Network Operating System
User Reference Manual

Phil Karn, KA9Q
and
Gerard van der Grinten, PA0GRI

This manual reflects version 911229 of NOS 2.0m
as released to the public by PA0GRI

1. The NOS.EXE Program

The MS-DOS executable file **NOS.EXE** (Further called **Nos**) provides Internet (TCP/IP), NET/ROM and AX.25 facilities. Because it has an internal multitasking operating system, **Nos** can act simultaneously as a client, a server and a packet switch for all three sets of protocols. That is, while a local user accesses remote services, the system can also provide those same services to remote users while also switching IP, NET/ROM and AX.25 packets and frames between other client and server nodes.

The keyboard and display is used by the local operator to control both host and gateway level functions, for which a number of commands are provided.

1.1. Startup

nos [-b] [-s <#sockets>] [-d </directory>] [-v] [<startup file>]

When **Nos** is executed without arguments, it attempts to open the file **autoexec.nos** in the root directory of the current drive. If it exists, it is read and executed as though its contents were typed on the console as commands. This feature is useful for attaching communication interfaces, configuring network addresses, and starting the various services.

Four command-line options are accepted:

1.1.1. -b

The **-b** option specifies the use of BIOS for console output; the default is to write directly to the video display buffer. Use this option if you are running under a windowing package and have trouble with output "bleeding through" on top of other windows.

1.1.2. -s <no_of_sockets>

The **-s** option specifies the size of the *socket* array to be allocated within **Nos**. This limits the number of network connections that may exist simultaneously; the default is 40.

1.1.3. -d </directory>

The **-d** option allows the user to specify a "root" directory for the configuration and spool files; it defaults to the root directory of the system.

1.1.4. -v

The **-v** option allows the user to view command execution during the startup of **Nos**. It lets the commands read from **autoexec.nos** echo before they are executed. This is a nice help if **Nos** stops (hangs) during initialization.

After all command line options, the name of a alternate startup file may be specified. This file is then opened and read instead of **autoexec.nos**.

1.2. DOS environement variables.

The following DOS environement variables can be used to specify things to **NOS**.

1.2.1. TZ

The TZ variable should be set to the local timezone. Default is UTC. This is used on the timestamp in smtp.

1.2.2. MAILER

The MAILER specifies what program should be started when the **mail** command is entered. Default is BM.EXE.

1.2.3. COMSPEC

The COMSPEC specifies what command shell will be used to *shell out* of **Nos**. This is normaly set by MS-DOS startup to COMMAND.COM. Default is also COMMAND.COM.

1.2.4. TMP

The TMP variable is used to create a spot where temporary files are created. Without TMP being set the temp files are created in the root directory. A sample is "set TMP=C:\tmp\".

1.2.5. USER

The USER variable is used by **ftp** and **Rlogin** to set the username for the rlogin daemon on the remote system. The default when not setting USER is guest. Guess you don't like it, but that's life. With **ftp** the user is suggested in the user name prompt. If a cr only is given the suggested name is used, otherwise the given name.

2. Console modes

The console may be in one of two modes: *command mode* and *converse mode*. In *command mode*, the prompt **net**> is displayed and any of the commands described in the **Commands** chapter may be entered. In *converse mode*, keyboard input is processed according to the *current session*.

Sessions come in many types: *Telnet, Ttylink, Rlogin, FTP, AX25, Finger, Command, NETROM, Ping, More, Dial, Dir, PPP PAP, Hopcheck* and *Tip*.

In a Telnet, Ttylink, AX25, NETROM, Rlogin, or Tip session, keyboard input is sent to the remote system and any output from the remote system is displayed on the console. In an FTP session, keyboard input is first examined to see if it is a known local command; if so it is executed locally. If not, it is "passed through" to the remote FTP server. (See the **FTP Subcommands** chapter). In a Ping session the user may test the path to a remote site, in a More session, the user may examine a local file. A Hopcheck session is used to trace the path taken by packets to reach a specified destination. A Finger session is used to peek at a remote system for its users (and what they are doing on some extended responses from UNIX systems). PPP PAP is used as a link setup like slip between two systems.

The keyboard also has *cooked* and *raw* states. In *cooked* state, input is line-at-a-time; the user may use the line editing characters ^U, ^R, ^B, ^W and backspace to erase the line, redisplay the line, redisplay the remainder of the previous line, erase last word and erase the last character, respectively. Hitting either return or line feed passes the complete line up to the application. In *raw* mode, each character is immediately passed to the application as it is typed. The keyboard is always in *cooked* state in command mode. It is also *cooked* in converse mode on an AX25, FTP or NET/ROM session. In a Telnet or Ttylink session it depends on whether the remote end has issued (and the local end has accepted) the Telnet WILL ECHO option. (See the **echo** command).

On the IBM-PC, the user may escape back to *command mode* by hitting the F10 key or the *escape* key. On other systems, the user must enter the *escape* character, which is by default control-] (hex 1d, ASCII GS). (Note that this is distinct from the ASCII character of the same name). The escape character can be changed (see the **escape** command). The F10 key can be redefined with the fkey command

so the user is now warned to leave one escape possibility open for himself. Setting both F10 and escape to unreachable codes renders a system unescapable and the user hung in a session.

In the IBM PC version, each session (including the command "session") has its own screen. When a new session is created, the command display is saved in memory and the screen is cleared. When the command escape key (usually F10 or ^]) is hit, the current session screen is saved and the command screen is restored. When a session is resumed, its screen is restored exactly as it appeared when it was last current. **NosFp expects that the driver NANSY.SYS is loaded to support the screen and terminal emulation routines. ANSY.SYS has many problems and should NOT been used. Version 24 of nansy.sys (nan24hyc.zip on download systems) should be used for best results. Users of DeskView should be ok with dvansy.sys as it has been reported that tab expansion is all right under DV.**

3. Commands

This section describes the commands recognized in command mode, or within a startup file such as **autoexec.nos**. These are given in the following notation:

command
command literal_parameter
command subcommand <parameter>
command [<optional_parameter>]
command a | b

Many commands take subcommands or parameters, which may be optional or required. In general, if a required subcommand or parameter is omitted, an error message will summarize the available subcommands or required parameters. (Giving a '?' in place of the subcommand will also generate the message. This is useful when the command word alone is a valid command.) If a command takes an optional value parameter, issuing the command without the parameter generally displays the current value of the variable. (Exceptions to this rule are noted in the individual command descriptions.)

Two or more parameters separated by vertical bar(s) denote a choice between the specified values. Optional parameters are shown enclosed in [brackets], and a parameter enclosed in <angle brackets> should be replaced with an actual value or string. For example, the notation <**hostid**> denotes an actual host or gateway, which may be specified in one of two ways: as a numeric IP address in dotted decimal notation (eg. 44.0.0.1.), or as a symbolic name listed in the file **domain.txt**,

All commands and many subcommands may be abbreviated. You only need type enough of a command's name to distinguish it from others that begin with the same series of letters. Parameters, however, must be typed in full.

Certain FTP subcommands (eg. **put, get, dir**, etc) are recognized only in converse mode with the appropriate FTP session; they are not recognized in command mode. (See the **FTP Subcommands** chapter.)

3.1. <CR>

Entering a carriage return (empty line) while in command mode puts you in converse mode with the current session. If there is no current session, **Nos** remains in command mode and reissues the **net**> prompt.

3.2. !

An alias for the **shell** command.

3.3.

Commands starting with the hash mark (#) are ignored. This is mainly useful for comments in the **autoexec.nos** file.

3.4. abort [<session #>]

Abort a FTP **get, put or dir** operation in progress. If issued without an argument, the current session is aborted. (This command works only on FTP sessions.) When receiving a file, **abort** simply resets the data connection; the next incoming data packet will generate a TCP RST (reset) response to clear the remote server. When sending a file, **abort** sends a premature end-of-file. Note that in both cases **abort** will leave a partial copy of the file on the destination machine, which must be removed manually if it is unwanted.

3.5. arp

Display the Address Resolution Protocol table that maps IP addresses to their subnet (link) addresses on subnetworks capable of broadcasting. For each IP address entry the subnet type (eg. Ethernet, AX.25), subnet address and time to expiration is shown. If the link address is currently unknown, the number of IP datagrams awaiting resolution is also shown.

3.5.1. arp add <hostid> ether | ax25 | netrom | arcnet <ether_addr> | <ax25_addr>

Add a permanent entry to the table. It will not time out as will an automatically-created entry, but must be removed with the **arp drop** command.

3.5.2. arp drop <hostid> ether | ax25 | netrom | arcnet

Delete a permanent entry from the arp table.

3.5.3. arp flush

Drop all automatically-created entries in the ARP table; permanent entries are not affected.

3.5.4. arp publish <hostid> ether | ax25 | netrom | arcnet <ether_addr> | <ax25_addr>

This command is similar to the **arp add** command, but the system will also respond to any ARP request it sees on the network that seeks the specified address. (Use this feature with great care.)

3.6. autoroute <yes|no>

Displays or sets the IP autorouting option. When set all AX25 IP packets are analysed and remembered.

3.7. asystat

Display statistics on attached asynchronous communications interfaces (8250 or 16550A), if any. The display for each port consists of three lines. The first line gives the port label and the configuration flags; these indicate whether the port is a 16550A chip, the *trigger character* if any, whether CTS flow control is enabled, whether RLSD (carrier detect) line control is enabled, and the speed in bits per second. (Receiving the *trigger character* causes the driver to signal upper layer software that data is ready; it is automatically set to the appropriate frame end character for SLIP, PPP and NRS lines.)

The second line of the status display shows receiver (RX) event counts: the total number of receive interrupts, received characters, receiver overruns (lost characters) and the receiver *high water mark*. The high water mark is the maximum number of characters ever read from the device during a single interrupt. This is useful for monitoring system interrupt latency margins as it shows how close the port hardware has come to overflowing due to the inability of the CPU to respond to a receiver interrupt in time. 8250 chips have no FIFO, so the high water mark cannot go higher than 2 before overruns occur. The 16550A chip, however, has a 16-byte receive FIFO which the software programs to interrupt the CPU when the FIFO is one-quarter full. The high water mark should typically be 4 or 5 when a 16550A is used; higher values indicate that the CPU has at least once been slow to respond to a receiver interrupt.

When the 16550A is used, a count of FIFO timeouts is also displayed on the RX status line. These are generated automatically by the 16550A when three character intervals go by with more than 0 but less than 4 characters in the FIFO. Since the characters that make up a SLIP or NRS frame are normally sent at full line speed, this count will usually be a lower bound on the number of frames received on

the port, as only the last fragment of a frame generally results in a timeout (and then only when the frame is not a multiple of 4 bytes long.)

Finally, the software fifo overruns and high water mark are displayed. These indicate whether the <buf-size> parameter on the attach command needs to be adjusted (see the **Attach Commands** chapter).

The third line shows transmit (TX) statistics, including a total count of transmit interrupts, transmitted characters, the length of the transmit queue in bytes, the number of status interrupts, and the number of THRE timeouts. The status interrupt count will be zero unless CTS flow control or RLSD line control has been enabled. The THRE timeout is a stopgap measure to catch lost transmit interrupts, which seem to happen when there is a lot of activity (ideally, this will be zero).

3.8. attach <hw type> ...

Configure and attach a hardware interface to the system. The details are highly interface dependent and dependent on configuration flags in the file config.h when the software is build. It can be that not all drivers listed below will be included in your copy of **Nos**. Detailed instructions for each driver are in the **Attach Commands** chapter. Drivers are available for the following hardware types:

3.8.1. attach 3c500

Don't use this one anymore. Use the packet driver instead. This driver is obsolete and not supported anymore.

3.8.2. attach asy

Standard PC asynchronous interface (com port) using the National 8250 or 16450 or 16550A chip or compatible equivalent.

3.8.3. attach axip

A "wormhole" ax25 digipeater device.

3.8.4. attach arcnet

A ARCnet driver via the PACKET driver.

3.8.5. attach drsi

N6TTO driver for the DRSI PCPA 8530 card.

3.8.6. attach eagle

WA3CVG/NG6Q driver for the Eagle Computer card (Zilog 8530).

3.8.7. attach hapn

KE3Z driver for the Hamilton Amateur Packet Network adapter board (Intel 8273).

3.8.8. attach hs

Special "high speed" 8530 driver for the WA4DSY 56kb/s modem.

3.8.9. attach kiss

This enables a multiplexed tnc type to be used for second channel. It is used to connect a second port to an already attached asy interface. It will copy most of the params of it's parent port.

3.8.10. attach netrom

This is a psuedo interface to enable NET/ROM operations.

3.8.11. attach packet

Driver for use with separate software "packet drivers" meeting the FTP Software, Inc, Software Packet Driver specification.

3.8.12. attach pc100

Driver for the PACCOMM PC-100 (Zilog 8530) card.

3.8.13. attach pi

Dma driven 8530 scc board from VE3IFB.

3.8.14. attach scc

PE1CHL driver for generic 8530 cards.

3.8.15. attach slfp

Serial Line Faming Protocol packet driver.

An easy way to obtain a summary of the parameters required for a given device is to issue a partial attach command (eg. **attach packet**.) This produces a usage message giving the complete command format.

3.9. attended [off | on]

Displays or sets the global "I am present" flag in **Nos**. This flag is used in the welcome header by incoming ttylink connections.

3.10. ax25 <subcommand>

These commands are for Ax25 interfaces.

3.10.1. ax25 bc <interface> [on | off]

The bc command displays or enables or disables broadcasts via interface **interface** when on | off are given.

3.10.2. ax25 bcinterval [<seconds>]

The bcinterval displays or sets the time in seconds between **bc** broadcasts. On display both the interval and the countdown values are shown.

3.10.3. ax25 bckick <interface>

The bckick command issues a direct broadcast to interface **interface** when so enabled by a *ax25 bc <interface> on* command.

3.10.4. ax25 bctext ["broadcast text"]

The bctext command displays or sets the text to be send for broadcast messages send out every **bcinterval** seconds.

3.10.5. ax25 blimit [<limit>]

Display or set the AX25 retransmission backoff limit. Normally each successive AX25 retransmission is delayed by twice the value of the previous interval; this is called *binary exponential backoff*. When the backoff reaches the blimit setting it is held at that value, which defaults to 30. To prevent the possibility of "congestive collapse" on a loaded channel, blimit should be set at least as high as the number of stations sharing the channel. Note that this is applicable only on actual AX25 connections; UI frames will never be retransmitted by the AX25 layer.

3.10.6. ax25 dest

See ax25 hearddest command. This is a shorthand version of it.

3.10.7. ax25 digipeat [on | off]

Display or set the digipeater enable flag. If the axip interface is used this flag MUST be on , otherwise the digipeat function will not work at all.

3.10.8. ax25 filter <0 | 1 | 2 | 3>

The filter commands enables or disables the logging in the heard lists of source and destination ax25_addresses. This is a bitwise or function where the 01 value is for source stations and the 02 value for destination stations. When the bit is off , logging is enabled, when on disabeled.

3.10.9. ax25 flush

Clear the AX.25 "heard" list (see **ax25 heard**).

3.10.10. ax25 heard [<interface>]

Display the AX.25 "heard" list. For each interface that is configured to use AX.25, a list of all ax25_addresses heard through that interface is shown, along with a count of the number of packets heard from each station and the interval, in hr:min:sec format, since each station was last heard. The local station always appears first in the listing; the packet count actually reflects the number of packets transmitted. This entry is always present even if no packets have been sent. If interface is given, only the heard list for that interface is displayed. Note that logging of heard stations is controlled with the **ax25 filter** command.

3.10.11. ax25 hearddest [<interface>]

Displays the destination list, i.e. the addressed to stations. Next to the time the last transmission to that station the time that station replied (if heard) is displayed. This gives a good reference to see if a station is reachable and responding.

3.10.12. ax25 irtt [<milliseconds>]

Display or set the initial value of smoothed round trip time to be used when a new AX25 connection is created. The value is in milliseconds. The actual round trip time will be learned by measurement once the connection has been established.

3.10.13. ax25 kick <axcb>

Force a retransmission on the specified AX.25 control block. The control block address can be found with the **ax25 status** command.

3.10.14. ax25 maxframe [<count>]

Establish the maximum number of frames that will be allowed to remain unacknowledged at one time on new AX.25 connections. This number cannot be greater than 7. Without **count** it will display the current setting. Note that the maximum outstaning frame count only works with virtual connections. UI frames are not affected.

3.10.15. ax25 mycall [<ax25_addr>]

Display or set the default local AX.25 address. The standard format is used, (eg. KA9Q-0 or WB6RQN-5). This command must be given before any **attach** commands using AX.25 mode are given.

3.10.16. ax25 paclen [<size>]

Limit the size of I-fields on new AX.25 connections. If IP datagrams or fragments larger than this are transmitted, they will be transparently fragmented at the AX.25 level, sent as a series of I frames, and reassembled back into a complete IP datagram or fragment at the other end of the link. To have any effect on IP datagrams, this parameter should be less than or equal to the MTU of the associated interface.

3.10.17. ax25 pthresh [<size>]

Display or set the poll threshold to be used for new AX.25 Version 2 connections. The poll threshold controls retransmission behavior as follows. If the oldest unacknowledged I-frame size is less than the poll threshold, it will be sent with the poll (P) bit set if a timeout occurs. If the oldest unacked I-frame size is equal to or greater than the threshold, then a RR or RNR frame, as appropriate, with the poll bit set will be sent if a timeout occurs.

The idea behind the poll threshold is that the extra time needed to send a "small" I-frame instead of a supervisory frame when polling after a timeout is small, and since there's a good chance the I-frame will have to be sent anyway (i.e., if it were lost previously) then you might as well send it as the poll. But if the I-frame is large, send a supervisory (RR/RNR) poll instead to determine first if retransmitting the oldest unacknowledged I-frame is necessary; the timeout might have been caused by a lost acknowledgement. This is obviously a tradeoff, so experiment with the poll threshold setting. The default is 128 bytes, one half the default value of **paclen**.

3.10.18. ax25 reset <axcb>

Delete the AX.25 connection control block at the specified address.

3.10.19. ax25 retry [<count>]

Limit the number of successive unsuccessful retransmission attempts on new AX.25 connections. If this limit is exceeded, link re-establishment is attempted. If this fails **retry** times, then the connection is abandoned and all queued data is deleted.

3.10.20. ax25 route

Display the AX.25 routing table that specifies the digipeaters to be used in reaching a given station.

3.10.20.1. ax25 route add <target> [digis ...]

Add an entry to the AX.25 routing table. An automatic **ax25 route add** is executed if digipeaters are specified in an AX25 **connect** command, or if a connection is received from a remote station via digipeaters. Such automatic routing table entries won't override locally created entries, however.

3.10.20.2. ax25 route drop <target>

Drop an entry for target from the AX.25 routing table.

3.10.20.3. ax25 route mode <target> [vc | datagram | interface]

Sets the mode to vc | datagram | interface for target. **Interface** is the default for that interface. **Vc** is a virtual circuit (ax25 connected mode) and **datagram** is unconnected mode, (AX25 UI frames).

3.10.21. ax25 status [<axcb>]

Without an argument, display a one-line summary of each AX.25 control block. If the address of a particular control block is specified, the contents of that control block are dumped in more detail. Note that the send queue units are frames, while the receive queue units are bytes.

3.10.22. ax25 t3 [<milliseconds>]

Display or set the AX.25 idle "keep alive" timer. Value is in milliseconds.

3.10.23. ax25 t4 [<seconds>]

Display or set the AX.25 Link "redundancy" timer. Value is in seconds. When no exchange has been had during this time the link is reset and closed.

3.10.24. ax25 timertype [l | e | o]

Sets or displays the type of timer used for retransmission and recovery: **linear, exponential** or **original**.

3.10.25. ax25 version [1 | 2]

Display or set the version of the AX.25 protocol to attempt to use on new connections. The default is 1 (the version that does not use the poll/final bits).

3.10.26. ax25 window [<size>]

Set the number of bytes that can be pending on an AX.25 receive queue beyond which I frames will be answered with RNR (Receiver Not Ready) responses. This presently applies only to suspended interactive AX.25 sessions, since incoming I-frames containing network (IP, NET/ROM) packets are always processed immediately and are not placed on the receive queue. However, when an AX.25 connection carries both interactive and network packet traffic, an RNR generated because of backlogged interactive traffic will also stop network packet traffic from being sent.

3.11. bbs

Enter the local bbs port (same as a telnet session to your own station).

```
Current commands are:
help       ? <command>
area       a [<area>]
bye        b
connect    c <node>
           c <interface> <ax25_addr> [[<ax25_addr>] .... ]
download   d <filename>
           du <filename>          (for uuencode bin file)
escape     e [<escape char.]
finger     f [<user>]
help       h [<command>]
info       i
jheard     j
kill       k <msg number>
list       l [<msg number>]
mbusers    m
nodes      n <node>
operator   o
ports      p
read       r [<msg number>]
send       s[f] <usename[@hostid] [<from_addr>] [$bul_id]
send repl sr [<mgs number>]
telnet     t <hostid>
upload     u <filename>
verbose    v <n>
what       w [<directory>]
expert     x
zap        z <filename>
sysop      @
```

3.12. bootp <subcommand>

This is a bootp server / client, included into **Nos**. It is picked up from the University of Michigan in Ann Arbor. It is included into the sources of **Nos** but is not used nor tested by me. (PA0GRI) Its usability is fague but the discussions on tcp-group drove me to include this for possible use of reusable addresses.

3.12.1. bootp start

Starts the bootp server.

3.12.2. bootp stop

Stops the bootp server.

3.12.3. bootp dns [<ipaddr>]

Display or set the list of domain name servers for **bootp**.

3.12.4. bootp dyip [<iface> | <iface> <ipaddr1> <ipaddr2> | <iface> off]

Display or set the interface address range. The range is between *ipaddr1* and *ipaddr2*, both in dot notation.

3.12.5. bootp host [<hostaddr> <hardware type> <hardware addr> <ip addr> [boot file]]

Display or activate a bootp process. *Hardware type* is netrom, ether, macappletalk or ax25. *Hardware addr* is the interface name. *Ip addr* must be in dot notation.

3.12.6. bootp rmhost <ipaddr>

Delete *ipaddr* from the host table.

3.12.7. bootp homedir [<directory> | default]

Display or set the directory where the **bootp** files reside. The default directory is bpfiles.

3.12.8. bootp defaultfile [<bootfile> | default]

Display or set the file name of the **bootp** file. The default is boot.

3.12.9. bootp logfile [<filename> | default] [on | off]

Starts or stops the logging of **bootp** requests to *filename* or the default file name *bootplog*.

3.12.10. bootp logscreen [on | off]

Enables or disables the logging of **bootp** to the screen.

3.13. bootpd

This starts the server daemon for **bootp**.

3.14. cd [<dirname>]

Change the current working directory, and display the new setting. Without an argument, **cd** simply displays the current directory without change. The **pwd** command is an alias for **cd**.

3.15. close [<session>]

Close the specified session; without an argument, close the current session. On an AX.25 session, this command initiates a disconnect. On a FTP or Telnet session, this command sends a FIN (i.e., initiates a close) on the session's TCP connection. This is an alternative to asking the remote server to initiate a close (**QUIT** to FTP, or the logout command appropriate for the remote system in the case of Telnet).

When either FTP or Telnet sees the incoming half of a TCP connection close, it automatically responds by closing the outgoing half of the connection. Close is more graceful than the **reset** command, in that it is less likely to leave the remote TCP in a "half-open" state.

3.16. cls

Clears the current session screen (command screen).

3.17. comm <**interface**> <**text-string**>

The **comm** command sends **text-string** via **interface**. This can be used to send straight text to an tnc still in TAPR command mode during **Nos** startup. Note that to preserve spaces tabs etc. to include the string between double quote characters. Aka: comm ax1 "start kiss"

3.18. connect <**iface**> <**ax25_addr**> [<**digipeater**> ...]

Initiate a "vanilla" AX.25 session to the specified ax25_addr using the specified interface. Data sent on this session goes out in conventional AX.25 packets with no upper layer protocol. The de-facto presentation standard format is used, in that each packet holds one line of text, terminated by a carriage return. A single AX.25 connection may be used for terminal-to-terminal, IP and NET/ROM traffic. The three types of data being automatically separated by their AX.25 Level 3 Protocol IDs.

Up to 7 optional digipeaters may be given; note that the word **via** is NOT needed. If digipeaters are specified, they are automatically added to the AX25 routing table as though the **ax25 route add** command had been given before issuing the **connect** command.

3.19. delete <**filename**>

The **filename** is removed from the file system.

3.20. detach <**iface**>

Detach a previously attached interface from the system. All IP routing table entries referring to this interface are deleted, and forwarding references by any other interface to this interface are removed.

3.21. dialer <**iface**> [<**file**> [<**seconds**> [<**pings**> [<**hostid**>]]]]

Setup an autodialer session for the interface. Whenever the interface is idle for the interval in <seconds>, the autodialer will ping the <hostid>. If there is no answer after <pings> attempts, the autodialer will execute the special commands contained in the <dialer-file>.

If the interval in <seconds> is zero, a previous dialer command process will be removed. If the number of <pings> is zero, the <dialer-file> will be executed without pinging the <hostid>.

The file may have any valid name, and must be located in the configuration root directory (see the **Installion** section). The commands in the file are described in the **Dialer Subcommands** chapter.
Commands in **file** are:

3.21.1. control <**up** | **down**>

3.21.2. send <**string**> [<**milliseconds**>]

Sends **string** to the interface. If **milliseconds** is given, inter character timing is **milliseconds** milliseconds.

3.21.3. speed <**bps**>

Displays or sets the current interface speed to **bps** baud.

3.21.4. status <up | down>

3.21.5. wait <milliseconds> [<string> [speed]]

Wait the amount of **milliseconds**. If **string** is given, incoming characters from the interface are compared with **string**. If an compare is found and **speed** is the string speed, the next numbers read from the interface is the new baudrate used. This works like HAYES response CONNECT 9600. The wait command could habe been "wait 10000 CONNECT speed". This waits 10 seconds for the CONNECT response from the modem.

3.22. dir [<dirname>]

List the contents of the specified directory on the console. If no argument is given, the current directory is listed. Note that this command works by first listing the directory into a temporary file, and then creating a **more** session to display it. After this completes, the temporary file is deleted.

3.23. disconnect [<session>]

An alias for the **close** command (for the benefit of AX.25 users).

3.24. domain <subcommand>

The domain commands control and show the working of the name to internet address mapping software. **NOS** currently only has a client with a simple file reading local server. A real server is needed to service the community for their growing needs.

3.24.1. domain addserver <hostid>

Add a domain name server to the list of name servers. Note that, when this command is given in the **autoexec.nos** file the **ip address** command should given be before this command is used. (If not, **Nos** will not kow how to resolve the address, and an answer will never be recognized, or worse: just plain hangs the system.)

3.24.2. domain cache <subcommand>

Following commands work on the domain cache. These are resource records (see RFC 1033/1034) held in memory.

3.24.2.1. domain cache clean [<yes | no>]

Displays or sets the discard of expired resource records. Expired records have their timeout value decremented to zero. Normaly resource records get a default timeout value of 1800 seconds. After this time they are considered "old" and if referenced again the domain name resolver should be enquired again. When clean is off (the default), expired records will be retained; if no replacement can be obtained from another domain name server, these records will continue to be used.

When clean is on, expired records will be removed from the file whenever any new record is added to the file.

3.24.2.2. domain cache list

This command shows the current content of the in memory cache for resource records.

3.24.2.3. domain cache size [<size>]

Display or set the nominal maximum size of the local memory cache. The default is 20.

(Note: The cache may be temporarily larger when waiting for new records to be written to the **domain.txt** file.)

3.24.2.4. domain cache wait [<seconds>]

Display or set the interval in seconds to wait for additional activity before updating the **domain.txt** file. The default is 300 seconds (5 minutes).

3.24.3. domain dropserver <hostid>

Remove a domain name server from the list of name servers. You are warned when you delete the last name server.

3.24.4. domain listservers

List the currently configured domain name servers, along with statistics on how many queries and replies have been exchanged with each one, response times, etc.

3.24.5. domain maxwait [<timeout>]

This sets a timeout value (1 to 255 seconds) to a query or domain name server. This is not set for a already defined server but will be used for a newly defined name server. Also the value is used for domain nslookups. Note that name servers can have (PC based) trouble finding records in an large database. The default is set to 30 seconds.

3.24.6. domain retry [<retries>]

The retry value (number) limits the number of queries send out to remote domain name resolvers before giving up and telling you that host xyzzy.ampr.org does not exist. The total time lost with a query is retries * timeout * number of domain servers defined.

3.24.7. domain suffix [<domain suffix> | none]

Display or specify the default domain name suffix to be appended to a host name when it contains no periods. For example, if the suffix is set to **ampr.org.** and the user enters **telnet ka9q**, the domain resolver will attempt to find **ka9q.ampr.org.**. If the host name being sought contains one or more periods, however, the default suffix is NOT applied if the last part of the name is less than 5 characters and contains only letters; e.g.,
telnet foo.bar would NOT be turned into **foo.bar.ampr.org.**.
telnet foo.ka9q will be turned into **foo.ka9q.ampr.org.**. Note that a trailing dot (.) is required for the suffix. If the suffix is the string **none** (without trailing period) the current suffix is cleared and forgotten.

3.24.8. domain trace [on | off]

Display or set the flag controlling the tracing of domain server requests and responses. Trace messages will be seen only if a domain name being sought is not found in the local cache file, **domain.txt**.

3.24.9. domain translate [off | on]

Display or set the flag that controls the translation of ip adress in dot notation into symbolic names. The translation process makes heavely use of reverse domain name lookups. Do not set this flag unless you have a good and fast connection to a domain name server or have a fast domain.txt handler and domain.txt contains all IN-ADDR.ARPA. records you ever wanted.

3.24.10. domain verbose [off | on]

Display or set the flag controlling the return of a full name (true) or only the first name (dot delimiter) (false). This is for IP address to name translation only.

3.24.11. domain xyzzy

This is (when spelled fully) a magic word to enable domain queries to outside domain servers when reading commands from the startup file. This should only be used by those who have reliable access to a domain name server.

3.25. drsistat

Shows the statistics for all configured drsi boards.

3.26. dump <hex-address | .> [decimal-range]

The dump command shows memory in hex and ascii. Hex-address is a 32 bit value for a PC split into page address and page offset. A splitting colon is not used nor accepted. If decimal-range is not given , 128 bytes are displayed. **dump .** displayes memory starting at the end of a previous dump command.

3.27. echo [accept | refuse]

Display or set the flag controlling client Telnet's response to a remote WILL ECHO offer.

The Telnet presentation protocol specifies that in the absence of a negotiated agreement to the contrary, neither end echoes data received from the other. In this mode, a Telnet client session echoes keyboard input locally and nothing is actually sent until a carriage return is typed. Local line editing is also performed: backspace deletes the last character typed, while control-U deletes the entire line.

When communicating from keyboard to keyboard the standard local echo mode is used, so the setting of this parameter has no effect. However, many timesharing systems (eg. UNIX) prefer to do their own echoing of typed input. (This makes screen editors work right, among other things). Such systems send a Telnet WILL ECHO offer immediately upon receiving an incoming Telnet connection request. If **echo accept** is in effect, a client Telnet session will automatically return a DO ECHO response. In this mode, local echoing and editing is turned off and each key stroke is sent immediately (subject to the Nagle tinygram algorithm in TCP). While this mode is just fine across an Ethernet, it is clearly inefficient and painful across slow paths like packet radio channels. Specifying **echo refuse** causes an incoming WILL ECHO offer to be answered with a DONT ECHO; the client Telnet session remains in the local echo mode. Sessions already in the remote echo mode are unaffected. (Note: Berkeley Unix has a bug in that it will still echo input even after the client has refused the WILL ECHO offer. To get around this problem, enter the **stty -echo** command to the shell once you have logged in.)

3.28. eol [unix | standard]

Display or set Telnet's end-of-line behavior when in remote echo mode. In standard mode, each key is sent as-is. In unix mode, carriage returns are translated to line feeds. This command is not necessary with all UNIX systems; use it only when you find that a particular system responds to line feeds but not carriage returns. Only SunOS release 3.2 seems to exhibit this behavior; later releases are fixed.

3.29. escape [<char>]

Display or set the current command-mode escape character in hex. On the PC, the escape character is default ^]. The alternate escape key is F10 unless F10 is redefined with **fkey**.

3.30. etherstat

Display 3-Com Ethernet controller statistics (if configured).

3.31. exit

Exit the *nos* program and return to MS-DOS.

3.32. finger <user@hostid> | <@hostid>

Issue a network finger request for user **user** at host **hostid**. This creates a client session which may be interrupted, resumed, reset, etc, just like a Telnet client session. If only @**hostid** is given, all users on that host are identified.

3.33. fkey [<number> [<string>]]

Fkey displays or sets values for the programmable keys on the PC keyboard. **fkey** alone gives a display of all remapable keys and their number. **fkey number** displays the current value for that key. **fkey number string** assigns string to that key. Control characters can be created by prefixing then with an ^ character. A cr is ^M. To insert an ^ in the string 2 ^'s next to each other are needed. Following is the map of keys and their number. F1 is function key 1. Sf1 is Shift function key 1. Cf1 is Control function key 1. Af1 is Alt function key 1. Etc. The right most row is the numeric keypad.

key	number	key	number	key	number	key	number	key	number
f1	59	sf1	84	cf1	94	af1	104	pgup	73
f2	60	sf2	85	cf2	95	af2	105	pgdn	81
f3	61	sf3	86	cf3	96	af3	106	home	71
f4	62	sf4	87	cf4	97	af4	107	end	79
f5	63	sf5	88	cf5	98	af5	108	arup	72
f6	64	sf6	89	cf6	99	af6	109	ardn	80
f7	65	sf7	90	cf7	100	af7	110	ar l	75
f8	66	sf8	91	cf8	101	af8	111	ar r	77
f9	67	sf9	92	cf9	102	af9	112	ins	82
f10	68	sf10	93	cf10	103	af10	113	del	83

The mapping is taken to look like a vt100 / ansi keyboard. Currently assigned value strings for following keys:

number	string	key
59	" 33OP"	/* F1 */
60	" 33OQ"	/* F2 */
61	" 33OR"	/* F3 */
62	" 33OS"	/* F4 */
71	" 10"	/* home*/
72	" 33[A"	/* up arrow*/
73	" 25"	/* pgup */
75	" 33[D"	/* left arrow */
77	" 33[C"	/* right arrow */
79	" 05"	/* end */
80	" 33[B"	/* down arrow */
81	" 12"	/* pgdn */
82	" 01"	/* ins */
83	" 177"	/* del */

3.34. ftp <hostid>

Open an FTP control channel to the specified remote host and enter converse mode on the new session. Responses from the remote server are displayed directly on the screen. See the**FTP Subcommands** chapter for descriptions of the commands available in a FTP session.

3.35. ftype [ascii | binary | image | logical <size>]

This command displays or sets the default start file mode (ascii or binary) for ftp transfers. If ftype binary or image is given the next ftp session started will be in binary type. No binary command is needed once the session is started. In case of logical, the "word" size is set to **size**.

3.36. help

Display a brief summary of top-level commands.

3.37. hop <subcommands>

These commands are used to test the connectivity of the network.

3.37.1. hop check <hostid>

Initiate a *hopcheck* session to the specified host. This uses a series of UDP "probe" packets with increasing IP TTL fields to determine the sequence of gateways in the path to the specified destination. This function is patterned after the UNIX **traceroute** facility.

ICMP message tracing should be turned off before this command is executed (see the **icmp trace** command).

3.37.2. hop maxttl [<hops>]

Display or set the maximum TTL value to be used in hop check sessions. This effectively bounds the radius of the search.

3.37.3. hop maxwait [<seconds>]

Display or set the maximum interval, in seconds, that a hopcheck session will wait for responses at each stage of the trace. The default is 5 seconds.

3.37.4. hop queries [<count>]

Display or set the number of UDP probes that will be sent at each stage of the trace. The default is 3.

3.37.5. hop trace [on | off]

Display or set the flag that controls the display of additional information during a hop check session.

3.38. hostname [<name>]

Display or set the local host's name. By convention this should be the same as the host's primary domain name. This string is used only in the greeting messages of the various network servers; note that it does NOT set the system's IP address.

If <name> is the same as an <iface> (see the **Attach commands** chapter), this command will search for a CNAME domain resource record which corresponds to the IP address of the <iface>.

3.39. hs

Display statistics about the HS high speed HDLC driver (if configured and active).

3.40. icmp <subcommand>

These commands are used for the Internet Control Message Protocol service.

3.40.1. icmp echo [on | off]

Display or set the flag controlling the asynchronous display of ICMP Echo Reply packets. This flag must be on for one-shot pings to work (see the **ping** command.)

3.40.2. icmp status

Display statistics about the Internet Control Message Protocol (ICMP), including the number of ICMP messages of each type sent or received.

3.40.3. icmp trace [on | off]

Display or set the flag controlling the display of ICMP error messages. These informational messages are generated by Internet routers in response to routing, protocol or congestion problems. This option should be turned off before using the **hop check** facility because it relies on ICMP Time Exceeded messages, and the asynchronous display of these messages will be mingled with **hop check** command output.

3.41. ifconfig

Display a list of interfaces, with a short status for each.

3.41.1. ifconfig [<iface> [[[<subcommand> <param>] <subcommand> <param>] ....]

When only iface is given, a extended interface status is displayed. Multiple subcommand / parameter can be put on one line.

3.41.2. ifconfig <iface> broadcast <addr>

Set the broadcast address of interface iface to **addr**. **Addr** can eighter be an ax25_addr or an ether_addr, depending on the interface type, with 1's in the host part of the address. This is related to the **netmask** sub-command. See also the **arp** command.

3.41.3. ifconfig <iface> description ["description"]

This command sets the interface description to the string specified. If no string is supplied, the current description is cleared. The description is displayed with the **ifconfig iface** command (no parameters) and with the mailbox commands.

3.41.4. ifconfig <iface> encapsulation <slip | ax25 | ether | encap | ppp>

Sets the encapsulation for interface iface to slip / ax25 / ether / encap / ppp.

3.41.5. ifconfig <iface> forward <iface-2>

When a forward is defined, all output for interface **iface** is redirected to the interface directed by **iface-2**. To remove the forward, set <iface-2> to <iface>.

3.41.6. ifconfig <iface> ipaddress <addr>

Set the IP address to **addr** for this interface. This might be nessesary when a system acts as a gateway. Like an system with IP address 44.137.1.8 has an Internet access via its ethernet. The Internet IP address could be 129.179.122.10. An **ifconfig ec0 ipaddress 129.179.122.10** sets the correct address for that interface. Now routing to that system will work. (Note that the 44.x.x.x address is NOT connected to the Internet.) See also the **hostname** and **ip address** commands.

3.41.7. ifconfig <iface> linkaddress <hardware-dependant>

Set the hardware dependant address for this interface. For AX.25 this can be the callsign, for ethernet a new ethernet address.

3.41.8. ifconfig <iface> mtu <param>

Set the maximum transfer unit to **param** octeds (bytes). See the **Setting ... MTU, MSS and Window** chapter for more information.

3.41.9. ifconfig <iface> netmask <address>

Set the sub-net mask for this interface. The <address> takes the form of an IP address with 1's in the network and subnet parts of the address, and 0's in the host part of the address. Sample: ifconfig ec0 netmask 0xffffff00 for a class C network (24 bits). This is related to the **broadcast** sub-command. See also the **route** command.

3.41.10. ifconfig <iface> rxbuf <size>

Set the receive buffer size.

3.42. info

Info gives information about the version of **Nos** currently running and its buildin configuration. The configuration info is build with defines in config.h in the source distribution. That way it gives automatically correct configuration information.

3.43. ip <subcommand>

These commands are used for the Internet Protocol service.

3.43.1. ip access <permit|deny|delete> <dest addr>[/<bits>] <ifname> [lowport [higport]]

Set or display access control for IP routing functions. This command implements router access functions to **NosFp. Permit enables dest-addr packets to be routed via ifname. Deny disables those. If lowport is not given, all ports are assumed. If only lowport is given, that port is only checked for permission If lowport and highport are given, that is the range of ports permitted/denied. Dest-addr can be the word all for all addresses possible. Lowport can be the word none for all ports. The ip access delete must match a previous defined permit or deny to be able to delete that definition. Some samples:**
ip access permit 44/8 ax0
ip access deny all ax0 1 1023
ip access permit all ax0 If no access list is created, all interfaces might carry all types. If an access control is defined for an interface there must be an permit defined for that interface to allow trafic. Thus an partial denyal without an permit is an complete denyal.

3.43.2. ip address [<hostid>]

Display or set the default local IP address. This command must be given before an **attach** command if it is to be used as the default IP address for the interface.

3.43.3. ip rtimer [<seconds>]

Display or set the IP reassembly timeout. The default is 30 seconds. **Value** is in seconds.

3.43.4. ip status

Display Internet Protocol (IP) statistics, such as total packet counts and error counters of various types.

3.43.5. ip ttl [<hops>]

Display or set the default time-to-live value placed in each outgoing IP datagram. This limits the number of switch hops the datagram will be allowed to take. The idea is to bound the lifetime of the packet should it become caught in a routing loop, so make the value slightly larger than the number of hops across the network you expect to transit packets. The default is set at compilation time to the official recommended value for the Internet.

3.44. isat [on | off]

Display or set the AT flag. Normaly this flag is set when an interface is attached with an interrupt of 8 or higher. This is to signal that the second interrupt controllerc in a AT also needs an return of interrupt signal. If an AT type clock is in use, this command will allow measurement of time in milliseconds, rather than clock ticks (55 milliseconds per clock tick). During I/O initialisation this flag is set if the monitor prom has the standard byte 0xfc at address f000:fffe.

3.45. kick [<session>]

Kick all sockets associated with a session; if no argument is given, kick the current session. Performs the same function as the **ax25 kick** and **tcp kick** commands, but is easier to type.

3.46. lock [password <"password string">]

Locks the keyboard or defines a password string. If **password** is given then <password> is saved as the unlock string. If no parameters are supplied, the keyboard becomes locked when a password was specified earlier. If the keyboard is locked, the password is requested. If a correct password is supplied, the keyboard becomes unlocked. The setting of the password and locking of the keyboard can only been done by the system console keyboard or **autoexec.nos** startup file. The password can not been shown.

3.47. log [stop | <filename>]

Display the current log filename or set the **filename** for logging server sessions. If **stop** is given as the argument, logging is terminated (the servers themselves are unaffected). If a file name is given as an argument, server session log entries will be appended to it.

3.48. lzw [<subcommand>]

Lzw is the data compression capability on some sockets. This command set defines or changes their definition. Normaly leave this set default.

3.48.1. lzw mode <fast|compact>

Displays or sets the compression methode used on data compression for specific sockets. Currently SMTP can use compression.

3.48.2. lzw bits <number>

Displays or sets the number of bits used for the compression size. The more bits defined the larger the table space needed. Range is 9 to 16.

3.49. mail

This command will start a shell escape command. The mailer used is defined with the DOS environment variable MAILER, wich defaults to **BM.EXE**.

3.50. mbox [<subcommand>]

Display the status of the mailbox server system (if configured).

3.50.1. mbox attend [yes | no]

Displays or sets the attended flag on. This is used to announce in the mailbox if the station mannager is willing to attend his station (chat).

3.50.2. mbox expert <on|off>

Displays or set the level expected of a mailbox user. When set (s)he gets a short (>) prompt. When not set a long prompt is send with the first letter of each mailbox command as prompt. The user can then set Xpert mode to get the short prompt.

3.50.3. mbox fwdinfo ["forward info"]

Displays or sets the mailbox forward info to be included in the R: line for forwarded BBS bulletins. An empty string ("") as info clears the info field.

Sample: netrom fwdinfo "HNLNET BBS"

Will show '[HNLNET BBS]' in the R: line.

3.50.4. mbox haddress ["home-address"]

Displays or sets the home address field to be included in the R: line for forwarded BBS bulletins. An empty string ("") clears the field.

Sample: netrom haddress "#CRV.OR.USA"

Will show '@WG7J#CRV.OR.USA' in the R: line. (when ax mycall is WG7J)

3.50.5. mbox jumpstart <on|off>

Displays or sets the mailbox jumpstart code. When set and an "known" node connects to the mailbox, no extra line needs to be send to activate the mailbox but the prompt is send directly. Warning: When set, it takes a while to hear all nodes and an AX25 connect could have an wrong start as there is no wait for a level 3 protocol check. Especialy RSPF in virtual mode could course serious problems.

3.50.6. mbox kick

To kick the mailbox back in activity after retry timeouts.

3.50.7. mbox maxmsg

To display or set the maximum number of messages per area when an notesfile is shown to a user. This reserves lots of memory for every mailbox session.

3.50.8. mbox motd ["message string"]

Display ot set the mailbox welcoming Message Of The Day.

3.50.9. mbox nrid <on|off>

DIsplays or sets the Netrom id flag. When set the node id is displayed on the prompt line.

3.50.10. mbox operator [<address>]

The operator command shows or sets an alternate address for the control operator. When set and mbox attend is off or the user is "shelled out" and an mbox user requests a Operator command, the ttylink session is redirected to the address specified. Else it works as before that the mbox user is notified with "Unattended". When the system pointed to by address has also set its attended to off, then an Unattended is replied with as well (but generated by the remote station). This is a "flaky" way of having an attended unattended system.

3.50.11. mbox password <"password string">

Sets a password string to be presented to Sysops entering that mode in the mailbox (using the @ command and having that priviledge set to their login name/password in **f/ftpusers**). When a password is defined (max 30 characters) then (s)he is prompted with 5 numbers before letting her/him in. The five numbers represent the 5 character locations in the string defined, whereby the first character is number 0. Multiple lines of 5 characters can be send to fool "snoopers". The end of password sending is signaled with an empty line. If there was a good response, sysop mode is entered. The setting of the password can only been done by the system console or **autoexec.nos** startup file. The password can not been shown.

3.50.12. mbox qth ["qth info"]

Sample: netrom qth "Driebruggen, NL"

Displays or sets the info for qth in the R: line for forwarded BBS messages.

3.50.13. mbox secure <yes|no>

Displays or sets the security option for mailbox gateway users. If set users comming in via telnet to the bbs are not allowed to use the gateway. If not set anybody can use the gateway. (Note: No check for Bozo's). Also the mailbox send command is disabled except for ax25 and netrom connects.

3.50.14. mbox smtptoo <yes|no>

Displays or sets the flag to include SMTP headers in BBS messages. When set SMTP headers are included in messages. When not set not included.

3.50.15. mbox status

An alias for just entering mbox on the prompt line.

3.50.16. mbox timer [<seconds>]

Display the current interval and time remainig or set the mailbox forwarding timer.

3.50.17. mbox tiptimeout

Displays or sets the timeout value for tip connection timeout. After **timeout** seconds of no activity the connection is closed.

3.50.18. mbox trace [yes | no]

Displays or sets mailbox forward trace code flag. The trace is very minimal but every one working on mailbox/forward code now has a common flag **Mtrace**.

3.50.19. mbox utc <offset>

Displays or sets the offset you are according to ZULU time. Positive and negative numbers are supported and full leap year and month calculation are done. This is used by the mailbox forwarding to give a standard time in the R: line.

3.50.20. mbox zipcode zip

Displays or sets the info for the zip field for R: line BBS header lines. This field is max 7 characters long as released. For USA the zip is only 6 numbers long. Netherlands has 4 figures a space and 2 letters. Guess every PTT wants something they "invented" their own.

Sample: netrom zip "3465 TJ" or: netrom zip 54551

3.51. memory <subcommand>

These commands are used for memory allocation.

3.51.1. memory debug [on|off]

Displays or sets the memory allocator debug flag. If set debug information is written to the log file containing most flags and parameters from the memory allocation routines.

3.51.2. memory efficient [yes | no]

Displays or sets the search algorithm for buffer memory. When set the search is always started from the beginning of the free list. This is slower but keeps memory fragmentation to a minimum. When clear search is started on the end, accounting for more memory fragmentation but keeping speed. Roumers

state that you should include this command as the first line in **autoexec.nos**. The current releases default sets efficient to yes.

3.51.3. memory freelist

Display the storage allocator free list. Each entry consists of a starting address, in hex, and a size, in decimal bytes.

3.51.4. memory ibufsize [<size>]

Display or set the size of the buffers in the interrupt buffer pool. The size should be set to the largest type of buffer plus a header size of 22. For example: If your ax25 is the omly interface and a packet length of 256 is defined, the ibufsize should be 256 + 10 * 6 + 22. The 10 * 6 is the ax25 header (source , destination and 8 digipeaters).

3.51.5. memory minheap [<number>]

Displays or sets the minimum heap size to be allocated before shell escaping to DOS. This assures a free heap so that **Nos** can run trough without gettig short on memory for a while.

3.51.6. memory nibufs [<number>]

Display or set the number of interrupt buffer pool buffers. If the number of buffers is set, the statistics in the **memory status** display are reset for number of interrupt buffer fails. The minimum available value is set to the requested number of buffers. A roule of thumb for the number of buffers is to watch the statistics and keep a minimum of 2 free buffers. Increase or decrease as required.

3.51.7. memory sizes

Display a histogram of storage allocator request sizes. Each histogram bin is a binary order of magnitude (i.e., a factor of 2).

3.51.8. memory status

Display a summary of storage allocator statistics. The first line shows the base address of the heap, its total size, the amount of heap memory available in bytes and as a percentage of the total heap size, and the amount of memory left over (i.e., not placed on the heap at startup) and therefore available for **shell** subcommands.

The second line shows the total number of calls to allocate and free blocks of memory, the difference of these two values (i.e., the number of allocated blocks outstanding), the number of allocation requests that were denied due to lack of memory, and the number of calls to free() that attempted to free garbage (eg. by freeing the same block twice or freeing a garbled pointer).

The third line shows the garbage collection status. Garbage collection is a separate process running every second. If memory is "low" garbage collection routines are called in "Red alert" mode, else in "Yellow alert".

The fourth line shows the number of calls to malloc and free that occurred with interrupts off. In normal situations these values should be zero. The fourth line shows statistics for the special pool of fixed-size buffers used to satisfy requests for memory at interrupt time. The variables shown are the number of buffers currently in the pool, their size, and the number of requests that failed due to exhaustion of the pool.

The fifth line shows the current setting of the interrupt buffer pool, its minimal value and the number of no buffers available. These statistics are reset when a **memory nibufs <number>** command is given.

3.51.9. memory thresh [<size>]

Displays or sets the memory threshold size in bytes. If memory gets below this value , no more new sessions are started or accepted.

3.52. mkdir <dirname>

Create a sub-directory in the current working directory.

3.53. mode <iface> [vc | datagram]

Control the default transmission mode on the specified AX.25 interface. In **datagram** mode, IP packets are encapsulated in AX.25 UI frames and transmitted without any other link level mechanisms, such as connections or acknowledgements.

In **vc** (virtual circuit) mode, IP packets are encapsulated in AX.25 I frames and are acknowledged at the link level according to the AX.25 protocol. Link level connections are opened if necessary.

In both modes, ARP is used to map IP to AX.25 addresses. The defaults can be overridden with the type-of-service (TOS) bits in the IP header. Turning on the "reliability" bit causes I frames to be used, while turning on the "low delay" bit uses UI frames. (The effect of turning on both bits is undefined and subject to change).

In both modes, IP-level fragmentation is done if the datagram is larger than the interface MTU. In virtual circuit mode, however, the resulting datagram (or fragments) is further fragmented at the AX.25 layer if it (or they) are still larger than the AX.25 **paclen** parameter. In AX.25 fragmentation, datagrams are broken into several I frames and reassembled at the receiving end before being passed to IP. This is preferable to IP fragmentation whenever possible because of decreased overhead (the IP header isn't repeated in each fragment) and increased robustness (a lost fragment is immediately retransmitted by the link layer).

3.54. more <file> [<file> ...]

Display the specified file(s) a screen at a time. To proceed to the next screen, press the space bar; to cancel the display, hit the 'q' key. A CR advances a line. The **more** command creates a session that you can suspend and resume just like any other session.

3.55. motd [<"message of the day">]

Display the current message or sets the welcome message to the defined string.

3.56. multitask [on | off]

Displays or sets the multi task flag. If set **Nos** continues to service its interrupts and handling of normal operation when put in the background by executing a shell escape. Some programs do not like to miss valuable time taken away this way and weard results can happen. Clearing the flag stops all activiity fron **Nos** when put in the background.

3.57. netrom <subcommand>

3.57.1. netrom acktime [<milliseconds>]

Displays or sets the ack delay timer, similary to ax25 t2.

3.57.2. netrom alias <alias>

This sets the **netrom** alias for this station. Other stations can connect to the ax25 callsign and to the Netrom alias (when set). The alias is broadcasted with a Netrom broadcast.

3.57.3. netrom bcnodes <iface>

Initiates an immediate broadcast of nodelist on **iface**.

3.57.4. netrom call <call>

Displays or sets the call to be used by the netrom interface. Note: this is a shortcut for the "ifconfig netrom linkaddress call" command.

3.57.5. netrom connect <node>

Start a connection to **node**.

3.57.6. netrom choketime [<milliseconds>]

Display or set the time breaking a send choke. Choke is the term netrom uses for flow control conditions.

3.57.7. netrom derate [on | off]

Display or set automatic derating of netrom routes on link failure.

3.57.8. netrom interface <iface> <quality>

Define a netrom interface **iface**. The quality can be between 1 and 255. For compatibility with other implementations set this to 192 normaly. A check is made that the interface is of type CL_AX25, meaning NETROM capable media. Note: the alias is not specified on this line anymore. Use netrom alias command.

3.57.9. netrom irtt [<milliseconds>]

Display or set the initial round trip time.

3.57.10. netrom kick <&nrcb>

Give the control block a kick to get activity going again.

3.57.11. netrom load [<filename>]

When fully implemented reloads a saved netrom node list to continue from where you left when you saved the list. All entry's read are time decremented as if normal time just went on. If it took some time (**obsotimer** * x) your list could be empty as they all timed out.

3.57.12. netrom minquality [<value>]

Display or set the minimum quality for recognizing a node entry. Entry's below this value are not considered valuable for usage.

3.57.13. netrom nodefilter <subcommand>

Node filtering.

3.57.13.1. netrom nodefilter add <neighbor> <iface>

Allow **neighbor** to give us node updates.

3.57.13.2. netrom nodefilter drop <neighbor> <iface>

Disallow **neighbor** to give us node updates.

3.57.13.3. netrom nodefilter mode [none | accept | reject]

Display or set the initial filter schema. **None** accepts all. **Accept**,accepts only from nodes defined in **netrom nodefilter add** lists. **Reject**,no accepts from nodes defined in **netrom nodefilter add** lists.

3.57.14. netrom nodetimer [<seconds>]

Display or set the interval our node list is transmitted.

3.57.15. netrom obsotimer [<seconds>]

Display or set the time a node quality gets decremented.

3.57.16. netrom promiscious [off | on]

Enables nodes with a higher quality than defined with minquality. If on, all nodes are received, independent od nodefilter.

3.57.17. netrom qlimit [<bytes>]

Display or set the maximum queue limit for chat. Similar to ax25 window.

3.57.18. netrom reset <&nrcb>

Remove the control block. You can find the control block with **socket**.

3.57.19. netrom retries [<value>]

Display or set the maximum number of retries on connect, disconnect or data.

3.57.20. netrom route <subcommand>

Netrom routing commands.

3.57.20.1. netrom route add <alias> <destination> <iface> <quality> <neighbor>

Add a netrom route.

3.57.20.2. netrom route drop <destination> <neighbor> <iface>

Delete a netrom route.

3.57.20.3. netrom route info <destination>

Display the route it would take to get to **destination**.

3.57.21. netrom status

Displays all netrom connections.

3.57.22. netrom save [<filename>]

When fully implemented, saves the current netrom node list in memory to /netrom.sav or filename, if given.

3.57.23. netrom timertype [exponential | linear]

Displays or sets the type of backoff.

3.57.24. netrom ttl [<hops>]

Display or set the maximum number of hops a frame might hop before discarded, if it did not reach its destination before.

3.57.25. netrom user [<username>]

Display or set our netrom user name. This is used on outgoing connects.

3.57.26. netrom verbose [off | on]

Display or set the verbose flag. If set all nodes known to us are transmitted each time the **nodetimer** expires. If off, only your stations node ident is send out.

3.57.27. netrom window [<frames>]

Display or set the size of the sliding window. This is the largest send and receive window we might negotiate.

3.58. nntp <subcommand>

Network news transfer protocol has the following subcommands:

3.58.1. nntp addserver <nntpserver> <seconds> [<range>] [<groups>]

Add a nntp news server to quiry every **seconds** interval for new articles specified by the nntp groups command. Range can be a limit in time to query like *nntp addserver someserver 600 22:00-23:00*, to only query the server between 10 and 11pm. Multiple nntp add can be used co concatenate groups (up to 512 bytes) The interval **seconds** can be set to 0, so that normaly the nntp client does not query the server regulary but an *nntp kick nntpserver* can be used to start a session.

3.58.2. nntp directory <directory>

Show or set default directory for spooling news. This is not the control directory but the (alternate) place to spool articles to.

3.58.3. nntp dropserver <nntpserver>

3.58.4. nntp groups <group> [<group> ...]

Set or display currently set newsgroups.

3.58.5. nntp kick <nntpserver>

Kick the client to get in touch with the named server.

3.58.6. nntp listservers

Lists the curently defined servers.

3.59. nntp quiet [yes | no]

Shows or sets the flag controlling the display of newly received nntp news messages. The smpt quiet fag controls the beep following the message. If both are off, nothing is shown, keeping the display from getting hung on -more- when running unattended.

3.60. nntp trace <level>

Sets or shows the current trace level for the nntp client. 0 No tracing 1 (default) displays serious errors. 2 displays as 1 and transient errors. 3 displays as 2 and session progress. 4 displays as 3 and actual received articles. 5 displays errors.

3.61. nrstat

Displays the netrom interface statistics.

3.62. param <iface> [<param> ...]

Invoke a device-specific control routine. On a KISS TNC interface, this sends control packets to the TNC. Data bytes are treated as decimal. For example, **param ax0 1 255** will set the keyup timer (type field = 1) on the KISS TNC configured as ax0 to 2.55 seconds (255 x .01 sec). On a SLIP interface, the **param** command allows the baud rate to be read (without arguments) or set. On asy lines DTR and RTS can be (and possibly should be) set with param iface dtr 1 and param iface rts 1. The implementation of this command for the various interface drivers is incomplete and subject to change.
Current commands defined are:

0 Data - Normal KISS data
1 TxDelay - set TX delay for TNC
2 Persist - set persistence delay for TNC
3 Slottime - set slottime delay for TNC
4 TxTail - set TX tail delay for TNC

```
  5  Fulldup     - set Full duplex for TNC
  6  Hardware    - Hardware dependent
  7  TxMute      -
  8  DTR         - 0 = down, 1 = up
  9  RTS         - 0 = down, 1 = up
 10  Speed       - baud rate
 11  Enddelay    -
 12  Group       -
 13  Idle        -
 14  Min         -
 15  Maxkey      -
 16  Wait        -
 17  Parity      - 0 = none, 1 = even, 2 = odd parity
129  Down        -
130  Up          -
254  Return2     - Some tnc's need this
255  Return      - Reset a tnc from KISS into command mode
```

3.63. ping <hostid> [<length> [<milliseconds> [<incflag>]]]

Ping (send ICMP Echo Request packets to) the specified host. By default the data field contains only a small timestamp to aid in determining round trip time; if the optional **length** argument is given, the appropriate number of data bytes (consisting of hex 55) are added to the ping packets.

If **milliseconds** is specified, pings will be repeated indefinitely at the specified number of milliseconds; otherwise a single, "one shot" ping is done. Responses to one-shot pings appear asynchronously on the command screen, while repeated pings create a session that may be suspended and resumed. Pinging continues until the session is manually reset.

The **incflag** option causes a repeated ping to increment the target IP address for each ping; it is an experimental feature for searching blocks of IP addresses for active hosts.

3.64. pop <subcommand>

The **pop** client provides an automatic interface to the **pop** server that is totaly transparen to the user; all that the user needs to do is to set a few key parameters and the client will take it from there, appart from registring with the server station manager for setting up the user data. (see *pop userdata* command. To read more about *pop*, read rfc937. The following paragraph describe the **pop** subcommands.

3.64.1. pop mailbox [<name>]

Show the defined name or sets the name to **name** for POP mail queries to the mailhost. The server keeps the mail in the mailbox **name** located in /spool/mail/**name**.txt. This is a mandatory parameter.

3.64.2. pop mailhost [<ipaddr>]

Show the currently defined host address for a **pop** server or sets the IP address to the system named **ipaddr** for POP connections. This is a mandatory parameter.

3.64.3. pop kick

Start a POP connection now. (one shot connection)

3.64.4. pop timer [<seconds>]

Show the current setting and time remaining untill next server enquiry or when **seconds** is given, defines the interval that a **pop** connection is started every <seconds> to retrieve mail from the server system. When the timer is not set the client is only started with a *pop kick* command.

3.64.5. pop userdata <name> <password>

Sets up the userdata used for validation on the POP server system. Note that <name> and <password> are case sensitive. When only pop userdata is entered to show the values only the <name> is shown. (ultimate in security, if you realy want to see the password look in the "ether"). Remember: the name and password should also be defined at the server site in the *popusers* file.

3.65. popmail <subcommand>

popmail is a newer implementation of the pop client/server. It can handle both pop2 and pop3 type pop client/servers. For functionality see pop but the subcommands are different.

3.65.1. popmail addserver <host> [<seconds>] [hh:mm-hh:mm] <protocol> <mailbox> <username> <password>

Add **hostP as a pop server. When seconds is given, a timer is started to query the host with that interval for mail. If not specified no quering to the pop host will be started. You have to do that manualy with a kick. When hh:mm-hh:mm is given then only in that exact timeframe are queries to the host made (allowed). Protocol is eighter POP2 or POP3, depending on the mail service the host is providing. Note: pop2 is superceded by pop3. Mailbox is the mailbox name on the host where fore mail has to be picked up. Username and password are this systems validation parameters for the host. Note that on entering this command the host name is looked up. If non exsistant an error message is presented.**

3.65.2. popmail dropserver <host>

Drops **host** from the list of pop servers to be queried. All references to the entry are deleted from the current system.

3.65.3. popmail kick <host>

Starts a pop session with **host** to retrieve mail. This command is needed when no interval is specified with the **popmail addserver** command.

3.65.4. popmail list

Lists the current popmail server table.

3.65.5. popmail quiet <yes|no>

Displays or sets the notification of new mail ariving via pop.

3.65.6. popmail trace <level>

Displays or sets the trace level of pop sessions. Current trace levels are:
0 - No tracing
1 - Serious errors reported
2 - Transient errors reported
3 session progress reported Note that tracing only goes to the log file.

3.66. ppp <subcommands>

These commands are used for Point to Point Protocol interfaces.

This implementation of PPP is designed to be as complete as possible. Because of this, the number of options can be rather daunting. However, a typical PPP configuration might include the following commands:

```
attach asy 0x3f8 4 ppp pp0 4096 1500 9600
dial pp0 30 <hostid> 3 dialer.pp0
#
ppp pp0 lcp local accm 0
ppp pp0 lcp local compress address on
ppp pp0 lcp local compress protocol on
ppp pp0 lcp local magic on
ppp pp0 lcp open active
#
ppp pp0 ipcp local compress tcp 16 1
ppp pp0 ipcp open active
#
route add default pp0
```

3.66.1. ppp <iface>

Display the status of the PPP interface.

3.66.2. ppp <iface> lcp ...

These commands are used for the LCP [Link Control Protocol] configuration.

3.66.2.1. ppp <iface> lcp close

Shutdown the PPP interface.

3.66.2.2. ppp <iface> lcp local ...

These commands control the configuration of the local side of the link. If an option is specified, the parameters will be used as the initial values in configuration requests. If not specified, that option will not be requested.

For each of these options, the **allow** parameter will permit the remote to include that option in its response, even when the option is not included in the request. By default, all options are allowed.

3.66.2.2.1. ppp <iface> lcp local accm [<bitmap> | allow [on | off]]

Display or set the Async Control Character Map. The default is 0xffffffff.

3.66.2.2.2. ppp <iface> lcp local authenticate [pap | none | allow [on | off]]

Display or set the authentication protocol. The default is **none**.

3.66.2.2.3. ppp <iface> lcp local compress address/control [on | off | allow [on | off]]

Display or set the option to compress the address and control fields of the PPP HLDC-like header. This is generally desirable for slow asynchronous links, and undesirable for fast or synchronous links. The default is off.

3.66.2.2.4. ppp <iface> lcp local compress protocol [on | off | allow [on | off]]

Display or set the option to compress the protocol field of the PPP HLDC-like header. This is generally desirable for slow asynchronous links, and undesirable for fast or synchronous links. The default is off.

3.66.2.2.5. ppp <iface> lcp local magic [on | off | <value> | allow [on | off]]

Display or set the initial Magic Number. The default is off (zero).

3.66.2.2.6. ppp <iface> lcp local mru [<size> | allow [on | off]]

Display or set the Maximum Receive Unit. The default is 1500.

3.66.2.2.7. ppp <iface> lcp local default

Reset the options to their default values.

3.66.2.3. ppp <iface> lcp open active | passive

Wait for the physical layer to come up. If **active**, initiate configuration negotiation. If **passive**, wait for configuration negotiation from the remote.

3.66.2.4. ppp <iface> lcp remote ...

These commands control the configuration of the remote side of the link. The options are identical to those of the local side. If an option is specified, the parameters will be used in responses to the remote's configuration requests. If not specified, that option will be accepted if it is allowed.

For each of these options, the **allow** parameter will permit the remote to specify that option in its request. By default, all options are allowed.

3.66.2.5. ppp <iface> lcp timeout [<seconds>]

Display or set the interval to wait between configuration or termination attempts. The default is 3 seconds.

3.66.2.6. ppp <iface> lcp try ...

These commands are used for the various counters.

3.66.2.6.1. ppp <iface> lcp try configure [<count>]

Display or set the number of configuration requests sent. The default is 10.

3.66.2.6.2. ppp <iface> lcp try failure [<count>]

Display or set the number of bad configuration requests allowed from the remote. The default is 5.

3.66.2.6.3. ppp <iface> lcp try terminate [<count>]

Display or set the number of termination requests sent before shutdown. The default is 2.

3.66.3. ppp <iface> ipcp ...

These commands are used for the IPCP [Internet Protocol Control Protocol] configuration.

The **close, open, timeout** and **try** sub-commands are identical to the LCP (described above).

3.66.3.1. ppp <iface> ipcp local ...

These commands control the configuration of the local side of the link. If an option is specified, the parameters will be used as the initial values in configuration requests. If not specified, that option will not be requested.

For each of these options, the **allow** parameter will permit the remote to include that option in its response, even when the option is not included in the request. By default, all options are allowed.

3.66.3.1.1. ppp <iface> ipcp local address [<hostid> | allow [on | off]]

Display or set the local address for negotiation purposes. If an address of 0 is specified, the other side of the link will supply the address. By default, no addresses are negotiated.

3.66.3.1.2. ppp <iface> ipcp local compress [tcp <slots> [<flag>] | none | allow [on | off]]

Display or set the compression protocol. The default is **none**.

The **tcp** <slots> specifies the number of "conversation" slots, which must be 1 to 255. (This may be limited at compilation time to a smaller number.) A good choice is in the range 4 to 16.

The **tcp** <flag> is 0 (don't compress the slot number) or 1 (OK to compress the slot number). KA9Q can handle compressed slot numbers, so the default is 1.

3.66.3.2. ppp <iface> ipcp remote ...

These commands control the configuration of the remote side of the link. The options are identical to those of the local side. If an option is specified, the parameters will be used in responses to the remote's configuration requests. If not specified, that option will be accepted if it is allowed.

For each of these options, the **allow** parameter will permit the remote to specify that option in its request. By default, all options are allowed.

3.66.4. ppp <iface> pap ...

These commands are used for the PAP [Password Authentication Protocol] configuration.

The **timeout** and **try** sub-commands are identical to the LCP (described above). However, the terminate counter is unused.

3.66.4.1. ppp <iface> pap user [<username> [<password>]]

Display or set the username (the password may be set, but not displayed). When the username is specified, but no password is supplied, the **ftpusers** file is searched for the password. When a username/password is unknown or rejected, a session will appear at the console to prompt for a new username/password.

3.66.5. ppp <iface> trace [<flags>]

Display or set the flags that control the logging of information during PPP link configuration.

The flag value is 0 for none, 1 for basic, and 2 for general. Values greater than 2 are usually not compiled, and are described in the appropriate source files where they are defined.

3.67. ps

Display all current processes in the system. The fields are as follows:

PID - Process ID (the address of the process descriptor).

SP - The current value of the process stack pointer.

stksize - The size of the stack allocated to the process.

maxstk - The apparent peak stack utilization of this process. This is done in a somewhat heuristic fashion, so the numbers should be treated as approximate. If this number reaches or exceeds the stksize figure, the system is almost certain to crash; the *nos.exe* program should be recompiled to give the process a larger allocation when it is started.

event - The event this task is waiting for, if it is not runnable.

fl - Process status flags. There are three: I (Interrupts enabled), W (Waiting for event) and S (suspended). The I flag is set whenever a task has executed a pwait() call (wait for event) without first disabling hardware interrupts. Only tasks that wait for hardware interrupt events will turn off this flag; this is done to avoid critical sections and missed interrupts. The W flag indicates that the process is waiting for an event; the **event** column will be non-blank. Note that although there may be several runnable processes at any time (shown in the **ps** listing as those without the W flag and with blank event fields) only one process is actually running at any one instant (The Refrigerator Light Effect says that the **ps** command is always the one running when this display is generated.)

3.68. pwd [<dirname>]

An alias for the **cd** command.

3.69. rarp <subcommand>

This uses the reverse address resolution protocol.

3.69.1. rarp query <iface> <callsign>

This starts a reverse resolution request via **iface** to find the IP address for **callsign**. It counts down for 10 seconds before giving up listening for an answer.

3.70. record [off | <filename>]

Append to **filename** all data received on the current session. Data sent on the current session is also written into the file except for Telnet sessions in remote echo mode. The command **record off** stops recording and closes the file.

3.71. remote [-p <port>] [-k <key>] [-a <kickaddr>] <hostid> exit | reset | kick

Send a UDP packet to the specified host commanding it to exit the **Nos** program, reset the processor, or force a retransmission on TCP connections. For this command to be accepted, the remote system must be running the **remote** server and the port number specified in the **remote** command must match the port number given when the server was started on the remote system. If the port numbers do not match, or if the remote server is not running on the target system, the command packet is ignored. Even if the command is accepted there is no acknowledgement.

The **kick** command forces a retransmission timeout on all TCP connections that the remote node may have with the local node. If the -a option is used, connections to the specified host are kicked instead. No key is required for the kick subcommand.

The **exit** and **reset** subcommands are mainly useful for restarting the nos program on a remote unattended system after the configuration file has been updated. The remote system should invoke the **Nos** program automatically upon booting, preferably in an infinite loop. For example, under MS-DOS the boot disk should contain the following in **autoexec.bat**:

```
:loop
nos
goto :loop
```

3.71.1. remote -s <key>

The exit and reset subcommands of remote require a password. The password is set on a given system with the **-s** option, and it is specified in a command to a remote system with the **-k** option. If no password is set with the **-s** option, then the exit and reset subcommands are disabled.

Note that **remote** is an experimental feature in NOS; it is *not* yet supported by any other TCP/IP implementation.

3.72. rename <oldfilename> <newfilename>

Rename **oldfilename** to **newfilename**.

3.73. reset [<session>]

Reset the specified session; if no argument is given, reset the current session. This command should be used with caution since it does not reliably inform the remote end that the connection no longer exists. (In TCP a reset (RST) message will be automatically generated should the remote TCP send anything after a local reset has been done. In AX.25 the DM message performs a similar role. Both are used to get rid of a lingering half-open connection after a remote system has crashed.)

3.74. rip <subcommand>

These commands are used for the RIP service.

3.74.1. rip accept <gateway>

Remove the specified gateway from the RIP filter table, allowing future broadcasts from that gateway to be accepted.

3.74.2. rip add <hostid> <seconds> [<flags>]

Add an entry to the RIP broadcast table. The IP routing table will be sent to **hostid** every interval of **seconds**. If **flags** is specified as 1, then "split horizon" processing will be performed for this destination. That is, any IP routing table entries pointing to the interface that will be used to send this update will be removed from the update. If split horizon processing is not specified, then all routing table entries except those marked "private" will be sent in each update. (Private entries are never sent in RIP packets).

Triggered updates are always done. That is, any change in the routing table that causes a previously reachable destination to become unreachable will trigger an update that advertises the destination with metric 15, defined to mean "infinity".

Note that for RIP packets to be sent properly to a broadcast address, there must exist correct IP routing and ARP table entries that will first steer the broadcast to the correct interface and then place the correct link-level broadcast address in the link-level destination field. If a standard IP broadcast address convention is used (eg. 128.96.0.0 or 128.96.255.255) then chances are you already have the necessary IP routing table entry, but unusual subnet or cluster-addressed networks may require special attention. However, an **arp add** command will be required to translate this address to the appropriate link level broadcast address. For example,

arp add 128.96.0.0 ethernet ff:ff:ff:ff:ff:ff

for an Ethernet network (ip address is a sample only!!!), and

arp add 44.255.255.255 ax25 qst-0

for an AX25 packet radio channel. (If multiple AX25 interfaces, make a unique address for each interface.)

3.74.3. rip drop <dest>

Remove an entry from the RIP broadcast table.

3.74.4. rip merge [on | off]

This flag controls an experimental feature for consolidating redundant entries in the IP routing table. When rip merging is enabled, the table is scanned after processing each RIP update. An entry is considered redundant if the target(s) it covers would be routed identically by a less "specific" entry already in the table. That is, the target address(es) specified by the entry in question must also match the target addresses of the less specific entry and the two entries must have the same interface and gateway fields. For example, if the routing table contains

Dest	Len	Interface	Gateway	Metric	P	Timer	Use
1.2.3.4	32	ethernet0	128.96.1.2	1	0	0	0
1.2.3	24	ethernet0	128.96.1.2	1	0	0	0

then the first entry would be deleted as redundant since packets sent to 1.2.3.4 will still be routed correctly by the second entry. Note that the relative metrics of the entries are ignored.

3.74.5. rip refuse <gateway>

Refuse to accept RIP updates from the specified gateway by adding the gateway to the RIP filter table. It may be later removed with the **rip accept** command.

3.74.6. rip request <gateway>

Send a RIP Request packet to the specified gateway, causing it to reply with a RIP Response packet containing its routing table.

3.74.7. rip status

Display RIP status, including a count of the number of packets sent and received, the number of requests and responses, the number of unknown RIP packet types, and the number of refused RIP updates from hosts in the filter table. A list of the addresses and intervals to which periodic RIP updates are being sent is also shown, along with the contents of the filter table.

3.74.8. rip trace [0 | 1 | 2]

This variable controls the tracing of incoming and outgoing RIP packets. Setting it to 0 disables all RIP tracing. A value of 1 causes changes in the routing table to be displayed, while packets that cause no changes cause no output. Setting the variable to 2 produces maximum output, including tracing of RIP packets that cause no change in the routing table.

3.74.9. rip ttl <seconds>

Displays or sets the time to live timer to **seconds**. Normal timeout value is 240 seconds. This is not the ttl in a rip broadcast (16 = infinite). Set this timer before starting rip. Change this timer only in cooperation with your surrounding nodes. Default is 240 seconds.

3.75. rlogin host

Sets up an rlogin session via port 511 to an *NIX compatible station. Defaut terminal is an ansi (as defined with the fkeys) compatible terminal. Default user name is guest. (Redefine with set USER= environement variable).

3.76. rmdir <dirname>

Remove a sub-directory from the current working directory.

3.77. route

With no arguments, **route** displays the IP routing table.

3.77.1. route add <dest_hostid>[/bits] | default <iface> [<gateway_hostid> | direct [<metric>]]

This command adds an entry to the routing table. It requires at least two more arguments, the hostid of the target destination and the name of the interface to which its packets should be sent. If the destination is not local, the gateway's hostid should also be specified. (If the interface is a point-to-point link, then **gateway_hostid** may be omitted even if the target is non-local because this field is only used to determine the gateway's link level address, if any. If the destination is directly reachable, **gateway_hostid** is also unnecessary since the destination address is used to determine the interface link address). If **rspf** is used and the system is a switch / router to multiple routes the keyword **direct** can be used i.s.o a **gateway_hostid** to set the metric higher as the default 1. This way routes advertized by other **rspf** stations can be cheaper and get selected. If **direct** is given but **metric** not, an new algoritm is used to set the metric dependent on the number of subnet mask bits.

The optional **/bits** suffix to the destination host id specifies how many leading bits in the host id are to be considered significant in the routing comparisons. If not specified, 32 bits (i.e., full significance) is

assumed. With this option, a single routing table entry may refer to many hosts all sharing a common bit string prefix in their IP addresses. For example, ARPA Class A, B and C networks would use suffixes of /8, /16 and /24 respectively; the command

route add 44/8 sl0 44.64.0.2

causes any IP addresses beginning with "44" in the first 8 bits to be routed to 44.64.0.2; the remaining 24 bits are "don't-cares".

When an IP address to be routed matches more than one entry in the routing table, the entry with largest **bits** parameter (i.e., the "best" match) is used. This allows individual hosts or blocks of hosts to be exceptions to a more general rule for a larger block of hosts.

The special destination **default** is used to route datagrams to addresses not matched by any other entries in the routing table; it is equivalent to specifying a **/bits** suffix of /0 to any destination hostid. Care must be taken with default entries since two nodes with default entries pointing at each other will route packets to unknown addresses back and forth in a loop until their time-to-live (TTL) fields expire. (Routing loops for specific addresses can also be created, but this is less likely to occur accidentally).

There are two build in interfaces: **loopback** and **encap**. **Loopback** is for internal purposes only. The **encap** is an IP encapsulator interface. This is used to encapsulate a complete IP datagram into a IP datagram so that it gets "piggy-backed". It is often used to carry ampr.org datagrams (net44) over the Internet. Note that the Internet is fully connected and that ampr.org is (at least) very loose. This way two sites can interchange net44 datagrams to each other. Some extra notes: A Internet gateway has 2 IP addresses: one on ampr.org and one on the Internet. You should make sure that the interface connected to the Internet has the ifconfig ipaddr set correctly. Note: This previously only worked as gateway for other stations. There was some guessing done in picking an IP address when encap is used localy. The guessing takes the worst guess. (always wrong with a 2.7 Murphy factor) The code now uses the local IP address as source when the route starts at the local station. If this is not what you want, you can overrule this by setting the IP address of the psuedo encap interface to what you want it to be.

Some extra notes on encap: I have 3 systems on an ethernet with network 129.179.122.128/25. In the office I have another net linked to the Internet. The addresses there are in the 129.179.122.0/25 range. Now i have a radio link with 44.137.0.2 and 44.137.1.2 on respective sites. On the 44.137.0.2 system i have *route add 44.0.0.0/8 encap 44.137.1.1 99*. On a next system on my local ethernet i have 44.137.0.1 / 129.179.122.129. To get from that system to say 44.62.0.1 i have to add an encap to my local gateway with *route add 44.0.0.0/8 encap 129.179.122.130*. A *route add default ec0 129.179.122.130* gives me access to the Internet. Otherwise it will lookup the address for the interface to be used to reach 129.179.122.130 and will use 129.179.122.129. Now 44.62.0.1 will NEVER know where it came from. So adding the encap on the second system solves the problem.

Here are some examples of the **route** command:

```
# Route datagrams to IP address 44.0.0.3 to SLIP line #0.
# No gateway is needed because SLIP is point-to point.
route add 44.0.0.3 sl0

# Route all default traffic to the gateway on the local Ethernet
# with IP address 44.0.0.1
route add default ec0 44.0.0.1

# The local Ethernet has an ARPA Class-C address assignment;
# route all IP addresses beginning with 192.4.8 to it
route add 192.4.8/24 ec0

# The station with IP address 44.0.0.10 is on the local AX.25 channel
route add 44.0.0.10 ax0

# An encapsulation link to 192.4.8.12 where the subnet 44.64.0.0 is
accessible. The internet does not know where we are but we just use them
with what they know:
route add 44.64.0.0/16 encap 192.4.8.12 4
```

3.77.2. route addprivate <dest hostid>[/bits] | default <iface> [<gateway hostid> [<metric>]]

This command is identical to **route add** except that it also marks the new entry as private; it will never be included in outgoing RIP updates.

3.77.3. route drop <dest hostid>

route drop deletes an entry from the table. If a packet arrives for the deleted address and a default route is in effect, it will be used.

3.78. rspf <subcmd>

RSPF is the Radio Shortest Path First protocol. Each station listens for RRH (Router to Router Hello) messages. When such a RRH message is received, **Nos** will figure out if the link is bi-directional by pinging the other station. The protocol is described in the RSPF 2.1 specification.

3.78.1. rspf interface <interface> <quality> <horizon>

<interface> is the required interface rspf should use. quality is from 1 to 127 horizon is between 1 to 255 End nodes should have the quality set to 1. Immediate nodes normaly set the quality to 8. The normaly used value for horizon is 32.

3.78.2. rspf mode [vc | datagram | none]

Without arguments, displays the preferred mode for RSPF. Modes are VC (Virtual Cirquit) and Datagram. none resets the preferred mode.

3.78.3. rspf rrhtimer [seconds]

Without arguments, displays the rrh timer value.

3.78.4. rspf suspecttimer [seconds]

Without arguments, displays the suspect timer value.

3.78.5. rspf timer [seconds]

Without arguments, displays the update timer value.

To activate RSPF, do the following. Set the broadcast address for the destination interface, example ax0.

ifconfig ax0 broadcast 44.255.255.255

This automatically creates a routing entry for 44.255.255.255 in the routing table. If intend to use RSPF on more than one interface, those interfaces must each have different broadcast adresses. Else the routing entries will be overwriten by the next definition.

Configure ax0 as an RSPF interface with horizon 32 and a quality of 1 (hops). This is typical for an end node. replace the 1 with an 8 for immediate nodes.

rspf interface ax0 1 32

Set the interval between RRH messages.

rspf rrhtimer 900

Define how long it takes until an idle link is suspected to be bad.

rspf suspecttimer 2000

Set the interval between routing updates.

rspf timer 900

3.79. sccstat

Display the defined PE1CHL scc driver statistics.

3.80. session [[<session>] [flowmode [on | off]]]

Without arguments, displays the list of current sessions, including session number, remote TCP or AX.25 address and the address of the TCP or AX.25 control block. An asterisk (*) is shown next to the current that session. Entering a session number as an argument to the **session** command will put you in *converse* mode with that session. If the Telnet session; entering a blank line at this point puts you in converse mode with server is enabled, the user is notified of an incoming request and a session number is automatically assigned. The user may then select the session normally to converse with the remote user as though the session had been locally initiated. Adding the flowmode displays or enables / disables setting of *more* handling for that session. This is handy for example by long directory listings comming from an ftp session. Escaping to command mode before issuing the dir command and entering "session # flowmode on" gives a page at a time to look at. At any time you can escape out again and switch flowmode off. Note that a ftp session has it's own flow command now buildin. See FTP commands later in this manual.

3.81. shell [

Suspends **Nos** and executes a sub-shell ("command processor" under MS-DOS). When the sub-shell exits, **Nos** resumes (under MS-DOS, enter the **exit** command). Note: see the **COMSPEC** environement variable. When Background activity (FTP servers, etc) is also suspended while the subshell executes. Note that this will fail unless there is sufficient unused memory for the sub-shell and whatever command the user tries to run. When shelled out, Mailbox Operator connects and ttylink incomming connections are refuced. A System unattended message is send to the "connector" of that socket.

3.82. skick #socket

This is a shorthand for the various kick subcommands. This one searches the socket for correct type and kicks the transport layer.

3.83. smtp <subcommand>

These commands are used for the Simple Message Transport Protocol service (that is, mail).

3.83.1. smtp batch [yes | no]

If set smtp will batch the commands into one frame. When off only one command is send and a response is waited for. Some old and flaky smtp servers cannot handle more than one command at a time. **Nos** can handle multiple. If you are not hindered by an old smpt server, setting batch reduces bandwith.

3.83.2. smtp gateway [<hostid>]

Displays or sets the host to be used as a "smart" mail relay. Any mail sent to a host not in the host table will instead be sent to the gateway for forwarding.

3.83.3. smtp kick

Run through the outgoing mail queue and attempt to deliver any pending mail. This command allows the user to "kick" the mail system manually. Normally, this command is periodically invoked by a timer whenever **Nos** is running.

3.83.4. smtp kill [<jobid>]

Kill the job and delete the message.

3.83.5. smtp list

List the current jobs. A "L" means locked and in progress. It is wise to add in autoexec.bat a "del /spool/mqueue/*.lck" command

3.83.6. smtp maxclients [<count>]

Displays or sets the maximum number of simultaneous outgoing SMTP sessions that will be allowed. The default is 10; reduce it if network congestion is a problem.

3.83.7. smtp mode [queue | route]

Sets the smtp delivery mode. If queue, all messages are left in /spool/rqueue for external forwarding and handling. If route, messages are handled, if for local, appended to a mailbox or if remote, they are forwarded.

3.83.8. smtp mxlookup [yes | no]

Displays or sets a flag enabling or disabling MX record lookups. This can be enabled if a domain server is available in the near distance (reachable). It should be disabled (default) if no domain server is in reach to satisfy the MX query. Note that MX record handling is very limited in NOS. If an answer from a domain name server comes in it is taken to be the destination.

3.83.9. smtp quiet [yes | no]

Enables or disables the message that new mail arrived at this system.

3.83.10. smtp timer [<seconds>]

Displays or sets the interval, between scans of the outbound mail queue. For example, **smtp timer 600** will cause the system to check for outgoing mail every 10 minutes and attempt to deliver anything it finds, subject of course to the **smtp maxclients** limit. Setting a value of zero disables queue scanning altogether, note that this is the default! This value is recommended for stand alone IP gateways that never handle mail, since it saves wear and tear on the floppy disk drive.

3.83.11. smtp trace [<value>]

Displays or sets the trace flag in the SMTP client, allowing you to watch SMTP’s conversations as it delivers mail. Zero (the default) disables tracing.

3.84. socket [[<socket #>] [flowmode [yes | no]]]

Without an argument, displays all active sockets, giving their index and type, the address of the associated protocol control block and the and owner process ID and name. If the index to an active socket is supplied, the status display for the appropriate protocol is called. For example, if the socket refers to a TCP connection, the display will be that given by the **tcp status** command with the protocol control block address. **Flow** mode can be set or cleared on a session if so wanted. This comes in handy when a large directory is displayed by ftp. Escaping from the session just before entering dir and setting flowmode and returning to that session, gives a screen full at a time with -more- processing.

3.85. source <filename>

Read subsequent commands from **filename** until EOF. Then resume reading commands from the previous stream. This can be used to keep routing statements in a separate file, wich can be read at some point in **autoexec.nos**.

3.86. start ax25 | convers | discard | echo | finger | ftp | lpd | netrom | nntp | pop | pop2 | pop3 | remote | rip | smtp | telnet | tip | ttylink

Start the specified Internet server, allowing remote connection requests.

3.87. status

Displays load information on **Nos**. When started, how long running, open files etc.

3.88. stop ax25 | convers | discard | echo | finger | ftp | lpd | netrom | nntp | pop | pop2 | pop3 | remote | rip | smtp | telnet | tip | ttylink

Stop the specified Internet server, rejecting any further remote connect requests. Existing connections are allowed to complete normally.

3.89. tail <filename>

This does a more on the last couple of lines of file **filename**.

3.90. tcp <subcommand>

These commands are used for the Transmission Control Protocol service.

3.90.1. tcp irtt [<milliseconds>]

Display or set the initial round trip time estimate, in milliseconds, to be used for new TCP connections until they can measure and adapt to the actual value. The default is 5000 milliseconds (5 seconds). Increasing this when operating over slow channels will avoid the flurry of retransmissions that would otherwise occur as the smoothed estimate settles down at the correct value. Note that this command should be given before servers are started in order for it to have effect on incoming connections.

TCP also keeps a *cache* of measured round trip times and mean deviations (MDEV) for current and recent destinations. Whenever a new TCP connection is opened, the system first looks in this cache. If the destination is found, the cached IRTT and MDEV values are used. If not, the default IRTT value mentioned above is used, along with a MDEV of 0. This feature is fully automatic, and it can improve performance greatly when a series of connections are opened and closed to a given destination (eg. a series of FTP file transfers or directory listings).

3.90.2. tcp kick <tcb_addr>

If there is unacknowledged data on the send queue of the specified TCB, this command forces an immediate retransmission.

3.90.3. tcp mss [<size>]

Display or set the TCP Maximum Segment Size in bytes that will be sent on all outgoing TCP connect request (SYN segments). This tells the remote end the size of the largest segment (packet) it may send. Changing MSS affects only future connections; existing connections are unaffected.

3.90.4. tcp reset <tcb_addr>

Deletes the TCP control block at the specified address.

3.90.5. tcp rtt <tcb_addr> <milliseconds>

Replaces the automatically computed round trip time in the specified TCB with the rtt in milliseconds. This command is useful to speed up recovery from a series of lost packets since it provides a manual bypass around the normal backoff retransmission timing mechanisms.

3.90.6. tcp status [<tcb_addr>]

Without arguments, displays several TCP-level statistics, plus a summary of all existing TCP connections, including TCB address, send and receive queue sizes, local and remote sockets, and connection state. If **tcb_addr** is specified, a more detailed dump of the specified TCB is generated, including send and receive sequence numbers and timer information.

3.90.7. tcp syndata [yes | no]

Display or set the tcp syn + data pigybacking flag. Some tcp systems cannot handle syn + data together.

3.90.8. tcp timertype [linear | exponential]

Displays the current setting or sets the timer type backoff algoritm.

3.90.9. tcp trace [yes | no]

Display or set the tcp trace flag on or off.

3.90.10. tcp window [<size>]

Displays or sets the default receive window size in bytes to be used by TCP when creating new connections. Existing connections are unaffected.

3.91. telnet <hostid> [<port>]

Creates a Telnet session to the specified host and enters converse mode. If <port> is given that number is used. Default port is 23.

3.92. test

Start an internal test for an overflow problem that might appear in the clock function of some AT computers.

3.93. thirh-party [yes | no]

This is a restriction setter for some countries where stations are not allowed to handle thirth party mail traffic.

3.94. ttylink <hostid> [<port>]

Creates a Telnet session to the specified host and enters converse mode. If <port> is given that number is used. Default port is 87. This uses a split screen interface for easy conversation.

3.95. tip <iface>

Creates a **tip** session that connects to the specified interface in "dumb terminal" mode. The interface must have already been attached with the **attach** command. Any packet traffic (IP datagrams, etc) routed to the interface while this session exists will be discarded. To close a **tip** session, use the **reset** command. It will then revert to normal **slip, nrs** or **kiss** mode operation.

This feature is primarily useful for manually establishing SLIP connections. At present, only the built-in "com" ports can be used with this command.

3.96. trace [<iface> [off | <btio> [<tracefile>]]]

Controls packet tracing by the interface drivers. Specific bits enable tracing of the various interfaces and the amount of information produced. Trace output to the screen is only send to the command screen, when in a session tracing to the screen is suspended. (From version 2.0l) Tracing is controlled on a per-interface basis; without arguments, **trace** gives a list of all defined interfaces and their tracing status. Output can be limited to a single interface by specifying it, and the control flags can be change by specifying them as well. The flags are given as a hexadecimal number which is interpreted as follows:

O - Enable tracing of output packets if 1, disable if 0
I - Enable tracing of input packets if 1, disable if 0
T - Controls type of tracing:
  0 - Protocol headers are decoded, but data is not displayed
  1 - Protocol headers are decoded, and data (but not the headers themselves) are displayed as ASCII characters, 64 characters/line. Unprintable characters are displayed as periods.
  2 - Protocol headers are decoded, and the entire packet (headers AND data) is also displayed in hexadecimal and ASCII, 16 characters per line.
B - Broadcast filter flag. If set, only packets specifically addressed to this node will be traced; broadcast packets will not be displayed.

If **tracefile** is not specified, tracing will be to the console.

3.97. udp status

3.98. verbose [0 | 1 | 2 | 3]

Set or display the level of message output in file transfers. **Verbose 0** gives the least output, and **verbose 3** the most, as follows:

0 - Display error messages only.
1 - Display error messages plus a one-line summary after each transfer giving the name of the file, its size, and the transfer time and rate.
2 - Display error and summary messages plus the progress messages generated by the remote FTP server. (This setting is the default.)
3 - Display all messages. In addition, a "hash mark" (#) is displayed for every 1,000 bytes sent or received.

If a command is sent to the remote server because it is not recognized locally, the response is always displayed, regardless of the setting of **verbose**. This is necessary for commands like **pwd** (display working directory), which would otherwise produce no message at all if **verbose** were set to 0 or 1.

Displays the status of all UDP receive queues.

3.99. upload [<filename>]

Opens **filename** and sends it on the current session as though it were typed on the terminal.

3.100. watch

Displays the current software stopwatch values, with min and max readings for each. This facility allows a programmer to measure the execution time of critical sections of code with microsecond resolution. This command is supported only on the IBM PC, and the meaning of each stopwatch value depends on where the calls have been inserted for test purposes; the distribution copy of *nos* usually has no stopwatch calls.

3.101. watchdog [on | off]

Enables or disables the watchdog timer. If internal operations cease for 300 second and watchdog is enabled, a system reset is performed. See the remote command for setting up autoexec.bat.

3.102. ?

Same as the **help** command.

4. Attach Commands

This chapter details the attach commands for the various hardware interface drivers. Not all of these drivers may be configured into every **Nos** binary; a list of the available types may be obtained by entering the command **attach ?**.

Some parameters are accepted by several drivers. They are:

4.0.1. <bufsize>

For asynchronous devices (eg. COM ports operating in SLIP or NRS mode) this parameter specifies the size of the receiver's ring buffer. It should be large enough to hold incoming data at full line speed for the longest time that the system may be busy in MS-DOS or the BIOS doing a slow I/O operation (eg. to a floppy disk). A kilobyte is usually more than sufficient.

For synchronous devices (eg. the **scc, hs, pc100, hapn** and **drsi** interfaces operating in HDLC mode), the bufsize parameter specifies the largest packet that may be received on the interface. This should be set by mutual agreement among stations sharing the channel. For standard AX.25 with a maximum I-frame data size of 256 bytes, a value of 325 should provide an adequate safety margin. On higher speed channels (eg. 56kb/s) larger values (eg. 2K bytes) will provide much better performance and allow full-sized Ethernet packets to be carried without fragmentation.

4.0.2. <ioaddr>

The base address of the interface's control registers.This might be specified
in hex as 0xnnn or decimal. (nnn is the hexadecimal number).

4.0.3. <vector>

The interface's hardware interrupt (IRQ) vector, in decimal. When a vector is followed by the character 'c' then the vector is added in a interrupt chain. This way multiple devices can share the same interrupt vector (hardware changes might be nessesary). A sample is the use of a 4 port comm board sharing the same vector. The first attach command has a plain vector and the following have the c appended. Note that the highest speed port should be defined last (as it is served first by the chaining). Do NOT specify the c with the first attach statement in that group as unpredictable results will occure.

```
attach asy 0x3f8 4 ax25 ax25 2048 256 1200
attach asy 0x3f0 4c ax25 ax25 2048 256 9600
```

4.0.4. <iface>

The name (an arbitrary character string) to be assigned to this interface. It is used to refer to the interface in **ifconfig** and **route** commands and in **trace** output.

4.0.5. <mtu>

The Maximum Transmission Unit size, in bytes. Datagrams larger than this limit will be fragmented at the IP layer into smaller pieces. For AX.25 UI frames, this limits the size of the information field. For AX.25 I frames, however, the **ax25 paclen** parameter is also relevant. If the datagram or fragment is still larger than **paclen**, it is also fragmented at the AX.25 level (as opposed to the IP level) before transmission. (See the **ax25 paclen** command for further information).

4.0.6. <speed>

The speed in bits per second (eg. 2400).

4.1. attach 3c500 <ioaddr> <vector> arpa <iface> <qlen> <mtu> [<ip_addr>]

Attach a 3Com 3C501 Ethernet interface. **qlen** is the maximum allowable transmit queue length. If the **ip_addr** parameter is not given, the value associated with a prior **ip address** command will be used.

The use of this driver is not recommended; use the packet driver interface with the loadable 3C501 packet driver instead.

4.2. attach asy <ioaddr> <vector> ax25 | nrs | ppp | slip | raw <iface> <bufsize> <mtu> <speed> [<vf>]

Attach a standard PC "com port" (asynchronous serial port), using the National 8250 or 16550A chip. Standard values on the IBM PC and clones for **ioaddr** and **vector** are 0x3f8 and 4 for COM1, and 0x2f8 and 3 for COM2. If the port uses a 16550A chip, it will be detected automatically and the FIFOs enabled.

4.2.1. ax25

Similar to **slip**, except that an AX.25 header and a KISS TNC control header are added to the front of the datagram before SLIP encoding. Either UI (connectionless) or I (connection-oriented) AX.25 frames can be used; see the **mode** command for details.

4.2.2. nrs

Use the NET/ROM asynchronous framing technique for communication with a local NET/ROM TNC.

4.2.3. ppp

Point-to-Point-Protocol. Encapsulates datagrams in an HDLC-like frame. This is a new Internet standard for point-to-point communication, compatible with CCITT standards.

4.2.4. slip

Serial Line Internet Protocol. Encapsulates IP datagrams directly in SLIP frames without a link header. This is for operation on point-to-point lines and is compatible with 4.2BSD UNIX SLIP.

4.2.5. raw

Raw serial line without protocol , special for lpd server.

4.2.6. <vf>

The optional flags are a string of the characters "vf": **v** enables Van Jacobson TCP/IP Header Compression, and is valid only for SLIP. **f** forces the fifo on for 16550AFN compatible chips not hindered by the original "design bug" and not needing the work-around. Unfortunate these good chips don't get their fifo enabled. Specifying **f** on the attach line forces the fifo to be used. (unpredictable results occure when specified on a non 16550 type chip).

4.3. attach axip <iface> <mtu> <ip_addr> <callsign>

This creates a RFC1226 compatible AX.25 frame encapsulator for transmission of AX.25 frames over the **Internet**. **Iface** will be the name of the interface, **ip_addr** the address of the remote system and **callsign** the AX.25 callsign this station is listening on for frames to digipeat. Note that each attached axip interface should have a different callsign to listen to and this should also be different from other callsign's used on this station.

4.4. attach drsi <ioaddr> <vector> ax25 <iface> <bufsize> <mtu> <ch_a_speed> <ch_b_speed>

N6TTO driver for the Digital Radio Systems PCPA 8530 card. Since there are two channels on the board, two interfaces are attached. They will be named **iface** with 'a' and 'b' appended. **bufsize** is the receiver buffer size in bytes; it must be larger than the largest frame to be received. **ch_a_speed** and **ch_b_speed** are the speeds, in bits/sec, for the A and B channels, respectively.

4.5. attach eagle <ioaddr> <vector> ax25 <iface> <bufsize> <mtu> <speed>

WA3CVG/NG6Q driver for the Eagle Computer card (Zilog 8530).

4.6. attach hapn <ioaddr> <vector> ax25 <iface> <bufsize> <mtu> csma | full

KE3Z driver for the Hamilton Amateur Packet Network adapter (Intel 8273). The **csma | full** parameter specifies whether the port should operate in carrier sense multiple access (CSMA) mode or in full duplex.

4.7. attach hs <ioaddr> <vector> ax25 <iface> <bufsize> <mtu> <keyup_delay> <p>

Attach a DRSI PCPA or Eagle Computer interface card using a special "high speed" 8530 driver. This driver uses busy-wait loops to send and receive each byte instead of interrupts, making it usable with high speed modems (such as the WA4DSY 56kb/s modem) on slow systems. This does have the side effect of "freezing" the system whenever the modem transmitter or receiver is active. This driver can operate only in CSMA mode, and it is recommended that no other interfaces requiring small interrupt latencies be attached to the same machine.

The **keyup_delay** parameter specifies the transmitter keyup delay in byte time intervals. The **p** value specifies the transmitter persistence value in the range 1-255; the corresponding slot time is fixed at one hardware clock tick, about 55 ms on the PC.

As with the other 8530 drivers, this driver actually attaches two interfaces, one for each 8530 channel.

4.8. attach packet <intvec> <iface> <txqlen> <mtu>

Driver for use with separate software "packet drivers" meeting the FTP Software, Inc, Software Packet Driver specification. The driver must have already been installed before the **attach** command is given. Packet drivers in the Ethernet, ARCNET, SLIP, SLFP, and KISS/AX25 classes are supported.

intvec is the software interrupt vector used for communication to the packet driver, and **txqlen** is the maximum number of packets that will be allowed on the transmit queue.

4.9. attach pc100 <ioaddr> <vector> ax25 <iface> <bufsize> <speed>

Driver for the PACCOMM PC-100 (Zilog 8530) card. Only AX.25 operation is supported.

4.10. attach scc <devices> init <addr> <spacing> <Aoff> <Boff> <Dataoff> <intack> <vec> [p | r]<clock> [<hdwe>] [<param>]

PE1CHL driver to initialize a generic SCC (8530) interface board prior to actually attaching it. The parameters are as follows:

4.10.1. <devices>

The number of SCC chips to support.

4.10.2. <addr>

The base address of the first SCC chip (hex).

4.10.3. <spacing>

The spacing between the SCC chip base addresses.

4.10.4. <Aoff>

The offset from a chip’s base address to its channel A control register.

4.10.5. <Boff>

The offset from a chip’s base address to its channel B control register.

4.10.6. <Dataoff>

The offset from each channel’s control register to its data register.

4.10.7. <intack>

The address of the INTACK/Read Vector port. If none, specify 0 to read from RR3A/RR2B.

4.10.8. <vec>

The CPU interrupt vector for all connected SCCs.

4.10.9. <clock>

The clock frequency (PCLK/RTxC) of all SCCs in hertz. Prefix with 'p' for PCLK, 'r' for RTxC clock (for baudrate gen).

4.10.10. <hdwe>

Optional hardware type. The following values are currently supported: 1 - Eagle card, 2 - PACCOMM PC-100, 4 - PRIMUS-PC card (DG9BL), 8 - DRSI PCPA card.

4.10.11. <param>

Optional extra parameter. At present, this is used only with the PC-100 and PRIMUS-PC cards to set the modem mode. The value 0x22 is used with the PC-100 and 0x2 is used with the PRIMUS-PC card.

The **attach scc ... init** command must be given before the interfaces are actually attached with the following command.

4.11. attach scc <chan> slip | kiss | nrs | ax25 <iface> <mtu> <speed> <bufsize> [<call>]

Attach an initialized SCC port to the system. The parameters are as follows:

4.11.1. <chan>

The SCC channel number to attach, 0 or 1 for the first chip's A or B port, 2 or 3 for the second chip's A or B port, etc.

4.11.2. slip | kiss | nrs | ax25

The operating mode of the interface. **slip, kiss** and **nrs** all operate the port hardware in asynchronous mode; **slip** is Internet-standard serial line IP mode, **kiss** generates SLIP frames containing KISS TNC commands and AX.25 packets and **nrs** uses NET/ROM local serial link framing conventions to carry NET/ROM packets. Selecting **ax25** mode puts the interface into synchronous HDLC mode that is suitable for direct connection to a half duplex radio modem.

4.11.3. <speed>

The interface speed in bits per second (eg. 1200). Prefix with 'd' when an external divider is available to generate the TX clock. When the clock source is PCLK, this can be a /32 divider between TRxC and RTxC. When the clock is at RTxC, the TX rate must be supplied at TRxC. This is needed only for full duplex synchronous operation. When this arg is given as 'ext', the transmit and receive clocks are external, and the internal baud rate generator (BRG) and digital phase locked loop (DPLL) are not used.

4.12. Attach Examples

Here are some examples of the attach command:

```
# Attach a 3Com Ethernet controller using the standard 3Com address and
# vector (i.e., as it comes out of the box) to use ARPA-standard encapsulation.
# The receive queue is limited to 5 packets, and outgoing packets larger
# than 1500 bytes will be fragmented
attach 3c500 0x300 3 arpa ec0 5 1500

# Attach the PC asynch card normally known as "com1" (the first controller)
# to operate in point-to-point slip mode at 9600 baud, calling it "sl0".
# A 1024 byte receiver ring buffer is allocated. Outgoing packets larger
# than 256 bytes are fragmented.
attach asy 0x3f8 4 slip sl0 1024 256 9600

# Attach the secondary PC asynch card ("com2") to operate in AX.25 mode
# with an MTU of 576 bytes at 9600 baud with a KISS TNC, calling it "ax0".
# By default, IP datagrams are sent in UI frames
attach asy 0x2f8 3 ax25 ax0 1024 576 9600

# Attach a axip wormhole
attach axip ai0 256 129.179.122.10 pa0gri-11
# on the other side of the wormhole
attach axip ai0 256 129.179.122.130 pa0gri-12
# Now assume 129.179.122.10 has a AX.25 interface with callsign pa0gri-10
# and 129.179.122.130 a interface with callsign pa0gri-8
# Now a AX.25 frame like:
# pe1chl->pa0gri-11->pa0gri-8->pe1dna [ data]
# Received by pa0gri-11, set the has-been-digipeated change the interface
# callsign with the one it realy came in from and ecapsulates it
# in a IP frame type 93 and ships it to 129.179.122.130
# pe1chl->pa0gri-10*->pa0gri-8->pe1dna [ data]
# Arived at 129.179.122.130 the next digi is searched for and found.
# The frame is changed into:
# pe1chl->pa0gri-10*->pa0gri-12*->pe1dna [ data]
# so that on the way back the frame will find the "right" interface.

# Attach the packet driver loaded at interrupt 0x7e
# The packet driver is for an Ethernet interface
attach packet 0x7e ethernet 8 1500
```

5. FTP Subcommands

During converse mode with an FTP server, everything typed on the console is first examined to see if it is a locally-known command. If not, the line is passed intact to the remote server on the control channel. If it is one of the following commands, however, it is executed locally. (Note that this generally involves other commands being sent to the remote server on the control channel.)

5.1. dir [<file> | <directory> [<local file>]]

Without arguments, **dir** requests that a full directory listing of the remote server's current directory be sent to the terminal. If one argument is given, this is passed along in the LIST command; this can be a specific file or subdirectory that is meaningful to the remote file system. If two arguments are given, the second is taken as the local file into which the directory listing should be put (instead of being sent to the console). The PORT command is used before the LIST command is sent.

5.2. flow [off|on]

Displays or sets the -more- processing on the current ftp session. When set to on, a -more- prompt is displayed after each screen full of data. You can do the same with a **session** # **flow** command. Note that this is a local extension to the standard ftp command set.

5.3. get <remote file> [<local file>]

Asks the remote server to send the file specified in the first argument. The second argument, if given, will be the name of the file on the local machine; otherwise it will have the same name as on the remote machine. The PORT and RETR commands are sent on the control channel.

5.4. hash

A synonym for the **verbose 3** command.

5.5. ls [<file> | <directory> [<local file>]]

ls is identical to the **dir** command except that the "NLST" command is sent to the server instead of the "LIST" command. This results in an abbreviated directory listing, i.e., one showing only the file names themselves without any other information.

5.6. mget <file> [<file> ...]

Fetch a collection of files from the server. File names may include wild card characters; they will be interpreted and expanded into a list of files by the remote system using the NLST command. The files will have the same name on the local system that they had on the server.

5.7. mkdir <remote directory>

Creates a directory on the remote machine.

5.8. mput <file> [<file> ...]

Send a collection of files to the server. File names may include wild card characters; they will be expanded locally into a list of files to be sent. The files will have the same name on the server as on the local system.

5.9. put <local file> [<remote file>]

Asks the remote server to accept data, creating the file named in the first argument. The second argument, if given, will be the name of the file on the remote machine; otherwise it will have the same name as on the local machine. The PORT and STOR commands are sent on the control channel.

5.10. rmdir <remote directory>

Deletes a directory on the remote machine.

5.11. type [a | i | l <bytesize>]

Tells both the local client and remote server the type of file that is to be transferred. The default, (wich can be changed with the **ftype** command) is 'a', which means ASCII (i.e., a text file). Type 'i' means *image*, i.e., binary. In ASCII mode, files are sent as varying length lines of text in ASCII separated by cr/lf sequences; in IMAGE mode, files are sent exactly as they appear in the file system. ASCII mode should be used whenever transferring text between dissimilar systems (e.g., UNIX and MS-DOS) because of their different end-of-line and/or end-of-file conventions. When exchanging text files between machines of the same type, either mode will work but IMAGE mode is usually faster. Naturally, when exchanging raw binary files (executables, compressed archives, etc) IMAGE mode must be used. Type 'l' (logical byte size) is used when exchanging binary files with remote servers having oddball word sizes (e.g., DECSYSTEM-10s and 20s). Locally it works exactly like IMAGE, except that it notifies the remote system how large the byte size is. **bytesize** is typically 8. The type command sets the local transfer mode and generates the TYPE command on the control channel.

5.12. verbose [0 | 1 | 2 | 3]

Set or display the level of message output in file transfers. **Verbose 0** gives the least output, and **verbose 3** the most, as follows:

0 - Display error messages only.
1 - Display error messages plus a one-line summary after each transfer
giving the name of the file, its size, and the transfer time and rate.
2 - Display error and summary messages plus the progress messages generated
by the remote FTP server. (This setting is the default.)
3 - Display all messages. In addition, a "hash mark" (#) is displayed for
every 1,000 bytes sent or received.

If a command is sent to the remote server because it is not recognized locally, the response is always displayed, regardless of the setting of **verbose**. This is necessary for commands like **pwd** (display working directory), which would otherwise produce no message at all if **verbose** were set to 0 or 1.

6. Dialer Subcommands

Each dialer command may (should) have a different dialer file. The file resides in the configuration directory, as specified in the **Installation** section (see chapter 1). A typical dialer file might be:

```
# Set the speed, and toggle DTR to ensure modem is in command mode.
control down
wait 3000
speed 2400
control up
wait 3000
# Dial, and wait for connection
send "atdt555-1212\r"
wait 45000 "CONNECT " speed
wait 2000
# PAD specific initialization
send "\r"
wait 15000 "Terminal ="
send "ppp\r"
wait 10000 "\r\n"
```

6.0.1. control down | up

Control **asy** interface. The **down** option drops DTR and RTS. The **up** option asserts DTR and RTS.

6.0.2. send "string"

This dialer command will write the specified string to the interface. The string quote marks are required, and the string may not contain embedded control characters. However, the standard C string escape sequences are recognized (\0 should not be used).

6.0.3. speed [9600 | 4800 | 2400 | 1200 | 300]

This dialer command will set the speed of the interface to one of the available speeds. If the speed is missing, the speed will be displayed in the dialer session window.

6.0.4. wait <milliseconds> ["test string"] [speed]

If only the time is specified, the dialer pauses for the desired number of milliseconds.

Otherwise, the dialer reads until the test string is detected on the interface. If the string is not detected within the desired time, the autodialer will reset. The string quote marks are required, and the string may not contain embedded control characters. However, the standard C string escape sequences are recognized (\0 should not be used).

Finally, if the *speed* parameter is specified, the dialer will continue to read characters until a non-digit is detected. The string read is converted to an integer, and used to set the interface speed. If the trailing non-digit is not detected within the desired time, or the integer value is not a valid speed, the autodialer will reset.

7. Installation

Nos uses the following file and directory structure:

~/alias
~/autoexec.nos
~/dialer
~/domain.txt
~/ftpusers
~/net.rc
~/netron.sav
~/popusers
~/finger/
~/etc/printcap
~/etc/lpdperms
~/etc/log
~/spool/areas
~/spool/mail.log
~/spool/net.log
~/spool/forward.bbs
~/spool/history
~/spool/rewrite
~/spool/help/
~/spool/mail/
~/spool/mqueue/
~/spool/news/
~/spool/news/active
~/spool/news/pointer
~/spool/news/info
~/spool/news/help
~/spool/news/history
~/spool/news/forward
~/spool/news/poll
~/spool/rqueue/
~/spool/signatur/
~/spool/lpd/

The ~ in front of all files is a directory offset definable with the **-d** option on the **Nos** command line. Any name may be chosen and is default empty. (eg. just /) If for example **-d /net** is given, the structure shifts to /net/... The **alias, autoexec.nos, dialer, domain.txt net.rc, popusers** and **ftpusers** configuration files are located here. The **netrom.sav** file will be created there.

The "/spool" directory and its sub-directories are used by the bbs, SMTP and NNTP services. The **areas, forward.bbs, history, mail.log** and **rewrite** configuration files are located here. The /spool/news directory can have many subdirectories and each subdirectory can have subdirectories. Newsgroups are split into heirarchical directory structures. A news article in newsgroup rec.amateur.radio.packet will end up in /spool/news/rec/amateur/radio/packet.txt.

7.1. The /ftpusers File

Since MS-DOS is a single-user operating system (some might say it is a glorified bootstrap loader), it provides no access control; all files can be read, written or deleted by the local user. It is usually undesirable to give such open access to a system to remote network users. Net therefore provides its own access control mechanisms.

The file **ftpusers** controls remote FTP and mailbox access. The FTP default is *no* access; if this file does not exist, the FTP server will be unusable. A remote user must first "log in" to the system with the USER and PASS commands, giving a valid name and password listed in **ftpusers**, before he or she

can transfer files.

Each entry in **ftpusers** consists of a single line of the form

username password /path permissions ip_address

There must be at least four fields, and there must be exactly one space between each field. Comments may be added after the last field. Comment lines begin with '#' in column one.

username is the user's login name.

password is the required password. Note that this is in plain text; therefore it is not a good idea to give general read permission to the root directory. A password of '*' (a single asterisk) means that any password is acceptable.

/path is the allowable prefix on accessible files. Before any file or directory operation, the current directory and the user- specified file name are joined to form an absolute path name in "canonical" form (i.e., a full path name starting at the root, with "./" and "../" references, as well as redundant /'s, recognized and removed). The result MUST begin with the allowable path prefix; if not, the operation is denied. This field must always begin with a "/", i.e., at the root directory. Multiple directories can be specified by separating them with a ";" caracter and no whitespace around them.

permissions is a decimal number granting permission for read, create and write operations. If the low order bit (0x1) is set, the user is allowed to read a file subject to the path name prefix restriction. If the next bit (0x2) is set, the user is allowed to create a new file if it does not overwrite an existing file. If the third bit (0x4) is set, the user is allowed to write a file even if it overwrites an existing file, and in addition he may delete files. Again, all operations are allowed subject to the path name prefix restrictions. Permissions may be combined by adding bits, for example, 0x3 (= 0x2 + 0x1) means that the user is given read and create permission, but not overwrite/delete permission.

Additional permission bits used by the mailbox and PPP are:

```
   1  Read files
   2  Create new files
   4  Overwrite and delete existing files
   8  AX.25 gateway operation allowed
  16  Telnet gateway operation allowed
  32  NET/ROM gateway operation allowed
  64  Remote sysop access allowed (DANGEROUS)
 128 This user is banned from BBS access (illegal user)
 256 Priv bit for PPP connection
 512 Priv bit for peerID/pass lookup
1024 disallow send commands (except to SYSOP)
2048 disallow read commands
4096 disallow 3rd party mail
8192 This station is a known BBS
```

ip_address is used for PPP only and is the remote IP address of the connected system.

A username of **univperm** has special meaning in the validation meganism. If **univperm** is included as a valid user in **ftpusers** then any unknown user (not in **ftpusers**) will be mapped into **univperm** and get its permission bits and file path. If **univperm** is not included in **ftpusers** unknown users are not permitted nor validated.

For example, suppose **ftpusers** on machine pc.ka9q.ampr.org contains the line

friendly test /testdir 7

A session using this account would look like this:

```
net> ftp pc.ka9q.ampr.org
Resolving pc.ka9q.ampr.org... Trying 128.96.160.1...
FTP session 1 connected to pc.ka9q.ampr.org
220 pc.ka9q.ampr.org FTP version 900418 ready at Mon May 7 16:27:18 1990
Enter user name: friendly
331 Enter PASS command
Password: test [not echoed]
230 Logged in
ftp>
```

The user now has read, write, overwrite and delete privileges for any file under /testdir; he may not access any other files.

Here are some more sample entries in **ftpusers**:

```
karn foobar / 7         # User "karn" with password "foobar" may read,
                        # write, overwrite and delete any file on the
                        # system.

guest bletch /g/bogus;/public 3
                        # User "guest" with password "bletch" may read
                        # any file under /g/bogus and its subdirectories,
                        # and /public and its subdirectories,
                        # and may create a new file as long as it does
                        # not overwrite an existing file. He may NOT
                        # delete any files.

anonymous * /public 1   # User "anonymous" (any password) may read files
                        # under /public and its subdirectories; he may
                        # not create, overwrite or delete any files.
```

This last entry is the standard convention for keeping a repository of public files; in particular, the username "anonymous" is an established ARPA convention.

7.2. The /popusers File

Here are the username / password combinations defined for the POP users. It has a simple convention:

user:password:

for every POP user such a line has to be added. The user and password fields sould match the **pop userdata** statement of the remote user. Both user and password have to be delimited with a colon character.

7.3. The /net.rc File

The net.rc file is a fast login file for known ftp stations. Each line starts with the name of the ftp station. Following are a user and password statemend to be send to the server for validation. The name, user and password are separated with a space. Not a tab or more than 1 space character. Following is a sample net.rc file.

```
ucsd.edu anonymous gvdg@gvdgpc.cdh.cdc.com
ka9q.ampr.org guest pa0gri
```

7.4. The /domain.txt File

Nos translates domain names (eg. "pc.ka9q.ampr.org") to IP addresses (eg. 128.96.160.3) through the use of an Internet Domain Name resolver and a local "cache" file, **domain.txt**. Whenever the user specifies a domain name, the local cache is searched for the desired entry. If it is present, it is used; if

not, and if domain name server(s) have been configured, a query is sent over the network to the current server. If the server responds, the answer is added to the **domain.txt** file for future use. If the server does not respond, any additional servers on the list are tried in a round-robin fashion until one responds, or the retry limit is reached (see the **domain retry** command). If **domain.txt** does not contain the desired entry and there are no configured domain name servers, then the request immediately fails.

If a domain name server is available, and if all references to host-ids in your /**autoexec.nos** file are in IP address format, then it is possible to start with a completely empty **domain.txt** file and have **Nos** build it for you. However, you may wish to add your own entries to **domain.txt**, either because you prefer to use symbolic domain names in your /**autoexec.nos** file or you don't have access to a domain server and you need to create entries for all of the hosts you may wish to access.

Each entry takes one line, and the fields are separated by tabs. For example:

pc.ka9q.ampr.org. IN A 128.96.160.3

IN is the *class* of the record. It means *Internet*, and it will be found in all entries. **A** is the *type* of the record, and it means that this is an *address* record. Domain name **pc.ka9q.ampr.org** therefore has Internet address 128.96.160.3.

Another possible entry is the **CNAME** (Canonical Name) record. For example:

ka9q.ampr.org. IN CNAME pc.ka9q.ampr.org.

This says that domain name "ka9q.ampr.org" is actually an alias for the system with (primary, or *canonical*) domain name "pc.ka9q.ampr.org." When a domain name having a **CNAME** record is given to **Nos**, the system automatically follows the reference to the canonical name and returns the IP address associated with that entry.

Entries added automatically by **Nos** will have an additional field between the domain name and the class (**IN**) field. For example:

pc.ka9q.ampr.org. 3600 IN A 128.96.160.3

This is the *time-to-live* value, in seconds, associated with the record received from the server. Clients (such as **Nos**) caching these records are supposed to delete them after the time-to-live interval has expired, allowing for the possibility that the information in the record may become out of date.

This implementation of **Nos** will decrement the TTL to zero, but will not delete the record unless the "clean" flag is on (see the **domain cache clean** command). When a remote server is not available, the old entry will be used.

When the *TTL* value is missing (as in the examples above), the record will never expire, and must be managed by hand. Since **domain.txt** is a plain text file, it may be easily edited by the user to add, change or delete records.

Additional types of records, include NS (name server) and SOA (start of authority) may appear in **domain.txt** from remote server responses. These are not currently used by **Nos** but are retained for future development (such as the incorporation of a domain name server into **Nos** itself).

7.5. The /alias file.

SMTP server ALIAS file. This is for resolving a given target address into a single or multiple entry mail list.

Format:

```
mail_list_name call_1@host_1 [call_2@host_2].....# comments

pa0gri gvdg@fridley.cdh.cdc.com
kelvin g1emm@g1emm.ampr.org
#
bob    gb3xp@gb3xp.ampr.org
ian    g3rra@g3rra.ampr.org
jim    g1wkk@g1wkk.ampr.org
john   g5ds%gb3kp@gb3xp.ampr.org
```

```
ted     gb3kp%gb3kp@gb3xp.ampr.org
ron     g6vug@g6vug.ampr.org
tim     g4uqe@g4uqe.ampr.org
gareth g6kvk%gb7spv@gb3xp.ampr.org
bolton gb7tcp%gb7crg@gb3xp.ampr.org
julian g7efe%gb7cfb@gb3xp.ampr.org
#
world  g3rra@g3rra.ampr.org gb3xp@gb3xp.ampr.org g1plt@g1plt.ampr.org
   g1wkk@g1wkk.ampr.org g8kwi@g8kwi.ampr.org g6kvk%gb7spv@gb3xp.ampr.org
   g1ttg@g1ttg.ampr.org g6vug@g6vug.ampr.org g6kqz@g6kqz.ampr.org
   g4uqe@g4uqe.ampr.org g8ogr@g8ogr.ampr.org g4xwv@g4xwv.ampr.org
#
locals g3rra@g3rra.ampr.org g1wkk@g1wkk.ampr.org g8kwi@g8kwi.ampr.org
   g4bio@g4bio.ampr.org g1plt@g1plt.ampr.org g4uqe@g4uqe.ampr.org
   g6kqz@g6kqz.ampr.org g4tnu@g4tnu.ampr.org g6xqb@g6xqb.ampr.org
```

Note that it is reaonable, and sometimes desiable, to have alias records in the form:

```
area  area dest1 dest2
```

As the /alias file is scanned only once, this does not result in an infinite recursion.

7.6. The /spool/areas file.

This file is a header file shown to a mailbox user when he requests the a display. It should show all public mailboxes to read. Here is a sample:

```
------------ Public -- Mail -- Areas ------------
                                                  |
General    --    Any old chit-chat that is clean.      |
TcpIp      --    General Tcp/Ip messages. NOS etc.     |
Bugs       --    Where to report bugs in Ka9q-NOS      |
Updates    --    Info - NOS version UPDATES by G1EMM.  |
                                                  |
-----------------------------------------------------------
```

7.7. The /spool/forward.bbs file.

The mailbox reads a forwarding file, **spool/forward.bbs**. Here is a sample file.

```
wb0ttw 0006
ax25 ax0 wb0ttw
wb0ttw
w0tn
mspbul
all
------
wb0gdb
netrom #msparh
..c msparh
all
```

------ The first word on the first line in a forwarding record is the name of the BBS to forward to. This should be the same type of name which is shown by the *mbox status* command. The second word is optional. It specifies a range of hours when forwarding may take place. **0006** means that there will only be forwarding to this station between midnight and 6am.

The second line specifies how to establish the connection. It should start with the protocol (*ax25, connect, tcp, telnet* or *netrom*) and be followed by all the parameters which are nessesary when **Nos** has to establish a connection.

Directly after the second line, there may be lines that start with a dot. What follows after the dot will be send to the remote BBS as soon as the connection is established.

Then follows the names of a number of message areas, publib or private. Finally, there should be a couple of '-' signs to separate one forwarding record from another. Also terminate the file with a last line of dashes.

7.8. The /spool/rewrite file.

Read the rewrite file for lines where the first word is a regular expression and the second word are rewriting rules. An optional third field, containing just the letter "r", when present, instructs **Nos** to restart the rewrite file, using the new destination address. The special character '$' followed by a digit denotes the string that matched a '*' character. The '*' characters are numbered from 1 to 9. Example: the line "*@*.* $2@$1.ampr.org" would rewrite the address "foo@bar.xxx" to "bar@foo.ampr.org".

```
#
*@g1emm.ampr.org $1
*@g1emm.ampr $1
*@g1emm $1
#
*!*!*!*!*!*!* $7%$6@$5@$4@$3@$2@$1
*!*!*!*!*!* $6%$5@$4@$3@$2@$1
*!*!*!*!* $5%$4@$3@$2@$1
*!*!*!* $4%$3@$2@$1
*!*!* $3%$2@$1
*!* $2@$1
!* $1 r
#
# The End
```

8. Setting Bufsize, Paclen, Maxframe, MTU, MSS and Window

Many **Nos** users are confused by these parameters and do not know how to set them properly. This chapter will first review these parameters and then discuss how to choose values for them. Special emphasis is given to avoiding interoperability problems that may appear when communicating with non-**Nos** implementations of AX.25.

8.1. Hardware Parameters

8.1.1. Bufsize

This parameter is required by most of **Nos**'s built-in HDLC drivers (eg. those for the DRSI PCPA and the Paccomm PC-100). It specifies the size of the buffer to be allocated for each receiver port. HDLC frames larger than this value cannot be received.

There is no default **bufsize**; it must be specified in the **attach** command for the interface.

8.2. AX25 Parameters

8.2.1. Paclen

Paclen limits the size of the data field in an AX.25 I-frame. This value does *not* include the AX.25 protocol header (source, destination and digipeater addresses).

Since unconnected-mode (datagram) AX.25 uses UI frames, this parameter has no effect in unconnected mode.

The default value of **paclen** is 256 bytes.

8.2.2. Maxframe

This parameter controls the number of I-frames that **Nos** may send on an AX.25 connection before it must stop and wait for an acknowledgement. Since the AX.25/LAPB sequence number field is 3 bits wide, this number cannot be larger than 7.

Since unconnected-mode (datagram) AX.25 uses UI frames that do not have sequence numbers, this parameter does *not* apply to unconnected mode.

The default value of **maxframe** in **Nos** is 1 frame.

8.3. IP and TCP Parameters

8.3.1. MTU

The MTU (Maximum Transmission Unit) is an interface parameter that limits the size of the largest IP datagram that it may handle. IP datagrams routed to an interface that are larger than its MTU are each split into two or more *fragments*. Each fragment has its own IP header and is handled by the network as if it were a distinct IP datagram, but when it arrives at the destination it is held by the IP layer until all of the other fragments belonging to the original datagram have arrived. Then they are reassembled back into the complete, original IP datagram. The minimum acceptable interface MTU is 28 bytes: 20 bytes for the IP (fragment) header, plus 8 bytes of data.

There is no default MTU in **Nos**; it must be explicitly specified for each interface as part of the **attach** command.

8.3.2. MSS

MSS (Maximum Segment Size) is a TCP-level parameter that limits the amount of data that the *remote* TCP will send in a single TCP packet. MSS values are exchanged in the SYN (connection request) packets that open a TCP connection. In the **Nos** implementation of TCP, the MSS actually used by TCP is further reduced in order to avoid fragmentation at the local IP interface. That is, the local TCP asks IP for the MTU of the interface that will be used to reach the destination. It then subtracts 40 from the MTU value to allow for the overhead of the TCP and IP headers. If the result is less than the MSS

received from the remote TCP, it is used instead.

The default value of **MSS** is 512 bytes.

8.3.3. Window

This is a TCP-level parameter that controls how much data the local TCP will allow the remote TCP to send before it must stop and wait for an acknowledgement. The actual window value used by TCP when deciding how much more data to send is referred to as the *effective window*. This is the smaller of two values: the window advertised by the remote TCP minus the unacknowledged data in flight, and the *congestion window*, an automatically computed time-varying estimate of how much data the network can handle.

The default value of **Window** is 2048 bytes.

8.4. Discussion

8.4.1. IP Fragmentation vs AX.25 Segmentation

IP-level fragmentation often makes it possible to interconnect two dissimilar networks, but it is best avoided whenever possible. One reason is that when a single IP fragment is lost, all other fragments belonging to the same datagram are effectively also lost and the entire datagram must be retransmitted by the source. Even without loss, fragments require the allocation of temporary buffer memory at the destination, and it is never easy to decide how long to wait for missing fragments before giving up and discarding those that have already arrived. A reassembly timer controls this process. In **Nos** it is (re)initialized with the **ip rtimer** parameter (default 30 seconds) whenever progress is made in reassembling a datagram (i.e., a new fragment is received). It is not necessary that all of the fragments belonging to a datagram arrive within a single timeout interval, only that the interval between fragments be less than the timeout.

Most subnetworks that carry IP have MTUs of 576 bytes or more, so interconnecting them with subnetworks having smaller values can result in considerable fragmentation. For this reason, IP implementors working with links or subnets having unusually small packet size limits are encouraged to use *transparent fragmentation*, that is, to devise schemes to break up large IP datagrams into a sequence of link or subnet frames that are immediately reassembled on the other end of the link or subnet into the original, whole IP datagram without the use of IP-level fragmentation. Such a scheme is provided in AX.25 Version 2.1. It can break a large IP or NET/ROM datagram into a series of **paclen**-sized AX.25 segments (not to be confused with TCP segments), one per AX.25 I-frame, for transmission and reassemble them into a single datagram at the other end of the link before handing it up to the IP or NET/ROM module. Unfortunately, the segmentation procedure is a new feature in AX.25 and is not yet widely implemented; in fact, **Nos** is so far the only known implementation. This creates some interoperability problems between **Nos** and non-**Nos** nodes, in particular, standard NET/ROM nodes being used to carry IP datagrams. This problem is discussed further in the section on setting the MTU.

8.4.2. Setting paclen and bufsize

The more data you put into an AX.25 I frame, the smaller the AX.25 headers are in relation to the total frame size. In other words, by increasing **paclen**, you lower the AX.25 protocol overhead. Also, large data packets reduce the overhead of keying up a transmitter, and this can be an important factor with higher speed modems. On the other hand, large frames make bigger targets for noise and interference. Each link has an optimum value of **paclen** that is best discovered by experiment.

Another thing to remember when setting **paclen** is that the AX.25 version 2.0 specification limits it to 256 bytes. Although **Nos** can handle much larger values, some other AX.25 implementations (including digipeaters) cannot and this may cause interoperability problems. Even **Nos** may have trouble with certain KISS TNCs because of fixed-size buffers. The original KISS TNC code for the TNC-2 by K3MC can handle frames limited in size only by the RAM in the TNC, but some other KISS TNCs cannot.

Nos's built-in HDLC drivers (SCC, PC-100, DRSI, etc) allocate receive buffers according to the maximum expected frame size, so it is important that these devices be configured with the correct **bufsize**.

To do this, you must know the size of the largest possible frame that can be received. The **paclen** parameter controls only the size of the data field in an I-frame and not the total size of the frame as it appears on the air. The AX.25 spec allows up to 8 digipeaters, so the largest possible frame is (**paclen** + 72) bytes. So you should make **bufsize** at least this large.

Another important consideration is that the more recent versions of NOS improve interrupt response by maintaining a special pool of buffers for use by the receive routines. These buffers are currently fixed in size to 2048 bytes and this can be changed only by editing config.h and recompiling NOS. This limits **bufsize**; in fact, attempting to set a larger value may cause the driver not to work at all. This situation can be detected by running the **memory status** command and looking for a non-zero count of **Ibuffail** events, although these events can also occur occasionally during normal operation.

One of the drawbacks of AX.25 that there is no way for one station to tell another how large a packet it is willing to accept. This requires the stations sharing a channel to agree beforehand on a maximum packet size. TCP is different, as we shall see.

8.4.3. Setting Maxframe

For best performance on a half-duplex radio channel, **maxframe** should always be set to 1. The reasons are explained in the paper *Link Level Protocols Revisited* by Brian Lloyd and Phil Karn, which appeared in the proceedings of the ARRL 5th Computer Networking Conference in 1986.

8.4.4. Setting MTU

TCP/IP header overhead considerations similar to those of the AX.25 layer when setting **paclen** apply when choosing an MTU. However, certain subnetwork types supported by **Nos** have well-established MTUs, and these should always be used unless you know what you're doing: 1500 bytes for Ethernet, and 508 bytes for ARCNET. The MTU for PPP is automatically negotiated, and defaults to 1500. Other subnet types, including SLIP and AX.25, are not as well standardized.

SLIP has no official MTU, but the most common implementation (for BSD UNIX) uses an MTU of 1006 bytes. Although **Nos** has no hard wired limit on the size of a received SLIP frame, this is not true for other systems. Interoperability problems may therefore result if larger MTUs are used in **Nos**.

Choosing an MTU for an AX.25 interface is more complex. When the interface operates in datagram (UI-frame) mode, the **paclen** parameter does not apply. The MTU effectively becomes the **paclen** of the link. However, as mentioned earlier, large packets sent on AX.25 *connections* are automatically segmented into I-frames no larger than **paclen** bytes. Unfortunately, as also mentioned earlier, **Nos** is so far the only known implementation of the new AX.25 segmentation procedure. This is fine as long as all of the NET/ROM nodes along a path are running **Nos**, but since the main reason **Nos** supports NET/ROM is to allow use of existing NET/ROM networks, this is unlikely.

So it is usually important to avoid AX.25 segmentation when running IP over NET/ROM. The way to do this is to make sure that packets larger than **paclen** are never handed to AX.25. A NET/ROM transport header is 5 bytes long and a NET/ROM network header takes 15 bytes, so 20 bytes must be added to the size of an IP datagram when figuring the size of the AX.25 I-frame data field. If **paclen** is 256, this leaves 236 bytes for the IP datagram. This is the default MTU of the **netrom** pseudo-interface, so as long as **paclen** is at least 256 bytes, AX.25 segmentation can't happen. But if smaller values of **paclen** are used, the **netrom** MTU must also be reduced with the **ifconfig** command.

On the other hand, if you're running IP directly on top of AX.25, chances are all of the nodes are running **Nos** and support AX.25 segmentation. In this case there is no reason not to use a larger MTU and let AX.25 segmentation do its thing. If you choose an MTU on the order of 1000-1500 bytes, you can largely avoid IP-level fragmentation and reduce TCP/IP-level header overhead on file transfers to a very low level. And you are still free to pick whatever **paclen** value is appropriate for the link.

8.4.5. Setting MSS

The setting of this TCP-level parameter is somewhat less critical than the IP and AX.25 level parameters already discussed, mainly because it is automatically lowered according to the MTU of the local

interface when a connection is created. Although this is, strictly speaking, a protocol layering violation (TCP is not supposed to have any knowledge of the workings of lower layers) this technique does work well in practice. However, it can be fooled; for example, if a routing change occurs after the connection has been opened and the new local interface has a smaller MTU than the previous one, IP fragmentation may occur in the local system.

The only drawback to setting a large MSS is that it might cause avoidable fragmentation at some other point within the network path if it includes a "bottleneck" subnet with an MTU smaller than that of the local interface. (Unfortunately, there is presently no way to know when this is the case. There is ongoing work within the Internet Engineering Task Force on a "MTU Discovery" procedure to determine the largest datagram that may be sent over a given path without fragmentation, but it is not yet complete.) Also, since the MSS you specify is sent to the remote system, and not all other TCPs do the MSS-lowering procedure yet, this might cause the remote system to generate IP fragments unnecessarily.

On the other hand, a too-small MSS can result in a considerable performance loss, especially when operating over fast LANs and networks that can handle larger packets. So the best value for MSS is probably 40 less than the largest MTU on your system, with the 40-byte margin allowing for the TCP and IP headers. For example, if you have a SLIP interface with a 1006 byte MTU and an Ethernet interface with a 1500 byte MTU, set MSS to 1460 bytes. This allows you to receive maximum-sized Ethernet packets, assuming the path to your system does not have any bottleneck subnets with smaller MTUs.

8.4.6. Setting Window

A sliding window protocol like TCP cannot transfer more than one window's worth of data per round trip time interval. So this TCP-level parameter controls the ability of the remote TCP to keep a long "pipe" full. That is, when operating over a path with many hops, offering a large TCP window will help keep all those hops busy when you're receiving data. On the other hand, offering too large a window can congest the network if it cannot buffer all that data. Fortunately, new algorithms for dynamic controlling the effective TCP flow control window have been developed over the past few years and are now widely deployed. **Nos** includes them, and you can watch them in action with the **tcp status** <**tcb**> or **socket** <**sockno**> commands. Look at the **cwind** (congestion window) value.

In most cases it is safe to set the TCP window to a small integer multiple of the MSS, (eg. 4times), or larger if necessary to fully utilize a high bandwidth*delay product path. One thing to keep in mind, however, is that advertising a certain TCP window value declares that the system has that much buffer space available for incoming data. **Nos** does not actually preallocate this space; it keeps it in a common pool and may well "overbook" it, exploiting the fact that many TCP connections are idle for long periods and gambling that most applications will read incoming data from an active connection as soon as it arrives, thereby quickly freeing the buffer memory. However, it is possible to run **Nos** out of memory if excessive TCP window sizes are advertised and either the applications go to sleep indefinitely (eg. suspended Telnet sessions) or a lot of out-of-sequence data arrives. It is wise to keep an eye on the amount of available memory and to decrease the TCP window size (or limit the number of simultaneous connections) if it gets too low.

Depending on the channel access method and link level protocol, the use of a window setting that exceeds the MSS may cause an increase in channel collisions. In particular, collisions between data packets and returning acknowledgements during a bulk file transfer may become common. Although this is, strictly speaking, not TCP's fault, it is possible to work around the problem at the TCP level by decreasing the window so that the protocol operates in stop-and-wait mode. This is done by making the window value equal to the MSS.

8.5. Summary

In most cases, the default values provided by **Nos** for each of these parameters will work correctly and give reasonable performance. Only in special circumstances such as operation over a very poor link or experimentation with high speed modems should it be necessary to change them.