COHERENT Operating System Release Notes

Mark Williams Company
Lake Bluff, Illinois 60044
Telephone: (708) 689-2300

Printed in U.S.A.

Section 1:
Release Notes

This document contains the installation notes for COHERENT version 3.0. It describes how to install COHERENT onto your hard disk. Be sure to read this section carefully before you begin the installation process.

Also attached is documentation for UUCP, the COHERENT remote communication utility. For a number of reasons, this documentation could not be included in the printed COHERENT manual (although it is in the online manual pages). You should keep it with your COHERENT manual.

Finally, these notes contain a set of errata — errors in the manual that were discovered after the manual went to press. Try as we might, in a manual more than 1,000 pages long, consisting of 16 lengthy tutorials and more than 700 individual Lexicon articles, mistakes are nearly inevitable. We suggest that you pencil the corrections into your manual, and we hope that you will accept our apologies for whatever inconvenience this may have caused you.

2 The COHERENT System

Section 2:
Installing COHERENT

This section describes how to install COHERENT onto your computer. Installation of COHERENT is straightforward, and Mark Williams Company has prepared a suite of programs that automate much of the work for you. However, installation requires that you make a few decisions regarding how you want your system to be configured. We strongly urge you to read these notes through at least once before you begin, so you can decide correctly whenever an installation program asks you to make a decision.

Before you begin, please note the following caveats:

First, the following conditions must be met if COHERENT is to work on your system:

1. COHERENT is designed for use on the IBM AT, or on a computer that is totally compatible with the IBM AT. It does not work on *any* MicroChannel computer, or on any computer that is not 100% compatible with the IBM AT.

2. Your system must have at least one high-density, 3.5-inch or 5.25-inch, floppy-disk drive. The COHERENT distribution disks cannot be read by a low-density floppy-disk drive.

3. Your system must have a hard disk, and the hard disk must have at least ten megabytes of space free on it. More is recommended, but seven megabytes is the minimum space required by COHERENT. If you do not have enough space on your hard disk, you will have to clear space by removing or compressing existing files.

4. COHERENT works with RLL and MFM hard-disk controllers. It does *not* work with SCSI or ESDI hard-disk controllers. If your system has SCSI or ESDI drives only, COHERENT will not work with your system.

If you are unsure whether your system meets any or all of these conditions, check the documentation that came with your system, or contact the dealer from whom you purchased your system. If you believe that your computer cannot run COHERENT, please contact Mark Williams Company.

With these caveats in mind, please continue — and we hope you enjoy working with COHERENT!

What Does Installation Do?

The installation procedure creates one or more *partitions* on your hard disk to contain COHERENT and its files.

When you (or your dealer) installed MS-DOS onto your computer, you (or the dealer) divided your computer's hard disk into *logical partitions* (often by using the program **fdisk**). A hard disk on the IBM AT can have from one to four logical partitions. Each logical partition works with a specific operating system, such as COHERENT or MS-DOS or Xenix. Not every partition has to be used — your hard

disk may be divided into four partitions, but have MS-DOS file systems in only three of them, with the fourth partition being unused. Note, too, that the four logical partitions do not necessarily have to encompass the entire hard disk — a disk may have unused space that is outside any logical partition.

It may well be that you do not know the configuration of your hard disk. COHERENT can figure this out, and the configuration information will be displayed for you at the appropriate point in installation. The partitioning procedure will let you change the location and size of any partition.

If you are using MS-DOS 3.3 or later, be aware that a single MS-DOS partition (called an extended MS-DOS partition) can contain multiple logical disk drives. For example, drive C: might be on one logical partition, and drives D:, E: and F: might all be on a second logical partition.

The file systems for COHERENT and MS-DOS are very different, so it is not possible to have both systems use the same logical partition — each must have one or more logical partitions completely to itself.

As you can see, installation must cope with a number of variables: the size of your disk, the number of partitions into which it is divided, and the number of partitions that are in use. Installation thus will follow any of a number of possible scenarios, depending on how your disk is organized and how much space you wish to give over to COHERENT. The installation process will walk you through these decisions, to make them as painless as possible.

You can abort the installation procedure at any time by typing **<ctrl-C>**.

Backing Up Existing Data

As explained above, you must assign a logical partition containing at least ten megabytes on your hard disk to COHERENT during installation. If you have from one to three logical partitions already in use, you can use the fourth partition for COHERENT. If you are already using all four partitions, you will have to reassign an existing partition to COHERENT.

The space allocated to the COHERENT partition must not be part of any already existing partition. If your disk has ten megabytes free in an unused partition or not assigned to any partition, you can use it for COHERENT, but free space on a MS-DOS file system is still assigned to the MS-DOS partition.

COHERENT does not change any data on your hard disk until you exit from the partitioning procedure. The entire COHERENT installation procedure will *not* change data in an existing partition if you leave the partition unchanged. However, data currently on your hard disk may be destroyed if you change the base or the size of a logical partition, or if you change the order of the partition table entries.

Before you begin the installation procedure, be sure to back up files from any partitions which you are going to change. If you are already using multiple disk partitions and need to allocate an existing partition for COHERENT, you might be able to simply copy all the existing data on one partition to another to free the partition. *If you do not back up your hard disk before you begin to install COHERENT, Mark Williams Company cannot accept responsibility for any files that are lost during installation!*

If you do not know how your hard disk is partitioned, follow the instructions below to begin the installation procedure anyway. The COHERENT installation will not change your hard disk contents until you complete the par-

titioning procedure. Continue the installation until COHERENT displays the partition information for your hard disk. Write down the information and type <Ctrl-C> to abort installation. Figure out which partition to allocate to COHERENT and whether you will have to change any existing partitions. If necessary, boot your existing operating system, back up to other hard disk partitions or to diskettes as required, and then begin COHERENT installation again.

The following sections describe the installation process in detail. Be sure to read them through before you begin.

Getting Started

To begin installation, insert disk 1 (the **Boot/Build** disk) from your installation kit into drive A on your system. Reboot the system by pressing the reset button or by typing **<ctrl><alt><del>**. In a moment, you will be prompted with a question mark '?'. Type:

```
begin
```

followed by the **<return>** or **<enter>** key. Be sure you type **begin** in lower case letters. Note that you *must* reboot before you type **begin** — do not type **begin** to the command prompt from your current operating system.

The installation program will clear the screen and display some copyright information. After you press **<return>**, you will see the following greeting:

```
Welcome to the COHERENT operating system!

Your computer is now running COHERENT from
the floppy disk.  This program will install
COHERENT onto your hard disk.  You can inter-
rupt installation at any time by typing
<Ctrl-C>; then reboot to begin the installa-
tion procedure again.  Please be patient and
read the instructions on the screen
carefully.
```

As the instructions say, you can interrupt installation at any point by typing **<ctrl-C>**.

Entering the Serial Number

The next screen will ask you to enter a nine-digit serial number. This number is included on a paper supplied with your copy of the COHERENT system. The installation process cannot continue until you enter this number correctly.

Setting the Date and Time

Next, you must set the date and time for COHERENT. Setting the date and time is vital to the correct operation of COHERENT. COHERENT records the date and time quite differently from the way MS-DOS does it.

Time under COHERENT is recorded as the number of seconds since January 1, 1970, at exactly midnight. Internally, COHERENT always stores time as Greenwich Mean Time. GMT is used to make it easy for COHERENT systems around the globe to coordinate time with each other. When COHERENT time-stamps a file or displays the time, it converts Greenwich Mean Time to your local time, depending on what time zone you live in and whether Daylight Savings Time is in effect. (For a detailed discussion of this topic, see the Lexicon's entry for **TIMEZONE**.)

The installation program will display the following text:

> It is important for the COHERENT system to know the correct date and time. You must provide information about your timezone and daylight savings time.
>
> According to your computer system clock, your current local date and time are: *date and time*

You will be asked if this is correct. If it is not correct, the installation program will prompt you for the correct date and time.

You will then be asked about daylight savings time:

> You can run COHERENT with or without daylight savings time conversion. You should normally run with daylight savings time conversion. However, if you are going to use both COHERENT and MS-DOS and you choose to run with daylight savings time conversion, your time will be wrong (by one hour) during daylight savings time while you are running under MS-DOS.

You will be asked if you want to run in daylight-savings mode. You should answer yes unless you have an overwhelming reason not to.

The installation program then describes the default daylight-savings settings:

```
By default, COHERENT assumes daylight
savings time begins on the first Sunday in
April and ends on the last Sunday in October.
If you want to change the defaults, edit the
file "/etc/timezone" after you finish in-
stalling COHERENT.
```

The default settings are those enacted by law for the United States. COHERENT will then ask you what time zone you live in:

```
Please choose one of the following timezones:
        1       Greenwich
        2       Newfoundland
        3       Atlantic
        4       Eastern
        5       Central
        6       Mountain
        7       Pacific
        8       Yukon
        9       Alaska
        10      Bering
        11      Hawaii
        12      Other
```

If you select 1 through 11, COHERENT will set your local time automatically. If you select "Other", you will be asked how many minutes of time you live east or west of Greenwich, and then asked to name your time zone. If you are unclear on these concepts, consult the Lexicon article on **TIMEZONE**. If you are unsure about how your local time relates to Greenwich time, consult an atlas, or check with your local library.

COHERENT will then display the corrected local time and ask if it is correct. If not, you can repeat the process until the time is correct.

Partitioning Information

After the time is set, installation moves on to its next phase, partitioning the hard disk. As you enter the partition phase of installation, you will see the following text:

```
This installation procedure allows you to
create one or more partitions on your hard
disk to contain the COHERENT system and its
files. Each disk drive may contain no more
than four logical partitions. If all four
partitions on your disk are already in use,
you will have to overwrite at least one of
them to install COHERENT. If your disk uses
fewer than four partitions and has enough
unused space for COHERENT (7 megabytes), you
can install COHERENT into the unused space.
If you intend to install MS-DOS after in-
stalling COHERENT, you must leave the first
physical partition free for MS-DOS.

The next part of the installation procedure
will let you change the partitions on your
hard disk. Data on unchanged hard disk par-
titions will not be changed. However, data
already on your hard disk may be destroyed if
you change the base or the size of a logical
partition, or if you change the order of the
partition table entries. If you need to back
up existing data from the hard disk, type
<Ctrl-C> now to interrupt COHERENT installa-
tion; then reboot your system and back up
your hard disk data onto diskettes.
```

If you have have not backed up the data on your hard disk in a partition which you are going to use for COHERENT,

you should interrupt COHERENT installation by typing <Ctrl-C>. Then back up your data and begin COHERENT installation again.

Use the COHERENT Bootstrap?

When you press **<return>**, you will see the following text:

```
COHERENT initialization normally writes a
new master bootstrap program onto your hard
disk. The COHERENT master boot allows you to
boot the operating system on one selected
disk partition (the active partition)
automatically; it also allows you to boot the
operating system on any disk partition by
typing a key when you reboot. Mark Williams
strongly recommends that you use the
COHERENT master boot. However, the COHERENT
master boot may not work with some operating
systems (for example, Xenix) if you make the
COHERENT partition active; instead, leave
the other partition (e.g. Xenix) active and
boot COHERENT by typing a key. If you do not
use the COHERENT bootstrap, you must under-
stand how to boot the COHERENT partition
using your existing bootstrap program.
```

As explained in the prompt text, a *bootstrap* is a program that pulls an operating system into memory and sets it to running — the name relates to the fact that the operating system "pulls itself up by its boot straps". The COHERENT master bootstrap can boot COHERENT as well as many other operating systems, including the many versions of MS-DOS that have been tested. If you choose not to use the COHERENT master bootstrap, you must consult the documentation that came with your system to see how you can use your operating system's current bootstrap routine to boot another operating system. If you

choose to use the COHERENT master bootstrap and find that it has trouble booting your current operating system, try marking the other operating system as the active partition; you will still be able to boot COHERENT by typing a key when you reboot. If all else fails, boot your current operating system by using a boot floppy disk; you will not be able to boot it off of the hard disk, but at least it will be available to you.

We strongly suggest that you use the COHERENT master bootstrap routine.

How a Disk Is Organized

Installation then moves into the next phase: selecting a disk partition for COHERENT. As described above, partitioning can vary greatly from disk to disk; how the disk is partitioned will determine how much space is allocated to COHERENT and how much to MS-DOS or other operating systems. This is the trickiest part of installation, so be sure to read carefully.

This phase begins by displaying the current layout of your hard disk: the number of partitions, the size of each partition, and the operating system mounted on each partition. The following gives the printout for a typical hard disk. This hard disk, called disk **0**, has approximately 33 megabytes on it organized into four approximately equal partitions, as follows:

```
Drive 0 Current has the following configuration:
                [In Cylinders] [  In Tracks  ]
Number  Type  Start End Size Start  End Size Megabytes  Name
0 Boot MS-DOS     0 149  150     1  899  899    7.82   /dev/at0a
1      MS-DOS   150 299  150   900 1799  900    7.83   /dev/at0b
2      MS-DOS   300 449  150  1800 2699  900    7.83   /dev/at0c
3      MS-DOS   450 613  164  2700 3683  984    8.56   /dev/at0d
1 cylinder (.05 megabytes) is unused starting at cylinder 613.
```

As you can see, this disk has four partitions, number 0

through 3. Partition 0 is marked as the boot partition; what this means is explained below. COHERENT has given each partition a name, **dev/at0a** through **/dev/at0d**. These are the COHERENT device names for the partitions; you will not be working with these during installation, however, so you can safely ignore them for now.

Note that the middle columns of the table give the size of each partition in three ways: in cylinders, in tracks, and in megabytes. How do these differ? Megabytes are the easiest to understand: that is the number of bytes that can be written into the partition. Cylinders and tracks relate to the way a hard disk is built. A moment spent here on background can make what is to come much easier to understand.

Consider a floppy disk. Its surface is organized into 80 concentric rings, or *tracks*, numbered 0 through 79. Each track holds a fixed amount of data, with the amount depending upon the size and density of the disk. When the disk is in your disk drive, a *head* moves back and forth, reading tracks as directed. Unlike a phonograph cartridge, however, the head jumps from track to track discretely — it does not spiral in. Thus, you can measure space on a disk simply by counting the tracks. Note, too, that the term "head" is often used to describe one surface of a multi-sided disk.

As you've probably noticed, a floppy disk has two surfaces: the top and the bottom. The top is usually refered to as side 0, the bottom as side 1. Each surface has its own system of tracks, each numbered 0 through 79, giving a floppy disk a total of 160 tracks. Also, a floppy disk drive actually has two heads to read the disk, one for each surface.

A *cylinder* is the set of identically numbered tracks from both surfaces of the disk. A floppy disk organizes its 160 tracks into 80 cylinders: side 0 track 0 plus side 1 track 0 form cylinder 0; side 0 track 1 plus side 1 track 1 form cylinder 1, etc. If you think of a track as being a ring on

the disk, then origin of the term "cylinder" should be obvious.

Now, consider a hard disk. The term "hard disk" is somewhat incorrect, because one hard disk actually contains many hard disks, or *platters*, inside itself. The platters are stacked on a spindle, much like a set of 45-rpm records stacked on a record changer — except that heads move between the platters, one head for each platter surface (or two per platter). The number of platters and the number of tracks on each platter determine both the number of cylinders and the amount of data that can be written to the disk. Consider the disk described in the above table, which is a fairly typical device. It has three platters (six heads). Each head has 615 tracks, each of which holds 8,704 bytes. Thus, the device has a total of 3,690 tracks (6 times 615), organized into 615 cylinders, with each cylinder holding 52,224 bytes (6 times 8,704).

Different operating systems organize disk partitions in different ways. MS-DOS, for historical reasons, organizes partitions along track boundaries; under this scheme, the tracks of a cylinder can be divided between two partitions. UNIX, COHERENT, and similar operating systems prefer to organize partitions along cylinder boundaries: all of the tracks of a cylinder belong to only one partition. This lessens movement of the heads, which in turn speeds up reading of the disk. Note that, strictly speaking, "megabytes" has no meaning when thinking about disk partitioning: partitioning must be done either in tracks, or in cylinders. Humans tend to think of partitions in terms of megabytes, that is, in terms of the amount of data we can write into a partition, but when organizing your disk it is much easier to think in terms of cylinders. However, it is simple to translate cylinders into megabytes, which gives you the best of both worlds; this will be discussed in the following sub-section.

Partitioning the Disk

When you enter the partitioning phase of installation, the installation programs will display the configuration of your hard disk for you, using a table like the one shown above. It then displays the following menu:

```
Possible actions:
        1 = Change active partition
        2 = Change one logical partition
        3 = Change all logical partitions
        4 = Display drive information
        5 = Quit
```

The following describes each option in detail.

1. The *active partition* is the partition that the bootstrap program reads by default. When a partition is the active partition, the operating system mounted on that partition is booted automatically when your turn on your computer. This option allows you to change the active partition. You will need this option only if your hard disk has more than one logical partition, and the partitions contain different operating systems. Note that if later you wish to change the active partition, you can use the COHERENT command **fdisk** to do so. See the Lexicon entry on **fdisk** for details.

2. This option lets you change one logical partition — in effect, it lets you select a logical partition for COHERENT. You should use this option if your hard disk has more than one partition and you wish to install COHERENT on only one of them. The partition you select must hold at least seven megabytes. Note that the previous contents of the partition will be deleted.

3. This reconfigures the entire disk. You can reset the number of partitions, and the size of each.

4. Give summary information about the disk.

Option 5 is self-explanatory.

Begin by entering option 4, to receive more information about your disk. The following display gives the display for the hard disk described in the above table:

```
Drive 0 has 615 cylinders, 6 heads, and 17
sectors per track.
It contains:
        615 cylinders of 52224 bytes each,
        3690 tracks of 8704 bytes each,
        62730 sectors of 512 bytes each
or a total of 32117760 bytes (32.12 MB).
```

If the owner of this hard disk wanted to organize his hard disk by megabytes, all he would have to do is divide 1 million by 52,224 to find that one megabyte is approximately equal to 20 cylinders; thus, to make a ten-megabyte partition, he would assign it 200 cylinders. The size of a cylinder may be different on your system, but the principle is the same.

The next step depends on two factors: the current organization of your hard disk, and the amount of space you wish to give to COHERENT. If your disk has only one logical partition, you must use option 3 to create at least one new logical partition, at least one for each operating system. If your disk already has more than one logical partition, you can use option 2 to assign one to COHERENT or use option 3 to assign more than one, reserving the rest for your current operating system. Of the partition(s) that you assign to COHERENT, one must hold at least seven megabytes — you cannot use two four-megabyte partitions; thus, if no partition on your disk holds seven megabytes, must use option 3.

How much space should you give COHERENT? COHERENT is a multi-user, multi-tasking operating system; the more space you assign to it on your disk, the more users and the more processes it can support. COHERENT, via UUCP and other communications programs, also gives you access to information on other COHERENT and UNIX systems throughout the world; you will want to exchange mail with other users and possibly download news and information. All of this takes up space. You must have one seven-megabyte partition to hold COHERENT's root file system (that is, the file system that holds the files that make COHERENT go), and you would be well advised to assign at least one more partition of at least the same size to hold users' accounts and their files — or one 15-megabyte partition to hold both.

On very large hard disks, try to keep COHERENT partitions at 30 megabytes or smaller for efficiency. For example, it would be better to create two 30 MB COHERENT partitions rather than a single 60 MB COHERENT partition.

The following two sub-sections describe what happen when you invoke options 2 or 3.

Changing One Logical Partition

You will first be asked which partition you wish to change. Reply by entering the partition you want. The system replies with the following text:

```
Existing data on a partition will be lost if
you change the base or the size of the parti-
tion.  Be sure you have backed up all data
from any partition which you are going to
change.

You may specify partition bases in cylinders
or in tracks.
```

Reply 'y', to use cylinders. The system then asks:

```
You may specify partition sizes in cylinders
or in megabytes.
```

Reply 'y', again to use cylinders. Next, the system says whether the partition is initialized or is unused. It then asks you whether you wish to install COHERENT into the partition, leave the partition unchanged, or mark the partition as unused. You must select one of these possibilities: install cannot install any operating system other than COHERENT into a partition. To install COHERENT into this partition, reply 'y' when asked if you want this to be a COHERENT partition.

The final two questions ask you to enter the new base cylinder for the partition and the size of the partition in cylinders. Each question will prompt you with the current value for the partition. Simply pressing **<return>** would leave this current value unchanged. It is possible to make the partition smaller, but this serves no practical purpose if you simply intend to install COHERENT into this partition.

If you have made a mistake during this process, the system will prompt you and ask you to correct it. Otherwise, you will proceed to the next phase of installation.

Changing All Logical Partitions

This process mirrors what occurs when only one partition is changed, except that it is iterated for every existing partition. If you have four partitions and wish to eliminate one, simply set its size to zero. If you have fewer than four partitions, you will be asked if you wish to create any additional ones.

Note one additional feature: the table that displays the layout of partitions (an example of which is shown above) is redisplayed after every partition, showing the changes you have made (if any). By looking at the table, you will find it easy to keep straight just what you have done — when you

work with this table, you will see the value of working in cylinders.

If you make a mistake, the system will prompt you to correct it. A common error is requesting overlapping partitions — that is, setting the base cylinder of a partition within an area already allocated to another partition. Another error would be to request an impossibly large partition.

This concludes the discussion of partitioning the disk. The system will then moves to the next phase of installation.

Scanning for Bad Blocks

When a partition is assigned to COHERENT, it must be scanned for bad blocks. (The terms *block* and *sector* are often used interchangeably.) Most hard disks have at least a few blocks in which the disk's surface is flawed and therefore cannot be trusted to hold data reliably.

COHERENT keeps a list of bad blocks for each partition, to ensure that it does not write data into an unreliable area. This checking is performed automatically, but takes a few minutes. Patience is recommended.

Creating COHERENT File Systems

Once COHERENT has created a list of bad blocks, it can generate a file system for each of the partitions that you have assigned to it. One partition must be assigned the root file system; the root file system is the one that holds the files owned by COHERENT itself, the files that make the system go. If you are assigning more than one partition to COHERENT, you will be asked which you want to hold the root file system.

Mounting Additional File Systems

If you have created only one COHERENT partition, you can skip to the next section below. If you have created more than one COHERENT partition, the partitions other than the root partition are usually *mounted* automatically by the system. The next step is to specify where you want to mount the additional file systems. The system will ask you to assign a name to each file system. For historical reasons, a file system is usually given the name of a single letter from the lower end of the alphabet, such as 'u' or 'v' or 'x', although there's no reason not to name a file system 'work' or 'new'. Each name must be preceded by a slash '/'.

Rebooting

Now that partitions have been allocated and file systems have been created and mounted, COHERENT must be booted from the hard disk. If you have elected to use the COHERENT bootstrap, and if you have marked the COHERENT boot partition as the active partition, all you have to do is remove disk 1 (the **Build/Boot** disk) from the floppy-disk drive when prompted, and then hit the reset button or type **<ctrl><alt><del>**.

If you have made an MS-DOS partition the active partition, you must perform one additional step: type the number of the partition that holds the COHERENT root file system as the system is attempting to read the floppy disk for the bootstrap program. The number must be typed from the numeric keys at the top of the keyboard, *not* from the numeric keypad at the right. Before it begins the rebooting process, the system will tell you which number to press.

Copying Files

If rebooting occurs correctly, you will then be running COHERENT off of the hard disk. Now comes the event for which all of this preparation has occurred: the system copies the COHERENT files onto your hard disk. The system will prompt you to insert the three remaining disks from the installation kit, one after the other.

The system will ask you whether you want the full set of manual pages on line in uncompressed form, and whether you want the dictionaries used by the COHERENT spelling checker also in uncompressed form. These files must be uncompressed before they can be used, but take up much more room on the disk when uncompressed. You must decide whether the extra convenience of having on-line manual pages and a spelling checker is worth the extra space they require.

Touring the COHERENT File System

Finally, for the last step in installation the system will ask you if you wish to take a tour of the COHERENT file system. We suggest you answer yes, for this is the best way to become familiar with the layout of your newly installed COHERENT system.

And with that, the installation of COHERENT is finished!

Where to Go From Here

The next step should be to become familiar with COHERENT. We suggest that you look at the tutorials in the first half your COHERENT manual. In particular, you should work through the first three sections: *Introduction*, *Using the COHERENT System*, and *Administering the COHERENT System*. Read them carefully, and work their exercises. They will help you become familiar with COHERENT, its features and its capabilities.

One last note: Numerous references have been made in these notes to the Lexicon. This is the dictionary-format reference that occupies the second half of your COHERENT manual. It describes every COHERENT command, library routine, operating-system call, and header file, and also contains numerous articles on technical aspects of the system, definitions, system maintenance, and general good advice. The Lexicon is in a tree format, and by following the chain of cross-references it is possible to work your way from any one article to any other. Although it will never replace a good novel for bed-time reading, we think that you will find it well organized and occasionally even enjoyable to work with.

Section 3:
UUCP Remote Communication

UUCP is a set of programs that together let you communicate in an unattended manner with remote COHERENT and UNIX sites. The term *UUCP* is an abbreviation for "UNIX to UNIX copy"; as its name implies, UUCP was developed under the UNIX operating system. Mark Williams Company has recreated UUCP for COHERENT.

UUCP allows your COHERENT system to talk over telephone lines to other computers that also run COHERENT or UNIX. It can transmit files and mail to other systems and receive material from them, without needing you to guide it by hand every step of the way. Moreover, you can instruct UUCP to telephone other computers at the same time each day; this permits regular, orderly exchange of mail, news, and files among computers, and allows you to take advantage of lower telephone rates during off-peak periods. In a similar fashion, UUCP allows other systems to log into your system, to exchange mail or other information, and otherwise perform useful tasks.

Numerous UUCP systems have linked together to create an informal network called the *Usenet*. Many megabytes of

source code, news, and technical information are available across the Usenet. Anyone who is connected to the Usenet can exchange mail with anyone else who is also connected to the Usenet. All that is required to hook into the Usenet is to obtain a UUCP connection to anyone else who is connected to the Usenet.

You can use UUCP only if you have telephone access to another computer that runs UUCP, and if your system and the remote system with which you wish to communicate have been described to each other. UUCP is standard with COHERENT and UNIX, and can be purchased for MS-DOS. If you wish to copy files from another system, you must arrange with the system administrator of that system before you can begin to use UUCP. Likewise, if you want someone else to dial into your system to upload or download files, you must first describe that system to your copy of UUCP.

UUCP, however, is a system that supports thousands of interlinked computers that exchange millions of bytes of data daily; as you can imagine, a subject of this scope is difficult to encompass in a document as brief as this. If you wish to explore the heights and depths of UUCP, we urge you to acquire the following books:

- O'Reilly T, Todino G: *Managing UUCP and Usenet.* Sepastopol, Calif, O'Reilly & Associates Inc., 1987.

- Seyer MD: *RS-232 Made Easy: Connecting Computers, Printers, Terminals, and Modems.* Englewood Cliffs, NJ, Prentice-Hall Inc., 1984.

Contents of This Manual

This tutorial describes UUCP and tells you how to set up and run your UUCP system. It briefly describes how to attach a modem to your computer and how to describe it to the COHERENT system, and it describes how to dial out to other systems to exchange files. Future editions of this manual will contain information on how to allow other sys-

tems to dial into your system, so it can act as the hub of a network.

An Overview of UUCP

UUCP is a set of programs that exchange files with other computers that run UUCP. You can set aside files or mail messages to be transferred to another computer; UUCP regularly checks to see if material has been set aside to be transferred, dials the remote system, and exchanges the files without requiring your assistance.

This appears to be a simple function, but it can be extremely useful to you. Suppose, for example, that you run a real-estate office that is a member of an organization with regional and national offices. You can tell UUCP to call your regional office each night, to send a file of your new listings and to accept a file of new listings in your district that had come from other local offices. Likewise, your association's regional office can telephone the national office each night to receive new listings in your region, which can then be passed on automatically to the appropriate neighborhood offices. All of this information can be transferred at night, when telephone rates are lowest, and without needing you to be at the console. When you come to work the next morning, you will have the latest listings instantly available on your terminal.

In brief, what UUCP offers is the ability to join a *network* of computers, in which every user of every computer can exchange information with every user on every other computer, automatically. What computer networks can do is limited only by your need to exchange information with other computer users, and by your imagination.

The Programs

UUCP consists of the following programs:

uucp The UUCP user interface. **uucp** copies files from one computer to another. Be sure not to confuse the **uucp** command with the UUCP system, despite their similar names.

uucico Call remote systems: log in to the remote system, and transfer files.

uudecode Translate files encoded by **uuencode** back into object code.

uuencode Translate binary files into printable ASCII characters for transmission to another system.

uuinstall This program displays a template on your screen, and helps you describe a system to UUCP relatively painlessly.

uulog Read the UUCP logs, which reds the processes that UUCP has initiated recently.

uumvlog Copy the current UUCP log files into backup files, named after the day on which they were generated. Throw away all log files older a requested number of days. UUCP logs everything that it does; and since it does a lot, its log files can grow very large very quickly. **uumvlog** helps to ensure that you have enough information on your system to see where UUCP has gone wrong, yet the UUCP log files do not grow large enough to overwhelm your system.

uuname List the systems that your computer can reach.

uutouch Create a file that triggers a call to a named remote system.

uuxqt Execute files with the prefix "X." in the directory **/usr/spool/uucp/***sitename*.

Three other programs, while not part of UUCP per se, are used by it:

ttystat Check the status of your asynchronous ports. If UUCP is not receiving files from other systems or not sending files to other systems, it may be because the appropriate ports have not been enabled.

mail Send "electronic mail" to another person, either on your system or on another system via UUCP.

uux Execute commands on remote systems.

Directories and Files

UUCP uses the following files and directories.

/bin/uulog
The **uulog** command.

/etc/domain
This file lists the UUCP domain. It is read by **mail**.

/etc/modemcap This file holds descriptions of modems that are understood by the COHERENT system.

/etc/uucpname
Holds the name of your system, as it is known to other UUCP sites.

/usr/bin/uucp
The **uucp** command. Copy a file to another system that runs UUCP.

/usr/bin/uuname
The **uuname** command.

/usr/bin/uudecode
The **uudecode** command.

/usr/bin/uuencode
The **uuencode** command.

/usr/lib/uucp
Contains UUCP commands and system data files.

/usr/lib/uucp/L-devices
Describe the outgoing lines. Note whether they are directly wired or modems; give the protocol needed to manipulate them.

/usr/lib/uucp/L.sys
Gives login data for remote sites. It gives the way to call remote sites and the sites that only call you.

/usr/lib/uucpp/Permissions
For each site, list the programs that that site has permission to execute on your system.

/usr/lib/uucp/ttystat
The **ttystat** command.

/usr/lib/uucp/uucico
The **uucico** command.

/usr/lib/uucp/uumvlog
The **uumvlog** command.

/usr/lib/uucp/uutouch
The **uutouch** command.

/usr/lib/uucp/uuxqt
The **uuxqt** command.

/usr/spool/logs/uucp
Log of UUCP activity.

/usr/spool/uucp/.Log
Directory containing UUCP logfiles, as follows:

/usr/spool/uucp/.Log/uucico/*sitename*
/usr/spool/uucp/.Log/uux/*sitename*
/usr/spool/uucp/.Log/uucp/*sitename*
/usr/spool/uucp/.Log/uuxqt/*sitename*

/usr/spool/uucp/*sitename***/C.***
Files that instruct the local system either to send or to receive files.

/usr/spool/uucp/*sitename***/D.***
Work files for outgoing and incoming files.

/usr/spool/uucp/LCK.*
The "lock" files UUCP uses to coordinate its resources. When UUCP invokes one of its subordinate programs, it writes a lock file; another version of that subordinate program cannot be invoked until the first exits and thereby erases its lock file. This prevents different invocations of the same program from cancelling out each other's work.

/usr/spool/uucp/.Sequence
This directory contains the sequence number of the last file handled by UUCP.

/usr/spool/uucp/*sitename***/X.***
Executed files. These files have been executed by the command **uuxqt**, and are generated by a remote system.

/usr/spool/uucppublic
Public directory accessible by all remote UUCP systems.

Attaching a Modem to Your Computer

It is straightforward to attach a modem to your computer.

First, read the documentation that comes with your modem, and look for the following: (1) the baud rate at which the modem operates, and (2) the command protocol that your modem uses.

Second, check the plug on the back of your modem. The modem will connect to your computer via a nine-pin or 25-pin D plug, also known as an RS-232 interface. Such a plug can be either *male* or *female*: the male plug has nine or 25 small pins projecting from it, whereas the female does not.

Third, obtain a cable to connect one of the serial ports on your computer to the modem. The serial ports on an IBM AT or AT compatible are almost always male; if your modem has a female plug, you will need a male-to-female cable, whereas if your modem's plug is male (which is very rare), you will need a female-to-female cable. Be sure to purchase a standard modem cable for an IBM AT; practically all computer dealers carry them. The cable you purchase should support "full modem control"; if it doesn't say on the package, be sure to ask your dealer before you buy it. If you are handy with a soldering iron you may be able to solder up such a cable for yourself, but unless you know precisely what you are doing it probably is not worth the trouble. When you plug in your cable, be sure to note whether you plugged it into port **com1** or **com2**.

Fourth, reconfigure the serial port to suit your modem. This involves the following steps:

1. Log in as the superuser **root**.

2. Edit the file **/etc/ttys**. This file normally has three lines in it, one that describes the console and one for each serial port. Each line has four fields: a one-character field that indicates whether a login prompt should be displayed (used only for devices from which people will be logging into your system); a one-character field that describes whether the device is local or remote (a local would be a modem from which you wished to dial out, a remote device would be a modem from which someone could dial in); a one-character field that describes the speed (or baud rate) at which the device operates; and a field of indefinite length that names the device being described. If you have plugged into serial port **com1** a 1200-baud modem that will allow remote logins, edit the line for **com1** to read as follows:

   ```
   0rIcom1
   ```

   If you have plugged into serial port **com2** a 2400-baud modem from which you are only going to dial out, edit the line for **com2** to read as follows:

   ```
   0lLcom2
   ```

   Note that the second character is a lower-case el, not a one. For more information, see the Lexicon entries for **getty** and **ttys**.

3. When you have finished editing **/etc/ttys**, type the following command:

   ```
   kill quit 1
   ```

   This will force COHERENT to read **/etc/ttys** and set up its ports in the manner that you have configured them.

Finally, test if you have connected your modem. Turn on your modem; then log in as the superuser **root** and type the following command:

```
echo "FOO" >/dev/port
```

where *port* is either **com1** or **com2**, depending on which serial port you have plugged your modem. If the systems are connected, the lights on your modem should blink briefly. For a more sophisticated test, try to communicate with your modem by using the command **kermit**. If you are not familiar with **kermit**, see its entry in the Lexicon for details.

If you continue to have problems making connections with your modem, see the volume *RS-232 Made Easy*, referenced above. It describes in lavish detail how to connect all manner of devices via the RS-232 interface.

Installing UUCP

Installing UUCP on your system means giving UUCP information about both your system and about the remote systems with which it will be making contact.

Before you can use UUCP to log into a remote system, you must find a remote system that will let you log in via UUCP. If you ask around among your friends or check local bulletin boards, you should have no trouble finding a UUCP system that will let you log in.

When you find a system that will let you log in, you must tell the administrator of that system what your system is called. If you have not yet selected a name for your system, do so now. The name must be eight characters or fewer, and must be unique — or unique, at least, to the system into which you will log in.

The administrator of the remote system, in turn, will give you the following information: (1) the name of his system; (2) the telephone number of his system and the speed of his modem; (3) your login name (this is apart from your system name); and (4) your password. He may also tell you when he would prefer for you to log into his system. This information should always be kept confidential, just as

you would keep confidential the combination to a friend's locker.

Once you have exchanged information with the adminstrator of the remote system, you must describe that remote system to your local UUCP system. This is an involved process that has many pitfalls; however, COHERENT includes a program to make this task easier for you: **uuinstall**. This program displays a template on the screen; you fill in the blanks, and it does the rest.

To begin, log in as the superuser **root**, and type **uuinstall**. In a moment, the screen will clear and the following menu will appear:

```
H - Help for screens
S - Sitename
L - Lsys
D - Devices
P - Permissions
```

You should type **H** first. **uuinstall** will show you the keystrokes it expects to move from one field to another. These keystrokes are the same as used by the MicroEMACS screen editor; if you're not familiar with MicroEMACS, the keystrokes are easy enough to learn.

Setting Up Your Local Site

If you are describing a system to UUCP for the first time, type **S** for site name. This will ask you for two bits of information: the name you have given your site, and your local domain. Type in the site name in the space indicated. The "domain" is used to organize groups of users on your system; the Mark Williams edition of UUCP does not yet use domains, but the **mail** program expects domain information to be included for use with mail headers. Therefore, enter something into the domain slot; you may wish to use your site name followed by **.UUCP** (e.g., **mwc.UUCP**), or something similar. When you have

finished entering information, type **<ctrl-Z>** to exit this screen and return to the main menu.

If you are working with UUCP for the first time, type **D** for devices. This will let you describe to UUCP the type of modem it will working with. You will see the following template:

```
Type:
Line:
Remote:
Baudrate:
brand:
```

Type **N** to go to the next entry. Keep pressing **N** until you see the entry for the communications port you wish to use. If you wish to add a device not already listed, type **A** to "add" an entry.

The first entry, **Type:** can have one of two entries: **DIR** or **ACU**. The former indicates devices that are directly wired into the computer, such as terminals; the latter is for remote devices like modems. Type **ACU**, then **<ctrl-N>** to move to the next field.

In the next field, **Line:**, enter the serial port into which you've plugged your modem: either **com1** or **com2**. Then type **<ctrl-N>**.

The next field, **Remote:**, gives the name of the port into which a remote device is connected. Enter the port into which you plugged your modem, followed by the letter 'r'. For example, if your modem is plugged into port **com2**, enter **com2r**. Type **<ctrl-N>** to move to the next field.

The next field, **Baudrate:**, is the speed at which your modem operates, e.g., 2400 or 9600. Enter it, then type **<ctrl-N>**.

Finally, enter the type of modem that you are using. The COHERENT system's file **/etc/modemcap** contains descriptions for a number of popular modems, to spare you

the trouble of typing control sequences for your modem. The following table gives the code name for each of the modems described in **/etc/modemcap**, plus a description of it:

hayes	Hayes Smartmodem 1200
tbfast	Trailblazer, 9600 baud
xtb2400	Trailblazer, 2400 baud
avatex	Hayes clone, 2400 baud
promodem	Prometheus Promodem 1200
mk12	Signalman Mark XII
dc300	Radio Shack Direct-Connect 300

Enter the code name for the appropriate modem. One hint: if you have a Hayes or Hayes-compatible modem that runs at 2400 baud, enter **avatex** instead of **hayes** — their **modem** descriptions are virtually identical except for the baud rate.

Please note that the dialing commands in **modemcap** assume that you have a Touch-Tone telephone. If you have a pulse telephone, you must modify your modem's entry in **modemcap**. First, consult the documentation for your modem and find the correct command for dialing a pulse telephone; on Hayes and Hayes-compatible modems, it is **DP**. Then open the file **/etc/modemcap** and locate the description of your modem; then change the characters that follow the string **ds=** to the command you just looked up. For example, to edit the **avatex** entry in **modemcap** so it will dial a pulse telephone, change the string **ds=DT** to **ds=DP**.

If you have described your modem correctly, there should be no need for you to do it again. Type **<ctrl-Z>** to save your changes and return you to the main menu.

Describing a Remote Site

Next, type **L** for **L.sys**. **L.sys** is a file that hold a description of every system to which you will make connection. You will see the following template:

```
System
Line
baud rate
phone number

Day to Call Time From   Time To
Any




Expect    ""
Send      ""
Expect    ogin -
Send
Expect    ssword:
Send
```

Type **M**, to **modify this entry.**

In the first entry, **System**, type the name of the system with which you will be connecting. For example, if the system for International Widget is named **intwidget**, type that followed by **<ctrl-N>**. If a system is described more than once in the file **L.sys**, UUCP will use the first description.

The next field, **Line**, names the line to which you have connected your modem, either **com1** or **com2**, followed by

<ctrl-N>. This may seem redundant with the description in the **device** file; however, it's not, because it's possible to connect to a remote system via more than one route or device.

In the next field **baud rate**, enter your modem's baud rate; then **<ctrl-N>**.

In the next field, **phone number**, enter the remote system's telephone number. If the remote system has an area code other than yours, be sure to include the '1' before the telephone number; also, do not include the hyphens in the telephone number, or it will not fit into the space allotted for it. Then type **<ctrl-N>**.

Day and Time of Connection

The next of fields let you set the days of the week and times at which you wish to dial the remote system. **Day to Call** recognizes the following values:

Wk	Every weekday, i.e., Mo - Fr
Su	Sunday
Mo	Monday
Tu	Tuesday
We	Wednesday
Th	Thursday
Fr	Friday
Sa	Saturday
Never	Don't call remote system
All	Call at any time

Time From and **Time To** set a "window" during which UUCP will attempt to contact the remote system. Both are set using a 2400-hour clock; for example, with the following setting

```
We      2100   2300
```

UUCP will try to contact this remote system between 9:00 PM and 11:00 PM. Likewise, with the following setting

```
We      2300  0200
```

UUCP will try to contact this remote system between 11:00 PM and 2:00 AM the following morning. If on the first try UUCP fails to make connection with the remote system (the line is busy, say), it will try again periodically until either it connects with the remote system or the time period for that system and day has ended. (The following section will tell you how to set when UUCP checks for newly queued files.) When the next "legal" time comes around, UUCP will then try again.

If you do not set the time for a given day, then UUCP will attempt to contact the remote system as soon as it discovers that a file for that system has been queued. The advantage of setting times is that you can force UUCP to work in the evening and on weekends, when telephone rates are cheaper, and you can spread UUCP's work around so it never overloads the system at any given point. After all, if you need your modem yourself during the day, you don't want to wait for UUCP to finish a call before you use it.

As you can see, the template for days and times has seven rows. This lets you established different times for each day of the week; for weekdays and weekends; for weekdays alone; or weekends alone. Note that you do not have to dial a remote site every day! Depending on the importance of the site, weekdays or weekends alone may be sufficient. Consider the following set of entries:

Day to Call	Time From	Time To
Wk	2300	0200
Sat	2300	0200
Sun	1300	1500

This scheme dials the remote site between 11 PM and 2 AM Monday through Saturday, and between 1 PM and 3 PM on Sunday. This takes advantage of the fact that on Sundays, lower telephone rates are in the afternoon rather

than the evening; and it also takes advantage of the fact that like most sensible people, you have better things to do on a Sunday afternoon besides work on your computer.

Note that the default setting, **Any** with no times, forces UUCP to transmit files as soon as the are queued. If you wish to change this, do so; in any case, you can move from field to field and from line to line by typing **<ctrl-N>**.

If you are interacting with a number of remote sites, be sure to stagger the times during which UUCP attempts to contact them. The more systems UUCP has to contact during a given time period, the fewer attempts it will be able to make to contact any of them.

The Chat Script

The last six fields on this template

```
Expect    ""
Send      ""
Expect    ogin:
Send
Expect    ssword:
Send
```

In the UUCP world, this is called the "chat script". Basically, it walks your UUCP system through the prompts and responses by which you actually log into the remote system.

To understand the structure of the chat script, consider the process by which a user logs into the system for Universal Widget. When he makes the connection, the phrase

```
Welcome to the Wonderful World of Widgets!
```

appears on the screen. What he really wants to see, however, is the prompt

```
Login:
```

so he hits the carriage return key to demonstrate his im-

patience to the system. The remote system then displays the **Login:** prompt, and the user replies by typing his user ID, say "frank". Finally, the system displays the prompt

```
Password:
```

and he replies by typing his password, say "bahHumBug". All then proceeds accordingly.

The chat script is designed to imitate this sequence of events. The first **Expect/Send** pair should hold the prompt that you need to log in and how you respond if you *don't* get it. In most cases, you should set the **Expect** field to 'ogin:" and leave the **Send** field as the pair of quotation marks, which sends a carriage return.

If we were to automate the "International Widget" example above, the first **Expect/Send** pair of our chat script should read as follows:

```
Expect:  ogin:
Send:    ""
```

Note the following: First, not all systems use the word **Login:** to prompt for logins. Be sure to check with the administrator of the remote system to make sure. Second, you do not have to enter the entire prompt, only what comes at the end of it; for example, for the prompt **Login:**, the fragment **n:** is sufficient. Third, you should set these fields even if the remote system displays the login prompt immediately; the prompt may be garbled through line noise, and setting the first **Expect/Send** pair will help UUCP to cope with this.

The second **Expect/Send** pair should hold your response when you *do* receive the prompt you expect. The **Expect** half of this pair usually (but not always) holds the same prompt as the **Expect** slot of the previous **Expect/Send** pair. The **Send** half, however, should hold your login ID, as established by the administrator of the remote system. In our Widget example, the second **Expect/Send** should read:

```
Expect:  ogin:
Send:    frank
```

The third **Expect/Send** pair should hold the prompt for the password, and the password itself. In our example, the third **Expect/Send** pair should read as follows:

```
Expect:  ssword:
Send:    bahHumBug
```

When you have finished writing the chat script, your description of the remote system is complete. Type **<ctrl-Z>** to indicate that you have finished editing, and then type **X** to exit from this screen and return to the main menu. Then type **P** to enter the last template needed for installation: the one that sets permissions on your system.

Granting Permissions

The last task in describing a remote site is setting its permissions. Unless you grant the remote system permissions, it can execute nothing on your system, not even the **mail** program to send you a letter. When you grant permissions, you name the remote system in question, name the programs it can execute on your system, name the directories into which it can write files, and the directories from which it can copy files. If permissions were not set rigidly, then every UUCP connection would be potentially a breach of system security.

uuinstall's "permissions" appears as follows:

```
Remote site name
Provide an entry for that site calling
in <y/n>
Provide an entry for calling that site <y/n>
      MACHINE            LOGNAME
Commands which can be executed at this
computer by this remote site:

Read directory list

Exceptions to read directory list

Write directory list

Exceptions to write directory list

Can the remote site request file transfers
from this computer <y/n>
Can this computer initiate file transfers
to the remote site <y/n>
```

The first slot in the template asks you to name the remote site. Enter the name of the site as you entered it in the **L.sys** template. Note that **uuinstall** will automatically translate that into entries in the **MACHINE** and **LOGNAME** slots, below " by default, the **MACHINE** slot is set to the remote site's name, and the **LOGNAME** is the site name with the letter 'u' appended to the beginning.

The second slot in the template asks if you want to provide an entry in **/usr/lib/uucp/Permissions** for that site to call you. Enter 'y' only if that site will be dialing into your system; otherwise, enter 'n'.

The third slot asks if you want to provide an entry in **Permissions** for calling the remote site. If your system will be dialing into that system, enter 'y'.

The next slot asks you to name the commands that the remote site can execute on your computer. Enter **rmail:rnews**; the former lets the remote system send electronic mail on your system, and **rnews** lets it transfer news files to you. Add other commands if you like, but remember that the short the list is, the less the chance an intruder will be able to mischief on your system.

In the next slot, enter the directories from which the remote system can copy files. Enter **/usr/spool/uucp/uucppublic:/tmp**, plus whatever directories are appropriate.

The next slot requests exceptions to the read list. When you enter a directory on the read list, that directory plus all of its children become available for reading. If you wish to place "off limits" a subdirectory of any directory named in the previous slot, enter it here.

The next slot asks you to name the directories into which the remote system can write files. Enter **/usr/spool/uucp/uucppublic:/tmp**, plus any others that you wish.

Next, enter the list of exceptions to the write list.

The next slot asks if the remote system can ask your system to transfer files to it. If you wish to grant this permission, enter "y", which is the usual order of things.

Finally, **uuinstall** asks if your computer can ask the remote system to transfer files to you. If you can do so, enter 'y'.

This concludes the process of describing a remote system to UUCP. Type **<ctrl-Z>** to end data entry and return to the main menu; then enter 'X' to exit. Type 'y' when asked if you wish to save your changes into the system's files. And that's all there is to it.

Every time you wish to make contact with a new system, you must use **uuinstall** as described above.

This description may need several revisions, as you attempt to make contact with the remote system. Writing these descriptions is something of a black art. Be patient and persistant: once contact is made, the connection should work without further maintenance being needed for months to come.

Setting a Polling Time

The next step in setting up your UUCP system is to edit the file **/usr/lib/crontab**. This file contains a description of programs that are to be executed periodically. The program **cron** reads this file once every minute, checks its contents against the system time, and executes the appropriate programs. By inserting descriptions of the UUCP commands into **/usr/lib/crontab**, you will ensure that UUCP will execute regularly to poll the remote sites you have described to it. If you do not insert entries into **/usr/lib/crontab**, UUCP will connect with a remote system only if it has a file to upload to it.

The format of **/usr/lib/crontab** is described in detail in the Lexicon entry for **cron**. Basically, a crontab entry has six fields:

1. The minute in the hour when a command is to be executed (0 through 59).
2. The hour of the day when the command is to be executed (0 through 23)
3. The day of the month (1 through 31).
4. The month of the year (1 through 12).
5. The day of the week (0 through 6, with 0 indicating Sunday).

6. The command to be executed.

Fields are delimited by space characters. Note that a command can be executed more than once in any given period; just separate the multiple entries with commas. For example, if you wish to print the date and time on your terminal every 15 minutes around the clock, insert the following entry into **/usr/lib/crontab**:

```
0,15,30,45 * * * *        date >/dev/console
```

An asterisk in a field indicates that every value of the field is to be used.

The command **uutouch** forces UUCP to schedule a poll to a remote site, regardless of whether any files are waiting to be uploaded to that system. To poll the site called **george** on a regular basis, insert the following lines into **/usr/lib/crontab**:

```
30 * * * * /usr/lib/uucp/uucico -sgeorge
0 22 * * * /usr/lib/uucp/uutouch george
```

The first line invokes the program **uucico** every hour on the half hour around the clock. **uucico** checks to see if there is a file to be sent to site **george**, and dispatches it if need be.

The second line invokes the program **uutouch** every night at 10 PM. **uutouch** will schedule a poll to site **george** to see if it has a file to send to you. The next time that **uucico** is invoked, it will then call site **george**.

Finally, the command **uumvlog** should be invoked daily by **cron**. **uumvlog** copies all of UUCP's log files into backup files that are named by the date they were saved. This command takes one argument, the number of days' worth of backup files to save. For example, the command

```
uumvlog 2
```

saves two days' worth of backup files; when a backup file becomes more than two days old, **uumvlog** throws it away.

UUCP is designed to log everything that it does; and since it does a great many things, the log files can grow very large, very quickly. On a small system especially, you should be ruthless in purging your UUCP log files, or you may find them overwhelming the available disk space on your system. For most users, two days' worth of log files is sufficient.

The following entry in file **/usr/lib/crontab** will purge your backup files at the stroke of midnight every day:

```
0 0 * * * /usr/lib/uucp/uumvlog 2
```

Sending files via UUCP

Suppose, for example that site **santa** has been described to your UUCP system, and everyone has permission to read from your current directory. Suppose, too, that you have permission to write into directory **/tmp/reports/parents** on machine **santa**. To send the file **good.kids** to **santa**, type the following command:

```
uucp good.kids santa!/tmp/reports/parents
```

The **uucp** command compels UUCP to copy one or more files from your site to a remote site. UUCP queues both files automatically and sends them at the next scheduled time.

Note, too, the use of the '!' in the above command. The '!' separates a site name from another site name, from a directory name, or from a user ID. In the above example, the '!' indicates that directory **/tmp/reports/parents** can be found at site **santa**. One feature of a UUCP network is that any member can send files to any other member. That does not mean that every member must have full permissions with every other member; rather, for the sake of efficiency it is possible to route files through one or more intermediate computers, to allow batch transmissions of files. For example, to send the file **visibility** to machine

blitzen via machines **santa** and **reindeer**, use the following command:

```
uucp visibility santa!reindeer!blitzen!/tmp
```

In this example, the string **santa!reindeer!blitzen!/tmp** indicates that directory **/tmp** can be contacted at site **blitzen**, which in turn can be contacted via site **reindeer**, which in turn can be contacted via site **santa**. This scenario assumes that site **reindeer** has permission to write into directory **/tmp** on site **blitzen**, and that site **santa** has permission to upload files to site **reindeer**. (And, of course, that you have permission to upload files to site **santa**.) If any of these are not true, the transaction will fail.

With UUCP networks growing to international dimensions, such path names can become quite complex. The command **mail** has an **alias** function that allows you to define a user's UUCP path name under a simpler name that serves as that user's alias. **mail** reads the file **/usr/lib/mail/aliases** for every user listed on its command line. If it finds a match, then it substitutes the description in **alises** for the user name. If the entry in aliases consists of two or more fields separated by exclamation points, **mail** then invokes the **uucp** command to copy the mail message to the named site. For example, if you regularly send mail to user **joe** at site **widget**, then insert the following entry into **/usr/lib/mail/aliases**:

```
joe:  widget!joe
```

Make sure, first, that you have described site **widget** to UUCP or this will not work. Second, make sure that your local system does not have a user named **joe**; if it does, his mail thereafter will be shipped to the other **joe** at the remote site.

UUCP Administration

Once you have written and debugged the descriptions of your devices, systems, and permissions, administering UUCP consists mainly of reviewing the log files periodically to ensure that all connections are being made, and all programs executed correctly. The command **uulog** will assist you in this. When you type the command

```
uulog widget
```

uulog will open all of the log files associated with site **widget**, and display them for you. Given that the log files for given site are kept in four different directories, this can be a great convenience.

Logfiles are organized as follows:

/usr/spool/uucp/.Log/uucico/*sitename*
/usr/spool/uucp/.Log/uucp/*sitename*
/usr/spool/uucp/.Log/uux/*sitename*
/usr/spool/uucp/.Log/uuxqt/*sitename*

As you can see, one logfile for each site is kept in a directory named after a given UUCP command. UUCP records every transaction; so by reading these files, you can see whether your UUCP commands are succeeding.

If you are having trouble with your UUCP connections, send files through UUCP and observe how they fail. You may need to use **uuinstall** a few times to tweak your description of the remote site. If all else fails, contact Mark Williams Company.

If all is going well, you should run **/usr/lib/uucp/uumvlog** every day. This keeps the log files from getting out of hand. The previous section on setting the polling time describes how to do this.

The main task of the UUCP administrator is to monitor the UUCP log files to see that hardware is functioning correctly, and that files are transferred correctly. For example, failure to connect with a remote site after several attempts may mean that the remote site is having problems with its modem, or that it is scheduling outgoing calls for when you were scheduled to call in. Likewise, failure to receive scheduled calls from several sites may indicate equipment failure on your end. You must also monitor the alias file, to see to it that mail is routed to the correct recipient.

Finally, the UUCP administrator must monitor the use of disk space on the system. Old mail and messages, multiple copies of files, and files automatically input by various subscription and network services can eat up disk space rapidly; extraneous material must be pruned ruthlessly.

Where to Go From Here

The following pages give the Lexicon entry for each UUCP command.

uucico — Command

Transmit data to or from a remote site
/usr/lib/uucp/uucico [-r1] [-s*site*] [-S*site*]

uucico is the UUCP command that actually transfers files to or from a remote *site.* Its syntax is as follows:

-r1 Poll *site* unconditionally.

-s*site*
The name of the *site* to be polled. *site* must name one of the entries in **/usr/lib/uucp/L.sys.**

-S*site*
The name of the *site* to be polled. *site* must name one of the entries in **/usr/lib/uucp/L.sys.** Unlike the **-s** option, force execution even if not the correct time.

The messages sent by **uucico** are differentiated by the first letter of the message.

Example

To poll the site *sys* at five minutes after the hour, each hour, put the following entry into **/usr/lib/crontab:**

```
05 * * * * /usr/lib/uucp/uucico -ssys -r1
```

Files

/usr/lib/uucp/L.sys — List of reachable systems
/usr/spool/uucp/.Log/uucico/*sitename* — **uucico** activities log file for *sitename*
/usr/spool/uucp/*sitename* — Spool directory for work

See Also

commands, cron, uucp, UUCP, uulog, uutouch, uuxqt

uucp — Command

Ready files for transmission to other systems
uucp [-bcCdm] *source1 ... sourceN dest*

uucp copies files *source1* through *sourceN* to the destination system *dest*. Either source or destination files can contain specifications for the remote system.

uucp recognizes the following options:

-c Instead of copying the *source* file to the spool directory, use the file itself. This is a default.

-C Copy the source file to the spool directory.

-d Make directories on *dest* if they are necessary for copying the files.

-f Do not make intermediate directories for the file copy.

-g*grade*
grade is a single ASCII character indicating the importance of the files being transmitted: the lower the value of *grade*, the more important the files.

-m Send mail to the requester when the file is sent.

-n*user*
Notify *user* on destination system that file was sent. Note that *user* may contain a path:

```
-nuser!site
```

-x*debug debug* is a single-digit number, 0 to 9. The higher the level, the more information yielded.

Examples

The first example copies file **foo** to directory **bar** on system **george**:

```
uucp foo george!bar
```

The next example copies file **foo** from system **george** into directory **tmp** on your system:

```
uucp george!foo /tmp
```

The next example copies file **foo** from system **george** into file or directory **bar** on system **ivan**:

```
uucp george!foo ivan!bar
```

Note that this assumes your system can talk to both **george** and **ivan**.

The next example downloads files **foo** and **bar** from remote systems **ivan** and **george** from two remote systems into directory **/tmp**:

```
uucp ivan!foo george!bar /tmp
```

Files

/usr/lib/uucp/L.sys — List of reachable systems
/usr/spool/.Log/uucp/*sitename* — **uucp** activities log file for *sitename*
/usr/spool/uucp/*sitename* — Spool directory for work

See Also

commands, mail, uucico, UUCP, uudecode, uuencode, uutouch, uuwatch, uuxqt

uudecode — Command

Decode a binary file sent from a remote system
uudecode [*file*]

uudecode takes a **file** encoded by **uuencode** and translates it back to binary. Any leading and trailing lines added by **uucp** are discarded.

If the *file* is not specified, standard input is read.

Example

Consider the file **tmp** consisting of (approximately):

```
begin 644 sys
M5&AE('%U:6-K(&)R;W=N(&90>"!J=
6UP<R!O=F5R('1H92!L87IY(&109RX*

end
```

Note that the third line is a space followed by a newline. To decode it, type:

```
uudecode tmp
```

The output contained in file **sys** will be:

```
The quick brown fox jumps over the lazy dog.
```

See Also

commands, uucp, UUCP, uuencode

Notes

The user on the remote system must be able to write the file.

uuencode — Command

Encode a binary file for transmission to a remote system
uuencode [*source*] *outputfile*

uuencode prepares a binary file for transmission to a remote destination via **uucp. uuencode** takes binary input and produces an encoded version, consisting of printable ASCII characters, on standard output, which may be redirected or piped to **uucp**. If *source* is not specified, the standard input is read.

The format of the encoded file is as follows:

1. A *header* line starting with the characters **begin** followed by a space. This is followed by the mode of the file in octal (see **chmod** for details) and the name of the output file specified on the command line. These last two fields are also separated by a space. The mode and the system name can be changed by directing the output into a file and editing it.

2. The *body* of the file, consisting of a number of lines, each no more than 62 characters long, including a newline character. Each line starts with a character count written as a single ASCII character, representing an integer value from 0 (octal 40) to 63 (octal 135) giving the number of characters in the rest of the line. This is followed by the encoded characters and a newline. The last line of the body is a line consisting of an ASCII space (octal 40).

3. The trailer line has just the characters **end** on a line by itself.

The encoding is done by taking three bytes and storing them in four characters, six bits per character.

Example

To encode the file **tmp** consisting of the line

```
The quick brown fox jumps over the lazy dog.
```

and place the output into a file called **sys**, enter:

```
uuencode tmp sys
```

The output will be (approximately):

```
begin 644 sys
M5&AE('%U:6-K(&)R;W=N(&9O>"!J=
6UP<R!O=F5R('1H92!L87IY(&1O9RX*

end
```

Note that the third line consists of a space followed by a newline.

See Also

commands, uucp, UUCP, uudecode

Notes

The file is expanded by more than one third, causing increased transmission time. This can be a factor when sending large files.

uuinstall – Command

Install UUCP
uuinstall

uuinstall assists with the installation of UUCP. It uses screen templates, help lines, and prompts to help walk you through the installation of devices, remote systems, site names, domains, and permissions. For a detailed description of its use, see the tutorial on UUCP in the front of this manual

See Also

commands, UUCP

Notes

Only the superuser **root** can execute **uuinstall.**

uulog – Command

Examine UUCP operations
uulog [-fx] [*system* **]**

uulog copies the last part of the file **/usr/spool/uucp/.Log/uucico/***system* to see what **uucico** has done recently. *system* names the remote system whose logfile will be examined. If it is not specified, logfiles for all systems are displayed.

uulog recognizes the following options:

-f Similar to the command **tail -f**: this forces **uulog** to display UUCP activity as it is written into the log file, until you interrupt it by typing **<ctrl-C>**.

-x Display the log files for the command **uuxqt** rather than **uucico.**

Files

/usr/spool/uucp/.Log/uucico/*system* — **uucico** activity log file for *system*
/usr/spool/uucp/.Log/uuxqt/*system* — **uuxqt** activity log file for *system*

See Also

commands, uucico, uucp, UUCP, uuxqt

uuname – Command

List uucp names of known systems
uuname [-l]

uuname lists the names of all systems reachable directly by **uucp.** When used with the **-l** option, it prints the name of the local system.

Files

/usr/lib/uucp/L.sys — Site and remote login data list

See Also

commands, uucico, uucp, UUCP, uulog

uutouch – Command

Touch a file to trigger uucico poll
uutouch *system*

uutouch creates an empty control file for *system* in the directory **/usr/spool/uucp/***system*. This forces UUCP to poll *system* when **uucico** is called with the option **-sany**. If the empty file for *system* aready exists, it is left alone.

There are three types of files in the spool directory **/usr/spool/uucp***system*:

C. Command file.

D. Data file.

X. Execute file.

Example

A typical usage is to put the following line into **/usr/lib/crontab:**

```
0 7 * * * /usr/lib/uucp/uutouch george
```

This forces UUCP to schedule a poll to the remote system **george** at 7 AM local time. The actual poll take

place when **uucico** is started.

Files

/usr/spool/uucp/*sitename* — Directory for uucp work files

See Also

commands, cron, uucico, uucp, UUCP, uuxqt

uuxqt — Command

Execute commands requested by a remote system
uuxqt

uuxqt takes the execute files, those marked with the prefix **X.** in the directory **/usr/spool/uucp,** and executes them. It will only execute programs for which the remote system has permission.

uuxqt may be called by either **uucp** or **uucico.** It is not generally considered a user-callable program.

Files

/usr/spool/uucp — Directory for execute files

See Also

commands, uucico, uucp, UUCP

Section 4:
Errata

The following sections are missing from the manual but are in the online manual pages used by the man command:

dcheck	Directory consistency check
floppy	General floppy-disk information
hp	LaserJet printer filter
learn	Computer-aided instruction

The following UUCP sections are included in the enclosed insert. They are also in the online manual pages:

uucico
uucp
uudecode
uuencode
uuinstall
uulog
uumvlog
uuname
uutouch
uuxqt

The following manual sections are incomplete in the printed Lexicon. See the online manual pages for later versions:

ATclock	Sets as well as reads the system clock
date	'-s' option suppresses DST conversion on setting
dos	No 'F' option
getty	Option '3' cycles 2400/1200/300, not 1200/300