The Linux Modem-HOWTO

David S.Lawyer <mailto:bf347@lafn.org> Contains much info from Serial-HOWTO by Greg Hankins
v0.05, October 1999

Help with selecting, connecting, configuring, trouble-shooting, and understanding modems for a PC. See Serial-HOWTO for multiport serial boards.

Contents

1 Introduction **6**

1.1 DSL, Cable, and ISDN Modems in other HOWTOs . . . 6
1.2 Also not covered: PCMCIA Modems, PPP . . . 6
1.3 Copyright, Disclaimer, Trademarks, & Credits . . . 6
1.3.1 Copyright . . . 6
1.3.2 Disclaimer . . . 7
1.3.3 Trademarks . . . 7
1.3.4 Credits . . . 7
1.4 Future Plans, You Can Help . . . 7
1.5 New Versions of this HOWTO . . . 7
1.6 What is a Modem ? . . . 7
1.7 Quick Install . . . 8
1.7.1 External Modem Install . . . 8
1.7.2 Internal Modems (on ISA bus) . . . 8
1.7.3 All Modems . . . 9

2 Modem & Serial Port Basics **9**

2.1 Modem Converts Digital to Analog (and conversely) . . . 9
2.2 What is a Serial Port ? . . . 9
2.2.1 Intro to Serial . . . 9
2.2.2 Pins and Wires . . . 10
2.2.3 Internal Modem Contains Serial Port . . . 10
2.3 IO Address & IRQ . . . 10
2.4 Names: ttyS0, ttyS1, etc. . . . 11
2.5 Interrupts . . . 11
2.6 Data Compression (by the Modem) . . . 12
2.7 Error Correction . . . 12
2.8 Data Flow (Speeds) . . . 12

2.9 Flow Control . . . 13
2.9.1 Example of Flow Control . . . 13
2.9.2 Hardware vs. Software Flow Control . . . 14
2.9.3 Symptoms of No Flow Control . . . 14
2.9.4 Modem-to-Modem Flow Control . . . 14
2.10 Data Flow Path; Buffers . . . 14
2.11 Modem Commands . . . 15
2.12 Serial Software: Device Driver Module . . . 15

3 Configuring Modems (including the serial port) 16
3.1 Configuring Overview . . . 16

4 Configuring the Serial Port 16
4.1 PCI Bus Not Yet Well Supported . . . 16
4.2 Configuring Overview . . . 16
4.3 Common mistakes made re low-level configuring . . . 17
4.4 I/O Address & IRQ: Boot-time messages . . . 18
4.5 What is the current IO address and IRQ of my Serial Port ? . . . 19
4.5.1 What does the device driver think? . . . 19
4.5.2 What is set in my serial port hardware ? . . . 19
4.5.3 What is set in my PnP serial port hardware ? . . . 20
4.6 Choosing Serial IRQs . . . 20
4.6.1 IRQ 0 is not an IRQ . . . 20
4.6.2 Interrupt sharing and Kernels 2.2+ . . . 20
4.6.3 What IRQs to choose? . . . 21
4.7 Choosing Addresses –Video card conflict with ttyS3 . . . 21
4.8 Set IO Address & IRQ in the hardware (mostly for PnP) . . . 22
4.8.1 Using a PnP BIOS to I0-IRQ Configure . . . 22
4.9 Giving the IRQ and IO Address to Setserial . . . 23
4.10 Other Configuring . . . 23
4.10.1 Configuring Hardware Flow Control (RTS/CTS) . . . 23

5 Modem Configuration (excluding serial port) 24
5.1 AT Commands . . . 24
5.2 Init Strings: Saving and Recalling . . . 24
5.3 Other Modem Commands . . . 25

6 Modems for a Linux PC **26**

6.1 External vs. Internal . . . 26

6.2 External Modems . . . 26

6.2.1 PnP External Modems . . . 26

6.2.2 Cabling & Installation . . . 27

6.2.3 What the Lights (LED's) Mean . . . 27

6.3 Internal Modems . . . 27

6.4 Internal Modems You Must Avoid (winmodems, etc.) . . . 28

6.5 Which Internal Modems are Best to Avoid ? . . . 28

6.5.1 MWave and DSP Modems . . . 29

6.5.2 Rockwell (RPI) Drivers . . . 29

6.5.3 PCI Modems . . . 29

7 Serial Port Devices /dev/ttyS2, etc. **29**

7.1 Serial Port Device Names & Numbers . . . 29

7.2 Link ttySN to /dev/modem ? . . . 30

7.3 The cua Device . . . 30

8 Interesting Programs You Should Know About **30**

8.1 What is setserial ? . . . 30

8.1.1 Introduction . . . 30

8.1.2 Probing . . . 31

8.1.3 Can Linux Configure The Serial Devices Automagically? . . . 31

8.1.4 Boot-time Configuration . . . 32

8.1.5 IRQs . . . 33

8.1.6 Where to Run Setserial ? . . . 33

8.2 What is isapnp ? . . . 34

8.3 What is wvdialconf ? . . . 34

8.4 What is stty ? . . . 34

9 Trying Out Your Modem (Dialing Out) . **35**

9.1 Are You Ready to Dial Out ? . . . 35

9.2 Dialing Out with Minicom . . . 35

9.3 Dialing Out with Kermit . . . 36

10 Dial-In **37**

10.1 Overview . . . 37

10.2 Getty . . . 37

10.2.1 About mgetty . . . 37

10.2.2 About uugetty . . . 38
10.2.3 About getty_em . . . 38
10.2.4 About agetty and mingetty . . . 38
10.3 What Happens when Someone Dials In ? . . . 38
10.4 Why Manual Answer is Best . . . 39
10.5 Callback . . . 39
10.6 Voice Mail . . . 39

11 Uugetty for Dial-In (from the old Serial-HOWTO) 40
11.1 Installing getty_ps . . . 40
11.2 Setting up uugetty . . . 40
11.2.1 Modern Modems . . . 41
11.2.2 Old slow modems . . . 41
11.2.3 Login Banner . . . 41
11.3 Customizing uugetty . . . 41

12 What Speed Should I Use with My Modem? 42
12.1 Speed and Data Compression . . . 43
12.2 Where do I Set Speed ? . . . 43
12.3 Can't Set a High Enough Speed . . . 43
12.3.1 How speed is set in hardware: the divisor and baud_base . . . 43
12.3.2 Work-arounds for setting speed . . . 44
12.3.3 Crystal frequency is not baud_base . . . 44
12.4 Speed Table . . . 44

13 Communications Programs And Utilities 44
13.1 Minicom vs. Kermit . . . 45
13.2 Lists of Programs . . . 45
13.2.1 Least Popular Dialout . . . 45
13.2.2 Most Popular Dialout . . . 45
13.2.3 Fax . . . 45
13.2.4 Voicemail Software . . . 45
13.2.5 Dial-in (uses getty) . . . 46
13.2.6 Other . . . 46
13.3 SLiRP and term . . . 46

14 What Are UARTs? How Do They Affect Performance? 46
14.1 Introduction to UARTS . . . 46
14.2 Two Types of UARTs . . . 47

14.3 FIFOs . . . 47
14.4 UART Model Numbers . . . 48

15 Troubleshooting **49**
15.1 I can't get 56K on my 56K modem . . . 49
15.2 Uploading (downloading) files is broken/slow . . . 49
15.3 For Dial-in I Keep Getting "line NNN of inittab invalid" . . . 49
15.4 When I Try To Dial Out, It Says "/dev/ttySN: Device or resource busy" . . . 49
15.5 I Keep Getting: "Id "S3" respawning too fast: disabled for 5 minutes" . . . 49
15.6 My Modem is Hosed after Someone Hangs Up, or uugetty doesn't respawn . . . 50
15.7 uugetty Still Doesn't Work . . . 50
15.8 My Modem is Physically There but Can't be Found . . . 50
15.9 The following subsections are in both the Serial and Modem HOWTOs: . . . 51
15.10My Serial Port is Physically There but Can't be Found . . . 51
15.11Slow: Text appears on the screen slowly after long delays . . . 51
15.12The Startup Screen Show Wrong IRQs for the Serial Ports. . . . 51
15.13"Cannot open /dev/ttyS?: Permission denied" . . . 52
15.14"Operation not supported by device" (error message) for ttySx . . . 52
15.15"Cannot create lockfile. Sorry" (error message) . . . 52
15.16"Device /dev/ttyS? is locked." . . . 52
15.17Software which may help . . . 52

16 Flash Upgrades **53**

17 Other Sources of Information **53**
17.1 Misc . . . 53
17.2 Books . . . 54
17.3 HOWTOs . . . 54
17.4 Usenet newsgroups . . . 54
17.5 Web Sites . . . 54

18 Appendix A: How Modems Work (technical) (unfinished) **55**
18.1 Modulation Details . . . 55
18.1.1 Intro to Modulation . . . 55
18.1.2 Frequency Modulation . . . 55
18.1.3 Amplitude Modulation . . . 55
18.1.4 Phase Modulation . . . 55
18.1.5 Combination Modulation . . . 56
18.1.6 56k Modems (v.90) . . . 56

19 Appendix B: "baud" vs. "bps" 57

19.1 A simple example 57

19.2 Real examples 58

20 Appendix C: Terminal Server Connection 58

21 Appendix D: Other Types of Modems 59

21.1 Digital-to-Digital "Modems" 59

21.2 ISDN "Modems" 59

21.3 Digital Subscriber Line (DSL) 59

1 Introduction

1.1 DSL, Cable, and ISDN Modems in other HOWTOs

This HOWTO covers conventional modems for PC's, mainly modems on the ISA bus (although much of this should also apply to the PCI bus).

- DSL modems: see the mini-howto: ADSL
- Cable-Modems-HOWTO (was once a LDP mini-Howto) `<http://www.cs.unm.edu/~vuksan/linux/Cable-Modem.html>`
- ISDN Howto (not a LDP Howto) `<http://sdb.suse.de/sdb/en/html/isdn.html>`: drivers for ISDN "Modems". Much related info on this is in German.

1.2 Also not covered: PCMCIA Modems, PPP

For modems on the PCMCIA bus see the PCMCIA-HOWTO: PCMCIA serial and modem devices. This HOWTO doesn't cover PPP (used to connect to the Internet via a modem) or communication programs. Except it does show how to use communication programs to test that your modem works OK and can make phone calls. If you want to use a modem to connect to the Internet then you need to set up PPP. There's a lot of documentation for PPP (including a PPP-HOWTO which is being revised). Some of it might be found in /usr/doc/ppp or the like.

1.3 Copyright, Disclaimer, Trademarks, & Credits

1.3.1 Copyright

Copyright (c) 1998-9 by David S. Lawyer Please freely copy and distribute (sell or give away) this document. For corrections and minor changes contact the maintainer. You may create a derivative work and distribute it provided:

1. For a translation just let the maintainer know about it.
2. For a non-translation, discuss it with the maintainer (if it is being maintained). If after due deliberation no agreement with the maintainer can be reached, then you may proceed with the derivative work without the maintainer's approval.

3. Send your derivative work to the LDP (or the like) for posting on the Internet. If not the LDP, then let the LDP know where it is available.

4. License the derivative work in the spirit of this license or use GPL.

5. Give due credit to previous authors and major contributors.

1.3.2 Disclaimer

While I haven't intentionally tried to mislead you, there are likely a number of errors in this document. Please let me know about them. Since this is free documentation, it should be obvious that I cannot be held legally responsible for any errors.

1.3.3 Trademarks

If certain words are trademarks, the context should make it clear to whom they belong. For example "MS Windows" (or just "Windows") implies that "Windows" belongs to Microsoft (MS). "Hayes" is a trademark of Microcomputer Products Inc. I use "winmodem" to mean any modem which requires MS-Windows and not in the trademark sense.

1.3.4 Credits

The following is only a rough approximation of how version 0.0 of this document was created: About 1/3 of the material here was lifted directly from Serial-HOWTO v. 1.11 by Greg Hankins. <mailto:gregh@cc.gatech.edu> (with his permission). About another 1/3 was taken from that Serial-HOWTO and revised. The remaining 1/3 is newly created by the author: David S. Lawyer <mailto:bf347@lafn.org>.

1.4 Future Plans, You Can Help

An explanation of how to set up modems for dial-in is lacking in this version but should be included in the next. Please let me know of any errors in facts, opinions, logic, spelling, grammar, clarity, links, etc. But first, if the date is over a months old, check to see that you have the latest version. Please send me any other info that you think belongs in this document. The French "Modems-HOWTO" needs to be somehow merged with this document (but I don't know French).

1.5 New Versions of this HOWTO

New versions of this Modem-HOWTO come out every month or so since modem situation is rapidly changing (and since I'm still learning). Your problem might be solved in the latest version. It will be available to browse and/or download at LDP mirror sites. For a list of such sites see: <http://metalab.unc.edu/LDP/mirrors.html> If you only want to quickly check the date of the latest version you may not want to use a mirror site so check out: <http://metalab.unc.edu/LDP/HOWTO/Modem-HOWTO.html>.

1.6 What is a Modem ?

A modem is a device that lets one send digital signals over ordinary telephone lines not designed for digital signals. If telephone lines were all digital then you wouldn't need a modem. It permits your computer to connect to and communicate with the rest of the world. When you use a modem, you normally use a communication program or web browser (which includes such a program) to utilize the modem and dial-out

on a telephone line. Advanced modem users can set things up so that others may phone in to them and use their computer. This is called "dial-in".

There are two basic types of modems for a PC: external and internal. The external sets on your desk outside the PC while the internal is not visible since it's inside the PC. The external modem plugs into a connector on the back of the PC known as a "serial port". The internal modem is a card that is inserted inside the computer and has an (invisible) serial port built into it. For a more detailed comparison see 6.1 (External vs. Internal). Thus when you get an internal modem, you also get a dedicated serial port (which can only be used with the modem and not with anything else such as another modem or a printer). In Linux, the serial ports are named ttyS0, ttyS1, etc. (usually corresponding respectively to COM1, COM2, etc. in Dos/Windows).

The serial port is not to be confused with the "Universal Serial Bus" (USB) which uses a special modular connector and may be used with modems in the future. See 2 (Modem & Serial Port Basics) for more details on modems and serial ports.

Modems often include the ability to send Faxes (Fax Modems). See 13.2.3 (Fax) for a list of fax software. "Voice" modems can work like an automatic answering machine and handle voicemail. See 10.6 (Voicemail).

1.7 Quick Install

1.7.1 External Modem Install

With a straight-thru or modem cable, connect the modem to an unused serial port on the PC. Make sure you know the name of the serial port: in most cases COM1 is ttyS0, COM2 is ttyS1, etc. You may need to check the BIOS setup menu to determine this. Plug in the power cord to provide power to the modem. See 1.7.3 (All Modems) for further instructions.

1.7.2 Internal Modems (on ISA bus)

(For the PCI bus see 4.1 (PCI Bus Not Yet Well Supported) and 6.5.3 (PCI Modems).) If the modem says it will only work under MS Windows, you are out of luck. If you already have 2 serial ports, make the modem the 3rd serial port (ttyS2 = COM3). Find an unused IRQ number to use. In the past IRQ 5 was often used but today IRQ 5 is also used for sound cards. Then set the jumpers (or the like) on the internal modem to the unused IRQ and IO address 3E8 (ttyS2) .

"Or the like" (in the previous sentence) may be a bit tricky. If the modem is a Plug and Play (PnP) for the ISA bus, the equivalent probably can be done using the "isapnp" program which comes with "isapnptools". See "man isapnp" or the FAQ for it. See also "Plug-and-Play-HOWTO. With a PnP-BIOS you may be able to tell the CMOS setup menu that you don't have a PnP OS and then the BIOS may set a suitable IRQ and IO address in the modem card. If you want to "force" the BIOS to set a certain IRQ and/or IO then you may be able to do this using Window9x on the same PC. It can set them into the PnP BIOS's flash memory where they will be used to configure for Linux as well as Windows. See "Plug-and-Play-HOWTO and search for "forced" (occurs in several places). For Windows 3.x you can do the same thing using the ICU under Windows 3.x. There may even be a way to disable PnP using software (under Windows) that came with the modem.

Finally you must also find the file where "setserial" is run and add a line something like: "setserial /dev/ttyS2 irq5". Except for setserial v2.15 and later you may (if you distribution lets you) just run "setserial" on the command line and the results are saved to a configuration file. See 8.1 (What is Setserial) for more info. See the next subsection 1.7.3 (All Modems) for further instructions on quick installation.

1.7.3 All Modems

Plug the modem into a telephone line. Then start up a communication program such as minicom and go to the configuration menu for the serial port. Assign it a high baud rate a few times higher than the bit rate of your modem. See 12.4 (Speed Table) for the "best" speeds to use. Tell it the full name of your serial port such as /dev/ttyS1. Set hardware flow control (RTS/CTS). Now you need to save these settings and exit minicom. Then start minicom again, type AT to see if your modem is there and responds with OK. Then go to the dial directory (or menu) and dial a number.

2 Modem & Serial Port Basics

You don't have to understand the basics to use and install a modem. But understanding it may help to determine what is wrong if you run into problems. After reading this section, if you want to understand it even better you may want to see 18 (How Modems Work) in this document (not yet complete). More details on the serial port (including much of this section) will be found in Serial-HOWTO.

2.1 Modem Converts Digital to Analog (and conversely)

Most all telephone main lines are digital already but the lines leading to your house (or business) are usually analog which means that they were designed to transmit a voltage wave which is an exact replica of the sound wave coming out of your mouth. Such a voltage wave is called "analog". If viewed on an oscilloscope it looks like a sine wave of varying frequency and amplitude. A digital signal is like a square wave. For example 3 v (volts) might be a 1-bit and 0 v could be a 0-bit. For most serial ports (used by external modems) +12 v is a 0-bit and -12 v is a 1-bit (some are + or - 5 v).

To send data from your computer over the phone line, the modem takes the digital signal from your computer and converts it to "analog". It does this by both creating an analog sine wave and then "MODulating" it. Since the result still represents digital data, it could also be called a digital signal instead of analog. But it looks something like an analog signal and almost everyone calls it analog. At the other end of the phone line another modem "DEModulates" this signal and the pure digital signal is recovered. Put together the "mod" and "dem" parts of the two words above and you get "modem" (if you drop one of the two d's). A "modem" is thus a MODulator-DEModulator. Just what modulation is may be found in the section 18.1 (Modulation Details).

2.2 What is a Serial Port ?

2.2.1 Intro to Serial

The serial port is an I/O (Input/Output) device. Since modems have a serial port between them and the computer, it's necessary to understand the serial port as well as the modem.

Most PC's have one or two serial ports. Each has a 9-pin connector (sometimes 25-pin) on the back of the computer. Computer programs can send data (bytes) to the transmit pin (output) and receive bytes from the receive pin (input). The other pins are for control purposes and ground.

The serial port is much more than just a connector. It converts the data from parallel to serial and changes the electrical representation of the data. Inside the computer, data bits flow in parallel (using many wires at the same time). Serial flow is a stream of bits over a single wire (such as on the transmit or receive pin of the serial connector). For the serial port to create such a flow, it must convert data from parallel (inside the computer) to serial on the transmit pin (and conversely).

Most of the electronics of the serial port is found in a computer chip (or a section of a chip) known as a UART. For more details on UARTs see the section 14 (What Are UARTs? How Do They Affect Performance?). But you may want to finish this section first so that you will hopefully understand how the UART fits into the overall scheme of things.

2.2.2 Pins and Wires

Old PC's used 25 pin connectors but only about 9 pins were actually used so today most connectors are only 9-pin. Each of the 9 pins usually connects to a wire. Besides the two wires used for transmitting and receiving data, another pin (wire) is signal ground. The voltage on any wire is measured with respect to this ground. Thus the minimum number of wires to use for 2-way transmission of data is 3. Except that it has been known to work with no signal ground wire but with degraded performance and sometimes with errors.

There are still more wires which are for control purposes (signalling) only and not for sending bytes. All of these signals could have been shared on a single wire, but instead, there is a separate dedicated wire for every type of signal. Some (or all) of these control wires are called "modem control lines". Modem control wires are either in the asserted state (on) of +12 volts or in the negated state (off) of -12 volts. One of these wires is to signal the computer to stop sending bytes out the serial port cable. Conversely, another wire signals the device attached to the serial port to stop sending bytes to the computer. If the attached device is a modem, other wires may tell the modem to hang up the telephone line or tell the computer that a connection has been made or that the telephone line is ringing (someone is attempting to call in). See the Serial-HOWTO: Pinout and Signals for more details.

2.2.3 Internal Modem Contains Serial Port

For an internal modem there is no 9-pin connector but the behavior is almost exactly as if the above mentioned cable wires existed. Instead of a a 12 volt signal in a wire giving the state of a modem control line, the internal modem may just use a status bit in its own memory (a register) to determine the state of this non-existent "wire". The internal modem's serial port looks just like a real serial port to the computer. It even includes the speed limits that one may set at ordinary serial ports such as 115200 bits/sec. Unfortunately for Linux, many internal modems today don't work exactly this way but instead use MS Windows software to help do their job. They will not work under Linux. See 6.4 (Avoid: winmodems).

2.3 IO Address & IRQ

Since the computer needs to communicate with each serial port, the operating system must know that each serial port exists and where it is (its I/O address). It also needs to know which wire (IRQ number) the serial port must use to request service from the computer's CPU. It requests service by sending an interrupt on this wire. Thus every serial port device must store in its non-volatile memory both its I/O address and its Interrupt ReQuest number: IRQ. See 2.5 (Interrupts). For the PCI bus it doesn't work exactly this way since the PCI bus has its own system of interrupts. But since the PCI-aware BIOS sets up chips to map these PCI interrupts to IRQs, it seemingly behaves just as described above except that sharing of interrupts is allowed (2 or more devices may use the same IRQ number).

I/O addresses are not the same as memory addresses. When an I/O addresses is put onto the computer's address bus, another wire is energized. This both tells main memory to ignore the address and tells all devices which have I/O addresses (such as the serial port) to listen to the address to see if it matches the device's. If the address matches, then the I/O device reads the data on the data bus.

2.4 Names: ttyS0, ttyS1, etc.

The serial ports are named ttyS0, ttyS1, etc. (and usually correspond respectively to COM1, COM2, etc. in DOS/Windows). The /dev directory has a special file for each port. Type "ls /dev/ttyS*" to see them. Just because there may be (for example) a ttyS3 file, doesn't necessarily mean that there exists a physical serial port there.

Which one of these names (ttyS0, ttyS1, etc.) refers to which physical serial port is determined as follows. The serial driver (software) maintains a table showing which I/O address corresponds to which ttyS. This mapping of names (such as ttyS1) to I/O addresses (and IRQ's) may be both set and viewed by the "setserial" command. See 8.1 (What is Setserial). This does `not` set the I/O address and IRQ in the hardware itself (which is set by jumpers or by plug-and-play software). Thus what physical port corresponds to say ttyS1 depends both on what the serial driver thinks (per setserial) and what is set in the hardware. If a mistake has been made, the physical port may not correspond to any name (such as ttyS2) and thus it can't be used. See 7 (Serial Port Devices /dev/ttyS2, etc.) for more details>

2.5 Interrupts

Bytes come in over the phone line to the modem, are converted from analog to digital by the modem and passed along to the serial port on their way to their destination inside your computer. When the serial port receives several bytes (may be set to 1, 4, 8, or 14) into its buffer, it signals the CPU to fetch it (them) by sending an electrical signal known as an interrupt on a certain wire. On the ISA bus, each device usually has it own private interrupt line (wire). This interrupt will also be sent if there is an unexpected delay while waiting for the next byte to arrive (known as a timeout). Thus if the bytes are being received slowly (such as someone typing on a terminal keyboard) there may be an interrupt issued for every byte received. For some UART chips if 4 bytes could have received, but none of these 4 show up, then the port gives up waiting for the next byte and issues an interrupt to fetch the 1 or more bytes waiting in the buffer.

Each interrupt conductor (inside the computer) has a number (IRQ) and the serial port must know which conductor to use to signal on. For example, ttyS0 normally uses IRQ number 4 known as IRQ4 (or IRQ 4). A list of them and more will be found in "man setserial" (search for "Configuring Serial Ports"). Interrupts are issued whenever the serial port needs to get the CPU's attention. It's important to do this in a timely manner since the buffer inside the serial port can hold only 16 (1 in old serial ports) incoming bytes. If the CPU fails to remove such received bytes promptly, then there will not be any space left for any more incoming bytes and the small buffer may overflow (overrun) resulting in a loss of data bytes.

For an external modem, there is no way (such as flow control) to stop the flow rapidly enough to prevent this. For an internal modem the 16-byte buffer is on the same card and a good modem will not write to it if it's full. Thus a good internal modem will not overrun the 16-byte buffers. This is one advantage of an internal modem over an external.

Interrupts are also issued when the serial port has just sent out all 16 of its bytes from its small transmit buffer out the external cable. It then has space for 16 more outgoing bytes. The interrupt is to notify the CPU of that fact so that it may put more bytes in the small transmit buffer to be transmitted. Also, when a modem control line changes state an interrupt is issued.

The buffers mentioned above are all hardware buffers. The serial port also has large buffers in main memory. This will be explained later

Interrupts convey a lot of information but only indirectly. The interrupt itself just tells a chip called the interrupt controller that a certain serial port needs attention. The interrupt controller then signals the CPU. The CPU runs a special program to service the serial port. That program is called an interrupt service routine (part of the serial driver software). It tries to find out what has happened at the serial port and then deals with the problem such a transferring bytes from (or to) the serial port's hardware buffer. This program

can easily find out what has happened since the serial port has registers at IO addresses known to the the serial driver software. These registers contain status information about the serial port. The software reads these registers and by inspecting the contents, finds out what has happened and takes appropriate action.

2.6 Data Compression (by the Modem)

Before continuing with the basics of the serial port, one needs to understand about something done by the modem: data compression. In some cases this task is actually done by software run on the computer's CPU but unfortunately at present, such software only works for MS Windows. The discussion here will be for the case where the modem itself does the compression since this is what must happen in order for the modem to work under Linux.

In order to send data faster over the phone line, one may compress (encode it) using a custom encoding scheme which itself depends on the data. The encoded data is smaller than the original (less bytes) and can be sent over the Internet in less time. This process is called "data compression".

If you download files from the Internet, they are likely already compressed and it is not feasible for the modem to try to compress them further. Your modem may sense that what is passing thru has already been compressed and refrain from trying a compress it any more. If you are receiving data which has been compressed by the other modem, your modem will decompress it and create many more bytes than were sent over the phone line. Thus the flow of data from your modem into your computer will be higher than the flow over the phone line to you. The ratio of this flow is called the compression ratio. Compression ratios as high as 4 are possible, but not very likely.

2.7 Error Correction

Similar to data compression, modems may be set to do error correction. While there is some overhead cost involved which slows down the byte/sec flow rate, the fact that error correction strips off start and stop bits actually increases the data byte/sec flow rate.

For the serial port's interface with the external world, each 8-bit byte has 2 extra bits added to it: a start-bit and a stop-bit. Without error correction, these extra start and stop bits usually go right thru the modem and out over the phone lines. But when error correction is enabled, these extra bits are stripped off and the 8-bit bytes are put into packets. This is more efficient and results in higher byte/sec flow in spite of the fact that there are a few more bytes added for packet headers and error correction purposes.

2.8 Data Flow (Speeds)

Data (bytes representing letters, pictures, etc.) flows from your computer to your modem and then out on the telephone line (and conversely). Flow rates (such as 56k (56000) bits/sec) are (incorrectly) called "speed". But almost everyone says "speed" instead of "flow rate". If there were no data compression the flow rate from the computer to the modem would be about the same as the flow rate over the telephone line.

Actually there are two different speeds to consider at your end of the phone line:

- The speed on the phone line itself (DCE speed) modem-to-modem
- The speed from your computer's serial port to your modem (DTE speed)

When you dial out and connect to another modem on the other end of the phone line, your modem often sends you a message like "CONNECT 28800" or "CONNECT 115200". What do these mean? Well, its either the DCE speed or the DTE speed. If it's higher than the advertised modem speed it must be the DTE

modem-to-computer speed. This is the case for the 115200 speed shown above. The 28800 must be a DCE (modem-to-modem) speed since the serial port has no such speed. One may configure the modem to report either speed. Some modems report both speeds and report the modem-to-modem speed as (for example): CARRIER 28800.

If you have an internal modem you would not expect that there would be any speed limit on the DTE speed from your modem to your computer since you modem is inside your computer and is almost part of your computer. But there is since the modem contains a dedicated serial port within it.

It's important to understand that the average speed is often less than the specified speed, especially on the short DTE computer-to-modem line. Waits (or idle time) result in a lower average speed. These waits may include long waits of perhaps a second due to 2.9 (Flow Control). At the other extreme there may be very short waits (idle time) of several micro-seconds separating the end of one byte and the start of the next byte. In addition, modems will fallback to lower speeds if the telephone line conditions are less than pristine.

For a discussion of what DTE speed is best to use see section 12 (What Speed Should I Use).

2.9 Flow Control

Flow control means the ability to stop the flow of bytes in a wire. It also includes provisions to restart the flow without any loss of bytes. Flow control is needed for modems to allow a jump in instantaneous flow rates.

2.9.1 Example of Flow Control

For example, consider the case where you connect a 36.6k modem to your serial port that simply sends and receives bytes over the phone line at exactly 36.6k bits per second (bps). It's not doing any data compression or error correction. You have set the serial port speed to 115,200 bits/sec (bps), and you are sending data from your computer to the phone line. Then the flow from the your computer to your modem is at 115.2k bps. However the flow from your modem out the phone line is at best only 33.6k bps. Since a faster flow (115.2k) is going into your modem than is coming out of it, the modem is storing the excess flow (115.2k -33.6k = 81.6k bps) in one of its buffers. This buffer would eventually overrun (run out of storage space) unless the 115.2k flow is stopped.

But now flow control comes to the rescue. When the modem's buffer is almost full, the modem sends a stop signal to the serial port. The serial port passes on the stop signal on to the device driver and the 115.2k bps flow is halted. Then the modem continues to send out data at 33.6k bps drawing on the data it previous accumulated in its (the modem's) buffer. Since nothing is coming into the buffer, the level of bytes in it starts to drop. When almost no bytes are left in the buffer, the modem sends a start signal to the serial port and the 115.2k flow from the computer to the modem resumes. In effect, flow control creates an average flow rate (in this case 33.6k) which is significantly less than the "on" flow rate of 115.2k bps. This is "start-stop" flow control.

The above is a simple example of flow control for flow from the computer to a modem , but there is also flow control which is used for the opposite direction of flow: from a modem (or other device) to a computer. Each direction of flow involve 3 buffers: 1. in the modem 2. in the uart (serial port) 3. in main memory managed by the serial driver. Flow control protects certain buffers from overflowing. The small serial hardware buffers are not protected in this way but rely instead on a fast response to the interrupts they issue. This is the essence of flow control but there are still some more details.

You don't often need flow control in the direction from the modem to a PC. For complex example of a case where it's needed see "Complex Flow Control Example" in the Serial-HOWTO. But if you don't have a high enough speed set between the modem and the computer (serial port speed) then you do need to slow down

the flow from the modem to the PC. To do this you must stop the incoming flow of bytes over the telephone line. Your modem must tell the other modem to stop sending. See 2.9.4 (Modem-to-Modem Flow Control)

2.9.2 Hardware vs. Software Flow Control

If feasible best to use "hardware" flow control that uses two dedicated "modem control" wires to send the "stop" and "start" signals. Software flow control uses the main receive and transmit wires to send the start and stop signals. It uses the ASCII control characters DC1 (start) and DC3 (stop) for this purpose. They are just inserted into the regular stream of data. Software flow control is not only slower in reacting but also does not allow the sending of binary data unless special precautions are taken. Since binary data will likely contain DC1 and DC3, special means must be taken to distinguish between a DC3 that means a flow control stop and a DC3 that is part of the binary code. Likewise for DC1. To get software flow control to work for binary data requires both modem (hardware) and software support.

2.9.3 Symptoms of No Flow Control

Understanding flow-control theory can be of practical use. For example I used my modem to access the Internet and it seemed to work fine. But after a few months I tried to send long files from my PC to an ISP and a huge amount of retries and errors resulted (but eventually Kermit could send a long file after many retries). Receiving in the other direction (from my ISP to me) worked fine. The problem turned out to be a hardware defect in my modem that had resulted in disabling flow control. My modem's buffer was overflowing (overrunning) on long outgoing files since no "stop" signal was ever sent to the computer to halt sending to the modem. There was no problem in the direction from the modem to my computer since the capacity (say 115.2k) was always higher than the flow over the telephone line. The fix was to enable flow control by putting an enable-flow-control command for the modem last in the init string.

2.9.4 Modem-to-Modem Flow Control

This is the flow control of the data sent over the telephone lines between two modems. Practically speaking, it only exists when you have error correction enabled. Actually, even without error correction it's possible to enable software flow control between modems but it may interfere with sending binary data so it's not often used.

2.10 Data Flow Path; Buffers

Although much has been explained about this including flow control, a pair of 16-byte serial port buffers (in the hardware), and a pair of larger buffers inside a modem there is still another pair of buffers. These are large buffers (perhaps 8k) in main memory also known as serial port buffers. When an application program sends bytes to the serial port (and modem) they first get stashed in the the transmit serial port buffer in main memory. The pair consists of both this transmit buffer and a receive buffer for the opposite direction of byte-flow.

The serial device driver takes out say 16 bytes from this transmit buffer, one byte at a time and puts them into the 16-byte transmit buffer in the serial hardware for transmission. Once in that transmit buffer, there is no way to stop them from being transmitted. They are then transmitted to the modem which also has a fair sized (say 1k) buffer. When the device driver (on orders from flow control) stops the flow of outgoing bytes from the computer, what it actually stops is the flow of outgoing bytes from the large transmit buffer in main memory. Even after this has happened and the flow to the modem has stopped, an application program may keep sending bytes to the 8k transmit buffer until it becomes fill.

When it gets fill, the application program can't send any more bytes to it (a "write" statement in a C_program blocks) and the application program temporarily stops running and waits until some buffer space becomes available. Thus a flow control "stop" is ultimately able to stop the program that is sending the bytes. Even though this program stops, the computer does not necessarily stop computing. It may switch to running other processes while it's waiting at a flow control stop. The above was a little oversimplified since there is another alternative of having the application program itself do something else while it is waiting to "write".

2.11 Modem Commands

Commands to the modem are sent to it from the communication software over the same conductor as used to send data. The commands are short ASCII strings. Examples are "AT&K3" for enabling hardware flow control (RTS/CTS) between your computer and modem; and "ATDT5393401 for Dialing the number 5393401. Note all commands are prefaced by "AT". Some commands such as enabling flow control help configure the modem. Other commands such as dialing a number actually do something. There are about a hundred or so different possible commands. When your communication software starts running, it first sends an "init" string of commands to the modem to configure it. All commands are sent on the ordinary data line before the modem dials (or receives a call).

Once the modem is connected to another modem (on-line mode), everything that is sent from your computer to your modem goes directly to the other modem and is not interpreted by the modem as a command. There is a way to "escape" from this mode of operation and go back to command mode where everything sent to the modem will be interpreted as a command. The computer just sends "+++" with a specified time spacing before and after it. If this time spacing is correct, the modem reverts to command mode. Another way to do this is by a signal on a certain modem control line.

There are a number of lists of modem commands on the Internet. The section 17.5 (Web Sites) has links to a couple of such web-sites. Different models and brands of modems do not use exactly the same set of such commands. So what works for one modem might not work for another. Some common command (not guaranteed to work on all modems) are listed in this HOWTO in the section 5 (Modem Configuration)

2.12 Serial Software: Device Driver Module

The device driver for the serial port is the software that operates the serial port. It is now provided as a serial module. This module will normally get loaded automatically if it's needed. The kernel 2.2 + will do this. In earlier kernels, you had to have `kerneld` running in order to do auto-load modules on demand. Otherwise the serial module needed to be explicitly listed in /etc/modules. Before modules became popular with Linux, the serial driver was usually built into the kernel.

When the serial module is loaded it displays a message on the screen about the existing serial ports (often showing a wrong IRQ). But once the module is used by `setserial` to tell the device driver the (hopefully) correct IRQ then you should see a second display with the correct IRQ, etc. See 8.1 (What is Setserial) for more info on `setserial`.)

One may modify the driver by editing the kernel source code. Much of the serial driver is found in the file serial.c. For details regarding writing of programs for the serial port see Serial-Programming-HOWTO (currently being revised by Vern Hoxie).

3 Configuring Modems (including the serial port)

If you want to use a modem only for MS Windows/Dos, then you can just install almost any modem and it will work OK. With a Linux PC it's not usually this easy unless you use an external modem. All external modems should work OK (even if they are labeled "Plug and Play") But most new internal modems are Plug-and-Play (PnP) and have PnP serial ports. You may need to use the Linux "isapnp" program to configure these PnP serial ports. See the Plug-and-Play-HOWTO for more information.

3.1 Configuring Overview

Since each modem has an associated serial port there are two parts to configuring a modem:

- Configuring the modem itself: Done by the communication program
- Configuring the modem's serial port: Done only partly by the communication program

Most of the above configuring (but not necessarily most of the effort) is done by the communication program that you use with the modem such as `minicom` or `seyon`, or by the PPP (usually needed for use of a web browser). If you use the modem for dial-in, then the `getty` program which you use to present outsiders with a login-prompt, will help configure. Thus to configure the modem (and much of the serial port) you need to configure the communication program (or PPP or `getty`).

But note that not all of the configuration of the serial port is done by the communication program (or getty). The remaining configuring is simple to state (but sometimes difficult to do). It mainly consists of setting the IO address of the port and its IRQ number. In fact, plug-and-play could set these without you doing a thing. But there's a serious problem: Linux (as of early 1999) doesn't support plug-and-play very well. This may create a difficult problem for you. The next section covers the configuring of the serial port.

4 Configuring the Serial Port

4.1 PCI Bus Not Yet Well Supported

The kernel 2.2 serial driver contains no support for the PCI bus. But current kernels 2.3 and (eventually 2.4) will support some PCI serial cards (and modem cards). Most PCI cards need special support in the driver. The driver reads the id number digitally stored on the card to determine how (or if) to support the card. If you have a PCI card which is not a winmodem but it will not work, then you may ask the maintainer of the serial driver, Theodore (Ted) Y. Ts'o, to add support for it. Email him a copy of the output of "lspci" for the port. His email: `<mailto:tytso@mit.edu>`.

PCI modems are not well standardized. Some use main memory for I/O with the PC (ISA ones use the I/O address space for this). Some require special enabling of the IRQ. The output of "lspci" can help determine if one can be supported. If the details are not too complex I may find them and put them here in a future revision.

4.2 Configuring Overview

In many cases, configuring will happen automatically and you have nothing to do. But sometimes you need to configure (or just want to check out the configuration). If so, first you need to know about the two parts to configuring the serial port under Linux:

The first part (low-level configuring) is assigning it an IO address, IRQ, and name (such as ttyS2). This IO-IRQ pair must be set in both the hardware and told to the serial driver. We might just call this "io-irq" configuring for short. The `setserial` is used to tell the driver. PnP methods, jumpers, etc, are used to set the hardware. Details will be supplied later. If you need to configure but don't understand certain details it's easy to get into trouble.

The second part (high-level configuring) is assigning it a speed (such as 38.4K bits/sec), selecting flow control, etc. This is often done by communication programs such as PPP, minicom, or by getty (which you may run on the port so that others may log into it). This high-level configuring may also be done with the `stty` program. `stty` is also useful to view the current status if you're having problems and to reset the terminal interface to close to "normal" by typing something like "stty sane". See also the Serial-HOWTO section: "Stty". When Linux starts, some effort is made to detect and configure (low-level) a few serial ports. Exactly what happens depends on your BIOS, hardware, Linux distribution, etc. If the serial ports work OK, there may be no need for you to do any configuring. Application programs often do the high-level configuring but you may need to supply them some information so that they can do it. With Plug-and-Play serial ports (often built into an internal modem), the situation has become more complex. Here are cases when you need to do low-level configuring (set IRQ and IO addresses):

- You want to use more than 2 serial ports
- Installing a new serial port (such as an internal modem)
- Having problems with serial port(s)

For kernel 2.2+ you may be able to use more that 2 serial ports without low-level configuring by sharing interrupts but this may be no easier than low-level configuring. See 4.6.2 (Interrupt sharing and Kernels 2.2+)

The low-level configuring (setting the IRQ and IO address) seems to cause people more trouble (than high-level), although for many it's fully automatic and there is no configuring to be done. Thus most all of this section is on that topic. Until the serial driver knows the correct IRQ and IO address the port will probably not work at all. It likely cannot even be found by Linux. Even if it can be found, it may work extremely slow if the IRQ is wrong. See 15.11 (Slow: Text appears on the screen slowly after long delays).

In the Wintel world, the IO address and IRQ are called "resources" and we are thus configuring certain resources. In review, the low-level configuring consists of putting two values (an IRQ number and IO address) into two places:

1. the device driver (often by running "`setserial`" at boot-time)
2. memory registers of the serial port hardware itself

You may watch the start-up (= boot-time) messages. They are usually correct. But if you're having problems, there's a good chance that some of these messages don't show the true configuration of the hardware (and they are not supposed to). See 4.4 (I/O Address & IRQ: Boot-time messages).

4.3 Common mistakes made re low-level configuring

Here are some common mistakes people make:

- setserial command: They run it (without the "autoconfig" option) and think it has checked out the hardware (it hasn't).

- setserial messages: They see them displayed on the screen at boot-time, and erroneously think that the result shows how their hardware is actually configured.
- /proc/interrupts: When their serial device isn't in use they don't see its interrupt there, and erroneously conclude that their serial port can't be found (or doesn't have an interrupt set).
- /proc/ioports: People think this shows the hardware configuration when it only shows about the same data (possibly erroneous) as setserial.

4.4 I/O Address & IRQ: Boot-time messages

In many cases your ports will automatically get low-level configured at boot-time (but not always correctly). To see what is happening, look at the start-up messages on the screen. Don't neglect to check the messages from the BIOS before Linux is loaded (no examples shown here). Use shift-PageUp to see the messages after they have flash by. The `dmesg` command may be used at any time to view some of the messages but it often misses important ones. Here's an example of the start-up messages (as of mid 1999). Note that ttyS00 is the same as /dev/ttyS0.

```
At first you see what was detected (but the irq is only a wild guess):

Serial driver version 4.27 with no serial options enabled
ttyS00 at 0x03f8 (irq = 4) is a 16550A
ttyS01 at 0x02f8 (irq = 3) is a 16550A
ttyS02 at 0x03e8 (irq = 4) is a 16550A

Later you see what was saved, but it's not necessarily correct either:

Loading the saved-state of the serial devices...
/dev/ttyS0 at 0x03f8 (irq = 4) is a 16550A
/dev/ttyS1 at 0x02f8 (irq = 3) is a 16550A
/dev/ttyS2 at 0x03e8 (irq = 5) is a 16550A
```

Note that there is a slight disagreement: The first message shows ttyS2 at irq=4 while the second shows it at irq=5. Your may only have the first message. In most cases the last message is the correct one. But if your having trouble it may be misleading. Before reading the explanation of all of this complexity in the rest of this section, you might just try using your serial port and see if it works OK. If so it may not be essential to read further.

The second message is from the `setserial` program being run at boot-time. It shows what the device driver thinks is the correct configuration. But this too could be wrong. For example, the irq could actually be set to irq=8 in the hardware (both messages wrong). The irq=5 is there because someone incorrectly put this into a configuration file (or the like). The fact that Linux sometimes gets IRQs wrong is because it doesn't probe for IRQs. It just assumes the "standard" ones (first message) or accepts what you told it when you configured it (second message). Neither of these is necessarily correct. If the serial driver has the wrong IRQ the serial port is very slow or doesn't seem to work at all.

The first message is a result of Linux probing the serial ports. If a port shows up here it exists but the irq may be wrong. Linux doesn't check IRQs because doing so is not foolproof. It just assumes the IRQs are as shown because they are the "standard" values. Your may check them manually with `setserial` using the `autoconfig` option but this isn't guaranteed to be correct.

The data shown by the BIOS messages (which you see at first) is what is set in the hardware. If your serial port is Plug-and-Play PnP then it's possible that the `isapnp` will run and change these settings. Look for messages about this after Linux starts. The last serial port message shown in the example above should

agree with the BIOS messages (as possibly modified by isapnp). If they don't agree then you either need to change the setting in the port hardware or use setserial to tell the driver what is actually set in the hardware.

Also, if you have Plug-and-Play (PnP) serial ports, Linux will not find them unless the IRQ and IO has been set inside the hardware by Plug-and-Play software. This is a common reason why the start-up messages do not show a serial port that physically exists. The PC hardware (a PnP BIOS) may automatically low-level configure this. PnP configuring will be explained later.

4.5 What is the current IO address and IRQ of my Serial Port ?

The previous section indicated how to attempt to do this by looking at the start-up messages. An alternate title for this section is "Other ways of finding out your IO Address and IRQ". If they give you sufficient info then you may not need to read this section. If they don't then there are some other ways to look into this.

There are really two answers to the question "What is my IO and IRQ?" 1. What the device driver thinks has been set (This is what setserial sets and shows). 2. What is actually set in the hardware. They both should be the same. If you're having trouble (including communication programs that can't communicate) it may mean that these two items are not set the same. In other words, this means that the driver has incorrect info on the physical serial port. If the driver has the wrong IO address it will try to send data to a non-existing serial port –or even worse, to an actual device that is not a serial port. If it has the wrong IRQ the driver will not get interrupt service requests from the serial port, resulting in the possible overflow of the serial port's buffer and in very slow response. See 15.11 (Slow: Text appears on the screen slowly after long delays). If it has the wrong model of UART there is also apt to be trouble. To determine if both I0-IRQ pairs are identical you must find out the answers to the questions posed in the next two subsections.

4.5.1 What does the device driver think?

This is easy to find out. Just look at the start-up messages or type "setserial -g /dev/ttyS*". If everything works OK then what it tells you is likely also set in the hardware. There are some other ways to find this info by looking at "files" in the /proc directory. An important reason for understanding these other ways is to warn you that they only show what the device driver thinks. Some people view certain "files" in the /proc directory and erroneously think that what they see is set in the hardware but "it ain't necessarily so".

/proc/ioports will show the IO addresses that the drivers are using. /proc/interrupts shows the IRQs that are used by drivers of currently running processes (that have devices open). Note that in both cases above you are only seeing what the driver thinks and not necessarily what is actually set in the hardware. /proc/interrupts also shows how many interrupts have been issued (often thousands) for each device. You can get a clue from this because if you see a large number of interrupts that have been issued it means that there is a piece of hardware somewhere that is using that interrupt. Sometimes a showing of just a few interrupts doesn't mean that that interrupt is actually being physically generated by any serial port. Thus if you see almost no interrupts for a port that you're trying to use, that interrupt might not be set in the hardware and it implies that the driver is using the wrong interrupt. To view /proc/interrupts to check on a program that you're currently running (such as "minicom") you need to keep the program running while you view it. To do this, try to jump to a shell without exiting the program.

4.5.2 What is set in my serial port hardware ?

How do you find out what IO address and IRQ are actually set in the device hardware? Perhaps the BIOS messages will tell you some info before Linux starts booting. Use the shift-PageUp key to step back thru the boot-time messages and look at the very first ones which are from the BIOS. This is how it was before Linux started. Setserial can't change it but isapnp or pciutils can.

One crude method is try probing with setserial using the "autoconfig" option. You'll need to guess the addresses to probe at. See 8.1 (What is Setserial). For a PCI serial port, look at /proc/pci or /proc/bus/pci/devices. If your serial port is is Plug-and-Play see the next two subsections.

For a port set with jumpers, its how the jumpers were set. If the port is not Plug-and-Play (PnP) but has been setup by using a DOS program then it's set at whatever the person who ran that program set it to.

4.5.3 What is set in my PnP serial port hardware ?

PnP ports don't store their configuration in the hardware when the power is turned off. This is in contrast to Jumpers (non-PnP) which remain the same with the power off. If you have an ISA PnP port, it can reach a state where it doesn't have any IO address or IRQ and is in effect disabled. It should still be possible to find the port using the `pnpdump` program.

For Plug-and-Play (PnP) on the ISA bus one may try the `pnpdump` program (part of `isapnptools`). If you use the –dumpregs option then it should tell you the actual IO address and IRQ set in the port. The address it "trys" is not the device's IO address, but a special

For PnP ports checking on how it's configured under DOS/Windows may not be of much help. Windows stores its configuration info in its Registry which is not used by Linux. It may supply the BIOS's non-volatile memory with some info but it may not be kept in sync with the current Window configuration in the Registry ?? If you let a PnP BIOS automatically do the configuring when you start Linux (and have told the BIOS that you don't have a PnP operating system when running Linux) then Linux should use whatever configuration is in the BIOS's non-volatile memory.

4.6 Choosing Serial IRQs

If you have a true Plug-and-Play set up where either the OS or a PnP BIOS configures all your devices, then you don't choose your IRQs. PnP determines what it thinks is best and assigns them. But if you use the tools in Linux for Plug-and-Play (isapnp and pcitools) then you have to choose. If you already know what IRQ you want to use you could skip this section except that you may want to know that IRQ 0 has a special use (see the following paragraph).

4.6.1 IRQ 0 is not an IRQ

While IRQ 0 is actually the timer (in hardware) it has a special meaning for setting a serial port with setserial. It tells the driver that there is no interrupt for the port and the driver then will use polling methods. This is quite inefficient but can be tried if there is an interrupt conflict or mis-set interrupt. The advantage of assigning this is that you don't need to know what interrupt is set in the hardware. It should be used only as a temporary expedient until you are able to find a real interrupt to use.

4.6.2 Interrupt sharing and Kernels 2.2+

The general rule is that every device should use a unique IRQ and not share them. But there are situations where sharing is permitted such as with most multi-port boards. Even when it is permitted, it may not be as efficient since every time a shared interrupt is given a check must be made to determine where it came from. Thus if it's feasible, it's nice to allocate every device it's own interrupt.

Prior to kernel 2.2, serial IRQs could be shared with each other only for most multiport boards. Starting with kernel 2.2 serial IRQs may be sometimes shared between all serial ports. In order for sharing to work in 2.2 the kernel must have been compiled with CONFIG_SERIAL_SHARE_IRQ, and the serial port hardware must support sharing (so that if two serial cards put different voltages on the same interrupt wire, only the

voltage that means "this is an interrupt" will prevail). Thus even if you have 2.2, it may be best to avoid sharing.

4.6.3 What IRQs to choose?

The serial hardware often has only a limited number of IRQs it can be set at. Also you don't want IRQ conflicts. So there may not be much of a choice. Your PC may normally come with `ttyS0` and `ttyS2` at IRQ 4, and `ttyS1` and `ttyS3` at IRQ 3. Looking at `/proc/interrupts` will show which IRQs are being used by programs currently running. You likely don't want to use one of these. Before IRQ 5 was used for sound cards, it was often used for a serial port.

Here is how Greg (original author of Serial-HOWTO) set his up in /etc/rc.d/rc.serial. rc.serial is a file (shell script) which runs at start-up (it may have a different name of location). For versions of "setserial" after 2.15 it's not always done this way anymore but this example does show the choice of IRQs.

```
/sbin/setserial /dev/ttyS0 irq 3        # my serial mouse
/sbin/setserial /dev/ttyS1 irq 4        # my Wyse dumb terminal
/sbin/setserial /dev/ttyS2 irq 5        # my Zoom modem
/sbin/setserial /dev/ttyS3 irq 9        # my USR modem
```

Standard IRQ assignments:

```
IRQ  0    Timer channel 0 (May mean "no interrupt".  See below.)
IRQ  1    Keyboard
IRQ  2    Cascade for controller 2
IRQ  3    Serial port 2
IRQ  4    Serial port 1
IRQ  5    Parallel port 2, Sound card
IRQ  6    Floppy diskette
IRQ  7    Parallel port 1
IRQ  8    Real-time clock
IRQ  9    Redirected to IRQ2
IRQ 10    not assigned
IRQ 11    not assigned
IRQ 12    not assigned
IRQ 13    Math coprocessor
IRQ 14    Hard disk controller 1
IRQ 15    Hard disk controller 2
```

There is really no Right Thing to do when choosing interrupts. Just make sure it isn't being used by the motherboard, or any other boards. 2, 3, 4, 5, 7, 10, 11, 12 or 15 are possible choices. Note that IRQ 2 is the same as IRQ 9. You can call it either 2 or 9, the serial driver is very understanding. If you have a very old serial board it may not be able to use IRQs 8 and above.

Make sure you don't use IRQs 1, 6, 8, 13 or 14! These are used by your motherboard. You will make her very unhappy by taking her IRQs. When you are done, double-check `/proc/interrupts` when programs that use interrupts are being run and make sure there are no conflicts.

4.7 Choosing Addresses –Video card conflict with ttyS3

The IO address of the IBM 8514 video board (and others like it) is allegedly 0x?2e8 where ? is 2, 4, 8, or 9. This may conflict (but shouldn't it the serial port is well designed) with the IO address of `ttyS3` at 0x02e8

if the serial port ignores the leading 0 hex digit (many do). That is bad news if you try to use `ttyS3` at this IO address.

In most cases you should use the default addresses if feasible. Addresses shown represent the first address of an 8-byte range. For example 3f8 is really the range 3f8-3ff. Each serial device (as well as other types of devices that use IO addresses) needs its own unique address range. There should be no overlaps (conflicts). Here are the default addresses for the serial ports:

```
ttyS0 address 0x3f8
ttyS1 address 0x2f8
ttyS2 address 0x3e8
ttyS3 address 0x2e8
```

4.8 Set IO Address & IRQ in the hardware (mostly for PnP)

After it's set in the hardware don't forget to insure that it also gets set in the driver by using `setserial`. For non-PnP serial ports they are either set in hardware by jumpers or by running a DOS program ("jumperless") to set them (it may disable PnP). The rest of this subsection is only for PnP serial ports. Here's a list of the possible methods of configuring PnP serial ports:

- Using a PnP BIOS CMOS setup menu (usually only for external modems on ttyS0 (Com1) and ttyS1 (Com2))
- Letting a PnP BIOS automatically configure a PnP serial port See 4.8.1 (Using a PnP BIOS to I0-IRQ Configure)
- Doing nothing if you have both a PnP serial port and a PnP Linux operating system (see Plug-and-Play-HOWTO).
- Using `isapnp` for a PnP serial port non-PCI)
- Using pciutils (pcitools) for the PCI bus

The IO address and IRQ must be set (by PnP) in their registers each time the system is powered on since PnP hardware doesn't remember how it was set when the power is shut off. A simple way to do this is to let a PnP BIOS know that you don't have a PnP OS and the BIOS will automatically do this each time you start. This might cause problems in Windows (which is a PnP OS) if you start Windows with the BIOS thinking that Windows is not a PnP OS. See Plug-and-Play-HOWTO.

Plug-and-Play was designed to automate this io-irq configuring, but for Linux at present, it has made life more complicated. The standard kernels for Linux don't support plug-and-play very well. If you use a patch to the Linux kernel to covert it to a plug-and-play operating system, then all of the above should be handled automatically by the OS. But when you want to use this to automate configuring devices other that the serial port, you may find that you'll still have to configure the drivers manually since many Linux drivers are not written to support a Linux PnP OS. If you use `isapnptools` or the BIOS for configuring plug-and-play this will only put the two values into the registers of the serial port section of the modem card and you will likely still need to set up setserial. None of this is easy or very well documented as of early 1999. See Plug-and-Play-HOWTO and the isapnptools FAQ.

4.8.1 Using a PnP BIOS to I0-IRQ Configure

While the explanation of how to use a PnP OS or isapnp for io-irq configuring should come with such software, this is not the case if you want to let a PnP BIOS do such configuring. Not all PnP BIOS can do

this. The BIOS usually has a CMOS menu for setting up the first two serial ports. There is often little to choose from. Unless otherwise indicated in menus, these first two ports get set at the standard IO addresses and IRQs. See 7.1 (Serial Port Device Names & Numbers)

Whether you like it or not, when you start up a PC a PnP BIOS starts to do PnP (io-irq) configuring of hardware devices. It may do the job partially and turn the rest over to a PnP OS (which you probably don't have) or if thinks you don't have a PnP OS it may fully configure all the PnP devices but not configure the device drivers. This is what you want but it's not always easy to figure out exactly what the PnP BIOS has done.

If you tell the BIOS that you don't have a PnP OS, then the PnP BIOS should do the configuring of all PnP serial ports –not just the first two. If you have MS Windows9x on the same PC, the BIOS might ?? have saved the io-irq configuration used for MS Windows in its non-volatile memory. If you're now reading the latest version of this HOWTO let me know if Windows puts it there. If so the BIOS will use this same configuration for Linux. In this case, if you can find out how MS Windows has set up io-irq then it should be the same under Linux.

If you add a new PnP device, the BIOS should change its PnP configuration to accommodate it. It could even change the io-irq of existing devices if required to avoid any conflicts. For this purpose, it keeps a list of non-PnP devices provided that you have told the BIOS how these non-PnP devices are io-irq configured. One way to tell the BIOS this is by running a program called ICU under DOS/Windows.

But how do you find out what the BIOS has done so that you set up the device drivers with this info? The BIOS itself may provide some info, either in it's setup menus of via messages on the screen when you turn on your computer. See 4.5.2 (What is set in my serial port hardware?)

4.9 Giving the IRQ and IO Address to Setserial

Once you've set the IRQ and IO address in the hardware (or arranged for it to be done by PnP) you also need to insure that the "setserial" command is run each time you start Linux. See the subsection 8.1.4 (Boot-time Configuration)

4.10 Other Configuring

4.10.1 Configuring Hardware Flow Control (RTS/CTS)

See 2.9 (Flow Control) for an explanation of it. You should always use hardware flow control if possible. Your communication program or "`getty`" should have an option for setting it (and if you're in luck it might be enabled by default). It needs to be set both inside your modem (by an init string or default) and in the device driver. Your communication program should set both of these (if you configure it right).

If none of the above will fully enable hardware flow control. Then you must do it yourself. For the modem, make sure that it's either done by the init string or is on by default. If you need to tell the device driver to do it is best done on startup by putting a file that runs at boot-time. See the subsection 8.1.4 (Boot-time Configuration) You need to add the following to such a file for each serial port (example is ttyS2) you want to enable hardware flow control on:

```
stty crtscts < /dev/ttyS2
```

If you want to see if flow control is enabled do the following: In minicom (or the like) type AT&V to see how the modem is configured and look for &K3 which means hardware flow control. Then see if the device driver knows about it by typing: stty -a < /dev/ttyS2 Look for "crtscts" (without a disabling minus sign).

5 Modem Configuration (excluding serial port)

5.1 AT Commands

While the serial port on which a modem resides requires configuring, so does the modem itself. The modem is configured by sending AT commands (or the like) to it on the same serial line that is used to send data.

Most modems use an AT command set. These are cryptic and short ASCII commands where all command strings are prefaced by the letters AT. For example: ATZ&K3 There are two commands here Z and &K3. Unfortunately there are many different variations of the AT command set so that what works for one modem may or may not work for another modem. Thus there is no guarantee that the AT commands given in this section will work on your modem. Another point is that to get the modem to act on the AT command string, a return character must be sent at the end of the string.

Such command strings are either automatically sent to the modem by communication programs or are sent directly by you. Most communication programs provide a screen where you can type commands directly to your modem. This is good for setting up the modem as you can have it remember how it was set even after its powered off.

If you have a manual for your modem you can likely look up the AT command set in it. Otherwise, you may try to find it on the Internet. One may use a search engine and include some actual commands in the search terms to avoid finding sites that just talk about such commands but fail to list them. You might also try a few of the sites listed in the subsection 17.5 (Web Sites)

5.2 Init Strings: Saving and Recalling

The examples given in this subsection are from the Hayes AT modem command set. All command strings must be prefaced by the two letters AT (for example: AT&C1&D3). When a modem is powered on, it automatically configures itself with one of the configurations it has stored in its non-volatile memory. If this configuration is satisfactory there is nothing further to do.

If it's not satisfactory, then one may either alter the stored configuration or configure the modem each time you use it by sending it a string of commands known as an "init string" (= initialization string). Normally a a communication program does this. What it sends will depend on how you configured the communications program or what script you wrote for it if you use Kermit. You can usually edit the init string your communication program uses and change it to whatever you want. Sometimes the communications program will let you select the model of your modem and then it will use an init string that it thinks is best for that modem.

The configuration of the modem uses when it's first powered on could be expressed by an init string. You might think of this as the default "string" (called a profile). If your communications program sends the modem another string (the init string), then this string will modify the default configuration. For example, if the init string only contains two commands, then only those two items will be changed. However, some commands will recall a stored profile from inside the modem so a single such command in the init string can thereby change everything in the configuration.

Modern modems have a few different stored profiles to choose from that are stored in the modem's non-volatile memory (it's still there when you turn it off). In my modem there are two factory profiles (0 and 1, neither of which you can change) and two user defined profiles (0 and 1) that the user may set and store. Your modem may have more. Which one of these user-defined profiles is used at power-up depends on another item stored in the profile. If the command &Y0 is given then in the future profile 0 will be used at power-on. If it's a 1 instead of a 0 then profile 1 will be used at power-on.

There are also commands to recall (use it now) any of the 4 stored profiles. One may put such a command

in an init string. Of course if it recalls the same profile as was automatically loaded at power-up, nothing is changed unless the active profile has been modified since power-up. Since it could have been modified It's a good idea to use some kind of an init string even if it does nothing more than recalling a stored profile.

Recalling a saved profile (use 1 instead of 0 for profile 1):
Z0 recalls user-defined profile 0 and resets (hangs up, etc.)
&F0 recalls factory profile 0

Once you have sent commands to the modem to configure it the way you want (including recalling a factory profile and modifying it a little) you may save this as a user-defined profile:
&W0 saves the current configuration to user-profile 0

Many people don't bother saving a good configuration in their modem, but instead, send the modem a longer init string each time the modem is used. Another method is to restore the factory default at the start of the init string and then modify it a little by adding a few other commands to the end of the init string. By doing it this way no one can cause problems by modifying the user-defined profile which is loaded at power-on.

You may pick an init string supplied by someone else that they think is right for your modem. Some communication programs have a library of init strings to select from. The most difficult method (and one which will teach you the most about modems) is to study the modem manual and write one yourself. You could save this configuration inside the modem so that you don't need an init string. A third alternative is to start with an init string someone else wrote, but modify it to suit your purposes.

Now if you look at init strings used by communication programs you may see symbols which are not valid modem commands. These symbols are commands to the communication program itself (such as meaning to pause briefly) and will not be sent to the modem.

5.3 Other Modem Commands

Future editions of Modem-HOWTO may contain more AT commands but the rest of this section is mostly what was in the old Serial-HOWTO. All strings must start with AT. Here's a few Hayes AT codes that should be in the string (if they are not set by using a factory default or by a saved configuration).

```
E1      command echo ON
Q0      result codes are reported
V1      verbose ON
S0=0    never answer (uugetty does this with the WAITFOR option)
```

Here's some more codes concerning modem control lines DCD and DSR:

```
&C1     DCD is on after connect only
&S0     DSR is always on
```

These affect what your modem does when calls start and end. What DTR does may also be set up but it's more complicated.

If your modem does not support a stored profile, you can set these through the `INIT` string in a config file (or the like). Some older modems come with DIP switches that affect register settings. Be sure these are set correctly, too.

Greg Hankins has a collection of modem setups for different types of modems. If you would like to send him your working configuration, please do so: `<mailto:gregh@cc.gatech.edu>` You can get these setups at *ftp://ftp.cc.gatech.edu/pub/people/gregh/modem-configs*.

Note: to get his USR Courier V.34 modem to reset correctly when DTR drops, Greg Hankins had to set `&D2` and `S13=1` (this sets bit 0 of register S13). This has been confirmed to work on USR Sportster V.34 modems as well.

Note: some Supra modems treat DCD differently than other modems. If you are using a Supra, try setting `&C0` and *not* `&C1`. You must also set `&D2` to handle DTR correctly.

6 Modems for a Linux PC

6.1 External vs. Internal

A modem for a PC may be either internal or external. The internal one is installed inside of your PC (you must remove screws, etc. to install it) and the external one just plugs into a serial port connector on a PC. Internal modems are less expensive, are less likely to to suffer data loss due to buffer overrun, usually use less electricity, and use up no space on your desk. External modems are much easier to install, require less configuration, and have lights which may give you a clue as to what is happening. External modems are easy to move to another computer. Most external modems have no switch to turn off the power supply when not in use and thus are likely to consume a little electricity even when turned off (unless you unplug the power supply from the wall). Each watt they draw costs you about $1/yr. Another possible disadvantage of an external is that you will be forced to use an existing serial port which may not support a speed of over 115,200 k (although as of late 1998 most new internal modems don't either –but some do). If a new internal modem had a 16650 UART it would put less load on the CPU (but almost none do as of late 1998).

Internal modems present a special problem for Linux, but will work just as well as external modems provided you avoid the high percentage of them that will work only for MS Windows, and also provided that you spend time (sometimes a lot of time) to configure them correctly. Some of the modems which will work only under MS Windows are, unfortunately, not labeled as such. If you buy a new one, make sure that you can return it for a refund if it will not work under Linux.

While most new modems are plug-and-play you have various ways to deal with them:

- Use the "isapnp" program
- Have a PnP BIOS do the configuring
- Patch the kernel to create a PnP Linux

Each of the above has shortcomings. Isapnp documentation is difficult to understand although reading the Plug-and-Play-HOWTO (at present incomplete) will aid in understanding it. If you want the PnP BIOS to do the configuring, all you need to do is to make sure that it knows you don't have a PnP operating system. But it may not do it correctly. To find out what it's done see 4.5.2 (What is set in my serial port hardware?). Patching the kernel can be complicated too and there might not be a patch available against recent kernels.

There are many Linux users that say that it's a lot simpler just to get an external modem and plug it in. But since new peripherals are mostly PnP today, you may eventually need to deal with it, so why delay it? Still, the most expedient (and expensive) solution is an external modem (if you have a free serial port).

6.2 External Modems

6.2.1 PnP External Modems

Many external modems are labeled "Plug and Play" (PnP) but they should all work fine as non-PnP modems. Since you usually plug the modem into a serial port which has its own IRQ number and IO address, the modem needs no PnP features to set these up. However, the serial port itself may need to be configured (IRQ number and IO address) unless the default configuration is OK.

How can an external modem be called PnP since it can't be configured by PnP? Well, it has a special PnP identification built into it that can be read (thru the serial port) by a PnP operating system. Such an operating system would then know that you have a modem on a certain port and would also know the model number. Then you might not need to configure application programs by telling them what port the modem is on (such as /dev/ttyS2 or COM3). But if you don't have such a PnP operating system you will need to configure your application program manually by giving it the /dev id (such as /dev/ttyS2).

6.2.2 Cabling & Installation

Connecting an external modem is simple compared to connecting most other devices to a serial port that require various types of "null modem" cables. Modems use straight through cable, with no pins crossed over. Most computer stores should have these. Make sure you get the correct gender. If you are using the DB9 or DB25 serial port at your computer, it will always be male which means that the connector on the cable should be female. Hook up your modem to one of your serial ports. If you are willing to accept the default IRQ and IO address of the port you connect it to, then you are ready to start your communication program and configure the modem itself.

6.2.3 What the Lights (LED's) Mean

- TM Test Modem
- AA Auto Answer (If on, your modem will answer an incoming call)
- RD Receive Data line = RxD
- SD Send Data line = TxD
- TR data Terminal Ready = DTR (set by your PC)
- RI Ring Indicator (If on, someone is "ringing" your modem)
- OH Off Hook (If off, your modem has hung up the phone line)
- MR Modem Ready = DSR ??
- EC Error Correction
- DC Data Compression
- HS High Speed (for this modem)

6.3 Internal Modems

An internal modem is installed in a PC by taking off the cover of the PC and inserting the modem card into a vacant slot on the motherboard. There are modems for the ISA slots and others for the PCI slots. While external modems plug into the serial port (via a short cable) the internal modems have the serial port built into the modem. In other words, the modem card is both a serial port and a modem.

Setting the IO address and IRQ for a serial port was formerly done by jumpers on the card. These are little black rectangular "cubes" about 5x4x2 mm in size which push in over pins on the card. Plug-and-Play modems (actually the serial port part of the modems) don't use jumpers for setting these but instead are configured by sending configuration commands to them (via IO address space on the ISA bus inside the computer). Such configuration commands can be sent by a PnP BIOS, the isapnp program (for the ISA bus only) or by a PnP operating system. The configuring of them is built into Windows 95/98 OSs. Under Linux you have a choice of ways (none of which is always easy) to io-irq configure them:

1. Use "isapnp" which may be run automatically at every boot-time
2. Use a PnP BIOS alone (which runs at every boot-time)
3. Patch Linux to make it a PnP operating system

6.4 Internal Modems You Must Avoid (winmodems, etc.)

A majority of internal modems made after about mid-1998 *don't* work with Linux since they are "winmodems" or the like. Names used include: HSP, HCF, and soft-... modem. Such modems turn over much (or even almost all) of the work of the modem to the main processor (CPU) chip of your computer (such as a Pentium chip). Since only Windows (and not Linux) software is provided to do this, it will not work under Linux. A list of modems which do and don't work under Linux is at *Linux modem list* <http://www.o2.net/~gromitkc/winmodem.html>

A better term for "winmodem" might be "software modem" or "soft-modem". But since this software is only for MS Windows, the term "winmodem" is now a good one since it also implies it's "Windows-only". The term "Winmodem" is a trademark for a certain type of "winmodem". Here is some more precise terminology regarding "winmodems":

HSP (Host Signal Processor) means that the host processor (your CPU chip) creates the code needed to produce the electrical signal on the phone line. The modem itself just creates whatever electrical waveshape the CPU tells it to. In contrast to this, a "controllerless" modem can create the waveshapes on its own (but can't control the modem). It contains no facilities to deal with bytes being sent and received. It can't compress strings of bytes; it can't check for errors; it can't put them into packets. In other words it can't control the modem but instead has the CPU do all this work using a program for Windows. The Rockwell HCF (Host Controlled Family) does this. If the software that does all this could be ported to Linux and then there wouldn't be this problem. Besides the above, a modem which doesn't simulate a serial port will not work under Linux.

How do you determine if an internal modem will work under Linux? If you don't know the model of the modem and you also have Windows on your Linux PC, click on the "Modem" icon in the "Control Panel". First check out the modem list on the Web mentioned 3 paragraphs above. If that doesn't work (or isn't feasible), you can look at the package it came in (or a manual) find the section on the package that says something like "Minimum System Requirements" or just "System Requirements". It may be in fine print. Read it closely. If Windows is listed as one of the requirements then it will likely not work under Linux.

Otherwise, it may work under Linux if it fails to state explicitly that you must have Windows. By saying it's "designed for Windows" it may only mean that it fully supports Microsoft's plug-and-play which is OK since Linux uses the same plug-and-play specs (but doesn't support them very well as yet). Being "designed for Windows" thus gives no clue as to whether or not it will work under Linux. You might check the Website of the manufacturer or inquire via email. I once saw a web-page that specifically stated that one model worked under Linux while implying that another model didn't.

As far as the author knows, there is no effort currently underway to support winmodems in Linux. You might request that modem manufacturers port their code to Linux (or the like).

6.5 Which Internal Modems are Best to Avoid ?

- 6.4 (Avoid: winmodems) or the like. They will NOT work at all
- 6.5.3 (PCI Modems) seldom work under Linux
- 6.5.1 (MWave and DSP Modems) might work, but only if you first start Windows/Dos each time you power on your PC

- Modems with 6.5.2 (RPI (Rockwell)) drivers work but with reduced performance

6.5.1 MWave and DSP Modems

Such modems use DSP's (Digital Signal Processors) which are programmed by algorithms which must be downloaded from the hard disk to the DSP's memory just before using the modem. Unfortunately, the downloading is done by Dos/Windows programs so one can't do it from Linux. Ordinary modems that work with Linux often have a DSP too (and may mention this on the packaging), but the program that runs it is stored inside the modem. This is not a "DSP modem" in the sense of this section and should work OK under Linux. An example of a DSP modem is IBM's Aptiva MWAVE.

If a DSP modem modem simulates a serial port, then it is usable with Linux which communicates with modems via the serial port. If you also have Dos/Windows on the same PC you may be able to use the modem: You first install the driver under DOS (using DOS and not Window drivers). Then start Dos/Windows (make sure the modem gets initialized) and without turning off the computer, go into Linux. One way to get to Linux is to use loadlin.exe which is a DOS program that will boot Linux from DOS (See Config-HOWTO). Another way is to just press CTRL-ALT-DEL. The modem remains on the same com port (same IO address) that it used under DOS.

6.5.2 Rockwell (RPI) Drivers

Modems that require Rockwell RPI drivers are not fully usable since the driver software doesn't work on Linux. The RPI does compression and error correction using MS Windows software on your computer's CPU. If you are willing to operate the modem without using the RPI (and have no compression nor error correction) then you may easily disable RPI by sending the modem (via the initialization string) a command to do so each time you power on your modem. On my modem this command is +H0. Not having data compression available may not be much of a handicap since most long files which you download from the Internet are already compressed and attempts at further compression may only slow things down a bit.

6.5.3 PCI Modems

A PCI modem card is one which inserts into a PCI-bus slot on the motherboard of a PC. Unfortunately, it seems that most PCI modems will not work under Linux but efforts are underway to support some of them. See 4.1 (PCI Bus Not Yet Well Supported)

7 Serial Port Devices /dev/ttyS2, etc.

For creating devices in the device directory see the Serial-HOWTO: "Creating Devices In the /dev directory".

7.1 Serial Port Device Names & Numbers

Devices in Linux have major and minor numbers. Each serial port may have 2 possible names in the /dev directory: ttyS and cua. Their drivers behave slightly differently. The cua device is deprecated and will not be used in the future. See 7.3 (The cua Device).

Dos/Windows use the COM name while the `setserial` program uses tty00, tty01, etc. Don't confuse these with dev/tty0, dev/tty1, etc. which are used for the console (your PC monitor) but are not serial ports. The dos names (COM1, etc.) and IO address is shown below for the "standard" case (but yours could be different).

```
       set-                                                  IO
dos    serial          major minor         major minor  address
COM1   ttyO0  /dev/ttyS0  4,  64;  /dev/cua0  5,  64      3F8
COM2   ttyO1  /dev/ttyS1  4,  65;  /dev/cua1  5,  65      2F8
COM3   ttyO2  /dev/ttyS2  4,  66;  /dev/cua2  5,  66      3E8
COM4   ttyO3  /dev/ttyS3  4,  67;  /dev/cua3  5,  67      2E8
```

Note that all distributions should come with ttyS devices (and cua devices until cua is finally abolished). You can verify this by typing:

```
linux% ls -l /dev/cua*
linux% ls -l /dev/ttyS*
```

7.2 Link ttySN to /dev/modem ?

On some installations, two extra devices will be created, `/dev/modem` for your modem and `/dev/mouse` for your mouse. Both of these are symbolic links to the appropriate device in `/dev` which you specified during the installation (unless you have a bus mouse, then `/dev/mouse` will point to the bus mouse device).

There has been some discussion on the merits of `/dev/mouse` and `/dev/modem`. The use of these links is discouraged. In particular, if you are planning on using your modem for dialin you may run into problems because the lock files may not work correctly if you use `/dev/modem`. Use them if you like, but *be sure they point to the right device*. However, if you change or remove this link, some applications might need reconfiguration.

7.3 The cua Device

Each ttyS device has a corresponding cua device. But the cua device is to be abolish so it's best to use ttyS (unless cua is required). There is a difference between cua and ttyS but a savvy programmer can make a ttyS port behave just like a cua port so there is no real need for the cua anymore. Except some older programs may need to use the cua.

What's the difference? The main difference between cua and ttyS has to do with what happens in a C-program when an ordinary "open" command tries to open the port. If a cua port has been set to check modem control signals, the port can be opened even if the DCD modem control signal says not to. Astute programming (by adding additional lines to the program) can force a ttyS port to behave this way also. But a cua port can be more easily programmed to open for dialing out on a modem even when the modem fails to assert DCD (since no one has called into it and there's no carrier). That's why cua was once used for dial-out and ttyS used for dial-in.

Starting with Linux kernel 2.2, a warning message will be put in the kernel log when one uses cua. This is an omen that cua is on the way out.

8 Interesting Programs You Should Know About

8.1 What is setserial ?

8.1.1 Introduction

Don't ever use `setserial` with Laptops (PCMCIA). `setserial` is a program which allows you to tell the device driver software the I/O address of the serial port, which interrupt (IRQ) is set in the port's hardware,

what type of UART you have, etc. It can also show how the driver is currently set. In addition, it can probe the hardware (if certain options are given). Besides the manual for `setserial`, check out info in `/usr/doc/setserial.../` or the like. It should tell you how setserial is handled in your distribution of Linux.

`Setserial` is often run automatically at boot-time by a start-up shell-script. It will only work if the serial module is loaded. If you should for some reason unload the serial module later on, the changes previously made by `setserial` will be forgotten. So `setserial` must be run again to reestablish them. In addition to running via a start-up script, something akin to `setserial` runs when the serial module is loaded. Thus when you watch the start-up messages on the screen it may look like it ran twice, and in fact it has.

With appropriate options, `setserial` can probe (at a given I/O address) for a serial port but you must guess the I/O address. If you ask it to probe for /dev/ttyS2 for example, it will only probe at the address it thinks ttyS2 is at. If you tell setserial that ttyS2 is at a different address, then it will probe at that address, etc. See 8.1.2 (Probing)

Setserial does not set either IRQ's nor I/O addresses in the serial port hardware itself. You must tell setserial the identical values that have been set in the hardware. It's set in the hardware either by jumpers or by plug-and-play. Do not just invent some values that you think would be nice to use and tell them to setserial. However, if you know the I/O address but don't know the IRQ you may command setserial to attempt to determine the IRQ.

You can see a list of possible commands to use (but not the one-letter options such as -v for verbose –which you should normally use when troubleshooting) by just typing `setserial` with no arguments. Note that setserial calls an I/O address a "port". If you type:

```
setserial-g /dev/ttyS*
```

you'll see some info about how that device driver is configured for your ports. Add a "v" to the option "-g" to see more. But this doesn't tell you if the hardware actually has these values set in it. If fact, you can run setserial and assign a purely fictitious I/O address, any IRQ, and whatever uart type you would like to have. Then the next time you type "setserial ..." it will display these bogus values without complaint. Note that assignments made by setserial are lost when the PC is powered down so it is usually run automatically somewhere each time that Linux is booted.

8.1.2 Probing

In order to try to find out if you have a certain piece of serial hardware you must first know (or guess) its I/O address (or the device driver must have an I/O address for it, likely previously set by setserial). To try to detect the physical hardware use the -v (verbose) and `autoconfig` command to `setserial`. If the resulting message shows a uart type such as 16550A, then you're OK. If instead it shows "`unknown`" for the uart type, then there is likely no serial port at all at that I/O address. Some cheap serial ports don't identify themselves correctly so if you see "`unknown`" you still might have something there.

Besides auto-probing for uart type, setserial can auto-probe for IRQ's but this doesn't always work right either. In versions of setserial >= 2.15, your last probe test may be saved and put into the configuration file `/etc/serial.conf` which will be used next time you start Linux. The script that runs `setserial` at boot-time does not usually probe, but you could change it so that it does. See the next section.

8.1.3 Can Linux Configure The Serial Devices Automagically?

Yes, but ... You may set Linux up to detect and set up the serial devices automatically on startup (probing). It's easy to do for setserial < 2.15. Just add some lines to the file that runs setserial on start-up. See 8.1.4 (Old configuration method: edit a script) For example, for ttyS3 you would add:

```
/sbin/setserial /dev/ttyS3 auto_irq skip_test autoconfig
```

to the file that runs setserial on startup. Do this for every serial port you want to auto configure. Be sure to give a device name that really does exist on your machine.

For versions >= 2.15 (provided your distribution implemented the change, Redhat didn't) it's much harder to do since the file that runs setserial on startup, /etc/init.d/setserial or the like was not intended to be edited by the user. There may be no helpful comments in it like there were in earlier versions.

8.1.4 Boot-time Configuration

When the kernel loads the serial module (or if the "module" is built into the kernel) then only `ttyS{0-3}` are auto-detected and the driver is set to IRQs 4 and 3 (regardless of what the hardware is actually set at). You see this as a boot-time message just like as if `setserial` had been run.. If you use 3 or more ports, this may result in IRQ conflicts.

To fix such conflicts by telling setserial the true IRQs (or for other reasons) there may be a file somewhere that runs `setserial` again. This happens early at boot-time before any process uses the serial port. In fact, your distribution may have set things up so that the setserial program runs automatically from a start-up script at boot-time. More info about how to handle this situation should be found in /usr/doc/setserial.../ or the like.

New configuration method using /etc/serial.conf Starting with version 2.15 (1999) of `setserial` the shell-script that runs from a file at boot-time gets its data from a configuration file: /etc/serial.conf. But serial.conf is normally never edited. Instead just use `setserial` on the command line. Normally, what you changed with the setserial command is saved to the configuration file (serial.conf) when you shutdown (normally) or reboot. This only works if "###AUTOSAVE### or the like is on the first line of serial.conf. If you should use `setserial` experimentally and it doesn't work out right, then don't forget to redo it so that the experimental settings don't get saved by mistake. The file most commonly used to run setserial at boot-time (in conformance with the configuration file) is now /etc/init.d/setserial (Debian) or /etc/init.d/serial (Redhat), or etc., but it also should not normally be edited.

To disable a port, use `setserial` to set it to "uart none". The format of /etc/serial.conf appears to be just like that of the parameters placed after "setserial" on the command line with one line for each port. If you don't use autosave, you may edit /etc/serial.conf manually. For 2.15, the Debian distribution installs the system with autosave enabled, but Redhat 6.0 just had a file /usr/doc/setserial-2.15/rc.serial which you have to move to /etc/init.d/.

BUG: As of July 1999 there is a bug/problem since with ###AUTOSAVE### only the setserial parameters displayed by "setserial -G /dev/ttyS?" (where ? = 0, 1, 2, ...) get saved but the other parameters don't get saved. This will only affect a small minority of users since the parameters not saved are seldom used anyway. It's been reported as a bug and may be fixed by now.

In order to force the current settings set by setserial to be saved to the configuration file (serial.conf) without shutting down, do what normally happens when you shutdown: Run the shell-script `/etc/init.d/{set}serial stop`. The "stop" command will save the current configuration but the serial ports still keep working OK.

Old configuration method: edit a script Prior to 2.15 (1999), or if your distribution chose not to use /etc/serial.conf then you need to look for a file that runs setserial at boot-time. If it doesn't exist, you need to create one (or place the commands in a file that runs early at boot-time). If such a file is currently being used it's likely somewhere in the /etc directory-tree. But Redhat has supplied it in /usr/doc/setserial/

although you need to move it to the /etc tree before using it. You might use "locate" to try to find such a file. For example, you could type: locate "*serial*".

What you are looking for could be named rc.serial, or 0setserial (Debian). If such a file is supplied, it should contain a number of commented-out examples. By uncommenting some of these and/or modifying them, you should be able to set things up correctly. Make sure that you are using a valid path for `setserial`, and a valid device name. You could do a test by executing this file manually (just type its name as the super-user) to see if it works right. Testing like this is a lot faster than doing repeated reboots to get it right. Of course you can also test a single `setserial` command by just typing it on the command line.

The script `/etc/rc.d/rc.serial` was commonly used in the past. The Debian distribution used `/etc/rc.boot/0setserial`. Another file once used is `/etc/rc.d/rc.local` but it's not a good idea to use this since it may not be run early enough. It's been reported that other processes may try to open the serial port before rc.local runs resulting in serial communication failure.

8.1.5 IRQs

By default, both ttyS0 and ttyS2 share IRQ 4, while ttyS0 and ttyS3 share IRQ 3. But sharing serial interrupts is not permitted unless you: 1. have kernel 2.2 or better, and 2. you've complied in support for this, and 3. your serial hardware supports it. See

4.6.2 (Interrupt sharing and Kernels 2.2+) If you only have two serial ports, ttyS0 and ttyS1, you're still OK since IRQ sharing conflicts don't exist for non-existent devices.

If you add an internal modem and retain ttyS0 and ttyS1, then you should attempt to find an unused IRQ and set it both on your serial port (or modem card) and then use setserial to assign it to your device driver. If IRQ 5 is not being used for a sound card, this may be one you can use for a modem. To set the IRQ in hardware you may need to use isapnp, a PnP BIOS, or patch Linux to make it PnP. To help you determine which spare IRQ's you might have, type "man setserial" and search for say: "IRQ 11".

————————From TT-HOWTO————————————-

Setserial tells the serial port driver both the interrupt (IRQ) number and the port IO address. It can tell you what type of UART chips you have if you use the autoconfig parameter. For the serial port to work, the serial module (for the Linux kernel) must be loaded. If you use more than 2 serial ports and want to assign unique IRQ numbers (or use more than 4 serial ports), then you need to use setserial. If your serial port is plug-and-play you should look in the Serial-HOWTO. In addition to the man page for it, but sure to look at /usr/doc/setserial/ if it exists. It may tell you how setserial is run automatically at boot time and how to alter the configuration. Don't ever use setserial with Laptops (PCMCIA bus).

starting with version 2.15 (1999) of `setserial` changes made to setserial on the command line are sometimes automatically saved so that they will remain after you shut off the computer.

8.1.6 Where to Run Setserial ?

Starting with version 2.15 of `setserial` the method of configuration changed. Before 2.15 you configured it by editing a shell-script that ran `setserial`. After 2.15 there's a configuration file /etc/serial.conf which is automatically edited to conform with whatever you do with the `setserial` command.

You may find you have both the old and new way of configuring installed. If the shell-script is named "...pre-2.15" or the like, it's probably not being used anymore so everything should be OK if you keep it.

Prior to version 2.15 Prior to 2.15 there was no configuration file for `setserial`. One needs to edit the shell-script that runs setserial at boot-time. You need to find it or create one if it doesn't exist. Where

it is depends on your distribution. It might be in a file named "rc.serial" . In the Debian distribution it was the 0setserial file in /etc/rc.boot. One way to try to find it is to search for it with the "locate" command.

Version 2.15 and later You neither edit the file that runs setserial nor the configuration file. Changes made by using `setserial` on the command line are *sometimes* saved in /etc/serial.conf during normal shutdown. This feature is called autosave. It only works if the first line of serial.conf contains ###AUTOSAVE### or the like.

There's a bug I've reported in 2.15 and it's been fixed (but you could still have it): With the bug, seldom used changes made by setserial don't get saved. Only the most important parameters that you see if you give the command "setserial -g /dev/ttyS?" do get saved.

If autosave is enabled, then each time you boot, setserial is run by /etc/init.d/setserial (or "serial", etc.) using the saved configuration.

8.2 What is isapnp ?

`isapnp` is a program to configure Plug-and-Play (PnP) devices on the ISA bus including internal modems. It comes in a package called "isapnptools" and includes another program, "pnpdump" which finds all your ISA PnP devices and shows you options for configuring them in a format which may be added to the PnP configuration file: /etc/isapnp.conf. It may also be used with the –dumpregs option to show the current IO address and IRQ of the modem's serial port. The isapnp command may be put into a startup file so that it runs each time you start the computer and thus will configure ISA PnP devices. It is able to do this even if your BIOS doesn't support PnP. See Plug-and-Play-HOWTO.

8.3 What is wvdialconf ?

`wvdialconf` will try to find which serial port has a modem on it. It also creates a configuration program for the wvdial program. `wvdial` is used for simplified dialing out using the PPP protocol to an ISP. But you don't need to install PPP in order to use `wvdialconf`. It will only find modems which are not in use. It will also automatically devise a "suitable" init strings but sometimes gets it wrong. Since this command has no options, it's simple to use but you must give it the name of a file to put the init string (and other data) into. For example type: wvdialconf my_file_name.

8.4 What is stty ?

`stty` is like setserial but it sets the baud rate and other parameters of a serial port. Typing "stty -a < /dev/ttyS2" should show you how ttyS2 is configured. Most of the settings are for things that you never need to use with modems (such as some used only for old terminals of the 1970s). Your communication package should automatically set up all the setting correctly for modems. But stty is sometimes useful for trouble-shooting.

Two items set by stty are: 1. Hardware flow control by "crtscts" and 2. Ignore the DCD signal from the modem: "clocal". If the modem is not sending a DCD signal and clocal is disabled (stty shows -clocal) then a program may not be able to open the serial port. If the port can't open, the program may just hang, waiting (often in vain) for a DCD signal from the modem.

Minicom sets clocal automatically when it starts up so there is no problem. But version 6.0.192 of Kermit hung when I set -clocal and tried to "set line ..." If -clocal is set and there is no DCD signal then even the "stty" command will hang and there is seemingly no way to set clocal (except by running minicom). But minicom will restore -clocal when it exits. One way to get out of this is to use minicom to send the "AT&C"

to the modem (to get the DCD signal) and then exit minicom with no reset so that the DCD signal remains on. Then you may use stty again.

9 Trying Out Your Modem (Dialing Out) .

9.1 Are You Ready to Dial Out ?

Once you've plugged in your modem and know which serial port it's on you're ready to try using it. Before you try to get the Internet on it or have people call in to you, first try something simpler like dialing out to some number to see if your modem is working OK. Find a phone number that is connected to a modem. It you don't know what number to call, ask at computer stores for such phone numbers of bulletin boards, etc. or see if a local library has a phone number for their on-line catalog.

Then make sure you are ready to phone. Do you know what serial port (such as ttyS2) your modem is on? You should have found this out when you io-irq configured your serial ports. Have you decided what speed you are going to use for this port? See 12.4 (Speed Table) for a quick selection or 12 (What Speed Should I Use) for more details. If you have no idea what speed to set, just set it a few times faster than the advertised speed of your modem. Also remember that if you see a menu where an option is "hardware flow control" and/or "RTS/CTS" or the like, select it. Is a live telephone cable plugged in to your modem? You may want to connect the cable to a real telephone to make sure that it can produce a dial tone.

Now you need to select a communication (dialing) program to use to dial out. Dialing programs include: minicom, seyon (X-windows), and kermit. See section 13 (Communications Programs) about some communications programs. Two examples are presented next: 9.2 (Dialing Out with Minicom) and 9.3 (Dialing Out with Kermit)

9.2 Dialing Out with Minicom

Minicom comes with most Linux distributions. To configure it you should be the root user. Type "minicom -s" to configure. This will take you directly to the configuration (set-up) menus. Alternatively you could just run "minicom" and then type ^A to see the bottom status line. This shows to type ^A Z for help (you've already typed the ^A so just type z). From the help menu go to the Configuration menu.

Most of the options don't need to be set for just simply dialing out. To configure you have to supply a few basic items: the name of the serial port your modem is on such as /dev/ttyS2 and the speed such as 115200. These are set at the serial port menu. Go to it and set them. Also (if possible) set hardware flow control (RTS/CTS). Then save them. When typing in the speed, you should also see something like "8N1" which you should leave alone. It means: 8-bit bytes, No parity, 1 stop-bit appended to each byte. If you can't find the speed you want, a lower speed will always work for a test. Exit (hit return) when done and save the configuration as default (dfl) using the menu. You may want to exit minicom and start it again so it can now find the serial port and initialize the modem, or you could go to help and tell minicom to initialize the modem.

Now you are ready to dial. But first at the main screen you get after you first type "minicom" make sure there's a modem there by typing AT and then hit the <enter> key. It should display OK. If it doesn't something is wrong and there is no point of trying to dial.

If you got the "OK" go back to help and select the dialing directory. You may edit it and type in a phone number, etc. into the directory and then select "dial" to dial it. Alternatively, you may just dial manually (by selecting "manual" and then type the number at the keyboard). If it doesn't work, carefully note any error messages and try to figure out what went wrong.

9.3 Dialing Out with Kermit

You can find the latest version of `kermit` at *http://www.columbia.edu/kermit/*. For example, say your modem was on `ttyS3`, and it's speed was 115200 bps. You would do the following:

```
linux# kermit
C-Kermit 6.0.192, 6 Sep 96, for Linux
 Copyright (C) 1985, 1996,
  Trustees of Columbia University in the City of New York.
Default file-transfer mode is BINARY
Type ? or HELP for help.
C-Kermit>set line /dev/ttyS3
C-Kermit>set carrier-watch off
C-Kermit>set speed 115200
/dev/ttyS3, 115200 bps
C-Kermit>c
Connecting to /dev/ttyS3, speed 115200.
The escape character is Ctrl-\ (ASCII 28, FS)
Type the escape character followed by C to get back,
or followed by ? to see other options.
ATE1Q0V1                          ; you type this and then the Enter key
OK                                ; modem should respond with this
```

If your modem responds to `AT` commands, you can assume your modem is working correctly on the Linux side. Now try calling another modem by typing:

```
ATDT7654321
```

where 7654321 is a phone number. Use ATDP instead of ATDT if you have a pulse line. If the call goes through, your modem is working.

To get back to the `kermit` prompt, hold down the Ctrl key, press the backslash key, then let go of the Ctrl key, then press the C key:

```
Ctrl-\-C
(Back at linux)
C-Kermit>quit
linux#
```

This was just a test using the primitive "by-hand" dialing method. The normal method is to let `kermit` do the dialing for you with its built-in modem database and automatic dialing features, for example using a US Robotics (USR) modem:

```
linux# kermit
C-Kermit 6.0.192, 6 Sep 1997, for Linux
 Copyright (C) 1985, 1996,
  Trustees of Columbia University in the City of New York.
Default file-transfer mode is BINARY
Type ? or HELP for help
C-Kermit>set modem type usr        ; Select modem type
C-Kermit>set line /dev/ttyS3       ; Select communication device
C-Kermit>set speed 115200          ; Set the dialing speed
C-Kermit>dial 7654321              ; Dial
 Number: 7654321
 Device=/dev/ttyS3, modem=usr, speed=115200
```

```
 Call completed.<BEEP>
Connecting to /dev/ttyS3, speed 115200
The escape character is Ctrl-\ (ASCII 28, FS).
Type the escape character followed by C to get back,
or followed by ? to see other options.

Welcome to ...

login:
```

10 Dial-In

10.1 Overview

Dial-in is where you set up your PC so that others may dial in to your phone number and use your PC. The "point of view" is your PC. When you dial out from your PC you are also dialing in to another computer (but not dialing in to your own computer)

Dialin works like this. Someone with a modem dials your telephone number. Your modem answers the call and connects. Once the caller is connected, your PC sends a login prompt to the caller. Then the caller logs in and uses your PC. Using your PC may mean that the caller has a shell account and can use your PC just as if they logged in at the console (or text-terminal). It could also mean that they get connected to the Internet. The program that you use at your PC to handle dialin is called `getty`.

Often, after login, another program is run, including programs to connect the caller to the Internet. This HOWTO doesn't cover such programs that run after your login. For connecting to the internet, the PPP program is usually run (see PPP-HOWTO, a new revision is expected soon). Login may be automated so that the communication program automatically types in a "name" and password each time.

10.2 Getty

`getty` is the program you run for dialin. You don't need it for dialout. In addition to presenting a login prompt, it also answers the telephone. Originally getty was used for logging in to a computer from a dumb terminal. It's currently used for logging in to a Linux console). There are a few different getty programs with slightly different names. Only certain ones work with modems for dialin. This `getty` program is usually started at boot-time. It must be called from the /etc/inittab file. You may find an example in this file of a call to getty which you will likely need to edit a bit. If you use a different getty program than the one shown in such an example, then you will need to edit it quite a bit since the options will have a different format.

There are four different getty programs to choose from that may be used with modems for dialin: `mgetty`, `uugetty`, `getty_em`, and `agetty`. A few details are given in the following subsections. `agetty` is the simplest (and weakest) of the four and some consider it mainly for use with directly connected text-terminals. `mgetty` has support for fax and voice mail but `Uugetty` doesn't. `mgetty` allegedly lacks a few of the features of `uugetty`. `getty_em` is a simplified version of `uugetty`. Thus `mgetty` is likely your best choice unless you are already familiar with `uugetty` (or find it difficult to get `mgetty`). The syntax for these getty programs differs, so be sure to check that you are using the correct syntax in `/etc/inittab` for whichever getty you use.

10.2.1 About mgetty

`mgetty` was written as a replacement for `uugetty` which was in existence long before `mgetty`. Both are for

use with modems. Although `mgetty` may be also used for directly connected terminals the documentation for this is hard to pinpoint and mgetty will not (as of mid 1999) support software flow control (used on many terminals) without recompiling. This defect is listed as a bug. In addition to allowing dialup logins, `mgetty` also provides FAX support and auto PPP detection. There is a supplemental program called `vgetty` which handles voicemail for some modems. `mgetty` documentation is good (except for voice mail), and does not need supplementing . Please refer to it for installation instructions. You can find the latest information on `mgetty` at *http://www.leo.org/~doering/mgetty/* and `<http://alpha.greenie.net/mgetty>`

10.2.2 About uugetty

`getty_ps` contains two programs: `getty` is used for console and terminal devices, and `uugetty` for modems. Greg Hankins (former author of Serial-HOWTO) used `uugetty` so his writings about it are included here. See 11 (Uugetty). The other gettys are well covered by the documentation that comes with them.

10.2.3 About getty_em

This is a simplified version of "uugetty". It was written by Vern Hoxie after he became fully confused with complex support files needed for getty_ps and uugetty.

It is part of the collection of serial port utilities and information by Vern Hoxie available via ftp from `<scicom.alphacdc.com>`. The name of the collection is "serial_suite.tgz". When logging into "scicom" as "anonymous", you must use your full e-mail address as the password. For example: greg.hankins@cc.gatech.edu

10.2.4 About agetty and mingetty

`agetty` is a simple, completely functional implementation of `getty` which is best suited for virtual consoles or terminals rather than modems. But it works fine with modems under favorable conditions (except you cannot dial out when agetty is running and waiting for a call). `agetty` in the Debian distribution is just named `getty`.

`mingetty` is a small getty that will work only for consoles (monitors) so you can't use it with modems for dialin.

10.3 What Happens when Someone Dials In ?

The caller runs some sort of communication program that dials your telephone number and your telephone rings. There are two different ways that your PC can answer the phone. One way is for the modem to automatically answer the call. The other way is for getty to sense the ringing and send a command to the modem to answer the call. Once the call is answered, your modem sends tones to the other modem (and conversely). The two modems negotiate how they will communicate and when this is done your modem sends a "CONNECTed" message (or the like) to `getty`. When `getty` gets this message, it sends a login prompt out the serial port. Sometimes `getty` just calls on a program named `login` to handle the logging in. `getty` usually starts running at boot-time but it must wait until a connection is made before sending out a "login" prompt.

Now for more details on the two methods of answering the call. By setting the S0 register of the modem to 3, the modem will automatically answer on the 3rd ring. If it's set to 0 then the modem will only answer the call if getty sends it an "A" (= Answer) command while the phone is ringing. Actually an "ATA" is sent since all modem commands are prefixed by "AT". You might think it best to utilize the ability of the

modem to automatically answer the call, but it's actually better if `getty` answers it. If the modem doesn't automatically answer, it's called manual answer (even though `getty` automatically handles it).

For the "manual" answer case, `getty` opens the port at boot-time and listens. When the phone rings, a "RING" message is sent to the listening `getty`. Then if `getty` wants to answer this ring, it sends the modem an "ATA" command. The modem then makes a connection and sends a "CONNECT ..." message to `getty` which then sends a login prompt to the caller.

The automatic answer case uses the CD (Carrier Detect) wire from the modem to the serial port to detect when a connection is made. It works like this. At boot-time `getty` tries to open the serial port but the attempt fails since there is normally no CD signal from the modem. Then the `getty` program waits at the open statement in the program until a CD signal appears. When a CD signal arrives (perhaps hours later) then the port is opened and `getty` sends the login prompt. While `getty` is waiting (sleeping) at the open statement, other processes can run since Linux is a multiprocessing operating system. What actually wakes `getty` up is an interrupt which is issued when the CD line from the modem changes state to on.

You may wonder how getty is able to open the serial port in the manual-answer case since there is no CD signal. Well, there's a way to write a program to force the port to open even if there is no CD signal present.

10.4 Why Manual Answer is Best

The difference between the two ways of answering will show itself when the computer happens to be down but the modem is still working. For the manual case, the "RING" message is sent to getty but since the computer is down, getty isn't there and the phone never gets answered. There are no telephone charges when there is no answer. For the automatic answer case, the phone is answered but no login message is ever sent since the computer is down. The phone bill runs up as the waiting continues. If the phone call is toll-free, it doesn't make much difference, although it may be frustrating waiting for a login prompt that never arrives. `mgetty` uses manual answer. `Uugetty` can do this too using a configuration script.

10.5 Callback

Callback is where someone first dials in to your modem. Then, you get a little info from the caller and then call it right back. Why would you want to do this? One reason is to save on telephone bills if you can call the caller cheaper than the caller can call you. Another is to make sure that the caller really is who it claims to be. If a caller calls you and claims to be calling from its usual phone number, then one way to verify this is to actually place a new call to that number.

There's a program for Linux called "callback" that works with mgetty. It's at `<ftp://ftp.icce.rug.nl/pub/unix/>`. Step-by-step instructions on how someone installed it (and PPP) is at `<http://www.stokely.com/unix.serial.port.resources/callback.html>`

10.6 Voice Mail

Voice mail is like an answering machine run by a computer. To do this you must have a modem that supports "voice" and supporting software. Instead of storing the messages on tape, they are stored in digital format on a disk. When a person phones you, they hear a "greeting" message and can then leave a message for you. More advanced systems would have caller-selectable mail boxes and caller-selectable messages to listen to. Free software is available in Linux for simple answering, but doesn't seem to be available yet for the more advanced stuff.

I know of two different voicemail packages for Linux. One is a very minimal package (see 13.2.4 (Voicemail Software)). The other, more advanced, but currently poorly documented, is `vgetty`. It's an optional addition

to the well documented and widely distributed `mgetty` program. It supports ZyXEL-like voice modem commands. In the Debian distribution, you must get the mgetty-voice package in addition to the mgetty package and mgetty-doc package. Obsolete documentation has been removed from mgetty but replacement documentation is lacking (except if you use the -h (help) option when running certain programs, etc.). But one sees postings about using it on the mgetty newsgroup. See 10.2.1 (About mgetty) and >. It seems that `vgetty` is currently not very stable but it's successfully being used and development of it continues. If this is the latest version of this HOWTO can someone who is familiar with vgetty please let me know its current status.

11 Uugetty for Dial-In (from the old Serial-HOWTO)

Be aware that you could use mgetty as a (better?) alternative to `uugetty`. `mgetty` is newer and more popular than uugetty. See 10.2 (What is getty?) for a brief comparison of these 2 gettys.

11.1 Installing getty_ps

Since uugetty is part of getty_ps you'll first have to install getty_ps. If you don't have it, get the latest version from *metalab.unc.edu:/pub/Linux/system/serial*. In particular, if you want to use high speeds (57600 and 115200 bps), you must get version 2.0.7j or later. You must also have libc 5.x or greater.

By default, `getty_ps` will be configured to be Linux FSSTND (File System Standard) compliant, which means that the binaries will be in `/sbin`, and the config files will be named `/etc/conf.{uu}getty.ttyS`*N*. This is not apparent from the documentation! It will also expect lock files to go in `/var/lock`. Make sure you have the `/var/lock` directory.

If you don't want FSSTND compliance, binaries will go in `/etc`, config files will go in `/etc/default/{uu}getty.ttyS`*N*, and lock files will go in `/usr/spool/uucp`. I recommend doing things this way if you are using UUCP, because UUCP will have problems if you move the lock files to where it isn't looking for them.

`getty_ps` can also use `syslogd` to log messages. See the man pages for `syslogd(1)` and `syslog.conf(5)` for setting up `syslogd`, if you don't have it running already. Messages are logged with priority LOG_AUTH, errors use LOG_ERR, and debugging uses LOG_DEBUG. If you don't want to use `syslogd` you can edit `tune.h` in the `getty_ps` source files to use a log file for messages instead, namely `/var/adm/getty.log` by default.

Decide on if you want FSSTND compliance and syslog capability. You can also choose a combination of the two. Edit the `Makefile`, `tune.h` and `config.h` to reflect your decisions. Then compile and install according to the instructions included with the package.

11.2 Setting up uugetty

With `uugetty` you may dial out with your modem while `uugetty` is watching the port for logins. The original author of this HOWTO said to use /dev/cuaN for dial-out (with uugetty running on /dev/ttyN). but others say you must use /dev/ttyN for both dial-in and dial-out ??

`uugetty` does important lock file checking. Update `/etc/gettydefs` to include an entry for your modem. When you are done editing `/etc/gettydefs`, you can verify that the syntax is correct by doing:

```
linux# getty -c /etc/gettydefs
```

11.2.1 Modern Modems

If you have a 9600 bps or faster modem with data compression, you can lock your serial port to one speed. For example:

```
# 115200 fixed speed
F115200# B115200 CS8 # B115200 SANE -ISTRIP HUPCL #@S @L @B login: #F115200
```

If you have your modem set up to do RTS/CTS hardware flow control, you can add `CRTSCTS` to the entries:

```
# 115200 fixed speed with hardware flow control
F115200# B115200 CS8 CRTSCTS # B115200 SANE -ISTRIP HUPCL CRTSCTS #@S @L @B login: #F115200
```

11.2.2 Old slow modems

If you have a slow modem (under 9600 bps) Then, instead of one line for a single speed, your need several lines to try a number of speeds. Note the these lines are linked to each other by the last "word" in the line such as #38400. Blank lines are needed between each entry.

```
# Modem entries
115200# B115200 CS8 # B115200 SANE -ISTRIP HUPCL #@S @L @B login: #57600

57600# B57600 CS8 # B57600 SANE -ISTRIP HUPCL #@S @L @B login: #38400

38400# B38400 CS8 # B38400 SANE -ISTRIP HUPCL #@S @L @B login: #19200

19200# B19200 CS8 # B19200 SANE -ISTRIP HUPCL #@S @L @B login: #9600

9600# B9600 CS8 # B9600 SANE -ISTRIP HUPCL #@S @L @B login: #2400

2400# B2400 CS8 # B2400 SANE -ISTRIP HUPCL #@S @L @B login: #115200
```

11.2.3 Login Banner

If you want, you can make `uugetty` print interesting things in the login banner. In Greg's examples, he has the system name, the serial line, and the current bps rate. You can add other things:

```
@B    The current (evaluated at the time the @B is seen) bps rate.
@D    The current date, in MM/DD/YY.
@L    The serial line to which uugetty is attached.
@S    The system name.
@T    The current time, in HH:MM:SS (24-hour).
@U    The number of currently signed-on users.  This is  a
      count of the number of entries in the /etc/utmp file
      that have a non-null ut_name field.
@V    The value of VERSION, as given in the defaults file.
To display a single '@' character, use either '\@' or '@@'.
```

11.3 Customizing uugetty

There are lots of parameters you can tweak for each port you have. These are implemented in separate config files for each port. The file `/etc/conf.uugetty` will be used by *all* instances of `uugetty`, and

/etc/conf.uugetty.ttyS*N* will only be used by that one port. Sample default config files can be found with the getty_ps source files, which come with most Linux distributions. Due to space concerns, they are not listed here. Note that if you are using older versions of uugetty (older than 2.0.7e), or aren't using FSSTND, then the default file will be /etc/default/uugetty.ttyS*N*. Greg's /etc/conf.uugetty.ttyS3 looked like this:

```
# sample uugetty configuration file for a Hayes compatible modem to allow
# incoming modem connections
#
# alternate lock file to check... if this lock file exists, then uugetty is
# restarted so that the modem is re-initialized
ALTLOCK=cua3
ALTLINE=cua3
# line to initialize
INITLINE=cua3
# timeout to disconnect if idle...
TIMEOUT=60
# modem initialization string...
# format: <expect> <send> ... (chat sequence)
INIT="" AT\r OK\r\n
WAITFOR=RING
CONNECT="" ATA\r CONNECT\s\A
# this line sets the time to delay before sending the login banner
DELAY=1
#DEBUG=010
```

Add the following line to your /etc/inittab, so that uugetty is run on your serial port, substituting in the correct information for your environment - run-levels (2345 or 345, etc.) config file location, port, speed, and default terminal type:

```
S3:2345:respawn:/sbin/uugetty -d /etc/default/uugetty.ttyS3 ttyS3 F115200 vt100
```

Restart init:

```
linux# init q
```

For the speed parameter in your /etc/inittab, you want to use the highest bps rate that your modem supports.

Now Linux will be watching your serial port for connections. Dial in from another machine and login to you Linux system.

uugetty has a lot more options, see the man page for uugetty) (often just called getty) for a full description. Among other things there is a scheduling feature, and a ringback feature.

12 What Speed Should I Use with My Modem?

By "speed" we really mean the "data flow rate" but almost everybody incorrectly calls it speed. For all modern modems you have no choice of the speed that the modem uses on the telephone line since it will automatically choose the highest possible speed that is possible under the circumstances. But you do have a choice as to what speed will be used between your modem and your computer. This is sometimes called "DTE speed" where "DTE" stands for Data Terminal Equipment (Your computer is a DTE.) You need to

set this speed high enough so this part of the signal path will not be a bottleneck. The setting for the DTE speed is the maximum speed of this link. Most of the time it will likely operate at lower speeds.

For an external modem, DTE speed is the speed (in bits/sec) of the flow over the cable between you modem and PC. For an internal modem, it's the same idea since the modem also emulates a serial port. It may seem ridiculous having a speed limit on communication between a computer and a modem card that is directly connected inside the computer to a much higher speed bus. But it's that way since the modem card probably includes a dedicated serial port which does have speed limits (and settable speeds).

12.1 Speed and Data Compression

What speed do you choose? If it were not for "data compression" one might try to choose a DTE speed exactly the same as the modem speed. Data compression takes the bytes sent to the modem from your computer and encodes them into a fewer number of bytes. For example, if the flow (speed) from the PC to the modem was 20,000 bytes/sec (bps) and the compression ratio was 2 to 1, then only 10,000 bytes/sec would flow over the telephone line. Thus for a 2:1 compression ratio you need to set the speed double the maximum modem speed on the phone line. If the compression ratio is 3 to 1 you need to set it 3 times faster.

12.2 Where do I Set Speed ?

This DTE speed is normally set by a menu in your communications program or by an option given to the getty command if someone is dialing in. You can't set the DCE modem-to-modem speed.

12.3 Can't Set a High Enough Speed

You need to find out the highest speed supported by your hardware. As of late 1998 most hardware only supported speeds up to 115.2K bps. A few 56K internal modems support 230.4K bps. Recent Linux kernels support high speeds (over 115.2K) but you might have difficulty using it because of one or both of the following reasons:

1. The application program (or stty) will not accept the high speed.
2. Setserial has a default speed of 115,200 (but this default is easy to change)

12.3.1 How speed is set in hardware: the divisor and baud_base

Here's a list of commonly used divisors and their corresponding speeds (assuming a maximum speed of 115,200): 1 (115.2K), 2 (57.6K), 3 (38.4K), 6 (19.2K), 12 (9.6K), 24 (4.8K), 48 (2.4K), 96 (1.2K), etc. The serial driver sets the speed in the hardware by sending the hardware only a "divisor" (a positive integer). This "divisor" divides the maximum speed of the hardware resulting in a slower speed (except a divisor of 1 obviously tells the hardware to run at maximum speed).

Normally, if you specify a speed of 115.2K (in your communication program or by stty) then the serial driver sets the port hardware to divisor 1 which obviously sets the highest speed. If you happen to have hardware with a maximum speed of say 230.4K, then specifying 115.2K will result in divisor 1 and will actually give you 230.4K. This is double the speed that you set. In fact, for any speed you set, the actual speed will be double. If you had hardware that could run at 460.8K then the actual speed would be quadruple what you set.

12.3.2 Work-arounds for setting speed

To correct this accounting (but not always fix the problem) you may use "setserial" to change the baud_base to the actual maximal speed of your port such as 230.4K. Then if you set the speed (by your application or by stty) to 230.4K, a divisor of 1 will be used and you'll get the same speed as you set. PROBLEM: stty and many communication programs (as of mid 1999) still have 115.2K as their maximum speed setting and will not let you set 230.4K, etc. So in these cases one solution is not to change anything with `setserial` but mentally keep in mind that the actual speed is always double what you set.

There's another work-around which is not much better. To use it you set the baud_base (with setserial) to the maximal speed of your hardware. This corrects the accounting so that if you set say 115.2K you actually get 115.2K. Now you still have to figure out how to set the highest speed if your communication program (or the like) will not let you do it. Fortunately, setserial has a way to do this: use the "spd_cust" parameter with "divisor 1". Then when you set the speed to 38400 in a communication program, the divisor will be set to 1 in the port and it will operate at maximum speed. For example:
setserial /dev/ttyS2 spd_cust baud_base 230400 divisor 1
Don't try using "divisor" for any other purpose other than the special use illustrated above (with spd_cust).

If there are two or more high speeds that you want to use that your communication program can't set, then it's not quite as easy as above. But the same principles apply. You could just keep the default baud_base and understand that when you set a speed you are really only setting a divisor. So your actual speed will always be your maximum speed divided by whatever divisor is set by the serial driver. See 12.3.1 (How speed is set in hardware: the divisor and baud_base)

12.3.3 Crystal frequency is not baud_base

Note that the baud_base setting is usually much lower than the frequency of the crystal oscillator in the hardware since the crystal frequency is often divided by 16 in the hardware to get the actual top speed. The reason the crystal frequency needs to be higher is so that this high crystal speed can be used to take a number of samples of each bit to determine if it's a 1 or a 0.

12.4 Speed Table

It's best to have at least a 16650 UART for a 56K modem but few modems support it. Second best is a 16550 that has been tweaked to give 230,400 bps. Here are some suggested speeds to set your serial line if your modem speed is:

- 56K (V.90) use 115200 bps or 230400 bps (a few % faster ?)
- 28.8K (V.34), 33.6K (V.34) use 115200 bps
- 14400 bps (V.32bis), with V.42bis data compression, use 57600 bps
- 9600 bps (V.32), with V.42bis data compression, use 38400 bps
- slower than a 9600 bps (V.32) modem, set your speed to the highest speed your modem supports.

13 Communications Programs And Utilities

PPP is by far the most widely used. It's used for Internet access. For dialing out to public libraries, bulletin boards, etc. `minicom` is the most popular followed by `Seyon` (X-Windows only) and `Kermit`.

13.1 Minicom vs. Kermit

Minicom is only a communications program while Kermit is both a communications program and a file transfer protocol. But one may use the Kermit protocol from within Minicom (provided one has Kermit installed on one's PC) . Minicom is menu based while Kermit is command line based (interactive at the special Kermit prompt). While the Kermit program is free software, the documentation is not all free. There is no detailed manual supplied and it is suggested that you purchase a book as the manual. However Kermit has interactive online help which tells all but lacks tutorial explanations for the beginner. Commands may be put in a script file so you don't have to type them over again each time. Kermit (as a communications program) is more powerful than Minicom.

Although all Minicom documentation is free, it's not as extensive as Kermit's. Since permission is required to include Kermit in a commercial distribution, and since the documentation is not entirely free, some distributions don't include Kermit. In my opinion it's easier to set up Minicom and there is less to learn.

13.2 Lists of Programs

Here is a list of some communication software you can choose from, available via FTP, if they didn't come with your distribution. I would like comparative comments on the dialout programs. Are the least popular ones obsolete?

13.2.1 Least Popular Dialout

- `ecu` - a communications program
- `pcomm` - `procomm`-like communications program with zmodem
- `xc` - xcomm communication package

13.2.2 Most Popular Dialout

- `minicom` - `telix`-like communications program. Supports scripts, zmodem, kermit
- *C-Kermit* <`http://www.columbia.edu/kermit/`> - portable, scriptable, serial and TCP/IP communications including file transfer, character-set translation, and zmodem support
- `seyon` - X based communication program

13.2.3 Fax

- `efax` a small fax program
- `hylafax` a large fax program based on the client-server model.
- `mgetty+fax` handles fax stuff and login for dial-ins

13.2.4 Voicemail Software

- *mvm* <`http://www-internal.alphabet.ch/~schaefer/mvm/`> is a Minimal VoiceMail for Linux
- vgetty is an extension to mgetty that handles voicemail for some modems. It should come with recent releases of mgetty.

13.2.5 Dial-in (uses getty)

- `mgetty+fax` is for modems and is well documented (except for voicemail as of early 1999). It also handles fax stuff and provides an alternative to `uugetty`. It's incorporating voicemail (using vgetty) features. See 10.2.1 (About mgetty)
- `uugetty` is also for modems. It comes as a part of the `ps_getty` package. See 11 (About getty_ps)

13.2.6 Other

- `callback` is where you dial out to a remote modem and then that modem hangs up and calls you back (to save on phone bills).
- `SLiRP` and `term` provide a PPP-like service that you can run in user space on a remote computer with a shell account. See 13.3 (term and SLiRP) for more details
- `ZyXEL` is a control program for ZyXEL U-1496 modems. It handles dialin, dialout, dial back security, FAXing, and voice mailbox functions.
- SLIP and PPP software can be found at *ftp://metalab.unc.edu/pub/Linux/system/network/serial*.
- Other things can be found on *ftp://metalab.unc.edu/pub/Linux/system/serial* and *ftp://metalab.unc.edu/pub/Linux/apps/serialcomm* or one of the many mirrors. These are the directories where serial programs are kept.

13.3 SLiRP and term

`SLiRP` and `term` are programs which are of use if you only have a dial-up shell account on a Unix-like machine and want to get the equivalent of a PPP account (or the like) without being authorized to have it (possibly because you don't want to pay extra for it, etc.). `SLiRP` is more popular than `term` which is almost obsolete.

To use `SLiRP` you install it in your shell account on the remote computer. Then you dial up the account and run SLiRP on the remote and PPP on your local PC. You now have a PPP connection over which you may run a web browser on your local PC such as Netscape, etc. There may be some problems as SLiRP is not as good as a real PPP account. Some accounts may provide SLiRP since it saves on IP addresses (You have no IP address while using SLiRP).

`term` is something like SLiRP only you need to run `term` on both the local and remote computer. There is no PPP on the phone line since `term` uses its own protocol. To use `term` from your PC you need to use a term-aware version of ftp to do ftp, etc. Thus it's easier to use SLiRP since the ordinary version of ftp works fine with SLiRP. There is an unmaintained Term HOWTO.

14 What Are UARTs? How Do They Affect Performance?

14.1 Introduction to UARTS

(This section is also in the Serial-HOWTO)

UARTs (**U**niversal **A**synchronous **R**eceiver **T**ransmitter) are serial chips on your PC motherboard (or on an internal modem card). The UART function may also be done on a chip that does other things as well. On older computers like many 486's, the chips were on the disk IO controller card. Still older computer have dedicated serial boards.

The UART's purpose is to convert bytes from the PC's parallel bus to a serial bit-stream. The cable going out of the serial port is serial and has only one wire for each direction of flow. The serial port sends out a stream of bits, one bit at a time. Conversely, the bit stream that enters the serial port via the external cable is converted to parallel bytes that the computer can understand. UARTs deal with data in byte sized pieces, which is conveniently also the size of ASCII characters.

Say you have a terminal hooked up to your PC. When you type a character, the terminal gives that character to it's transmitter (also a UART). The transmitter sends that byte out onto the serial line, one bit at a time, at a specific rate. On the PC end, the receiving UART takes all the bits and rebuilds the (parallel) byte and puts it in a buffer.

Along with converting between serial and parallel, the UART does some other things as a byproduct (side effect) of it's primary task. The voltage used to represent bits is also converted (changed). Extra bits (called start and stop bits) are added to each byte before it is transmitted. See the Serial-HOWTO section, "Voltage Waveshapes" for details. Also, while the flow rate (in bytes/sec) on the parallel bus inside the computer is very high, the flow rate out the UART on the serial port side of it is much lower. The UART has a fixed set of rates (speeds) which it can use at its serial port interface.

14.2 Two Types of UARTs

There are two basic types of UARTs: dumb UARTS and FIFO UARTS. Dumb UARTs are the 8250, 16450, early 16550, and early 16650. They are obsolete but if you understand how they work it's easy to understand how the modern ones work with FIFO UARTS (late 16550, 16550A, 16c552, late 16650, 16750, and 16C950).

There is some confusion regarding 16550. Early models had a bug and worked properly only as 16450's (no FIFO). Later models with the bug fixed were named 16550A but many manufacturers did not accept the name change and continued calling it a 16550. Most all 16550's in use today are like 16550A's. Linux will report it as being a 16550A even though your hardware manual (or a label note) says it's a 16550. A similar situation exists for the 16650 (only it's worse since the manufacturer allegedly didn't admit anything was wrong). Linux will report a late 16650 as being a 16650V2. If it reports it as 16650 it is bad news and only is used as if it had a one-byte buffer.

14.3 FIFOs

To understand the differences between dumb and FIFO (First In, First Out queue discipline) first let's examine what happens when a UART has sent or received a byte. The UART itself can't do anything with the data passing thru it, it just receives and sends it. For the original dumb UARTS, the CPU gets an interrupt from the serial device every time a byte has been sent or received. The CPU then moves the received byte out of the UART's buffer and into memory somewhere, or gives the UART another byte to send. The 8250 and 16450 UARTs only have a 1 byte buffer. That means, that every time 1 byte is sent or received, the CPU is interrupted. At low transfer rates, this is OK. But, at high transfer rates, the CPU gets so busy dealing with the UART, that is doesn't have time to adequately tend to other tasks. In some cases, the CPU does not get around to servicing the interrupt in time, and the byte is overwritten, because they are coming in so fast. This is called an "overrun" or "overflow".

That's where the FIFO UARTs are useful. The 16550A (or 16550) FIFO chip comes with 16 byte FIFO buffers. This means that it can receive up to 14 bytes (or send 16 bytes) before it has to interrupt the CPU. Not only can it wait for more bytes, but the CPU then can transfer all 14 (or more) bytes at a time. Although the interrupt threshold (trigger level) may be set at 8 instead of 14, this is still a significant advantage over the other UARTs, which only have 1 byte buffers. The CPU receives less interrupts, and is free to do other things. Data is not lost, and everyone is happy.

While most PC's only have a 16550 with 16-byte buffers, better UARTS have even larger buffers. Note that the interrupt is issued slightly before the buffer get full (at say a "trigger level" of 14 bytes for a 16-byte buffer). This allows room for a few more bytes to be received during the time that the interrupt is being serviced. The trigger level may be set to various permitted values by kernel software. A trigger level of 1 will be almost like a dumb UART (except that it still has room for 15 more bytes after it issues the interrupt).

If you type something while visiting a BBS, the characters you type go out thru the serial port. Your typed characters that you see on the screen are what was echoed back thru the telephone line thru your modem and then thru your serial port to the screen. If you had a 16-byte buffer on the serial port which held back characters until it had 14 of them, you would need to type many characters before you could see what you typed (before they appeared on the screen). This would be very confusing but there is a "timeout" to prevent this. Thus you normally see a character on the screen just as soon as you type it.

The "timeout" works like this for the receive UART buffer: If characters arrive one after another, then an interrupt is issued only when say the 14th character reaches the buffer. But if a character arrives and the next character doesn't arrive soon thereafter, then an interrupt is issued. This happens even though there are not 14 characters in the buffer (there may only be one character in it). Thus when what you type goes thru this buffer, it acts almost like a 1-byte buffer even though it is actually a 16-byte buffer (unless your typing speed is a hundred times faster than normal). There is also "timeout" for the transmit buffer as well.

14.4 UART Model Numbers

Here's a list of UARTs. *TL* is *T*rigger *L*evel

- 8250, 16450, early 16550: Obsolete with 1-byte buffers
- 16550, 16550A, 16c552: 16-byte buffers, TL=1,4,8,14
- 16650: 32-byte buffers. Speed up to 460.8 Kbps
- 16750: 64-byte buffer for send, 56-byte for receive. Speed up to 921.6 Kbps
- Hayes ESP: 1K-byte buffers.

The obsolete ones are only good for modems no higher than 14.4k (DTE speeds up to 38400 bps). For modern modems you need at least a 16550 (and not an early 16550). For V.90 56k modems, it may be a several percent faster with a 16650 (especially if you are downloading uncompressed files). The main advantage of the 16650 is its larger buffer size as the extra speed isn't needed unless the modem compression ratio is high. Some 56k internal modems may come with a 16650 ??

Non-UART, and intelligent multiport boards use DSP chips to do additional buffering and control, thus relieving the CPU even more. For example, the Cyclades Cyclom, and Stallion EasyIO boards use a Cirrus Logic CD1400 RISC UART, and many boards use 80186 CPUs or even special RISC CPUs, to handle the serial IO.

Most newer PC's (486's, Pentiums, or better) come with 16550A's (usually called just 16550's). If you have something really old the chip may unplug so that you may be able to upgrade by buying a 16550A chip and replacing your existing 16450 UART. If the functionality has been put on another type of chip, you are out of luck. If the UART is socketed, then upgrading is easy (if you can find a replacement). The new and old are pin-to-pin compatible. It may be more feasible to just buy a new serial board on the Internet (few retail stores stock them today).

15 Troubleshooting

15.1 I can't get 56K on my 56K modem

There must be very low noise on the line for it to work at even close to 56K. Some phone lines are so bad that the speeds obtainable are much slower than 56K (like 28.8K or even slower). Sometimes extension phones connected to the same line can cause problems. To check this you might try connecting your modem directly to the point where the telephone line enters the building with the feed for everything else on that line disconnected (if others can tolerate that).

15.2 Uploading (downloading) files is broken/slow

Flow control (both at your PC and/or modem) may not be enabled. If you have set a high DTE speed (like 115.2K) then flow from your modem to your PC may work OK but a lot of flow in the other direction will not all get thru due to the telephone line bottleneck. This will result in many errors and the resending of packets. It may thus take far too long to send a file. In some cases, files don't make it thru at all. If you're downloading long uncompressed files or web pages (and your modem uses data compression) or you've set a low DTE speed, then downloading may also be broken due to no flow control.

15.3 For Dial-in I Keep Getting "line NNN of inittab invalid"

Make sure you are using the correct syntax for your version of `init`. The different `init`'s that are out there use different syntax in the `/etc/inittab` file. Make sure you are using the correct syntax for your version of `getty`.

15.4 When I Try To Dial Out, It Says "/dev/ttySN: Device or resource busy"

This problem can arise when DCD or DTR are not implemented correctly. DCD should only be on (asserted) when there is an actual connection (ie someone has dialed in), not when `getty` is watching the port. Check to make sure that your modem is configured to only assert DCD when there is a connection. DTR should be on (asserted) whenever something is using, or watching the line, like `getty`, `kermit`, or some other comm program.

Another common cause of "device busy" errors, is that you set up your serial port with an interrupt already taken by something else. As each device initializes, it asks Linux for permission to use its hardware interrupt. Linux keeps track of which interrupt is assigned to whom, and if your interrupt is already taken, your device won't be able to initialize properly. The device really doesn't have much of any way to tell you that this happened, except that when you try to use it, it will return a "device-busy" error. Check the interrupts on all of your boards (serial, ethernet, SCSI, etc.). Look for IRQ conflicts.

15.5 I Keep Getting: "Id "S3" respawning too fast: disabled for 5 minutes"

Id "S3" is just an example. In this case look on the line which starts with "S3" in `/etc/inittab` . This is causing the problem. Make sure the syntax for this line is correct and that the device (ttyS3) exists and can be found.

Make sure your modem is configured correctly. Look at registers `E` and `Q`. This can occur when your modem is chatting with `getty`.

If you use uugetty, verify that your `/etc/gettydefs` syntax is correct by doing the following:

```
linux# getty -c /etc/gettydefs
```

This can also happen when the `uugetty` initialization is failing. See section 15.7 (uugetty Still Doesn't Work).

15.6 My Modem is Hosed after Someone Hangs Up, or uugetty doesn't respawn

This can happen when your modem doesn't reset when DTR is dropped. Greg Hankins saw his RD and SD LEDs go crazy when this happened. You need to have your modem reset. Most Hayes compatible modems do this with `&D3`, but on his USR Courier, he had to set `&D2` and `S13=1`. Check your modem manual (if you have one).

15.7 uugetty Still Doesn't Work

There is a `DEBUG` option that comes with `getty_ps`. Edit your config file `/etc/conf.{uu}getty.ttyS`*N* and add `DEBUG=`*NNN*. Where *NNN* is one of the following combination of numbers according to what you are trying to debug:

```
D_OPT   001         option settings
D_DEF   002         defaults file processing
D_UTMP  004         utmp/wtmp processing
D_INIT  010         line initialization (INIT)
D_GTAB  020         gettytab file processing
D_RUN   040         other runtime diagnostics
D_RB    100         ringback debugging
D_LOCK  200         uugetty lockfile processing
D_SCH   400         schedule processing
D_ALL   777         everything
```

Setting `DEBUG=010` is a good place to start.

If you are running `syslogd`, debugging info will appear in your log files. If you aren't running `syslogd` info will appear in `/tmp/getty:ttyS`*N* for debugging `getty` and `/tmp/uugetty:ttyS`*N* for `uugetty`, and in `/var/adm/getty.log`. Look at the debugging info and see what is going on. Most likely, you will need to tune some of the parameters in your config file, and reconfigure your modem.

You could also try `mgetty`. Some people have better luck with it.

15.8 My Modem is Physically There but Can't be Found

If you know about the existing serial ports before you installed an internal modem, then the problem is to find the new serial port. This is covered in the next section. This section is about finding out which serial port has the modem on it.

There's a program that looks for modems on commonly used serial ports called "wvdialconf". Just type "wvdialconf <a-new-file-name>". It will create the new file as a configuration file but you don't need this file unless you are going to use "wvdial" for dialing. See 8.3 (What is wvdialconf ?)

Your problem could be due to a winmodem (or the like) which can't be used with Linux. See 6.4 (Avoid: winmodems). The "setserial program may be used to detect serial ports but will not detect modems on them. Thus "wvdialconf" is best to try first.

Another way to see if there's a modem on a port is to start "minicom" on the port (go to the setup menus with ^AO). Then type "AT" and you should see OK (or 0 if it's set for "digit result codes"). If it takes

many seconds to get a response (including only the cursor moving down one line) then see 15.11 (Slow: Text appears on the screen slowly after long delays)

15.9 The following subsections are in both the Serial and Modem HOWTOs:

15.10 My Serial Port is Physically There but Can't be Found

Check the BIOS menus and BIOS messages. If it's an ISA bus PnP serial port, try "pnpdump –dumpregs" and/or see Plug-and-Play-HOWTO. For the PCI bus look at /proc/pci. You may try probing with setserial. See 8.1.2 (Probing) If nothing seems to get thru the port it may be there but have a bad interrupt. See 15.11 (Slow: Text appears on the screen slowly after long delays)

15.11 Slow: Text appears on the screen slowly after long delays

It's likely mis-set/conflicting interrupts. Here are some of the symptoms which will happen the first time you try to use a modem, terminal, or printer. In some cases you type something but nothing appears on the screen until many seconds later. Only the last character typed may show up (and it may be just an invisible <return> character but you notice the cursor jumps down one line). In other cases where a lot of data should appear on the screen, only a batch of about 16 characters appear. Then there is a long wait of many seconds for the next batch of characters. You might also get "input overrun" error messages (or find them in logs).

For more details on the symptoms and why this happens see the Serial-HOWTO section: "Interrupt Problem Details".

As a final check to see if it really is an interrupt problem, set the IRQ to zero with "setserial". This will tell the driver not to use interrupts but to use inefficient (but faster) polling methods. If this seemingly fixes the "slow" problem then you had an interrupt problem but should still try to fix it since polling uses excessive computer resources.

Checking to find the interrupt conflict may not be easy since looking in the /proc directory may be misleading. Make sure there are no IRQs being illegally shared. Check all your boards (serial, ethernet, SCSI, etc...). Make sure the jumper (or PnP) settings, and the `setserial` parameters are correct for all your serial devices. Also check `/proc/ioports` and `/proc/interrupts` and `/proc/pci` for conflicts. For more details see Serial-HOWTO: Interrupt Problem Details. If it involves Plug-and-Play devices, see also Plug-and-Play-HOWTO.

15.12 The Startup Screen Show Wrong IRQs for the Serial Ports.

Linux does not do any IRQ detection on startup. When the serial module loads it only does serial device detection. Thus, disregard what it says about the IRQ, because it's just assuming the standard IRQs. This is done, because IRQ detection is unreliable, and can be fooled. But if and when setserial runs from a start-up script, it changes the IRQ's and displays the new (and hopefully correct) state on on the startup screen. If the wrong IRQ is not corrected by a later display on the screen, then you've got a problem.

So, even though I have my `ttyS2` set at IRQ 5, I still see

```
ttyS02 at 0x03e8 (irq = 4) is a 16550A
```

at first when Linux boots. (Older kernels may show "ttyS02" as "tty02") You have to use `setserial` to tell Linux the IRQ you are using.

15.13 "Cannot open /dev/ttyS?: Permission denied"

Check the file permissions on this port with "ls -l /dev/ttyS?"_ If you own the ttyS? then you need read and write permissions: crw with the c (Character device) in col. 1. It you don't own it then it should show rw- in cols. 8 & 9 which means that everyone has read and write permission on it. Use "chmod" to change permissions. There are more complicated ways to get access like belonging to a "group" that has group permission.

15.14 "Operation not supported by device" (error message) for ttySx

This means that an operation requested by setserial, stty, etc. couldn't be done because the kernel doesn't support doing it. Formerly this was often due to the "serial" module not being loaded. But with the advent of PnP, it may likely mean that there is no modem (or other serial device) at the address where the driver (and setserial) thinks it is. If there is no modem there, commands (for operations) sent to that address obviously don't get done. See 4.5.2 (What is set in my serial port hardware?)

If the "serial" module wasn't loaded but "lsmod" shows you it's now loaded it might be the case that it's loaded now but wasn't loaded when you got the error message. In many cases the module will automatically loaded when needed (if it can be found). To force loading of the "serial" module it may be listed in the file: /etc/modules.conf or /etc/modules. The actual module should reside in: /lib/modules/.../misc/serial.o.

15.15 "Cannot create lockfile. Sorry" (error message)

When a port is "opened" by a program a lockfile is created in /var/lock/. Wrong permissions for the lock directory will not allow a lockfile to be created there. Use "ls -ld /var/lock" to see if the permissions are OK: usually rwx for everyone (repeated 3 times). If it's wrong, use "chmod" to fix it. Of course, if there is no "lock" directory no lockfile can be created there. For more info on lockfiles see the Serial-HOWTO subsection: "What Are Lock Files".

15.16 "Device /dev/ttyS? is locked."

This means that someone else (or some other process) is supposedly using the serial port. Use "ps -alx" to find out who. Look for any program that uses the serial port. You may also check the lockfile found in /var/lock such as /var/lock/LCK..ttyS2 (contains the id number of the process that is using the port). It could even be a rogue process that you need to kill. If the process doesn't exist, then you may remove the lockfile. Most applications will do this automatically so you seldom need to remove lockfiles manually anymore.

15.17 Software which may help

- `modemstat` and `statserial` show the current state of various modem signal lines (such as DTR, CTS, etc.)
- `irqtune` will give serial port interrupts higher priority to improve performance.
- `hdparm` for hard-disk tuning may help some more.

16 Flash Upgrades

Many modems can be upgraded by reprogramming their flash memories with an upgrade program which you get from the Internet. By sending this "program" from the PC via the serial port to the modem, the modem will store this program in its non-volatile memory (it's still there when the power is turned off). The instructions on installing it are usually on how to do in under Windows so you'll need to figure out how to do the equivalent under Linux (unless you want to install the upgrade under Windows). Sending the program to the modem is often called a download.

If the latest version of this HOWTO still contains this request (see 1.5 (New Versions of this HOWTO)) please send me your experiences with installing such upgrades that will be helpful to others.

Here's the general idea of doing an upgrade. First, there may be a command that you need to send your modem to tell it that what follows is a flash ROM upgrade. In one case this was AT** You can do this by starting a communications program (such as minicom) and type. First type AT <enter> to see if your modem is there and answers "OK".

Next, you need to send an file (sometimes two files) directly to the modem. Communication programs (such as minicom) often use zmodem or kermit to send files to the modem (and beyond) but these put the file into packets which append headers and you want the exact file sent to the modem, not a modified one. But the kermit communications program has a "transmit" command that will send the file directly (without using the kermit packets) so this is one way to send a file directly. Minicom didn't have this feature in 1998.

Another way to send the file(s) would be to escape from the communications program to the shell (in minicom this is ^AJ) and then: `cat upgrade_file_name > /dev/ttyS2` (if your serial port is ttyS2). Then go back to the communication program (type fg at the command line prompt in minicom) to see what happened.

Here's an example session for a certain Rockwell modem (C-a is ^A):

```
- Run minicom
- Type AT** : see "Download initiated .."
- C-a J
- cat FLASH.S37 > /dev/modem
- fg : see "Download flash code .."
- C-a J
- cat 283P1722.S37 > /dev/modem
- fg : see "Device successfully programmed"
```

17 Other Sources of Information

17.1 Misc

- man pages for: `agetty(8)`, `getty(1m)`, `gettydefs(5)`, `init(1)`, `isapnp(8)`, `login(1)`, `mgetty(8)`, `setserial(8)`
- Your modem manual (if it exists). Some modems come without manuals.
- *Serial Suite* <ftp://scicom.alphadec.com/pub/linux> by Vern Hoxie is a collection of blurbs about the care and feeding of the Linux serial port plus some simple programs.
- The Linux serial mailing list. To join, send email to *majordomo@vger.rutgers.edu*, with "`subscribe linux-serial`" in the message body. If you send "`help`" in the message body, you get a help message. The server also serves many other Linux lists. Send the "`lists`" command for a list of mailing lists.

17.2 Books

I've been unable to find a good up-to-date book on modems.

- The Complete Modem Reference by Gilbert Held, 1997. Contains too much info about obsolete topics. More up-to-date info may be found on the Internet.
- Modems For Dummies by Tina Rathbone, 1996. (Have never seen it.)
- Ultimate Modem Handbook by Cass R. Lewart, 1998.

17.3 HOWTOs

- Cable-Modem mini-howto
- ISDN Howto (not a LDP Howto) <http://www.suse.de/Support/sdb_e/isdn.html>: drivers for ISDN "Modems". Much related info on this is in German.
- Modems-HOWTO: In French (Not used in creating this Modem-HOWTO)
- NET-3-HOWTO: all about networking, including SLIP, CSLIP, and PPP
- PPP-HOWTO: help with PPP including modem set-up
- Serial-HOWTO has info on Multiport Serial Cards used for both terminals and banks of modems. Covers the serial port in more detail than in the HOWTO.
- Serial-Programming-HOWTO: for some aspects of serial-port programming
- Text-Terminal-HOWTO: (including connecting up with modems)
- UUCP-HOWTO: for information on setting up UUCP

17.4 Usenet newsgroups

- comp.os.linux.answers FAQs, How-To's, READMEs, etc. about Linux.
- comp.os.linux.hardware Hardware compatibility with the Linux operating system.
- comp.os.linux.setup Linux installation and system administration.
- comp.dcom.modems Modems for all OS's

17.5 Web Sites

- Hayes AT modem commands *Technical Reference for Hayes (tm) Modem Users* <http://www.hayes.com/TechSupport/techref/>
- *Rockwell-based modem commands* <http://www.rss.rockwell.com/techinfo/>
- A white paper discussing serial communications and multiport serial boards is available from Cyclades at *http://www.cyclades.com*
- Modem FAQs:
  Navas 28800 Modem FAQ <http://web.aimnet.com/jnavas/modem/faq.html>
- *Curt's High Speed Modem Page* <http://www.teleport.com/curt/modems.html>

- Much info on 56k modems *56k Modem = v.Unreliable* <http://808hi.com/56k/>
- *Links to modem manufacturers* <http://www.56k.com/links/Modem_Manufacturers/>
- *Identifying modems by FCC ID* <http://www.sbsdirect.com/fccenter.html>
- Partial list of modems which work/don't_work under Linux *modem list* <http://www.o2.net/~gromitkc/winmodem.html>

18 Appendix A: How Modems Work (technical) (unfinished)

18.1 Modulation Details

18.1.1 Intro to Modulation

Modulation is the conversion of a digital signal represented by binary bits (0 or 1) into an analog signal something like a sine wave. The modulated signal consists pure sine wave "carrier" signal which is modified to convey information. A pure carrier sine wave, unchanging in frequency and voltage, provides no flow of information at all (except that a carrier is present). To make it convey information we modify (or modulate) this carrier. There are 3 basic types of modulation: frequency, amplitude, and phase. They will be explained next.

18.1.2 Frequency Modulation

The simplest modulation method is frequency modulation. Frequency is measured in cycles per second (of a sine wave). It's the count of the number of times the sine wave shape repeats itself in a second. This is the same as the number of times it reaches it peak value in a second. The word "Hertz" (abbreviated Hz) is used to mean "cycles per second".

A simple example of frequency modulation is where one frequency means a 0 and another means a 1. For example, for some obsolete 300 baud modems 1070 Hz meant a binary 0 while 1270 Hz meant a binary 1. This was called "frequency shift keying". Instead of just two possible frequencies, more could be used to allow more information to be transmitted. If we had 4 different frequencies (call them A, B, C, and D) then each frequency could stand for a pair of bits. For example, to send 00 one would use frequency A. To send 01, use frequency B; for 10 use C; for 11 use D. In like manner, by using 8 different frequencies we could send 3 bits with each shift in frequency. Each time we double the number of possible frequencies we increase the number of bits it can represent by 1.

18.1.3 Amplitude Modulation

Once one understands frequency modulation example above including the possibilities of representing a few bits by a single shift in frequency, it's easier to understand both amplitude modulation and phase modulation. For amplitude modulation, one just changes the height (voltage) of the sine wave analogous to changing the frequency of the sine wave. For a simple case there could only be 2 allowed amplitude levels, one representing a 0-bit and another representing a 1-bit. As explained for the case of frequency modulation, having more possible amplitudes will result in more information being transmitted.

18.1.4 Phase Modulation

To change the phase of a sine wave at a certain instant of time, we stop sending this old sine wave and immediately begin sending a new sine wave of the same frequency and amplitude. If we started sending the

new sine wave at the same voltage level (and slope) as existed when we stopped sending the old sine wave, there would be no change in phase (and no detectable change at all). But suppose that we started up the new sine wave at a different point on the sine wave curve. Then there would likely be a sudden voltage jump at the point in time where the old sine wave stopped and the new sine wave began. This is a phase shift and it's measured in degrees (deg.) A 0 deg. (or a 360 deg.) phase shift means no change at all while a 180 deg. phase shift just reverses the voltage (and slope) of the sine wave. Put another way, a 180 deg. phase shift just skips over a half-period (180 deg.) at the point of transition. Of course we could just skip over say 90 deg. or 135 deg. etc. As in the example for frequency modulation, the more possible phase shifts, the more bits a single shift in phase can represent.

18.1.5 Combination Modulation

Instead of just selecting either frequency, amplitude, or phase modulation, we may chose to combine modulation methods. Suppose that we have 256 possible frequencies and thus can send a byte (8 bits) for each shift in frequency (since 2 to the 8 power is 256). Suppose also that we have another 256 different amplitudes so that each shift in amplitude represents a byte. Also suppose there are 256 possible phase shifts. Then a certain points in time we may make a shift in all 3 things: frequency, amplitude and phase. This would send out 3 bytes for each such transition.

No modulation method in use today actually does this. It's not practical due to the relatively long time it would take to detect all 3 types of changes. But what is quite common is the simultaneous change in both phase and amplitude. This is called phase-amplitude modulation (sometimes also called quadrature amplitude modulation = QAM). This method is used for the common modem speeds of 14.4k, 28.8k, and 33.6k. The only significant case where this modulation method is not used today is for 56k modems. But even 56k modems exclusively use QAM (phase-amplitude modulation) in the direction from your PC out the telephone line. Sometimes even the other direction will also fall back to QAM when line conditions are not good enough. Thus QAM (phase-amplitude modulation) still remains the most widely used method on ordinary telephone lines.

18.1.6 56k Modems (v.90)

The modulation method used above 33.6k is entirely different than the common phase-amplitude modulation. Since ordinary telephone calls are converted to digital signals at the local offices of the telephone company, the fastest speed that you can send digital data by an ordinary telephone call is the same speed that the telephone company uses over its digital portion of the phone call transmission. What is this speed? Well, it's close to 64Kbps. It would be 64k but sometimes bits are "stolen" for the purposes of signalling. But if the phone Co. knows that the link is digital, bits may not get stolen. The case of 64k will be presented and then it will be explained why the actual speed is lower (56k or less –usually significantly less).

Thus 64k is the absolute top speed possible for an ordinary telephone call using the digital portion of the circuit that was designed to send digital encodings of the human voice. In order to use 64k, the modem must know exactly how the telephone company is doing its digital encoding of the analog signals. This task is far too complicated if both sides of a telephone call have only an analog interface to the telephone company. But if one side has a digital interface, then it's possible (at least in one direction). Thus if your ISP has a digital interface to the phone company, the ISP may send out a certain digital signal over the phone lines toward your PC. The digital signal from the ISP gets converted to analog at the local telephone office near your PC's location (perhaps near your home). Then it's your modem's task to try to figure out exactly what that digital signal was. If it can do this then transmission at 64k (the speed of the telephone company's digital signal) is possible in this direction.

What method does the telephone company use to digitally encode analog signals? It uses a method of sampling the amplitude of the analog signal at a rate of 8000 samples per second. Each sample amplitude

is encoded as a 8-bit (ASCII-like) byte. (Note: 8 x 8000 = 64k) This is called "Pulse Code Modulation" = PCM. These bytes are then sent digitally on the telephone company's digital circuits where many calls share a single circuit using a time-sharing scheme known as "time division multiplexing". Then finally at the local telephone office near your home, the digital signal is de-multiplexed resulting in the same digital signal as was originally created by PCM. This signal is then converted back to analog and sent to your home. Each 8-bit byte creates a certain amplitude of the analog signal. Your modem's task is to determine just what that PCM 8-bit byte was based on the amplitude it detects.

This is (sort of) "amplitude demodulation" but not really. It's not amplitude demodulation because there is no carrier. Actually, it's called "modulus conversion" which is the inverse of PCM. In order to determine the digital codes the telephone Co. used to create the analog signal, the modem must sample this analog signal amplitude at exactly the same points in time the phone Co. used when it created the analog signal. To do this a timing signal is generated from a residual 4k Hz signal on the phone line. The setting of amplitudes to go out to your home/office at 8k settings/sec sort of creates a 4k signal. Suppose every other amplitude was of opposite polarity. Then there would be a 4k Hz sine-like wave created. Each amplitude is in a sense a 8-bit symbol and when to sample amplitudes is known as "symbol timing".

Now the encoding of amplitudes in PCM is not linear and at low amplitudes the differences between possible amplitude is quite small. So to make the distinction of amplitudes that are close to each other easier, certain amplitudes are not used. This give a larger delta between possible amplitudes and makes correct detection of them by your modem easier. If half the amplitude levels were not used by v.90 and it would be tantamount to each symbol (allowed amplitude level) representing 7 bits instead of 8. This is where 56k comes from: 7 bits/symbol x 8k symbols/sec = 56k bps. Of course each symbol is actually generated by 8-bits but only 128 bytes of the possible 256 bytes are actually used. There is a code table mapping these 128 8-bit bytes to 128 7-bit bytes.

But it's a little more complicated that this. If the line conditions are not nearly perfect, then even fewer possible levels (symbols) are used resulting in speeds under 56k. Also due to government rules prohibiting high power levels on phone lines, certain high amplitudes levels can't be used resulting in only about 53.3k at best for "56k" modems.

Note that the digital part of the telephone network is bi-directional. Two such circuits are used for a phone call, one in each direction. The direction from your home/office to the ISP uses the conventional phase-amplitude modulation scheme with a maximum of 36.6k (and not 53.3k).

19 Appendix B: "baud" vs. "bps"

19.1 A simple example

"baud" and "bps" are perhaps one of the most misused terms in the computing and telecommunications field. Many people use these terms interchangeably, when in fact they are not! bps is simply the number of bits transmitted per second. The baud rate is a measure of how many times per second a signal changes (or could change). For a typical serial port a 1-bit is -12 volts and a 0-bit is +12 v (volts). If the bps is 38,400 a sequence of 010101... would also be 38,400 baud since the voltage shifts back and forth from positive to negative to positive ... and there are 38,400 shifts per second. For another sequence say 111000111... there will be fewer shifts of voltage since for three 1's in sequence the voltage just stays at -12 volts yet we say that its still 38,400 baud since there is a possibility that the number of changes per second will be that high.

Looked at another way, put an imaginary tic mark separating each bit (even thought the voltage may not change). 38,400 baud then means 38,400 tic marks per second. The tic marks at at the instants of permitted change and are actually marked by a synchronized clock signal generated in the hardware but not sent over the external cable.

Suppose that a "change" may have more than the two possible outcomes of the previous example (of +- 12 v). Suppose it has 4 possible outcomes, each represented by a unique voltage level. Each level may represent a pair of bits (such as 01). For example, -12v could be 00, -6v 01, +6v 10 and +12v 11. Here the bit rate is double the baud rate. For example, 3000 changes per second will generate 2 bits for each change resulting in 6000 bits per second (bps). In other words 3000 baud results in 6000 bps.

19.2 Real examples

The above example is overly simple. Real examples are more complicated but based on the same idea. This explains how a modem running at 2400 baud, can send 14400 bps (or higher). The modem achieves a bps rate greater than baud rate by encoding many bits in each signal change (or transition). Thus, when 2 or more bits are encoded per baud, the bps rate exceeds the baud rate. If your modem-to-modem connection is at 14400 bps, it's going to be sending 6 bits per signal transition (or symbol) at 2400 baud. A speed of 28800 bps is obtained by 3200 baud at 9 bits/baud. When people misuse the word baud, they may mean the modem speed (such as 33.6K).

Common modem bps rates were formerly 50, 75, 110, 300, 1200, 2400, 9600. These were also the bps rates over the serial_port-to-modem cables. Today the bps modem-to-modem rates are 14.4K, 28.8K, 33.6K, and 56K, but the common rates over the serialPort-to-modem cables are not the same but are: 19.2K, 38.4K, 57.6K, 115.2K. Using modems with V.42bis compression (max 4:1 compression), rates up to 115.2K bps are possible for 33.6K modems (230.4K is possible for 56K modems).

Except for the 56k modems, most modems run at 2400, 3000, or 3200 baud. Because of the bandwidth limitations on voice-grade phone lines, baud rates greater than 2400 are harder to achieve, and only work under conditions of pristine phone line quality.

How did this confusion between bps and baud start? Well, back when antique low speed modems were high speed modems, the bps rate actually did equal the baud rate. One bit would be encoded per phase change. People would use bps and baud interchangeably, because they were the same number. For example, a 300 bps modem also had a baud rate of 300. This all changed when faster modems came around, and the bit rate exceeded the baud rate. “baud” is named after Emile Baudot, the inventor of the asynchronous telegraph printer. One way this problem gets resolved is to use the term "symbol rate" instead of "baud" and thus avoid using the term "baud".

20 Appendix C: Terminal Server Connection

This section was adapted from Text-Terminal-HOWTO.

A terminal server is something like an intelligent switch that can connect many modems (or terminals) to one or more computers. It's not a mechanical switch so it may change the speeds and protocols of the streams of data that go thru it. A number of companies make terminal servers: Xyplex, Cisco, 3Com, Computone, Livingston, etc. There are many different types and capabilities. Another HOWTO is needed to compare and describe them (including the possibility of creating your own terminal server with a Linux PC). Most are used for modem connections rather than directly connected terminals.

One use for them is to connect many modems (or terminals) to a high speed network which connects to host computers. Of course the terminal server must have the computing power and software to run network protocols so it is in some ways like a computer. The terminal server may interact with the user and ask what computer to connect to, etc. or it may connect without asking. One may sometimes send jobs to a printer thru a terminal server.

A PC today has enough computing power to act like a terminal server except that each serial port should have its own hardware interrupt. PC's only have a few spare interrupts for this purpose and since they

are hard-wired you can't create more by software. A solution is to use an advanced multiport serial card which has its own system of interrupts (or on lower cost models, shares one of the PC's interrupts between a number of ports). See Serial-HOWTO for more info. If such a PC runs Linux with getty running on many serial ports it might be thought of as a terminal server. It is in effect a terminal server if it's linked to other PC's over a network and if its job is mainly to pass thru data and handle the serial port interrupts every 14 (or so) bytes. Software called "radius" is sometimes used.

Today real terminal servers serve more than just terminals. They also serve PC's which emulate terminals, and are sometimes connected to a bank of modems connected to phone lines. Some even include built-in modems. If a terminal (or PC emulating one) is connected directly to a modem, the modem at the other end of the line could be connected to a terminal server. In some cases the terminal server by default expects the callers to use PPP packets, something that real text terminals don't generate.

21 Appendix D: Other Types of Modems

This HOWTO currently only deals with the common type of modem used to connect PC's to ordinary analog telephone lines. There are various other types of modems, including devices called modems that are not really modems.

21.1 Digital-to-Digital "Modems"

The standard definition of a modem is sometimes broadened to include "digital" modems. Today direct digital service is now being provided to many homes and offices so a computer there sends out digital signals directly (well almost) into the telephone lines. But a device is still needed to convert the computer digital signal into type allowed on telephone circuits and this device is sometimes called a modem. The next 2 sections: ISDN and DSL concern digital-to-digital "modems".

21.2 ISDN "Modems"

The "modem" is really a Terminal Adapter (TA). A Debian package "isdnutils" is available. There is a ISDN Howto in German with an English translation: `<http://www.suse.de/Support/sdb_e/isdn.html>`. It's put out by the SuSE distribution of Linux and likely is about drivers available in that distribution. There is an isdn4linux package and a newsgroup: de.alt.comm.isdn4linux. Many of the postings are in German. You might try using a search engine (such as DejaNews) to find "isdn4linux".

21.3 Digital Subscriber Line (DSL)

DSL uses the existing twisted pair line from your home (etc.) to the local telephone office. This can be used if your telephone line can accept higher speeds than an ordinary modem (say 56k) sends over it. It replaces the analog-to-digital converter at the local telephone office with a converter which can accept a much faster flow of data (in a different format of course). The device which converts the digital signals from your computer to the signal used to represent digital data on the local telephone line is also called a modem. This document presently does not cover the special aspects of these modems.

END OF Modem-HOWTO