sh
.
:x
"passwd.bak" 24 lines, 602 characters
# []
```

When you are satisfied with the edited **passwd.bak** file, use it to replace the old copy of **passwd**:

```
# chmod 644 passwd
# mv passwd.bak passwd
```

Before replacing the old passwd file with our edited version, passwd.bak, it is necessary to change the file access permission for the passwd file to read by all and write for owner only (644). Then, passwd.bak can replace passwd. The passwd command strips away write permission from /etc/passwd every time the command is used.

At this point, a new user name, ralph, has been added to the system. Ralph still needs to have his HOME directory created before he can log in. He can also be added to the list of users in his group.

Before continuing, it is a good idea to assign a password to ralph. Entries without passwords in the /etc/passwd file pose threats to system security. Having a password protects you from anyone, whether it is your 10 year old whizz kid, or the employee you just fired. Only the superuser and the user his self can change (or add) a password. Here is the way the superuser adds a password:

```
# passwd ralph
Password:
Retype password:
```

Notice that the password does not appear while it is typed. This prevents someone from seeing the password on the screen. If the superuser or the user forgets the password, the superuser can always add a new password.

Also, the `passwd` command automatically strips write permission from the `passwd` file. This includes write permission even for the superuser. This means that the superuser must take the extra step of adding write permission (with `chmod`) before replacing the old password file.

4.3.3 Editing `/etc/group`

The `/etc/group` file contains entries similar to the entries in `/etc/passwd`. Editing this file is not essential. The purpose for this file is to allow programs like `ls` to substitute group id numbers for names.

Each item of information is separated by colons, and a password may be present. There are four types of information in `/etc/group`:

- o The group name,
- o An encrypted password,
- o The group id number and
- o The list of user names that can use the group id number separated by commas.

When `/etc/group` is edited, the superuser adds the user name to the entry for the selected groups. The user name is always added to the entry for the group id number used in the user's entry in `/etc/passwd`. In Ralph's entry,

```
ralph::223:100:Ralph Minkler:/users/ralph:/bin/sh,
```

the group id number is 100. This is the group id normally associated with ralph. Ralph may have access to other group id numbers if his user name is added to the list following the group id number. The `newgrp` command is used to change group id number. In the example that follows, ralph is added to his own group list, and to an additional group id number:

```

# ex group
"group" 3 lines, 169 characters
:%
users::100:admin,guest,guest1,kevin,rik:
sys::1:root,drivers,stand:
staff::200:bob,kevin,splen,len,rik,kevinb,bill:
:s/bill/bill,ralph/
staff::200:bob,kevin,splen,len,rik,kevinb,bill,ralph:
:ls/rik/rik,ralph/
users::100:admin,guest,guest1,kevin,rik,ralph:
:x
"group" 3 lines, 181 characters
# []

```

The /etc/group file allows selected users to be members of multiple groups. When ralph logs in, he has the group id number 100. Because ralph's name is included in the "staff" entry, ralph can also have the group id 200 assigned to her:

newgrp guest

The password field of /etc/group entries is usually null. If a password is present, a user whose name is not in the entry can attempt to access the group id number by using the **newgrp** command and typing the password when requested. The user attempting to change to a group id without a password or the user name on the list gets the message:

Sorry.

The /etc/group file is not required for the creation of a new user name. /etc/group is provided as a means for switching group id's. There is no easy way to add a password to the entries in /etc/group.

4.3.4 Creating the New HOME Directory

The superuser makes a directory for the new user with the pathname in the HOME directory field of the /etc/passwd entry. The HOME directory should be part of the same file system as the rest of the user's group.

There are three things to be done in setting up a new HOME directory:

- o Making the directory,
- o Copying a login profile to it and
- o Changing the user and group ownership.

Here are the commands used to create a new HOME directory for ralph in the directory /users, the directory that has been set up for users' HOME directories:

```

# mkdir /users/ralph
# cp /.profile /users/ralph
# chown ralph /users/ralph /user/ralph/.profile
# chgrp users /users/ralph /user/ralph/.profile
# []

```

Changing the ownership (**chown**) of the new directory and the files within it is an essential task. Otherwise, ralph may log in to find that he does not have permission to write in his own directory, or edit any files.

The file **.profile** is used as the configuration file by the shell. It is not essential to have this file, but it does provide an example of a configuration file that can be edited by the new user. The file **.login** is used by the C shell.

Now, if everything has been done correctly, you should be able to log in as ralph. New user names are always tried by the superuser before they are given to the new user to check for mistakes.

```

# su ralph
Password:

$ ls -al
drwxr-x--- ralp root          96 Oct  3, 1984 .
drwxr-x--- bin  root         214 Aug 18, 1984 ..
-rwxr-x--- ralp root          27  Oct  3, 1984 .profile
$ []

```

Notice that the owner of "." and the **.profile** file is now ralph.

4.4 Removing Users

In the course of ordinary events, people move on. You may find that your list of users has changed requiring the removal of some names from **/etc/passwd**. We actually advise against the immediate removal of user names from **/etc/passwd**. What you can do instead is to give them an impossible password to prevent them from logging in.

The reason for keeping old user names around is that files are identified by their user numbers. Then, programs look in **/etc/passwd** and cross reference the user number to a user name. If you have removed the user name from **/etc/passwd**, you will be faced with a mysterious number instead of a recognizable name in long directory listings (**ls -l**), for example. So, here is the way to edit **/etc/passwd** to prevent old users from logging in:

```

# cp passwd passwd.bak
# ex passwd.bak
"passwd.bak" 23 lines 567 characters
/ralph
ralph:hd64kghnd3hgG:223:200:Ralph Minkler:/users/ralph:bin/sh
C
ralph:deleted:223:200:Ralph Minkler:/users/ralph:/bin/sh
.
:x
"passwd.bak" 24 lines, 601 characters
# []

```

The word "deleted" in the password field makes logging into the user "ralph" impossible. Programs can still identify files using ralph's name because the entry still exists. This is why the /etc/passwd file is referred to as a user database: some information about users is kept here.

You might also want to remove ralph's files from the file system. Presumably, you will want to backup the files first, in case there is something of value there. The following commands back up the files beginning in ralph's HOME directory, then delete all the files:

```

# tar cvfB /dev/dj0a 798 /users/ralph
# rm -fr /users/ralph

```

The tar command requires that you have inserted a formatted floppy diskette, and have enough formatted diskettes ready to complete the backup. This was explained earlier in the section on First Time Use.

4.5 etc/termcap: Terminal Capabilities

Some of the UNIX commands take advantage of intelligent terminal capabilities, like cursor addressing and line insertion. The editors **ex** and **vi** both use this information, as do some games. The information about terminals is kept in the file /etc/termcap, **terminal capabilities**.

The /etc/termcap file that you have was edited so that it only supports terminals manufactured by Morrow. However, the full termcap file exists as /etc/termcap.all, and you can add entries from this file that support other terminals.

The reason for not including the entire file is that much time is wasted by vi or ex scanning the lengthy, original file. If your terminal is not included in the edited version of /etc/termcap, edit the file /etc/termcap.all, find the lines that support your terminal and add them to the /etc/termcap file. The script that follows shows an editing session that "clips" out the entry for a Soroc 120 and adds it to the termcap file.

```

ex /etc/tercap.all
"/etc/termcap.all" 1098 lines, 48870 characters
:/soroc
MI|soroc|Soroc 120:\
:nu
  293 MI|soroc|Soroc 120:\
:294,296nu
  294      :cd=\EY:ce=\ET:cl=2\E*:ma=^K^P^R^L^L:\
  295      :kl=^H:ku=^K:kr=^L:kd=^J:tc=adm3a:
  296 Ma|aa|annarbor|ann arbor:\
:293,295w /etc/soroc
"/etc/soroc" 3 lines, 56 characters
:q
cat soroc >> /etc/termcap

```

Entries for terminals in /etc/termcap.all begin with a line of possible names. This line of names always starts in the leftmost column. The line beginning with Ma| marks the start of the next entry. What you need to do is find your entry, discover the line numbers for your entry (293 through 295 in this example), and write these lines to a file. Then, you can append this little file to the /etc/termcap file.

4.6 Connecting Printers

The TRICEP system supports both parallel and serial printers. Since having a serial printer uses one of the possible login ports, you may wish to use a parallel printer instead. Other than this, there is no particular advantage to having one type of interface or the other.

The difficulty in adding either type of printer is usually in the cable. We can tell you exactly what your TRICEP system produces in the way of signals on the serial and parallel ports. You must match this up with what your printer needs and produces.

4.6.1 Serial Printers

Serial printers use the bit at a time method for communicating. The ends of a serial cable have DB-25 connectors on them, just like terminal cables. The required pinout on the cables AT THE TRICEP END is exactly the same as for terminals:

pin 2	Transmit data
pin 3	Received data
pin 7	Signal ground
pin 20	Data terminal ready

Pins 2, 3 and 7 will be the same at the printer end. However, pin 20 might need to be connected to some other pin at the printer end.

The various different printer manufacturers use any of a number of pins to signal when their printer is not ready. For example, NEC Spinwriters use pin 19. If you hook up the wire from pin 20 at the TRICEP end to the correct one of these pins, everything will work fine. Your task is to read your printer manual and decipher which pin holds the printer ready signal. Most printers (for example, Epson, TI, DEC, etc.) use pin 20.

The SIO-4 board has a built-in interlock, not controlled by software, that prevents it from sending characters unless it sees Data Terminal Ready on pin 20. The Data Terminal Ready signal is logic true, meaning that when the voltage on pin 20 is high (between 12 and 20 volts), Data Terminal Ready is true, and transmission can proceed.

Your printer may use X-ON or X-OFF handshaking to turn off and on transmission. This will work fine with the UNIX software, but pin 20 must still be true for the SIO-4 board to work correctly.

When you have connected your printer, you can test it with a couple of simple commands. First, you need to prevent init from trying to login your printer. This was explained in the section called Modified Inittab Example. Essentially, you need to edit the /etc/inittab file, change the entry that corresponds to your printer from "respawn" to "off", and signal init that you have edited the /etc/inittab file with

```
telinit Q
```

Next, you need to set up the correct baud rate for the port. Suppose, you have connected your printer to port tty2 (and have edited this line in inittab). Let's assume that your printer is set up to work at 1200 baud, a nice fast baud rate. You set up the TRICEP to work at 1200 baud on port tty2 by typing

```
stty B1200 < /dev/tty2
```

If you get back the message "Permission denied", it is because you must be logged in as superuser (the # prompt) to edit inittab, use telinit, or to change the baud rate of a port that you aren't connected to. You may also need to change the file access permission mode with chmod, as in

```
chmod 600 /dev/tty2
```

With the baud rate correctly established, try this command to send a test message to your printer:

```
echo This is a test of the... > /dev/tty2
```

The line up to the greater than sign should appear at your printer. If it doesn't, you need to check the power to the printer, that the printer is ON-LINE and that it has paper ready. When all these things are true, but your test fails, you have a problem in your cable. Make certain you are using the correct port and device name. (Look at the drawing in the earlier section entitled Adding Additional Terminals). If everything else checks out okay, your problem is probably with pin 20 and the printer ready pin.

When you have your printer working correctly, go on to LP, the Print Spooling Facility.

4.6.2 Parallel Printers

The TRICEP has a Centronics style printer interface for parallel printers. The connection for this printer is located just above the connection for tty2 (see Back Panel drawing in the section labeled Adding Additional Terminals). The cable enclosed with the TRICEP is connected here **ONLY WHILE THE TRICEP IS OFF**. Connect the cable so that the cable drops down without being folded over the connector. The red stripe will be on the left (when viewed from the back). The other end of the cable can be connected to printers with Centronics style interfaces.

The name of the parallel port is /dev/cent. Connect your printer, turn the power on to it, get its paper ready and make it ON-LINE. Then, you can test out your hookup with

```
echo This is a test... > /dev/cent
```

This should print the line up to the greater than sign on your printer. If it didn't, check your printer's setup, and your cable connections. If these things don't seem to help, compare the signals your printer requires (as outlined in your printer's manual) with the signals your TRICEP is sending. The connector pin 11, BUSY, may be the culprit.

4.7 LP, Print Spooling Facility

The lp command arranges background printing of files in the UNIX system. There is considerable flexibility allowed with lp, meaning that you must configure it so that it will work with your system. Some of the things lp can do are:

- o collect files for later printing,
- o add banners and other information to each job printed,
- o send files to a default printer,
- o send files to the first available printer of a group of printers, and
- o be stopped and restarted with simple commands.

Before you can use lp, you must configure your system with the lpadmin command. The lpadmin command can only be used by the superuser, and only when the scheduler for lp, lpsched, is not running. The next section explains how to configure your system with lpadmin and add lpsched to your /etc/rc file. The section that follows outlines the basic use of lp and the enable and disable printer commands. The information presented here is based on the Section in the System Administrator's Guide labeled LP.

4.7.1 LPADMIN, the lp Configuration Command

Your TRICEP system comes configured to use lp with a centronics interface printer connected to /dev/cent. This printer has been set up as the default printer, and is called mp300. You still need to read the instructions in this section to understand how lp works. You also need to add several lines to the /etc/rc file to make starting of the lp printer daemon automatic with going multi-user.

You can follow the instructions in this section if you want to change how lp works. If you want to change your configuration after lpsched has been started, you must halt lpsched with lpshut.

The lpadmin command allows you to

- o create printer names and associate them with ports,
- o select the interface, a program that controls the printing of characters,
- o group printers into classes, so that files may be sent to any one of a group of printers, and
- o establish the name of a default printer destination.

The first bit of funny business about lp is that the commands lpadmin, lpshut and lpsched are all located in the /usr/lib directory. This is because these commands will only work for the system administrator, not ordinary users. Practically speaking, these commands won't work unless you add the /usr/lib directory to your search path. This may have already been done for you, but you can check with the command

```
echo $PATH
/etc:/bin:/usr/bin:/usr/lib
```

Each directory name in the path is searched whenever you type a command to the shell. In our example, /usr/lib is the directory that is searched last. If /usr/lib does not appear in your PATH, you can add it temporarily by typing

```
PATH=/etc:/bin:/usr/bin:/usr/lib
```

and permanently by adding this line to the /.profile file. You can also prepend the directory name, /usr/lib, to all your commands, as in /usr/lib/lpadmin, but this is awkward. Or, change directory to /usr/lib and prepend ./ to all your lp commands.

Okay, before setting up lp with lpadmin, you need to select a name and an interface for your printer, or printers. For a name, select something short but easily remembered. For example, mp300 or daisy are good printer names and spinwriter or pl are not so good (too long and unclear).

The interface program can be a shell script or a C-program. The task of the interface is to prepare the printer, print an identification of the job being printed (if desired) and to send the characters from the files being printed to the standard output. lp handles routing the characters to their correct destination. The directory /usr/spool/lp/model has several model interface programs. These programs are shell scripts, and can be modified to fit your own requirements. There is a separate directory used for modified models, /usr/spool/lp/interface, that you can copy your modified models to.

The interface program or shell script is passed a list of arguments. These arguments can be used to create the identifying title for the print job. The filenames to be printed are included in this list of arguments and MUST be used by the interface program. The arguments are:

```
arg[0]    the request id given when lp receives a request,
arg[1]    the user name of the user who made the request,
arg[2]    the (optional) title of the request,
arg[3]    the number of copies of the request to print,
arg[4]    the options that can be used by the interface
           program, specified by the user, and
arg[5-n]  the full pathnames of the files to be printed.
```

If all that your interface program does is print the files without an identifying page, you could use the simple shell script:

```
shift 5
files="$*"
for file in $files
do
/usr/bin/xtab "$file" 2>&1
echo "\n\014"
done
echo "\014"
exit 0
```

This script skips the first five arguments (request id, user name, title, number of copies and options) with the five shifts. Then, it prints the file (xtab "\$file"), sending any error messages to the error device. Next, the echo "\n\014" sends a newline and a formfeed, and the next file, if any, is printed. Then, a final formfeed is sent. This script is named /usr/spool/lp/model/mp300. The xtab program, in /usr/bin, expands tabs into the correct number of spaces to reach the next tab stop. Tab stops are every 8 spaces.

If you are using a serial printer, you need to add a command to the interface script (like mp300) to set the correct baud rate, like

```
stty B1200
```

before sending any characters to the printer. In the mp300 script, add the stty command after the shift 5.

The model named /usr/spool/lp/model/dumb provides a much more complex example of an interface program. Dumb provides an identifying page before each request is started, complete with the user name in a banner. This model also expects a 132 character wide printer, so you may want to edit it before using it.

If you want to filter (modify) the characters sent to your printer, you need to replace the xtab command (which is a filter that replaces tabs) with the name of the appropriate filter program. For example, if your printer can expand tabs, you don't need xtab, and can use cat instead.

Okay, let's suppose you have selected the name of your printer and the interface program to be "mp300" and "mp300", respectively. The printer is connected to the parallel device, /dev/cent. The lpadmin command that we used to set this printer up is

```
lpadmin -pmp300 -v/dev/cent -mmp300
```

This creates a printer destination (-p) named mp300 connected to port (-v) /dev/cent, using the model (-m) interface file /usr/spool/lp/model/mp300. Now, if you want this printer to be the default destination, that is, the printer that lp uses unless some other printer is specified, then you would use the command

```
lpadmin -dmp300
```

This makes printer mp300 the default (-d) destination of lp requests. Now, imagine that you have a second printer, and that you want this printer to be included also as a default destination. You can make this second printer, that we will name daisy, and the original printer into a class of printers, and change the default to this class. For example, let's create the new printer, daisy, a serial printer connected to tty2:

```
lpadmin -pdaisy -v/dev/tty2 -mmp300 -cany
```

The -cany argument creates the class "any" with a new printer "daisy". To add the original printer to this class, you type

```
lpadmin -pmp300 -cany
```

Now, both printers are a member of the class any. To make the class the default, type

```
lpsched  
accept any  
enable any
```

This sets up lp so that requests go to the first available printer in the class any. If all that you have is a serial printer connected to /dev/tty2, just use these commands:

```
lpsched  
accept any  
enable any  
lpadmin -pname -v/dev/tty2 -mmp300  
lpadmin -dname
```

and you have set up lp for a printer named "name", connected to /dev/tty2 and using the interface program mp300. The printer "name" is also the default printer.

Great going! You've set up lp to work with your printer (we hope). But before you can test it, you must do three things: start the scheduler, enable the printer and allow it to accept requests. This is done with the following three commands:

```
lpsched  
accept mp300  
enable mp300
```

If you created a class, like "any", you must also use the accept command before it can accept requests, as in

```
accept any
```

The lpsched command starts the lp scheduler. The lp scheduler initiates the printing of each file. The accept command allows lp to request printing by a particular printer. The enable command enables the printer. The accept and enable commands are only necessary when you change or modify lp with lpadmin, or after you have used disable or reject.

Is everything working? You can try the lp command and see:

```
lp /etc/passwd
```

should send the passwd file to the default printer. This will take some thrashing about, as lp creates a new file with the request to print /etc/passwd (and other information) in it, and the scheduler actually starts printing the file. If nothing happens after a thirty seconds or so, you've made a mistake somewhere. Have you enabled the printer? Will it accept requests? Is lpsched running? lpsched will refuse to run if the file /usr/spool/lp/SCHEDLOCK exists, so you can remove this file and try again.

If you want or need to modify lp, you must first halt lpsched with lpshut. Any job that was being printed will be halted, and will be restarted at the beginning when lpsched starts running again.

Making lpsched a regular part of your system is done by adding the following lines to the file /etc/rc:

```
rm -f /usr/spool/lp/SCHEDLOCK
/usr/lib/lpsched
echo "LP scheduler started"
```

These lines will automatically start the lp scheduler every time your system goes multiuser. You don't need to enable or accept requests except after modifying a printer or class with lpadmin.

4.7.2 Using lp

Here are four ways to use lp:

```
lp /etc/passwd
lp < /etc/passwd
cat /etc/passwd | lp
lp -c /etc/passwd
```

Only the first of these four ways does not make a copy of the file /etc/passwd. This means that if the passwd file is modified between the time of the request and the actual printing of the file, the modified passwd file will be printed. Making a copy of the file to be printed will also take lp longer than using the original.

Every user can use the lp command. Several other commands are also available to all users:

lpstat	reports on the status of a printer or request
enable	enables a printer
disable	disables a printer
cancel	stops a request from printing

The lpstat command allows you to discover what has happened to either a particular printer or to a request for printing a file. For example,

```
lpstat -pmp300
```

returns the status of the printer named mp300, and

```
lpstat mp300-002
```

reports the status of the request named mp300-002. The request names are reported when you invoke the lp command.

The disable command allows anyone to halt a printer. For example, if you notice that a printer (say, daisy) is about to run out of paper, or has jammed, you can stop it by typing


```
disable -r"paper jam" daisy
```

This not only stops output to daisy, but also adds a reason (the -r argument) that can be printed by the lpstat command. After the printer is ready to be restarted, simply type

```
enable daisy
```

Output of the last file to be printed will be restarted.

The cancel command is used to remove any file from the list of requests. Suppose, for example, some request is sending endless formfeeds to the printer. You can quickly stop this by typing

```
disable daisy  
cancel daisy-005  
enable daisy
```

The disable command halts the printer, the cancel command removes the request from the queue and the enable command restarts the printer. If the request belongs to a different user than the one that canceled it, the owner of the request is sent mail by cancel.

4.8 Adding Disks to Your System

You can add additional disks, both hard and floppy, to your root file system with the mount command. The mount command makes a file system on an additional disk part of the root file system.

There are only three requirements for mounting additional disks:

1. There must be a file system on the disk;
2. The file system must not already be mounted; and
3. A directory in a mounted file system must be available.

If you want to add a new hard or floppy disk, you must first format it. Floppy disks are formatted by typing

```
diskformat /dev/rdj0
```

To format additional hard disk, the `fmw` program is used. This program is interactive, and will ask you to supply the drive number and the type of the hard disk. Although the format program checks the hard disks, we recommend that the `badblock` program also be run. For the second hard disk, use

```
badblk /dev/mwlhw0
```

After formatting the new disk, a file system must be made on it with the `mkfs` (make file system) command. Remember please, that this is only for new disks that you are adding. Both formatting and `mkfs` will destroy information already on a disk.

To make a file system on a 5 1/4" floppy disk, type

```
mkfs /dev/dj0 400 2 5
```

To make a file system on a hard disk, type

```
mkfs /dev/mwlc `devsize -k /dev/mwlc` 13 17      (1024 byte)
mkfs1B /dev/mwlc `devsize /dev/mwlc` 5 17      (512 byte)
```

These commands assume that you have configured the disk as the second drive (drive 1), with the device name `/dev/mwlc`. The ``devsize -k /dev/mwlc`` provides the number of (1024 byte) blocks available for making the file system. The interleave factor for the free list is 13 and the number of sectors per track is 17. The `mkfs1B` makes 512 byte per block file systems. The `1k` block sized file systems are faster and more efficient.

After making a file system, you need to check it with

```
fsck /dev/mwlc      hard disks
fsck /dev/dj0       floppy disks
```

Now, you are ready to mount the new file system.

4.8.1 Mounting File Systems

The mount command requires two arguments, the device name and a directory name. The /t directory in the root directory has been set aside for mounting temporary file systems. If you intend to mount a hard disk file system regularly, you may wish to create a directory just for it, as in

```
mkdir /a
```

This makes a new directory in root for semi-permanent hard disk file systems.

Mounting file systems is done with this command:

```
mount /dev/mw1c /a          for hard disks
mount /dev/dj0 /t          for floppies
```

The decision to use directory /a or /t is an arbitrary one. You can use either, or make your own directory name with mkdir.

The mount command currently produces an errant warning message:

```
WARNING!! - mounting: <> as </t>
```

This is a bug. Mount will correctly complain if there isn't a file system on the disk, or if the device is already mounted. The list of mounted devices is kept in the human readable file /etc/mnttab. You can display this list with the command without any arguments.

```
mount
/ on /dev/mw0a  read/write on Fri Aug 24 02:31:59 1984
/t on /dev/dj0  read/write on Fri Aug 24 16:44:23 1984
```

You can routinely mount hard disks by adding a line with the appropriate mount command in the file /etc/rc. If you are routinely mounting a file system, you should also add its name (for example, /dev/mw1a) on a line by itself in the file /etc/checklist.

```
/dev/mw0a
/dev/mw1a
```

/etc/checklist is used by fsck as the default list of file systems to check.

4.8.2 Unmounting File Systems

Disks are unmounted using the `umount` command. Please observe that this is not `un-mount`, but `umount` (no first n). The `umount` command requires the device name as its argument, and that the file system be idle before it is unmounted.

```
umount /dev/dj0
Device busy.
```

A device is busy when there is a file open in the file system or someone's current directory is in that file system. The system administrator can use the `fuser` (find user) command to discover the identity of a user preventing a file system from being unmounted.

4.9 Connecting Modems

Your TRICEP can be set up for remote login by connecting a modem to one of the tty ports. This way a person possessing a modem and a terminal can access your TRICEP via phone lines. But, before you can use a modem with TRICEP, certain configuration changes must be made. The cable used between the TRICEP and the modem must also meet very specific conditions.

You need to make configuration changes to two files: `/etc/inittab` and `/etc/gettydefs`. Back in the section "Modified inittab Example" we explained how to change `inittab` to include a modem on port `tty0`. Please follow the example in that section for including a modem on port `tty0`. (We assume that you are using port `tty0` for your modem. This is not required; you can use a different port. Just remember to make the correct substitution for the port name in your modifications.)

The `/etc/gettydefs` file needs to have two entries added to it for security reasons. These entries differ from the others by including the line condition `HUPCL`. `HUPCL` means Hang UP on Close, or disconnect the phone line whenever a user logs out. When the user logs off, the TRICEP hangs up the modem from its end. The lines that you need to add to `/etc/gettydefs` are:

```
mo_1200# B1200 # B1200 SANE TAB$3 HUPCL #
        \r\n\nTRICEP System V Remote\r\n\nlogin: #mo_300

mo_300# B300 # B300 SANE TAB$3 HUPCL #
        \r\n\nTRICEP System V Remote\r\n\nlogin: #mo_1200
```

This sets up the tty line specified in `inittab` for proper operation with a modem. You may specify either 1200 or 300 baud as the initial baud rate. Getty will switch to the alternate baud rate if it detects a framing error caused by non-matching baud rates.

You also need to change the minor device number of the tty line that you wish to use. Bit 7 must be set to indicate to the device driver that pin 5 (CLEAR TO SEND) will be used to enable the modem. As superuser, type

```
rm /dev/tty0
mknod /dev/tty0 c 0 129
```

If you are using some other tty line than tty0, simply add the number of the line to the minor device number used in this example (129). For example, to change the minor device number for a modem connected to /dev/tty7, use

```
mknod /dev/tty7 c 0 136
```

Once you have completed these configuration changes, you need to construct a cable.

4.9.1 Modem to TRICEP Cable

The modem to TRICEP cable differs radically from a terminal to TRICEP cable. This is because more signals are necessary to properly carry out modem communication. Two additional signals are used at the TRICEP end: CLEAR TO SEND and REQUEST TO SEND. In addition, the chassis ground is included. Here is a diagram showing the way the cables is built:

TRICEP	MODEM
1 (CHASSIS GROUND)	1 (CHASSIS GROUND)
2 (RECEIVE DATA)	3 (TRANSMIT DATA)
3 (TRANSMIT DATA)	2 (RECEIVE DATA)
4 (REQUEST TO SEND)	8 (CARRIER DETECT)
5 (CLEAR TO SEND)	20 (DATA TERMINAL READY)
7 (SIGNAL GROUND)	7 (SIGNAL GROUND)
20 (DATA TERMINAL READY)	6 (DATA SET READY)

Computers are normally set up to be DCE's, that is, Data Communication Equipment, so they can be used with terminals (DTE's). Of course, modems are set up the same way (as computers). So, when modems are connected to computers (or computers to computers), the transmit and receive lines must be crossed (pin2 to pin3, and pin 3 to pin 2). This is called a null-modem connection.

The modem's pin 6, DATA SET REAY line, drives the TRICEP's pin 20, DATA TERMINAL READY. Remember, pin 20 must be logic TRUE (high) before the SIO4 will allow communication, the same as for printers.

After you have activated a getty for the modem line (via inittab and telinit), getty will open the modem line, /dev/tty0 in this example. This brings the TRICEP's pin 5 (CLEAR TO SEND) high, enabling the modem to receive calls by activating pin 20 (DATA TERMINAL READY) on the modem.

The modem brings its CARRIER DETECT line (pin 8) high when it answers a call from another modem. This is connected to the TRICEP's pin 4, (REQUEST TO SEND), waking up getty so it can collect a login name. From here on, the process is identical to an ordinary user login.

The TRICEP drops its CLEAR TO SEND line, pin 5, when the user logs out or exits. This forces the modem to drop its carrier (disconnecting the remote modem) and hang up. This is done by the HUPCL flag in the gettydefs line for the modem.

If the carrier is lost inadvertently, the TRICEP will kill all process associated with the modem line and hang up. Once the all the processes are killed, the TRICEP prepares for the next caller.

NOTE: Additional voltages are provided at limited amperages on pins 8, 9 and 10 of each RS-232 connector on the TRICEP. These are intended for use with short haul modems and the like.

pin 8	+12V	(Limited to 4 ma)
pin 9	+12V	(Limited to 20 ma)
pin 10	-12V	(Limited to 40 ma)

The cu command, use to call up remote modems, has its baud rate configured by the file /usr/lib/uucp/L-Device. This means that trying to set the baud rate with the command

```
cu -s 1200 /dev/tty0
```

doesn't work because it is overridden by the L-Device file. This is intended to insulate users from needing to know the correct baud rate.

5. MORROW ENHANCEMENTS

Your TRICEP system includes software that is not provided by other System V Unix ports. This software is part of the bridge between your TRICEP and the world of single-user microcomputers. Currently, this bridge includes file transfer programs. The ability to run MS-DOS programs on TRICEP slaves will be added soon, with the release of the 80188 slave.

Other programs have been added to your TRICEP system. The `ddt` command (disk debugging tool) allows the editing of non-ascii files. The `cptree` command copies subdirectories from one part of a file system to another. Documentation on these and other programs (as they are added) is found in the Morrow appendix to the User's Manual (also `/usr/man/local/man1` directory).

5.1 Diskette Specialists: far and dar

The two most popular operating systems for microcomputers are CP/M and MS-DOS. Each of these operating systems uses its own specific disk organization, making disks containing files from one operating system unusable by the other. These disks are also unusable to almost everything else but a particular brand of microcomputer and its operating system.

There is a way of making such disks usable, by combining hardware that can read different formats with software that understands different operating systems. The TRICEP programs, `far` and `dar`, understand CP/M and MS-DOS disk organization. And, the TRICEP floppy disk controller (and device drivers) can read and write disks formatted by:

- Micro Decisions using CP/M 2.2 or 3.0
- IBM PC's and other compatibles using MS-DOS 1.0, 2.0, 2.1

The driver software attempts to decide whether the disk is double or single sided, so that this is not a concern of the user.

Both `far` and `dar` use a simple command line interface to replace, extract, show a directory and delete files. The arguments for both programs are similar, making them easier to learn and use. The similar arguments will be explained first, followed by the differences in each program.

5.1.1 far and dar: Command Line Syntax

As mentioned, both far and dar share most of their syntax. The format for these commands is

```
far device_name {arguments} {files}
dar device_name {arguments} {files}
```

The device_name is the name of a special file in the device directory, usually /dev/dj0. /dev/dj0 is the name of the first 5 1/4" floppy disk drive.

You do not have to include the /dev directory name when using far and dar. These commands prepend the directory name /dev to device_names if necessary. For example,

```
far /dev/dj0
far dj0
```

are equivalent.

The default argument when no other arguments are present is t, for show directory (table of contents). Thus, the previous far examples would display the directory of a CP/M diskette in the first 5 1/4" drive (/dev/dj0).

There are four other common arguments to far and dar: create, delete, extract and replace. Only one of r, d, or x arguments (or t) may be used at a time. That is, you cannot extract and replace a file with the same command. The actions of these four arguments follows:

- c create an empty directory on the diskette; this deletes all the files listed in the diskette directory; the c option works together with the d and r options to create a new CP/M or MS-DOS directory where (possibly) no directory existed before. This should be used with caution because it ignores the current diskette format. The c option used without d or r has no effect.
- d delete the named files.
- x extract the named files from the diskette.
- r replace the named files to the diskette.

The verbose argument can be used in combination with any of the other five arguments. Verbose means just what it implies: be verbose (talky) while performing the action. The verbose option also provides a summary of the space remaining on the diskette.

The {files} argument means all the files on the diskette (or in the directory) when no filenames are present. In other words,

```
far dj0 x
```

means copy all the files from a CP/M diskette to the current directory, and

```
dar dj0 r
```

means copy all the files from the current directory to a MS-DOS diskette. Otherwise, far and dar operate only on the filenames given, as in

```
dar dj0 d doc.bak
```

that removes the file doc.bak on a MS-DOS diskette.

Expansion metacharacters, * and ?, must be used with caution, because the shell will attempt to expand them before passing them to the far or dar commands. If you want far or dar to expand a filename, for example, *.com, you must protect it from the shell by quoting it.

```
far dj0 x "*.com"
```

copies all the .com files from a CP/M diskette, and

```
far dj0 x *.com
```

copies all the .com files that already exist in the current directory AND on the floppy to the current directory.

5.1.2 far Differences

Since far deals with CP/M diskettes, it works somewhat differently than dar. The major difference is that CP/M has 16 user areas for subdividing diskettes.

The user option sets up the far command so that it works in a particular user area. For example,

```
far dj0 xu3
```

extracts all the files in user area 3 of a CP/M diskette. The default user area is 0.

CP/M filename conventions specify that filenames can be up to 8 characters long, with a three character extension. Small characters are converted to capitol letters, so the filenames

```
CHAP31.BAK  
chap31.BAK
```

are identical to CP/M. Besides capitol letters and digits, CP/M allows the following punctuation characters:

! @ # \$ % ^ & () _ - + ~ ` ' " | { } \ /

far will automatically convert lower case letters to upper case, truncate filenames longer than 8 characters and extensions longer than three, and remove directory names. For example, /user/rik/enhancements.text would become ENHANCEM.TEX in a CP/M directory.

5.1.3 dar Differences

MS-DOS uses a different strategy for organizing files on a disk. MS-DOS uses a system of subdirectories very much like the UNIX file system. The dar command does not provide a special argument for creating subdirectories, but will create them as needed.

As an example, if we copy a file to a DOS diskette with the command

```
dar dj0 r /users/rik/enhancement.text
```

dar creates two directories on the diskette: \USERS and \USERS\RIK. Notice that MS-DOS uses a backslash (\) to separate directories. dar handles the conversion between a slash and a backslash, so you will always use a slash.

Then, the file ENHANCEM.TEX is created in the \USERS\RIK directory. MS-DOS has the same limits on filenames as CP/M (8 characters for the name, three for the extension), and allows the use of the following characters in filenames:

A-Z	0-9	\$	&	#	%	'	()	-
@	^	{	}	~	`	!			

5.2 ddt, Disk Debugging Tool

Although this program takes its name from an old CP/M program, its use is somewhat different. ddt allows the viewing and editing of non-text files. This is useful when you need to change several bytes in an already existing file that cannot be handled by a text editor.

6. TRICEP SYSTEM SPECIFIC VALUES

Anyone who gets a System V UNIX port gets a standardized kernel with one exception. The exception is the device drivers, the software that controls and communicates with devices: disks, terminals, memory, printers and the clock. The implementation of device drivers has much to do with the performance of a UNIX port. Poorly written device drivers slow down UNIX, because the device drivers are among the most often used parts of the kernel.

The device drivers for the TRICEP were rewritten by Morrow to increase speed and add flexibility. Generally speaking, device drivers are based on a generic device driver that is edited to fit the hardware. The TRICEP device drivers were written specifically to take advantage of the built-in intelligence of the TRICEP devices.

The hard disk controller (HDC-DMA), the floppy disk controller (DJ-DMA) and the terminal controller (SIO4-DMA) all have separate processors that unburden the main CPU. Each peripheral CPU handles device I/O in parallel with the main CPU, uses DMA to transfer information to/from main memory, and signals completion to the main CPU by interrupts.

The connection between the rest of the kernel and the device drivers is made through the device switches. There are two device switches, one for block devices (disks) and one for character devices (terminals, printers, etc.). The device switch structure definitions are contained in the file /usr/src/sys/conf.c, for those of you who have purchased reconfiguration licenses.

6.1 Major and Minor Device Numbers

The UNIX system accesses the devices through the switch table by means of special files in the /dev directory. These files are block or character special files. Block or character special files are indicated in long directory listings by having a b (block) or c (character) at the beginning of the permissions:

```
ls -l /dev/mw0a /dev/console
crw----- 2 root    0,  0 Oct 3 13:45 /dev/console
brw----- 1 root    0,  0 Oct 3 13:47 /dev/mw0a
```

The two numbers preceding the date are the major and minor device numbers. Special files do not have any data blocks. Instead, they are used to hold (in the inode) the major and minor device numbers. The major device number selects one particular device driver from the block or the character device switch. The minor device number is passed to the device driver and contains information used by the driver, for example, which of several drives or terminals to access, or which disk partition to use.

Special files are created with the `mknod` command. Your system comes with most of the special files that you will need in the `/dev` directory. The next several sections define the major and minor device numbers for the different devices.

If you ever lose a special file, or need to add one, the `mknod` command is used to create special device files. The syntax for this command is

```
mknod device_name b-c major_device_number minor_device_number
```

where `b` stands for block special and `c` for character special. If you needed to create a new block device, for example, the fourth hard disk used as a file system without a swap space, you would use the command

```
mknod /dev/mw3c b 0 26
```

The definition and meanings of the major and minor device numbers are explained in the next several sections. The table at the end of this chapter provides most of the information you are likely to need for creating special files.

6.1.1 TTY Devices (SIO4-DMA)

The `tty` devices are character devices and have major device number 0. The minor device number has two purposes: to select one of eight possible ports and to enable/disable pin 5 (Request to Send). The least significant three bits select the port. Bit seven enables Request to Send when it is a 1, and disables it when it is zero. (See Adding Modems). The command

```
mknod /dev/console c 0 0
```

was used to create the console device, and

```
mknod /dev/modem c 0 129
```


creates a modem device using port `tty0` and Request to Send enabled (129 = 1000 0001). The console is the first port (number 0), `tty0` the second port (number 1), and `tty6` the seventh port (number 7).

6.1.2 Mini Winchester Devices (HDC-DMA)

The hard disks (mini winchesters) have a block major device number of zero. This does not conflict with the `tty` device, since the `tty` major device number refers to the character device switch. The minor device number for hard disks has two purposes: to select one of four possible hard disks and select a partition on the disk.

Hard disks are traditionally partitioned on cylinder boundaries. A partition is a logical (non-physical) division of a disk. A cylinder is the collection of tracks on a disk that may be accessed by one of the two to eight read/write heads without stepping.

There are six partitions built-into the hard disk device driver. They are described in the following diagram:



The designation of partition "a" as a file system is purely arbitrary: there is nothing special in the device driver requiring this partition to be a file system. Likewise, the "swap", "alt" (alternate blocks) and "boot" (boot loader) are just names. This is the way your root hard disk is currently set up. The following table is based on the illustration of partitions, and includes the number that is built into the minor device number for each partition.

a	0	File system with swap space reserved
b	1	Swap space (at least two megabytes)
c	2	File system without a reserved swap space
d	3	reserved for future use
e	4	reserved for future use
f	5	Boot loader space (at least 50 blocks)
g	6	Alternate blocks space (at least 50 blocks)
h	7	Whole drive - used for formatting, copying

The location of the partitions is written in the boot block, block 0, when a disk is formatted. This is not typical: most UNIX systems compile the partitions into the kernel. Having the partitions written in the boot block allows the device driver to determine the partitions dependent upon the size of the hard disk. Otherwise, a special kernel (or minor device number) would be needed for each different hard disk size.

The structure that defines the boot block is found in the file /usr/src/sys/mw.h (reconfigurable systems only).

The minor device number for the hard disk drivers is built as follows:

BIT	7	6	5	4	3	2	1	0
	Write-Block 0	X	X	Drive number			Disk partition	

Most mw devices do not allow writing to block 0 (the boot and configuration block). Setting bit 7 to 1, as in the device /dev/mw0hw0, allows writing in block 0. This protects block 0 from accidental overwriting. The disk formatting programs (diskformat and fmw) have private access to block 0.

Bits 3 and 4 determine which of four possible drives to access. Bits 0 through 2 select one of the 5 possible partitions to use. Bits 5 and 6 are not used.

There is also a raw hard disk device, major character device number 4. The raw, as opposed to cooked, device transfers data directly between the disk and user memory, bypassing the block buffers. This is useful for system programs that perform block reads and writes (such as dcopy and dd). The same minor device numbers apply to the raw device. Use of the raw device locks a process into memory during disk transfers, and should only be used for copying disks.

6.1.3 Floppy Disks (DJ-DMA)

The block major device number for floppies is 1. The minor device number has two purposes: to select one of eight possible drives and select one of two partitions. The least significant three bits select one of eight possible drives: 0 to 3, the 5 1/4" drives, and 4 to 7, the 8" drives.

The dj device has a special partition, indicated by setting bit 7 to 1, used specifically with the tar program. The tar program was designed for tape archiving and doesn't know that it shouldn't overwrite the boot block of 5 1/4" floppies. The boot block is special because it contains information used to determine single or double sidedness for 5 1/4 inch CP/M and MS-DOS format floppies. The special tar dj device, indicated by the letter "a", (e.g., /dev/dj0a), is ONLY for use with tar and 5 1/4" floppies. The regular partition, including the entire diskette, is used for all other applications.

There is also a character special floppy device with major device number 5. The raw floppy device is used mainly for formatting (diskformat). The minor device numbers work the same as for the cooked (block) device.

Master Device Table

There are several other device drivers used in the TRICEP. These are for accessing memory, kernel memory, the clock and the parallel port. There is also the null device, a bit bucket used for throwing away unwanted data. The table that follows defines these devices, along with the devices already mentioned.

Device	Type	Major	Minor	Description
console	c	0	0	system console
<u>dj0</u>	b	1	0	<u>first 5 1/4" floppy</u>
dj0a	b	1	128	first 5 1/4" floppy for tar only
dj1	b	1	1	second 5 1/4" floppy
dj1a	b	1	129	second 5 1/4" floppy for tar only
dj2	b	1	2	third 5 1/4" floppy
dj2a	b	1	130	third 5 1/4" floppy for tar only
dj3	b	1	3	fourth 5 1/4" floppy
dj3a	b	1	131	fourth 5 1/4" floppy for tar only
dj4	b	1	4	first 8" floppy
dj5	b	1	5	second 8" floppy
dj6	b	1	6	third 8" floppy
dj7	b	1	7	fourth 8" floppy
error	c	3	0	used by kernel
kmem	c	2	1	kernel memory
mem	c	2	0	memory
mw0a	b	0	0	root file system
<u>mw0h</u>	b	0	7	<u>whole first hard disk</u>
mw0hw0	b	0	135	like mw0h & enable write boot block
mw1a	b	0	8	second hard disk file system
mw1b	b	0	9	second hard disk swap space
(mw1c)	b	0	10	second hard disk, f.s. w/o swap
mw1f	b	0	13	second hard disk, boot cylinder
mw1g	b	0	14	second hard disk, alternate blocks
<u>mw1h</u>	b	0	15	<u>whole second hard disk</u>
mw1hw0	b	0	143	like mw1h & enable write boot block
mw2a	b	0	16	third hard disk file system
mw2b	b	0	17	third hard disk swap space
mw2c	b	0	18	third hard disk, f.s. w/o swap
mw2f	b	0	21	third hard disk, boot cylinder
mw2g	b	0	22	third hard disk, alternate blocks
mw2h	b	0	23	whole third hard disk
mw2hw0	b	0	151	like mw2h & enable write boot block
mw3a	b	0	24	fourth hard disk file system
mw3b	b	0	25	fourth hard disk swap space
mw3c	b	0	26	fourth hard disk, f.s. w/o swap
mw3f	b	0	29	fourth hard disk, boot cylinder
mw3g	b	0	30	fourth hard disk, alternate blocks
mw3h	b	0	31	whole fourth hard disk
mw3hw0	b	0	159	like mw3h & enable write boot block
null	c	2	2	null device, throw things away here
<u>rdj0</u>	c	5	0	<u>raw first 5 1/4" floppy (format)</u>
<u>rdj4</u>	c	5	4	<u>raw first 8" floppy (for format)</u>
rmw0a	c	4	0	first h.d. file system (copy f.s.)
rmw0h	c	4	7	whole first hd (copy disk)
rmw0hw0	c	4	135	whole first hd & enable boot write

Reserve mounted as →

mkfs 15119 89

*Reserve (Anper) Info 83
Disk Format /dev/rmw1h - size 1024 - cyl 0-319 - head 0-5
- discs 256 - il 4*

rmw1a	c	4	8	second hd file system (copy f.s.)
<u>rmw1h</u>	c	4	15	second whole hd (copy disk)
rmw1hw0	c	4	143	whole second hd & enable boot wrt
rmw2a	c	4	16	third hd file system (copy f.s.)
rmw2h	c	4	23	whole third hd (copy disk)
rmw2hw0	c	4	151	whole third hd & enable boot wrt
rmw3a	c	4	24	fourth hd file system (copy f.s.)
rmw3h	c	4	31	whole fourth hd (disk copy)
rmw3hw0	c	4	159	whole fourth hd & enable boot wrt
swap	b	0	1	swap device (only use one)
syscon	c	0	0	linked to console
tty0	c	0	1	second serial port
tty0m	c	0	129	second serial port as modem
tty1	c	0	2	third serial port
tty2	c	0	3	fourth serial port
tty3	c	0	4	fifth serial port
tty4	c	0	5	sixth serial port
tty6	c	0	6	seventh serial port
tty7	c	0	7	eighth serial port

Some of these devices do not currently exist in your device directory. And, you don't necessarily need them. If you add hard or floppy disks, for example, you will probably need to add a new device name. Simply find the name (or description) that matches your new device, change directory to /dev (cd /dev), and use the first four columns of this table for arguments to the mknod command. For example, to add a third hard disk device (that doesn't need a swap space), you would type

```
cd /dev
mknod mw2c b 0 18
```

to create the new device name.

All file system devices (device names beginning with "mw" or "rmw") must be owned by check for security reasons. The file access permissions for these devices must be read and write by owner only. To make these changes to our example, use these commands:

```
chown check mw2c
chmod 600 mw2c
```

These commands make the file system on mw2c accessible through the operating system only. Otherwise, a user could accidently (or maliciously) access the file system directly.

6.2 Every Day Procedures

This section is intended to guide you through every day activities, like

booting

going multi-user, running fsck and entering the date

checking for free disk space, policing the disk

going single-user and backing up

system shutdown

Each section that follows covers one of these topics. Your daily activities may vary from this; this is intended to cover a wide variety of installations.

Some systems may be running 24 hours a day; others may be shut down every evening and restarted in the morning. The things that everyone will have in common is a need to: check for free disk space and a regular backup routine.

6.2.1 Booting

Booting was covered in considerable detail in Step 7 of the Installation section. The only difference you will notice is that the system automatically resets itself when you turn the power on. To refresh your memory, you

turn on the console, and any external hard disks

turn on the TRICEP (automatically resetting it)

press RETURN to continue Standalone boot

press RETURN to begin execution

If you are rebooting the TRICEP after a system crash, you will not be turning on the power, but turning the keyswitch to RESET. And, you MUST run fsck everytime the system is reset after a crash. Your system may need to be reset because of a hardware failure. The consequences will probably be minor as long as you remember to use fsck.

6.2.2 Going Multi-User

TRICEP systems are delivered so that they are in single-user mode after booting. This can be changed by modifying the /etc/inittab file (see Modified Inittab Example).

You change from single-user mode to multi-user mode by typing

```
telinit 2
```

This sends a message to the `init` process to change to run level 2. Now, in the unmodified `/etc/inittab` file, this run level matches every entry. After you have modified `inittab`, this run level will probably match the multi-user mode, although this is merely custom, and not a default. If `telinit 2` doesn't wake up the other terminals (with a sign-on prompt from `gettydefs`), look at your `/etc/inittab` file and read the section named `Inittab Controls Init`. Run level 2 starts the daemons in `/etc/rc`.

6.2.3 Setting the Date

Going multi-user starts the `/etc/bcheckrc` script. This script displays the date (or what the TRICEP believes to be the date) and allows you to agree or disagree. Please check this date carefully, as it is used in backups and other important system functions.

The date is displayed as the day of the week, the first three letters of the month and the date, and the time in 24 hour notation. In 24 hour notation, you subtract 12 from the hour to convert to "ordinary" time. For example, in

```
date
Fri Aug 24 16:44:23 1984
```

16:44:23 is 4:44 and 23 seconds, P.M. Subtract 12 from 16 to get 4 P.M.

Entering a date can be a little confusing too. The date command likes the new date in one particular format, a series of pairs of digits. Each pair represents one part of the date, as in

```
Aug 24 16:44
 \ | / /
  08241644
```

The month comes first (January is 01, February is 02, and December is 12), followed by the date, the hour (in 24 hour notation) and the minutes. You must represent each part of the date with a pair of digits, and you must have all four pairs. Optionally, you can include the last two digits of the year after the other four pairs.

You will need to reset the date everytime the power is turned back on. It is considered bad practice to reset the date while in multi-user mode, so the `/etc/bcheckrc` script is the perfect time to do it.

6.2.4 Running FSCK

The `/etc/bcheckrc` script also asks you if you wish to check your file systems. The answer is YES, YES, YES! This is so important, a little exaggeration is permissible.

`fsck` can detect minor deficiencies in the file system before they grow into major ones. There are two reasons for running `fsck` before going multi-user: first, the file system that you are checking must be idle, and second, because damage to the file system occurs most often from incorrect shutdown procedures.

If you run `fsck` on an active file system, files that are modified while `fsck` is running will make `fsck` think that things are amiss. This can have unpleasant results, as `fsck` fixes things that aren't broken. `fsck` also may request that you halt the system and reset without sync. This means exactly that. Turn the keyswitch to RESET and reboot. Otherwise you will undo the work that `fsck` has just done. You can see that this could easily upset people if you were in multi-user mode and needed to reboot.

Damage to the file system occurs from two causes: improper shutdown and hardware failures. Improper shutdown means that the human in charge was lazy or forgetful, and shut the system off without using `/etc/shutdown` while in multi-user, or sync; sync if in single-user mode. Hardware failure either happens or it doesn't, but proper system shutdown can be controlled.

If `fsck` asks you for a yes or no (y or n) response, always respond y, unless you KNOW better. The worst that `fsck` will do to you is to remove two files when only one is incorrect (when both files share a block, that is, a duplicate block). The file that was last updated probably is correct, but `fsck` can't be sure. The best thing to do is replace both files from backups.

Is there ever a time when you shouldn't run `fsck`? Yes. If you suspect that your system has hardware problems with the hard disk (you keep getting disk error messages on the console), and `fsck` discovers lots of problems, stop `fsck` and have your hard disk checked. `fsck` can make things worse if the problem is in the hardware and not in the file system, by trying to fix non-existent problems.

6.2.5 Checking for Free Disk Space

One of the most common problems with any computer systems is running out of disk space. UNIX has lots of ways to create files. The accounting programs, if activated, create lots of files automatically. Users, if unchecked, will quickly overrun the system with things like `a.out`, `test`, `tempfile`, and a host of other files with names that look much more important. How's a poor system administrator to keep up?

Well, there are several tools at your disposal. Since prevention is the best medicine, you can keep track of the amount of free space available on the hard disk with the `df` (disk free) command. This simply reports the number of free disk blocks remaining.

```
df /dev/mw0a
2844
```

2844 blocks translates into 2.844 megabytes remaining. (512 byte systems cut this in half.) Bell Labs recommends that the root device have at least 2 megabytes free for temporary files, and that the file system with user directories have at least 1 megabyte per user free. This is pretty unreasonable, especially if you have a single 16 megabyte drive with both the root and user directories resident. You can strive for at least .5 megabytes free space (500 blocks) for each user, with a minimum of 1 megabyte total, and you will remain trouble free. If you have less space than this, check for free space often (hourly).

The second set of tools are `du` and `find`. `du` stands for disk usage, and reports on the number of blocks in files or directories. Keep a list of user home directories in a file, for example, `/etc/homes`. You can use this list with `du` to see how much space each user has consumed:

```
du -s `cat /etc/homes`
445 /user/bob
223 /user/rik
1022 /user/len
345 /user/kevin
701 /user/norm
122 /user/gayle
```

The file `/etc/homes` has the list of home directories in it, one per line. The `du -s` command produces a summary of blocks used by each user. You can talk to users who have taken up more than their share and ask them which files they want you to back up and remove. If they are truculent (and unyielding), you can be obnoxious, too. You, as system administrator, can search out and remove (after backing up, of course) old files.

The `find` command is ideally suited for uncovering old and unused files. For example,

```
find /user/joe -atime +30 -print > joes.oldfiles
```

creates a file containing the names of files in the `/user/joe` subtree that haven't been accessed in the last month. You can be nice and point this out to Joe, and you can be tyrannical, and backup and remove the files (using `tar` and `rm`) if nothing else works.

The find example in the User's Manual is a classic. This clever command searches the entire file system for files named either a.out (created by cc) or core (created by program crashes) that haven't been accessed in a week. The names a.out and core have certainly lost their meaning after a week, so they are true space wasters. Use

```
find / \(-name a.out -o -name core\) -atime +7 -exec rm {}\;
```

to find and remove them.

Find can also be used to find large files (with -size), files owned by particular user (-user) and files with dangerous permission bits set (like set user-id to root).

6.2.6 Going Single-User and Backing Up

Backing up is the second most important duty of the system administrator. (Running fsck is the first). Running any computer without a complete set of backups is asking for trouble. Even if the hardware never fouls up, some human will accidentally erase a file and come crying for help. Most computer centers discourage this tendency (humans making mistakes) by punishing such behavior with stiff fees. But, whether you are running a business or a school, backups are a must.

Before you can make good backups, you must make the TRICEP go into Single-User mode. Active file systems are changing file systems. If the system is multi-user when you backup, you can easily be backing up a file while it is changing. A backup of a file that changed is useless. At some time convenient to almost everyone (probably inconvenient to you), make the system single user with the shutdown command (used from the root directory while logged in as root):

```
shutdown
```

This takes about a minute to complete. The init program tries to be polite and gives everyone 20 seconds to finish their business. If you consider this extreme, use a wall (write all) to warn folks amount the imminent shutdown:

```
wall System going down for backup in 5 minutes!!<control-D>  
System going down for backup in 5 minutes!!
```

The <control-D> terminates the message.

When you see the message SINGLE USER MODE on the console, you are ready to begin backing up. This is a two stage process: selecting the files to backup and copying the selected files.

The find command selects the files on the basis of the last time that the file was modified (written to). There are at least two ways to do this.

```
find / -mtime -1 -print > tar.8.24
find / -newer tar.8.23 > tar.8.24
```

The first find command selects any file that has been modified within the last 24 hours. The -mtime -1 means modified in less than 1 day. If you want to increase the margin, substitute larger values for the 1.

The second find command selects files modified since the last time you used the find command to select files. This obviously won't work the first time you do an incremental backup, but is better for subsequent backups. The -newer option goes back to exactly the time of last backup, whereas -mtime only goes back exactly a day at a time.

Once you have your list of files, you need to copy them to your backup medium. For floppies, use

```
tar cvfB /dev/dj0a 798 `cat tar.8.24`
```

This creates tar archive floppies. Date these diskettes, please.

After you have been backing up your system for a while with tar, you need to consolidate your backups. You do this by using find to collect the names of all files modified in a longer period, say a week, two weeks or a month. For example,

```
find / -newer tar.8.23 -print > bigtar.9.23
```

collects all the files modified since tar.8.23 was created. This consolidates all of the backed up tar diskettes by including the latest version of all backed up files in one tar sequence. Now you can start reusing your daily tar backup diskettes. Remember to label diskettes with the date the backup was made, the format (tar) and the volume number.

The original backup, your monthly backups and your daily backups comprise your entire file system in most recently backed up form. If you need to restore your system from scratch, you start with the original tar diskettes (all 40), then use the latest months backup, and finally, the daily backups since the last monthly backup. This will restore your file system to the state it was in at the point of the last daily backup.

```
cd /
tar xvf /dev/dj0a
```

restores files to their position in the file system.

6.2.7 Turning the System Off

There are two ways to go about this, depending on whether you are in multi-user or single-user mode. In single-user mode, you can simply type

sync;sync

wait about ten seconds (the hard disk should be quiet), and turn off the power with the keyswitch, turn off other peripherals (hard disks, printers, terminals, etc.) after turning off the TRICEP.

When you are in multi-user mode, you need to gracefully shutdown the system before you can type your syncs. The /etc/shutdown script will do this for you. Simply type

shutdown

while logged in as root and in the root directory, answer "y" when requested and wait for the SINGLE USER MODE message. Then you can either backup the system (as in the previous section) or type

sync;sync

and turn the power off.

7. USING THE REST OF THE MANUALS

We don't really have to tell you that you now have access to the somewhat overwhelming resources of UNIX. You already know that. What you probably would like to know is how to learn about what you have. The TRICEP portion of the manuals provides only a brief glimpse of what is available.

This section of the TRICEP manual is a guide to the rest of the documentation. It is intended to get you started on the road to becoming a UNIX guru (one who really understands UNIX). But, you must be warned first that the material presented in the UniSoft documentation is not sufficient for learning UNIX.

As mentioned previously, the UniSoft documentation was written mostly by people at Bell Labs and UC Berkeley. This documentation was written by members of programming staffs for other programmers familiar with the UNIX system. Thus, there is a type of "catch 22" involved: this documentation can help you learn about UNIX if you already know about UNIX.

The documentation takes two forms: manual entries and papers that expand on manual entries. There is not a wealth of examples. There are instead concise definitions of commands, files and system calls. Once you have a grasp of how things work, this becomes useful. At first, and even later, the manual entries can be headache provoking.

We are not apologising for this. If the documentation were written in a more didactic style, it would expand to encyclopedic size. Instead of 6 inches of dense documentation, you would have 6 feet of readable documentation piled up on your shelf. This wouldn't help you very much. It's still too much to read.

The TRICEP manual and the UniSoft User Guide are both written in a terse but generally easy to understand form. We recommend that you start by reading both of these manuals, and trying the examples provided.

The TRICEP manual tries to follow a "cookbook" style approach: we give you the exact commands needed to perform a specific task, backup, for example. There is also some tutorial material designed to quickly give you a feel for the UNIX system.

The User Guide explains many of the basic commands used in UNIX. But only the section on `ed` provides you with hands-on working examples. If you don't find working with the User Guide enough help, buy one of the many books written about the UNIX system. Some of these books, those listed later in this section, provide many examples of how to do things. This hands-on approach makes UNIX easier to learn.

Later, you will find that the manual entries and the supporting papers sufficient for your use.

7.1 Divisions in the Documentation

There are eight divisions in the documentation:

- o TRICEP Installation and Maintenance Guide
- o User's Guide
- o User's Manual
- o Administrator's Guide
- o Administrators Manual
- o Programming Guide
- o Support Tools Guide
- o Document Processing Guide

The Guides are papers written about various programs or features of the UNIX system. They are generally more comprehensive than the manual entries and often provide some examples. The Manuals (User and Administration) are a printed form of the information produced by the man command. They are generally terse, but sometimes provide interesting and useful examples. The manual pages are especially helpful for discovering the correct options and syntax for the commands.

Physically speaking, you have received three pieces of documentation: two large binders and the TRICEP Installation and Maintenance Manual. One of the large binders contains the User's Guide and Manual. The Guide is the first section, the Manual pages form the second section.

The second binder contains the other five divisions of the documentation. These are presented in this order: Document Processing Guide, Programmers Guide, Support Tools Guide, Administrators Guide and Administrators Manual.

The material that follows explains what is defined in each division of the documentation.

7.1.1 TRICEP Installation and Maintenance Guide

This division of the manuals explains how to install your TRICEP. Included in the installation section is information on how to diagnose and correct hardware problems. The rest of this manual provides you with specific information on how to manage your TRICEP system. Much of this material is in the form of tutorials intended to give you a basic understanding of your system. The rest is in the form of commands that can be followed exactly, or with some file name substitution, so that the system can be managed without needing to understand it.

7.1.2 User Guide and Manuals

This is your next stop. The User Guide contains moderately simple descriptions of the basics of UNIX: the two main editors (ed and vi) and the two shells (sh and csh). Reading these papers gives you the information that you need to learn more. Remember, though, that reading alone won't teach you anything. It is best to read these sections while sitting in front of your terminal, logged in. Try everything that is suggested, even if an example is not provided. This may be a little frustrating, but produces the fastest results.

The User Manual consists of the manual pages for all of the User commands, system calls and subroutines, file formats and games. The first section on commands is the one that you will use the most. A permuted index is provided at the beginning of the commands section that can help you locate commands. Look in the middle column of the index for whatever you need. The last column gives you the name of the manual entry. The first column is actually a continuation of the middle column.

Entries in the Manual section often have a number appended after them. This number refers to the subsection of the manual that the entry belongs to. There are six sections to the User Manual:

1. Commands - programs generally accessible and used as tools, with definitions of the options and arguments required, and the type of results produced;
2. System calls - these are the subroutine calls for accessing and controlling the kernel, used by the commands and other application programs;
3. Subroutines - these are the subroutine calls available to the 'C' language, similar to system calls except they do not provide access to the kernel but provide other functions (string manipulation, math functions, i/o processing);
4. File formats - here are the definitions of the file formats used with some of the commands, such as the structure for the file system, gettydefs, and executable files;

5. Miscellaneous files - some information on networking, ascii characters, and, most important, a description of the termcaps file (terminal capabilities); and
6. Games - this is exactly what it sounds like, a listing of games that may be available on your system (unless some sourpuss removed them).

Following this scheme, you would know to look in section 4 (File formats) when you see something like inode(4) or a.out(4). Commands, like cp, ls or tar, are in the Commands Section, so they have a 1 appended to them (cp(1), ls(1) and tar(1)). Commands found in the Administrators Manual have 1M attached, as in fsck(1M) or mkfs(1M).

7.1.3 Document Processing Guide

Here you will find directions for using nroff and troff, the text formatting programs. These programs are intended to be used with particular printers and terminals for typesetting. In addition to the sections on nroff and troff, three special purpose macros are discussed: tbl, for tables, eqn, for equations and mm, for memorandum formatting.

7.1.4 Programming Guide

This division expounds mainly on the 'C' language. The 'C' language and its interface for the 68000 are defined here. The libraries, defined in the Subroutines section of the User Manual, are mentioned again. Lint, a clever program that checks 'C' syntax, is described. EFL, a preprocessor for FORTRAN, is defined also. (This must be used with a FORTRAN compiler that is not provided).

The information on the UNIX operating system in the back of this division will give you a quick overview of how UNIX works (as viewed from 'C' programs).

7.1.5 Support Tools Guide

There are nine tools described in this division: two calculator programs, four supplements to 'C' (or RATFOR) programs, two program development support tools and the communication program. Of these, I would rate dc (desk calculator), make (for compiling large programs) and awk (pattern scanning) as the most useful. The make program, for example, is used to define interrelationships between files and how to link modules together.

7.1.6 Administration Guide and Manual

This is actually two divisions: a guide section and manual entries. The guide section provides generalized information on administering to your system. This information is in addition to the more specific information in the TRICEP manual. However, the sections on fsck, lp, uucp and accounting go into much more detail than does the TRICEP manual. These four sections suffer from the typically dense writing style, but are important enough to read.

The Manual section is similar to the User Manual in that it contains man style entries. The entries in this section were chosen because they are NOT generally available to users. These programs, special files and procedures are limited to the superuser and other special users (like the lp administrator).

An M in a reference to a command, such as fsck(1M), means that the command's entry is in the Administrator's Manual.

NAME

udos - micro-DOS, MS/DOS 2.0 emulator for Tricep with SP-188 board(s)

SYNOPSIS

udos [options] file

DESCRIPTION

udos provides a system call interface consistent with Microsoft's DOS 2.0 operating system. udos also allows the use of some of the built-in commands that are a part of DOS. The intent of udos is to enable users to run software that is written for the popular MS/DOS operating system, while still enjoying the power and flexibility of the UNIX operating system.

Any program that uses the DOS system call interface will work with udos. Programs that are designed to work specifically with the IBM PC, that is, that make references to the ROM or the video-mapped memory for graphics, will NOT work with udos.

There are three options that are recognized by udos. Any other options (a single character preceded by a dash (-)) will be passed to the DOS program being invoked by udos.

- t This is the most important option. The "t" stands for terminal emulation. Including this option in the udos command line causes terminal control characters to be translated from those required by ADM31 terminals to the terminal specified in the TERM or TERMCAP environment variables.
- c This is a debugging flag, intended for creating a summary of DOS system calls processed by the simulator. A report is sent to the standard output when udos is terminating that lists the number of the system call (in hex) and the number of times the system call was made.
- v Another debugging flag, this one produces a display similar to the DOS debugger register display. One "-v" produces a display every system call, two -v's or three -v's result in more frequent displays. This may be useful to program developers, but is really a remnant from the development of udos.

If you wish to pass an option to a DOS program that uses the letters t, c or v, merely capitalize the letter and udos will ignore the option flag and pass the option to the DOS program as a command line option. For example,

udos conf -T

passes the option -T to the DOS program conf.com without activating terminal emulation. If "T" hadn't been capitalized, terminal emulation would be turned on, even though the option follows the DOS program name.

TERMINAL EMULATION

The terminal emulation ability of udos works most acceptably with programs that produce a relatively low volume of character output. Spreadsheet programs, like SuperCalc II, are low volume users of character I/O. Wordprocessing programs, like NewWord or WordStar, are examples of high volume users of character I/O. The terminal emulation ability built-into udos uses the library routines. These routines can result in a faster character output in the low volume I/O programs, and faster execution. In high volume I/O programs, the overhead in using terminal emulation will outstrip the performance gained by using the library routines.

There are two preconditions for using terminal emulation. The first is that the DOS program must be installed for an ADM31 terminal. This is a terminal type with most of the intelligent capabilities provided by newer terminals. The second requirement is that the TERM or TERMCAP environment variable be present. Morrow TRICEP systems are factory configured (in .setup or .cshrc files) to invoke the tset command and set the environment variables TERM and TERMCAP. If you have edited or removed these setup files from your home directory, you can restore them by copying the prototype files from /usr/lib.

DRIVE NAME ASSIGNMENTS

The DOS operating system provides for a hierarchical directory structure similar to UNIX. Things are complicated, however, because DOS is designed for use with floppy disk drives, and uses a combination of drive letters (like A:) and pathname specifications.

udos handles drive name assignments by mapping the drive letter to the pathname of a UNIX directory. For example, typing

```
udos
```

assigns drive A: to be your current directory. When you enter a drive name while using udos the real pathname transparently replaces the DOS drive name.

You can assign drive names to UNIX directories explicitly with the assignment operator (=). For example,


```
A>c:=\users\len\sc
```

assigns drive C: as the directory /users/len/sc. Notice that backslashes (also called "slants") were used in between directory names. This is a MS/DOS convention, but ordinary slashes will also work with udos.

udos will also set up drive letters A:, and possibly B:, automatically during invocation. If you invoke udos without specifying a program to run (interactive mode), drive A: will be set to your current directory. However, if you invoke udos with the name of a MS/DOS program that is not in your current directory, the directory containing the MS/DOS program will be assigned to drive A:, and your current directory will be assigned to drive B:. For example, typing

```
udos /usr/dos/sc/sc2
```

from your home directory will cause drive A: to be /usr/dos/sc, and drive B: to be your home directory. You will be logged onto drive B:. This allows program files (and overlays) that have the SYS attribute to be accessed while logged onto another drive, a feature of many MS/DOS programs.

The drive assignments are displayed automatically when you enter udos without a DOS program name (interactively). Current drive assignments can be displayed by typing an equal sign (=) while in interactive mode. Drive assignments may also be contained in dosrc files (see below).

BUILT-IN COMMANDS

udos implements some of the MS/DOS built-in commands. These commands are:

date	displays the MS/DOS date
time	displays the MS/DOS time
type	send a file to the standard output
del	remove a file
dir	display directory
cd	change to a subdirectory of a drive
ren	rename a file

There are several possible causes for confusion in this list. Date and time are displayed in MS/DOS format. But, you cannot set either the date or the time from udos. The "cd" command sets both the current directory AND the subdirectory of a drive name to be used. Thus, the following sequence

```
A>cd b:\other
A>
```

establishes the subdirectory `the B: drive` to be the current directory for THAT drive. Any references to drive B: then refer to the subdirectory (not the root directory).

The wild card metacharacters "*" and "?" are not implemented, so that the command

```
dir *.c
```

will not produce the expected result (a listing of the files in the current directory that end in .c).

The copy command is missing, but can be implemented using an alias. All of the commands associated with the processing of batch files are also missing. Users wishing to emulate MS/DOS batch files are advised to use the more powerful shells provided by UNIX. Redirection of output is also lacking from udos, but, once again, can be accomplished by invoking udos from a shell. For example,

```
udos program > file
```

redirects the output of "program" to "file".

ALIASES

An alias feature, similar to C-shell's, has been included with udos. Aliasing essentially works by replacing one string with another. In practice, this allows users to define short, mnemonic, strings to be replaced with longer and more complex ones. Argument substitution, a feature of aliasing under the C-shell, is not supported. The following list contains several examples of useful aliases:

```
alias ls !ls7 -F
alias copy !cp
alias ren !mv
alias mkdir !mkdir
alias rmdir !rmdir
alias n nw myfile
```

The first alias, `ls`, is replaced with the string `!ls7 -F` anytime `ls` is typed after the DOS prompt. The alias `"copy"` causes the UNIX command `"cp"` to be executed in a sub-shell. The next three examples show how you can replace DOS built-ins with UNIX commands. A warning is necessary here: when you invoke a shell (with an exclamation point), you can't use drive letters like `a:` as part of a filename.

The last example allows the execution of a 8 character command with two keystrokes. Aliases may be included in `dosrc` files.

DOSRC FILES

The udos program works in a manner similar to the other UNIX shells, and shares a feature common to both shells: runtime files. There are two runtime files that are used with udos: /usr/lib/dosrc and HOME/.dosrc. dosrc files may contain anything that can be typed after the command prompt (A>). The dosrc files may contain aliases, drive assignments and names of programs to execute. The commands in the file /usr/lib/dosrc are executed first. The commands in the user's own .dosrc file in his home directory are executed last, possibly overriding drive assignments or aliases made by the system-wide /usr/lib/dosrc file. These files are optional; there is no penalty if they don't exist.

COM, EXE and BAT EXTENSIONS

MS/DOS recognizes four filename extensions (three characters following a dot in a filename) as special. udos will recognize three of these four. The fourth extension, SYS, refers to special files that modify the operating system and won't work with udos.

When you issue a command to udos, you don't include the file extension. udos will automatically try appending the extensions .COM, .EXE and .BAT to the name of the command, in that order. If you specify a complete pathname (beginning with a /), only the specified directory will be searched. If a relative pathname is used, (one not beginning with a /), udos will search all the directories contained in the DOSPATH environment variable. Format of the DOSPATH environment variable is just like the PATH environment variable, a series of UNIX directory path names separated by colons.

INSTALLING SOFTWARE

Tricep systems built after January of 1985 include the /usr/dos directory. This is intended as the repository for directories containing MS/DOS software. The /usr/dos/sc subdirectory, for example, contains the programs and files used by the MS/DOS version of SuperCalc II. It is suggested that other software packages are also installed in subdirectories of /usr/dos, although this is not a requirement. The following script outlines the installation of NewWord.

```

cd /usr/dos
mkdir nw
cd nw
dar dj0 x
udos nwinstal
(install NewWord by selecting a terminal and printer)...
chmod 4444 *; chown bin *; chgrp bin *
chown bin . ; chgrp bin . ; chmod 555 .

```

The suggested procedure is to create a new subdirectory (/usr/dos/nw), change directory to it, extract the files from the MS/DOS diskette containing the programs (with dar), run the installation program (under udos) and change the mode and ownership of the files in this new directory.

The mode "4444" has a special significance to udos. Permission mode 4444 means "Read permission for all, and set the user-id". udos interprets the set user-id bit as designating a file as a "SYSTEM" file, which allows some programs to access files in the A: drive while logged onto another drive.

If re-installation of a program is necessary, you will need to change the permission mode to "Read and Write" (4644) before udos will allow alteration of the file. Although udos is owned by "root" and has the set user-id bit, it checks file permissions and does not allow writing to files that normally are off limits to the "real" user. In a file set to be Read-only for all, this even includes the superuser, root.

As noted in the section on Terminal Emulation, you must install software that will be used with the emulator for the ADM31 terminal. This allows one copy of a program to be used with several different terminals. If you are not using terminal emulation, a copy of the program must be installed for each type of terminal connected to the system.

Printer interfaces should be installed as the LST: device. If you are allowed graphics as an option, disable the graphics. There is a special case where you have connected an intelligent graphics device to your Tricep, and the program drives this particular device as if it were a printer. In this special case, graphics will work.

PRINTERS AND UDOS

MS/DOS programs that wish to direct output to printers should use the LST: device. The LST: device currently defaults as the /dev/cent device under UNIX. This is acceptable if you are the only user of a system that has a printer connected to the /dev/cent port. Otherwise, there will be trouble.

When multiple users direct output to /dev/cent, the output will be combined in an unpredictable fashion (garbage, in other words). Or, if you are using a serial printer, you will never see any printed output. The solution to this is to use the DOSPRN and DOSPRNBAUD environment variables.

The DOSPRN variable is set to the name of the file that you wish to redirect your LST: (printer output) to. This may be

the name of a serial port with a printer attached to it. In this case, you must also have set the environment variable DOSPRNBAUD to the correct baud rate for your printer. For example,

```
setenv DOSPRN=/dev/tty2
setenv DOSPRNBAUD=1200
```

establishes /dev/tty2 as the file to redirect the LST: output to, and attempts to set the baud rate of this port to 1200 baud for C-shell users. Bourne shell users need to use

```
DOSPRN=/dev/tty2
DOSPRNBAUD=1200
export DOSPRN DOSPRNBAUD
```

to accomplish the same effect. The C-shell commands can be added to the user's .login file, and the Bourne shell commands to the user's .profile file in their home directory.

If you do not want to send output directly to a printer, substitute a file name instead of a device name for DOSPRN. You can also add an alias to your .dosrc file to send the file to lpr and clear the file for new input. For example, if you have set DOSPRN to be dosprint in your home directory, the alias

```
alias print !(lp -c $HOME/dosprint; echo > $HOME/dosprint)
```

copies dosprint to the spool directory and then readies the file for more LST: output. Notice that this is not used for printing any file while under udos, but only for directing that the output sent to the LST: device be sent to the lp daemon for queuing and printing.

FILE NAMES

MS/DOS is a single case system. udos compensates for this by changing uppercase letters in pathnames to lowercase. This can result in unexpected difficulties if you normally mix upper and lower case letters in your file and directory names. For example, you wish to assign your directory /users/rik/TRI to be drive C: with the command

```
A>c=/users/rik/TRI
Select a directory for Drive C:
```

udos translated the letters "TRI" to "tri", so the assignment failed.

MS/DOS places other restrictions on file names. Only the first eight characters (up to the first .) are used, then

the first three characters (if any) after the dot. Longer filenames are truncated. As an example, suppose you wanted to access the file enhancements.text while using udos. The name gets truncated (twice!) and becomes enhancem.tex, which won't correspond to the original file name.

The set of characters permitted in MS/DOS filenames is:

```
A-Z      0-9      $      &      #      !      (
-      @      ^      {      }      ~      ,
```

Since some of these characters have a special meaning to the UNIX shells, care must be taken in creating filenames from udos. If you stick with names using only letters and numbers, you won't have any difficulties.

EMERGENCY ESCAPE

The interrupt character (normally DELETE) that is used under UNIX will not work under udos. MS/DOS programs use a variety of characters for special functions other than interrupting a runaway program. However, there is a special mechanism built into udos that will allow you to regain control of your terminal if a program has crashed.

Four tildes (~~~~) will terminate udos after a program crash. The four tildes may also terminate a running program if you can type enough of them fast enough. The catch is that udos waits until there are four tildes in its input buffer before terminating. If the program is still alive and calling for input, it may be difficult to get the four tildes into the buffer before the program reads (and clears) the buffer. Sometimes holding the ~ has the desired effect.

EXAMPLES

We will assume that you have correctly installed the MS/DOS program Super Calc II for an ADM31 terminal. Then, if you type the following from your home directory,

```
udos -t /usr/dos/sc/sc2
```

terminal emulation will be on, drive A: will be assigned to /usr/dos/sc, drive B: to your home directory, and you will be logged onto drive B:. Being "logged onto drive B:" in this case is equivalent to saying that your home directory is also your current directory.

FILES

```
/bin/udos
/usr/dos/lib/sp      (gets loaded into the SP188)
/usr/lib/dosrc      (system wide DOS rc file)
HOME/.dosrc        (user's own DOS rc file)
```

SEE ALSO

dar(1), csh(1), sh(1), chmod(1), mkdir(1), lp(1)

DIAGNOSTICS

If you don't have an SP188 slave and try to use udos you will get the diagnostic

udos requires an SP188 slave for operation.

If there are more instances of udos than there are sp-188 boards, then the message

No MSDOS cells available

appears.

BUGS

Perhaps calling these "bugs" is going too far. Wildcards and command line editing are missing from this version of udos. It is suggested that the user make use of the power of UNIX instead of trying to get a poor overworked slave to try to duplicate it.

```
echo /.profile
SHELL=/bin/sh
export SHELL
PATH=/bin:/usr/bin:/etc:/usr/lib:/usr/dos:
export PATH
TERM=vt1925
export TERM
DOSPRN=/dev/cent
export DOSPRN
. /.setup
cd /usr/dos/mba
udos -t hello
```

Shell for udos

PRINTER SET-UP MODE

FIRMWARE LEVEL: 3D

#01 CHARACTER SET (1-254)..... 101
<READY>=MODIFY <LF>=SCROLL <FF>=JUMP <FORMS>=SAVE ALL <RESET>=EXIT
<TEST>=PRINT PARAMETER LIST

EXIT FROM SET-UP MODE

16line
1995
2795

From len Tue Feb 12 06:49 PST 1985
To: belllabs benelux bestlimi bridgeco computer drakemic justcomp microbaz
micromas mrstevens pacificb pestcont promicro software
Subject: Caution / WARNING
Cc: gayl shirley tony

Caution: (I almost forgot)

Do not attempt to replace your booter (cp boot /dev/mw0f)
unless you have Tricep CPU Prom version 2.3 or greater.
The current version is 2.4 and it is not likely to change
again for quite some time.
The new prom can be obtained from our dealer support dept.
Ask for "triboot 2.4"

Len Edmondson
Morrow, Inc.

? usage
q quit
x exit without changing mail
p print
s [file] save (default mbox)
w [file] same without header
- print previous
d delete
+ next (no delete)
m [user] mail to user
! cmd execute cmd

? From len Tue Feb 12 06:34 PST 1985
To: belllabs benelux bestlimi bridgeco computer drakemic justcomp microbaz
micromas mrstevens pacificb pestcont promicro software
Subject: dj 1.11

We'd like to announce the release of a new floppy driver.
This driver utilizes track-at-a-time I/O and offers significant
performance improvement in many situations.

Numbers:
35 seconds to read all of a 5 inch. diskette.
That is 53% of the theoretical maximum (as
I figured it.)

~~Note:~~

To find out which version of the drivers you are using
issue the command

```
what /unix
```

A list of version numbers will be displayed.

To get your copy:

You will need 2 files:

```
/unix and /boot
```

I recommend the following series of commands:

(copy the files over)

```
take /boot /tmp/boot  
take /unix /tmp/unix
```

(copy the new booter to the "boot track")

```
cp /tmp/boot /dev/mw0f
```

(save your old operating system just in case)

```
mv /unix /oldunix
```

(move the new one into place)

```
mv /tmp/unix /unix
```

Keep a copy of the booter in a reasonable place

```
mv /tmp/boot /boot
```

Len Edmondson
Morrow, Inc.

? From len Tue Feb 12 06:18 PST 1985

To: belllabs benelux bestlimi bridgeco computer drakemic justcomp microbaz
micromas mrsteves pacificb pestcont promicro software

You may direct any inquiries regarding service/support of Morrow products
to gayl or tony via UNIX mail.

Len Edmondson
Morrow, Inc.

?

```
echo QUANTSAVE
echo formatting ...
/bin/diskformat /dev/rmw0h -size 1024 -cyl 0-511 -head 0-7 -dens 256 -il 4
echo determining bad blocks ...
/bin/badblk -w /dev/rmw0hw0
echo making file system ...
/etc/mkfs /dev/mw0a 34631 8 9
echo copying booter to hard disk ...
cp /dev/dj0f /dev/mw0f
echo mounting hard disk and making directories ...
mount /dev/mw0a /a
mkdir /a/bin
mkdir /a/etc
mkdir /a/dev
mkdir /a/a
mkdir /a/tmp
echo copying floppy files to hard disk ...
cp /unix /a
sync
cp /etc/init /a/etc
cp /etc/fsck /a/etc
cp /etc/mkfs /a/etc
sync
cp /bin/sh /a/bin
cp /bin/badblk /a/bin
cp /bin/cat /a/bin
cp /bin/tar /a/bin
cp /bin/patch /a/bin
cp /bin/ls /a/bin
sync
cp /bin/cp /a/bin
ln /a/bin/cp /a/bin/ln
ln /a/bin/cp /a/bin/mv
cp /bin/echo /a/bin
cp /bin/diskformat /a/bin
cp /bin/mkdir /a/bin
cp /bin/mknod /a/bin
sync
cp /dev/MAKEDEV /a/dev
echo making device files ...
cd /a/dev
sh MAKEDEV
sync
cd /
echo "using patch to change (presumably from dj-) to a mw-based kernel ..."
patch -w /a/unix _rootdev 0000
patch -w /a/unix _pipedev 0000
patch -w /a/unix _dumpdev 0001
patch -w /a/unix _swapdev 0001
patch -l /a/unix _swplo 0000
patch -l /a/unix _nswap 0800
sync
sync
umount /dev/mw0a
echo QUANTSAVE DONE!
```

```
echo CMI16SAVE
echo formatting ...
/bin/diskformat /dev/rmw0h -size 1024 -cyl 0-305 -head 0-5 -dens 128 -il 4
echo determining bad blocks ...
/bin/badblk -w /dev/rmw0hw0
echo making file system ...
/etc/mkfs /dev/mw0a 14363 8 9
echo copying booter to hard disk ...
cp /dev/dj0f /dev/mw0f
echo mounting hard disk and making directories ...
mount /dev/mw0a /a
mkdir /a/bin
mkdir /a/etc
mkdir /a/dev
mkdir /a/a
mkdir /a/tmp
echo copying floppy files to hard disk ...
cp /unix /a
sync
cp /etc/init /a/etc
cp /etc/fsck /a/etc
cp /etc/mkfs /a/etc
sync
cp /bin/sh /a/bin
cp /bin/badblk /a/bin
cp /bin/cat /a/bin
cp /bin/tar /a/bin
cp /bin/patch /a/bin
cp /bin/ls /a/bin
sync
cp /bin/cp /a/bin
ln /a/bin/cp /a/bin/ln
ln /a/bin/cp /a/bin/mv
cp /bin/echo /a/bin
cp /bin/diskformat /a/bin
cp /bin/mkdir /a/bin
cp /bin/mknod /a/bin
sync
cp /dev/MAKEDEV /a/dev
echo making device files ...
cd /a/dev
sh MAKEDEV
sync
cd /
echo "using patch to change (presumably from dj-) to a mw-based kernel ..."
patch -w /a/unix _rootdev 0000
patch -w /a/unix _pipedev 0000
patch -w /a/unix _dumpdev 0001
patch -w /a/unix _swapdev 0001
patch -l /a/unix _swplo 0000
patch -l /a/unix _nswap 0800
sync
sync
umount /dev/mw0a
echo CMI16SAVE DONE!
```

- /, 28
- /etc/passwd
 - The User Name File, 50
- 24 hour notation, 89
- accept (lp requests), 66
- access, 27
 - hard disk, 18
 - memory, 34
- accounting, 99
- action, 36
- active file system, 92
- adding file systems, 70
- address of memory, 17
- Administration Guide and Manual, 99
- alternate block area, 34
- alternate blocks, 84
- alternate sectors, 22
- amount of free space, 91
- anti-static rug, 8
- arguments
 - interface program, 64
- automatic RESET on power-on, 16
- background
 - printing, 62
 - processes, 44
- backspace, 45
- backup
 - entire system, 24
- backups
 - daily, 93
- badblocks, 22, 70
- baud rate, 48, 61
 - change, 44
 - console, 12
 - modem, 72
 - sign of wrong, 12
 - while booting, 17
- bcheckrc, 40, 89, 90
- Bell Labs, 33
- Berkeley UNIX, 33
- bin, 26, 28, 29
- binary, 29
- block 0, 26, 84
- block device
 - floppy, 85
- block
 - 0, 22
 - device, 29
 - devices, 83
 - numbers, 22
 - special files, 82
- blocks, 70, 91
- boards, 6
 - insertion, 6
 - loose, 6
- boot, 15
- boot action
 - init, 37
- BOOT UNIX (NO SYNC!), 23
- boot
 - code, 34
 - from floppy, 15
 - loader, 84
 - state, 37
- Booter code, 34
- Booting, 88
 - the process of, 16
- bootwait, 37
- Bourne shell, 47, 51
- brain damage, 8
- break key, 36, 44, 48
- brown-outs, 10
- C-shell, 47
- cable, 4
- cables, 6
- can't boot, 18
- cancel lp request, 67
- carriage returns, 45
- cat, 65
 - example, 91
- cd, 28
 - command, 52
 - example, 29, 30
- centronics, 62
 - cable, 4
- change baud rate, 44
- change baud rate on login, 36
- character
 - device, 83
 - special files, 82
- charactor
 - device, 29
- check
 - for bootable floppy disks, 34
 - memory, 34
 - of hard disk, 22
- Checking for Bad Blocks, 22
- Checking for Free Disk Space, 90

- checklist, 23
- child, 43
 - death, 43
 - of getty-login, 41
 - process, 44
- chmod
 - example, 53, 61, 87
- chown, 50
 - example, 56, 87
- class, 63
 - of printers, 62, 65
- CLEAR TO SEND, 73
- colon prompt, 34
- command
 - execution, 44
 - searchpath, 63
- commands
 - manual entries, 98
- computer to computer connection, 73
- configuration files, 33
- connections
 - terminal, 12
- console, 11, 39
 - device, 83
- contents of TRICEP carton, 4
- Controlling Shells: /etc/profile
 - /etc/cshrc, 47
- cooked device, 84, 85
- cooling, 8
- CP/M, 77, 85
 - legal filenames, 79
- CPU, 82
- crash, 19, 88
 - power problems, 10
- creating new directories, 71
- Creating the New HOME Directory, 55
- cshrc, 47
 - file, 36
- cu command, 74
- cursor, 14
- cylinder, 70, 84

- daemons, 40
 - starting, 89
- daily backups, 93
- damage to the file system, 90
- dar Differences, 79
- data bits
 - 8, 12
- DATA SET READY, 73
- Data Terminal Ready, 16, 60, 73
- database
 - user, 50

- date
 - example, 89
- death, 44
- DEC, 60
- default printer, 65
- dev, 26, 28, 30
- dev/console, 83
- dev/dj0, 77
- dev/dj0a, 25
- dev/tty0, 83
- device driver, 85
- device name, 70
- device
 - connections, 30
 - driver, 83, 84
 - drivers, 82
 - name table, 87
 - numbers, 82
- devices, 29, 83
- df
 - example, 91
- directory, 28, 55, 63, 70, 71, 91
 - home, 46, 47
 - names, 25, 27
- disable
 - printers, 67
- disabling lpsched, 66
- disk errors
 - fsck, 90
- disk
 - drive noises, 16
 - free space, 91
 - partitions, 82
 - usage, 91
- diskette
 - labels, 25
- diskettes, 24
- dj0, 77
- dj0a, 25
- DMA, 82
- Document Processing Guide, 98
- double-characters, 12
- double-sided, 24, 85
- DRAM, 17
 - memory settings, 17
- drive
 - partitions, 84
 - time-out, 34
- DTR, 60
- du
 - example, 32, 91
- duplex
 - full, 12
- duplicate block, 90

echo, 45, 62
 Editing /etc/group, 54
 Editing /etc/passwd, 53
 eight terminals, 41
 enable
 printer, 66
 printers, 67
 encrypted, 51
 environment, 47
 user, 46, 52
 Epson, 60
 erase temporary files, 40
 erratic performance, 8, 10
 errno 2, 41
 error
 read, 22
 etc, 26, 28, 30
 directory, 33
 etc/bcheckrc, 40, 89, 90
 etc/checklist, 23
 etc/cshrc, 47
 etc/gettydefs, 36, 44, 89
 etc/group, 50
 etc/initab, 89
 etc/inittab, 36, 88
 etc/mnttab, 71
 etc/passwd, 46, 51
 etc/passwd Controls login, 46
 etc/profile, 47
 etc/rc, 40, 67, 89
 etc/utmp data, 36
 etc/vchk, 23
 exec, 36, 43, 44, 46
 executable files, 29
 EXTA, 45
 extracting files from tar, 25

 fan, 10
 not working, 11
 far and dar: Command Line Syntax, 77
 far Differences, 78
 file system, 26, 28, 30, 84
 active, 90, 92
 adding, 70
 backup, 24
 security, 87
 file
 access permissions, 47
 descriptor, 27, 34
 sizes, 32
 filenames, 27
 CP/M, 79
 filter, 65
 final flags, 44

 find, 50
 example, 92, 93
 examples, 91
 uses etc/passwd, 50
 finding files, 91
 floppy
 8", 85
 floppy device driver, 85
 Floppy Disk Backup, 24
 floppy disk
 /dev/dj0, 77
 devices, 29
 during boot, 16
 floppy diskettes, 24
 Floppy Disks (DJ-DMA), 85
 foreign voltages, 10
 fork, 36, 43
 format, 24
 formats
 other floppy, 77
 formatting disks, 70
 formatting
 hard disk, 84
 framing error, 36
 free list, 70
 free space on disk, 91
 fsck, 23, 30, 88, 90, 99
 example, 70
 response, 90
 fsdb, 30
 fuser (find user), 72

 games, 32
 garbled characters, 12, 17
 getty, 36, 43, 44, 46
 error, 41
 getty-login-shell, 39, 41, 42
 gettydefs, 30, 33, 36, 44, 48, 72, 89
 gettydefs: Sets Up Baud Rate, 44
 gibberish on screen, 14
 Going Multi-User, 88
 Going Single-User and Backing Up, 92
 grounds, 9
 group, 50, 51, 54
 group id, 51
 guest user, 46

 handshaking, 60
 hard disk, 27, 84
 devices, 29
 dropped, 8
 failure, 18
 location, 8

- mounting, 71
- power-up, 18
- hardware failure, 88
- hardware problems, 90
- heads, 84
- HOME, 47, 52, 55
- HOME directory, 28, 36
- home
 - directories, 91
 - directory, 46, 50, 51, 56
- human readable files, 33
- humidity, 8
- HUPCL, 72

- incorrectly addressed, 17
- init, 36, 43
- init states, 37
- initdefault, 39, 41
- initial baud rate, 17, 48
- initial flags, 44
- inittab, 30, 33, 36, 40, 41, 42, 48, 61, 72, 89
- Inittab Controls INIT, 36
- Inittab Example, 40
- inode, 27, 34
- insertion of pc boards, 6
- Installation
 - summary, 3
- insufficient power, 18
- insufficient RAM, 17
- interface, 63, 64
 - programs, 65
- interleave, 70
- interlock
 - print, 60
- interrupts, 19, 34, 82

- jumper settings, 17

- kernel, 15, 30, 34, 37, 43, 84, 85
- keyboard doesn't work, 18
- kill, 36, 39, 42, 43, 74
 - character, 45

- L-Device file, 74
- lib, 26
- light
 - power on, 10
- linefeed, 45
- loader
 - boot, 84
- loading, 15
- loading fails, 18
- log off, 39

- login, 36, 41, 44, 46, 50, 56
 - file, 36
 - message, 44, 45
 - names, 30
- lost+found, 26
- lp, 62, 99
 - accept requests, 66
 - daemon, 40
- lpadmin, 63
 - example, 65
 - the lp Configuration Command, 62
- lpsched, 40, 66
- lpshut, 66
- lpstat, 67
- ls
 - uses etc/passwd, 50
- ls7, 28
 - example, 29

- M, 99
 - after command names, 98
- major device number, 82, 83, 84, 85
- make file system, 70
- manual pages, 32
- memory, 15
- memory tester, 34
- memory
 - board settings, 17
 - first address 0, 17
 - minimal requirements, 17
- message
 - login, 45
- Mini Winchester Devices (HDC-DMA), 83
- minor device number, 73, 82, 83, 84, 85
- missing ram, 17
- mkdir
 - example, 55, 56, 71
- mkfs
 - examples, 70
- mknod
 - example, 73, 83, 87
- modem, 41, 42, 72, 73
- Modem to TRICEP Cable, 73
- Modified inittab Example, 41
- modify, 27
- modifying lp, 66
- moisture, 8
- monthly backups, 93
- motd, 30
- motherboard, 6
- mount table, 40

- mount
 - example, 71
 - file system, 70
- Mounting File Systems, 71
- Moving Through the File System (Made Easy), 27
- mp300 script, 65
- MS-DOS, 77, 85
- multi-user, 11, 34, 39, 89, 90
- mw0, 19
- mw0a, 29
- mw0a error, 22

- NEC, 60
- new file creation, 47
- newgrp, 54, 55
- noises
 - boot, 16
- normal state, 37
- NULL, 44
- null modem, 73
- numbers after command names, 98

- obnoxious
 - system administrators, 91
- off, 61, 94
 - init action, 39, 42
- offset
 - tar device, 25
- ON-LINE, 62
- once
 - init action, 39
- overheat, 11

- parallel, 62
- parallel I/O processing, 82
- Parallel Printers, 61
- parity
 - no, 12
- partition, 84
 - floppy, 85
 - hard disk, 84
- partitions, 82
- passwd, 30, 33, 50, 51
 - command, 51, 53, 54
- passwords, 36, 46, 51
- PATH, 63
- path way to directory, 28
- pathname, 51
- pc boards
 - loose, 6
- permission denied
 - passwd, 53
- pilot light, 10

- pin 20, 73
 - RS-232, 16
- pin 5 of SIO4, 73
- ports, 63
- power off, 94
- power
 - cord, 5
 - requirements, 9, 10
- powerfail, 37
- powerfailure, 39
- powerwait, 39
- printer, 41
- printer ready, 60
- printer
 - baud rate, 65
 - default, 65
 - enable, 66
- printers
 - class of, 65
 - group of, 63
- process, 36, 41, 42, 43, 44
- process status data, 36
- Processes: A Few Words, 43
- profile, 33, 36, 47, 56, 63
- Programming Guide, 98
- PROM
 - code, 34
 - program, 15
 - program's duties, 16
- prompt
 - login, 45
- pseudo-run level, 42
- pwd, 28
 - example, 29

- ram, 17
- raw device, 84, 85
- rc, 40
- read all files, 23
- read error, 22
- reboot without sync, 90
- received data, 60
- recovering lost files
 - tar, 25
- repair file system, 23
- request, 65
- reset, 11, 14, 16, 88
 - light, 10
 - without sync, 90
- respawn, 36, 37, 39, 41, 48
- restoring the file system, 93
- restricted user, 52
- restricted users, 50, 51
- RETURN fails, 18
- returning a TRICEP, 5

- rm example, 40
- root
 - device, 34
 - directory, 28
 - file system, 30
 - password, 46
 - user, 46
- RS-232, 60
 - connector, 12
 - minimum cable, 5
 - missing cable, 4
 - voltages, 74
- run commands, 40
- run level, 36, 37, 39, 40, 42
- run level 2, 89
- Running FSCK, 90
- SANE, 45
- scheduler, 34
- security, 50, 51, 52, 54
 - file systems, 87
 - umask, 47
- Sending Signals to init, 42
- serial printer, 41
- Serial Printers, 60
- Setting the Date, 89
- sh file, 36
- sharing files, 54
- shell, 44
- shell script, 64
- shell
 - script example, 65
- shells, 36
- shift, 65
- shutdown, 90, 94
- sign-on prompt, 89
- signals, 41, 42
- single-sided, 85
- single-user, 34, 39
- single-user mode, 11, 92
- single-user
 - mode, 88
 - run level, 37
- SIO-4, 60
 - board, 60
- SIO4
 - second board, 41
- slash "/", 28
- soft-sectored, 24
- sort
 - example, 32
- sounds
 - booting, 16
- source files, 34
- space
 - user disk, 91
- spawn, 41
- special files, 82, 83
- spelling dictionary, 32
- spooler, 67
- standalone
 - boot, 15
 - boot program, 16
 - program, 34
- standard output, 64
- static, 8
- status
 - printer, 67, 68
- stepper noise, 16
- stop bits
 - 2, 12
- strobe, 62
- stty, 65
- superblock, 23
- superuser, 34, 42, 46, 50, 61
 - prompt, 15
- Support Tools Guide, 98
- surface scan of hard disk, 22
- swap space, 84
- swapping, 34
- switched outlets, 9
- sync, 90, 94
- syscon, 12, 40
- system calls, 36
- system
 - administrator, 11
 - calls, 34, 43
 - console, 39
 - doesn't respond, 18
 - ram, 17
- systty, 40
- take, 23
- tar and 5 1/4" floppies, 85
- tar device
 - dj0a, 25
- tar
 - example, 24, 57, 93
 - failure, 25
- telinit, 89
 - example, 61
 - examples, 42
- temperature
 - high, 8
- Terminal Settings, 11

- terminal
 - checkout, 14
 - connections, 12
 - devices, 29
 - first, 11
 - line conditioning, 44
 - no display, 16
 - setup, 12
 - stops working, 18
 - trouble, 12
- terminals
 - login, 41
- terminator, 45
- testing memory, 14
- TI, 60
- time, 89
- time-out, 18, 34
- tmp, 26, 40
- toggle, 62
- tools, 29
- tracks, 84
- transmitted data, 60
- trees, 26, 27
- Trees: the Basic Philosophy, 26
- TRICEP Installation and Maintenance Guide, 97
- Troubleshooting Booting Problems, 15
- truculent users, 91
- true, 60
- tty, 36
- TTY Devices (SIO4-DMA), 83
- tty
 - devices, 83
- turn on power, 88
- turning off, 94
- Turning the System Off, 94

- umask, 47
- umount
 - example, 72
- UniSoft, 33
- unix, 26, 28, 34
 - System V, 33
- Unmounting File Systems, 72
- unpacking, 4
- unreadable block, 23
- unused files, 91
- upage, 43, 46
- upper case, 36
- User Guide and Manuals, 97
- user's prompt, 33

- user
 - database, 50, 57
 - directories, 91
 - environment, 52
 - id number, 51
 - name, 51
 - names, 51
 - Using lp, 67
 - usr, 26
 - usr directory, 91
 - usr/spool/lp, 64
 - uucp, 99
 - uushell, 50
 - vchk example, 23
 - vchk tree, 23
 - ventilation, 8
 - voltage, 10
 - voltage switch
 - 115-230, 9

- wait
 - init action, 39
- wall
 - example, 92
- WARNING, 71
- who, 36
- who login, 52
- whole drive, 84
- Why Trees?, 27
- write all, 92
- WUNDERBUS I/O, 6

- X-ON-X-OFF, 60

