Serial HOWTO

Table of Contents

Serial HOWTO....1
David S.Lawyer dave@lafn.org original by Greg Hankins....1
1. Introduction....1
2. How the Hardware Transfers Bytes....1
3. Serial Port Basics....1
4. Is the Serial Port Obsolete?....1
5. Multiport Serial Boards/Cards/Adapters....2
6. Servers for Serial Ports....2
7. Configuring Overview....2
8. Locating the Serial Port: IO address, IRQs....2
9. Configuring the Serial Driver (high−level) "stty"....2
10. Serial Port Devices /dev/tts/2 = /dev/ttyS2, etc....2
11. Interesting Programs You Should Know About....3
12. Speed (Flow Rate)....3
13. Locking Out Others....3
14. Communications Programs And Utilities....3
15. Serial Tips And Miscellany....3
16. Troubleshooting....3
17. Interrupt Problem Details....4
18. What Are UARTs? How Do They Affect Performance?....4
19. Pinout and Signals....4
20. Voltage Waveshapes....4
21. Other Serial Devices (not async EIA−232)....4
22. Other Sources of Information....5
23. Appendix: Obsolete Hardware (prior to 1990) Info....5
1. Introduction....5
1.1 Copyright, Disclaimer, & Credits....5
Copyright....5
Disclaimer....6
Trademarks....6
Credits....6
1.2 New Versions of this Serial−HOWTO....6
1.3 New in Recent Versions....6
1.4 Related HOWTO's re the Serial Port....7
1.5 Feedback....7
1.6 What is a Serial Port?....7
2. How the Hardware Transfers Bytes....8
2.1 Transmitting....8
2.2 Receiving....9
2.3 The Large Serial Buffers....9
3. Serial Port Basics....10
3.1 What is a Serial Port ?....10
Intro to Serial....10
Pins and Wires....10
RS−232 or EIA−232, etc....11
3.2 IO Address & IRQ....11
3.3 Names: ttyS0, ttyS1, etc....11
3.4 Interrupts....11

Table of Contents

Serial HOWTO
3.5 Data Flow (Speeds) 12
3.6 Flow Control 12
Example of Flow Control 13
Symptoms of No Flow Control 13
Hardware vs. Software Flow Control 14
3.7 Data Flow Path; Buffers 14
3.8 Complex Flow Control Example 15
3.9 Serial Driver Module 16
4. Is the Serial Port Obsolete? 17
4.1 Introduction 17
4.2 EIA−232 Cable Is Low Speed & Short Distance 17
4.3 Inefficient Interface to the Computer 17
5. Multiport Serial Boards/Cards/Adapters 18
5.1 Intro to Multiport Serial 18
5.2 Modem Limitations 18
5.3 Dumb vs. Smart Cards 19
5.4 Getting/Enabling a Driver 19
Introduction 19
If you need driver built into the kernel (mostly dumb boards) 19
If you use a module (mostly for smart boards) 19
Getting info on multiport boards 19
5.5 Multiport Devices in the /dev Directory 20
5.6 Making Legacy Multiport Devices in the /dev Directory 20
5.7 Standard PC Serial Cards 20
5.8 Dumb Multiport Serial Boards (with standard UART chips) 21
5.9 Intelligent Multiport Serial Boards 22
5.10 Unsupported Multiport Boards 24
6. Servers for Serial Ports 24
7. Configuring Overview 25
8. Locating the Serial Port: IO address, IRQs 25
8.1 IO & IRQ Overview 25
8.2 PCI Bus Support 27
Introduction 27
More info on PCI 27
8.3 Common mistakes made re low−level configuring 27
8.4 IRQ & IO Address Must be Correct 28
8.5 What is the IO Address and IRQ per the driver ? 28
Introduction 28
I/O Address & IRQ: Boot−time messages 28
The /proc directory and setserial 29
8.6 What is the IO Address & IRQ of my Serial Port Hardware? 30
Introduction 30
PCI: What IOs and IRQs have been set? 30
PCI: Enabling a disabled port 30
ISA PnP ports 31
Finding a port that is not disabled (ISA, PCI, PnP, non−PnP) 31
Exploring via MS Windows (a last resort) 31

Table of Contents

Serial HOWTO

8.7 Choosing Serial IRQs....31
IRQ 0 is not an IRQ....31
Interrupt sharing, Kernels 2.2+....32
What IRQs to choose?....32
8.8 Choosing Addresses −−Video card conflict with ttyS3....33
8.9 Set IO Address & IRQ in the hardware (mostly for PnP)....33
Using a PnP BIOS to I0−IRQ Configure....34
8.10 Giving the IRQ and IO Address to Setserial....34
9. Configuring the Serial Driver (high−level) "stty"....35
9.1 Overview....35
9.2 Flow Control....35
10. Serial Port Devices /dev/tts/2 = /dev/ttyS2, etc....35
10.1 Serial Port Names: ttyS2, tts/1, etc....35
10.2 Devfs (The Device File System)....36
10.3 Legacy Serial Port Device Names & Numbers....36
10.4 More on Serial Port Names....36
10.5 USB (Universal Serial Bus) Serial Ports....37
10.6 Link ttySN to /dev/modem....37
10.7 Which Connector on the Back of my PC is ttyS1, etc?....37
Inspect the connectors....37
Send bytes to the port....37
Connect a device to the connector....38
Missing connectors....38
10.8 Creating Devices In the /dev directory....38
11. Interesting Programs You Should Know About....39
11.1 Serial Monitoring/Diagnostics Programs....39
11.2 Changing Interrupt Priority....39
11.3 What is Setserial ?....39
Important information....39
Introduction....40
Serial module unload....40
Slow baud rates of 1200 or less....40
Giving the setserial command....41
Configuration file....41
Probing....42
Boot−time Configuration....43
Edit a script (required prior to version 2.15)....43
Configuration method using /etc/serial.conf, etc....44
IRQs....45
Laptops: PCMCIA....45
11.4 Stty....46
Introduction....46
Flow control options....46
Using stty at a "foreign" terminal....47
Old redirection method....47
Two interfaces at a terminal....47
Where to put the stty command ?....48

Table of Contents

Serial HOWTO

11.5 What is isapnp ?....48
11.6 What is slattach?....49
12. Speed (Flow Rate)....49
12.1 Very High Speeds....49
Speeds over 115.2k....49
How speed is set in hardware: the divisor and baud_base....49
Setting the divisor, speed accounting....50
Crystal frequency is higher than baud_base....50
12.2 Higher Serial Throughput....51
13. Locking Out Others....51
13.1 Introduction....51
13.2 Lock−Files....51
13.3 Lock−Files and devfs Problems....52
13.4 Change Owners, Groups, and/or Permissions of Device Files....52
14. Communications Programs And Utilities....53
14.1 List of Software....53
14.2 kermit and zmodem....53
15. Serial Tips And Miscellany....53
15.1 Serial Module....53
15.2 Serial Console (console on the serial port)....54
15.3 Line Drivers....54
15.4 Stopping the Data Flow when Printing, etc....54
15.5 Known IO Address Conflicts....54
Avoiding IO Address Conflicts with Certain Video Boards....54
IO address conflict with ide2 hard drive....54
15.6 Known Defective Hardware....55
Problem with AMD Elan SC400 CPU (PC−on−a−chip)....55
16. Troubleshooting....55
16.1 Serial Electrical Test Equipment....55
Breakout Gadgets, etc....55
Measuring voltages....55
Taste voltage....56
16.2 Serial Monitoring/Diagnostics....56
16.3 (The following subsections are in both the Serial and Modem HOWTOs)....56
16.4 My Serial Port is Physically There but Can't be Found....56
16.5 Extremely Slow: Text appears on the screen slowly after long delays....57
16.6 Somewhat Slow: I expected it to be a few times faster....57
16.7 The Startup Screen Show Wrong IRQs for the Serial Ports....58
16.8 "Cannot open /dev/ttyS?: Permission denied"....58
16.9 "Operation not supported by device" for ttyS?....58
16.10 "Cannot create lockfile. Sorry"....58
16.11 "Device /dev/ttyS? is locked."....59
16.12 "/dev/tty? Device or resource busy"....59
16.13 "Input/output error" from setserial, stty, pppd, etc....60
16.14 "LSR safety check engaged"....60
16.15 Overrun errors on serial port....60
16.16 Port gets characters only sporadically....60

Table of Contents

Serial HOWTO

16.17 Troubleshooting Tools....60
16.18 Almost all characters are wrong; Many missing or many extras....61
17. Interrupt Problem Details....61
17.1 Types of interrupt problems....61
17.2 Symptoms of Mis−set or Conflicting Interrupts....62
17.3 Mis−set Interrupts....62
17.4 Interrupt Conflicts....63
17.5 Resolving Interrupt Problems....64
18. What Are UARTs? How Do They Affect Performance?....64
18.1 Introduction to UARTS....64
18.2 Two Types of UARTs....64
18.3 FIFOs....65
18.4 Why FIFO Buffers are Small....66
18.5 UART Model Numbers....66
19. Pinout and Signals....66
19.1 Pinout of 9−pin and 25−pin serial connectors....66
19.2 Signals May Have No Fixed Meaning....67
19.3 Cabling Between Serial Ports....67
19.4 RTS/CTS and DTR/DSR Flow Control....68
The DTR and DSR Pins....69
19.5 Preventing a Port From Opening....69
20. Voltage Waveshapes....69
20.1 Voltage for a Bit....69
20.2 Voltage Sequence for a Byte....70
20.3 Parity Explained....70
20.4 Forming a Byte (Framing)....70
20.5 How "Asynchronous" is Synchronized....71
21. Other Serial Devices (not async EIA−232)....71
21.1 Successors to EIA−232....71
21.2 EIA−422−A (balanced) and EIA−423−A (unbalanced)....71
21.3 EIA−485....72
21.4 EIA−530....72
21.5 EIA−612/613....72
21.6 The Universal Serial Bus (USB)....73
21.7 Firewire....73
21.8 MIDI....73
21.9 Synchronization & Synchronous....73
Defining Asynchronous vs Synchronous....74
Synchronous Communication....74
22. Other Sources of Information....74
22.1 Books....74
22.2 Serial Software....75
22.3 Related Linux Documents....75
22.4 Usenet newsgroups:....75
22.5 Serial Mailing List....76
22.6 Internet....76
23. Appendix: Obsolete Hardware (prior to 1990) Info....76

Table of Contents

Serial HOWTO
23.1 Replacing obsolete UARTS.....76

Serial HOWTO

David S.Lawyer dave@lafn.org original by Greg Hankins

v2.22 December 2003

This document describes the UART serial port features other than those which should be covered by Modem−HOWTO, PPP−HOWTO, Serial−Programming−HOWTO, or Text−Terminal−HOWTO. It lists info on multiport serial cards. It contains technical info about the serial port itself in more detail than found in the above HOWTOs and should be best for troubleshooting when the problem is the serial port itself. If you are dealing with a Modem, PPP (used for Internet access on a phone line), or a Text−Terminal, those HOWTOs should be consulted first.

1. Introduction

- 1.1 Copyright, Disclaimer, & Credits
- 1.2 New Versions of this Serial−HOWTO
- 1.3 New in Recent Versions
- 1.4 Related HOWTO's re the Serial Port
- 1.5 Feedback
- 1.6 What is a Serial Port?

2. How the Hardware Transfers Bytes

- 2.1 Transmitting
- 2.2 Receiving
- 2.3 The Large Serial Buffers

3. Serial Port Basics

- 3.1 What is a Serial Port ?
- 3.2 IO Address & IRQ
- 3.3 Names: ttyS0, ttyS1, etc.
- 3.4 Interrupts
- 3.5 Data Flow (Speeds)
- 3.6 Flow Control
- 3.7 Data Flow Path; Buffers
- 3.8 Complex Flow Control Example
- 3.9 Serial Driver Module

4. Is the Serial Port Obsolete?

- 4.1 Introduction
- 4.2 EIA−232 Cable Is Low Speed & Short Distance
- 4.3 Inefficient Interface to the Computer

5. Multiport Serial Boards/Cards/Adapters

- 5.1 Intro to Multiport Serial
- 5.2 Modem Limitations
- 5.3 Dumb vs. Smart Cards
- 5.4 Getting/Enabling a Driver
- 5.5 Multiport Devices in the /dev Directory,
- 5.6 Making Legacy Multiport Devices in the /dev Directory
- 5.7 Standard PC Serial Cards
- 5.8 Dumb Multiport Serial Boards (with standard UART chips)
- 5.9 Intelligent Multiport Serial Boards
- 5.10 Unsupported Multiport Boards

6. Servers for Serial Ports

7. Configuring Overview

8. Locating the Serial Port: IO address, IRQs

- 8.1 IO & IRQ Overview
- 8.2 PCI Bus Support
- 8.3 Common mistakes made re low−level configuring
- 8.4 IRQ & IO Address Must be Correct
- 8.5 What is the IO Address and IRQ per the driver ?
- 8.6 What is the IO Address & IRQ of my Serial Port Hardware?
- 8.7 Choosing Serial IRQs
- 8.8 Choosing Addresses −−Video card conflict with ttyS3
- 8.9 Set IO Address & IRQ in the hardware (mostly for PnP)
- 8.10 Giving the IRQ and IO Address to Setserial

9. Configuring the Serial Driver (high−level) "stty"

- 9.1 Overview
- 9.2 Flow Control

10. Serial Port Devices /dev/tts/2 = /dev/ttyS2, etc.

- 10.1 Serial Port Names: ttyS2, tts/1, etc.
- 10.2 Devfs (The Device File System)
- 10.3 Legacy Serial Port Device Names & Numbers
- 10.4 More on Serial Port Names
- 10.5 USB (Universal Serial Bus) Serial Ports
- 10.6 Link ttySN to /dev/modem
- 10.7 Which Connector on the Back of my PC is ttyS1, etc?
- 10.8 Creating Devices In the /dev directory

11. Interesting Programs You Should Know About

- 11.1 Serial Monitoring/Diagnostics Programs
- 11.2 Changing Interrupt Priority
- 11.3 What is Setserial ?
- 11.4 Stty
- 11.5 What is isapnp ?
- 11.6 What is slattach?

12. Speed (Flow Rate)

- 12.1 Very High Speeds
- 12.2 Higher Serial Throughput

13. Locking Out Others

- 13.1 Introduction
- 13.2 Lock−Files
- 13.3 Lock−Files and devfs Problems
- 13.4 Change Owners, Groups, and/or Permissions of Device Files

14. Communications Programs And Utilities

- 14.1 List of Software
- 14.2 kermit and zmodem

15. Serial Tips And Miscellany

- 15.1 Serial Module
- 15.2 Serial Console (console on the serial port)
- 15.3 Line Drivers
- 15.4 Stopping the Data Flow when Printing, etc.
- 15.5 Known IO Address Conflicts
- 15.6 Known Defective Hardware

16. Troubleshooting

- 16.1 Serial Electrical Test Equipment
- 16.2 Serial Monitoring/Diagnostics
- 16.3 (The following subsections are in both the Serial and Modem HOWTOs)
- 16.4 My Serial Port is Physically There but Can't be Found
- 16.5 Extremely Slow: Text appears on the screen slowly after long delays
- 16.6 Somewhat Slow: I expected it to be a few times faster
- 16.7 The Startup Screen Show Wrong IRQs for the Serial Ports.
- 16.8 "Cannot open /dev/ttyS?: Permission denied"
- 16.9 "Operation not supported by device" for ttyS?
- 16.10 "Cannot create lockfile. Sorry"

- 16.11 "Device /dev/ttyS? is locked."
- 16.12 "/dev/tty? Device or resource busy"
- 16.13 "Input/output error" from setserial, stty, pppd, etc.
- 16.14 "LSR safety check engaged"
- 16.15 Overrun errors on serial port
- 16.16 Port gets characters only sporadically
- 16.17 Troubleshooting Tools
- 16.18 Almost all characters are wrong; Many missing or many extras

17. Interrupt Problem Details

- 17.1 Types of interrupt problems
- 17.2 Symptoms of Mis−set or Conflicting Interrupts
- 17.3 Mis−set Interrupts
- 17.4 Interrupt Conflicts
- 17.5 Resolving Interrupt Problems

18. What Are UARTs? How Do They Affect Performance?

- 18.1 Introduction to UARTS
- 18.2 Two Types of UARTs
- 18.3 FIFOs
- 18.4 Why FIFO Buffers are Small
- 18.5 UART Model Numbers

19. Pinout and Signals

- 19.1 Pinout of 9−pin and 25−pin serial connectors
- 19.2 Signals May Have No Fixed Meaning
- 19.3 Cabling Between Serial Ports
- 19.4 RTS/CTS and DTR/DSR Flow Control
- 19.5 Preventing a Port From Opening

20. Voltage Waveshapes

- 20.1 Voltage for a Bit
- 20.2 Voltage Sequence for a Byte
- 20.3 Parity Explained
- 20.4 Forming a Byte (Framing)
- 20.5 How "Asynchronous" is Synchronized

21. Other Serial Devices (not async EIA−232)

- 21.1 Successors to EIA−232
- 21.2 EIA−422−A (balanced) and EIA−423−A (unbalanced)
- 21.3 EIA−485
- 21.4 EIA−530
- 21.5 EIA−612/613

- 21.6 The Universal Serial Bus (USB)
- 21.7 Firewire
- 21.8 MIDI
- 21.9 Synchronization & Synchronous

22. Other Sources of Information

- 22.1 Books
- 22.2 Serial Software
- 22.3 Related Linux Documents
- 22.4 Usenet newsgroups:
- 22.5 Serial Mailing List
- 22.6 Internet

23. Appendix: Obsolete Hardware (prior to 1990) Info

- 23.1 Replacing obsolete UARTS

---

1. Introduction

This HOWTO covers basic info on the Serial Port and multiport serial cards. It contains much more information in it than most people need to know and most people are able to use it without reading this HOWTO. But if you're having problems or just want to understand how it works, this is one place to find out about it.

This HOWTO is about the original serial port which uses a UART chip and is sometimes called a "UART serial port" to differentiate it from the newer Universal Serial Bus. Information specific to modems and text−terminals is found in Modem−HOWTO and Text−Terminal−HOWTO. Info on getty (the program that runs the login process or the like) has been also moved to these HOWTOs since mgetty and uugetty are best for modems while agetty is best for text−terminals. If you are dealing with a modem, text terminal, or printer, then you may not need to consult this HOWTO. But if you are using the serial port for some other device, using a multiport serial card, trouble−shooting the serial port itself, or want to understand more technical details of the serial port, then you may want to use this HOWTO as well as some of the other HOWTOs. (See Related HOWTO's) This HOWTO lists info on various multiport serial cards since they may be used for either modems or text−terminals. This HOWTO addresses Linux running on PCs (ISA or PCI buses), although it might be valid for other architectures.

1.1 Copyright, Disclaimer, & Credits

Copyright

Copyright (c) 1993−1997 by Greg Hankins, (c) 1998−2003 by David S. Lawyer mailto:dave@lafn.org

Please freely copy and distribute (sell or give away) this document in any format. Send any corrections and comments to the document maintainer. You may create a derivative work and distribute it provided that you:

1. If it's not a translation: Email a copy of your derivative work (in a format LDP accepts) to the author(s) and maintainer (could be the same person). If you don't get a response then email the LDP

(Linux Documentation Project): submit@en.tldp.org.
2. License the derivative work in the spirit of this license or use GPL. Include a copyright notice and at least a pointer to the license used.
3. Give due credit to previous authors and major contributors.

If you're considering making a derived work other than a translation, it's requested that you discuss your plans with the current maintainer.

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.

Trademarks.

Any brand names (starts with a capital letter such as MS Windows) should be assumed to be a trademark). Such trademarks belong to their respective owners.

Credits

Most of the original Serial−HOWTO was written by Greg Hankins. mailto:gregh@twoguys.org He also rewrote many contributions by others in order to maintain continuity in the writing style and flow. He wrote: ``Thanks to everyone who has contributed or commented, the list of people has gotten too long to list (somewhere over one hundred). Special thanks to Ted Ts'o for answering questions about the serial drivers." Approximately half of v2.00 was from Greg Hankins HOWTO and the other half is by David Lawyer. Ted Ts'o has continued to be helpful.

1.2 New Versions of this Serial−HOWTO

New versions of the Serial−HOWTO will be available to browse and/or download at LDP mirror sites. For a list of mirror sites see: http://www.tldp.org/mirrors.html. Various formats are available. If you only want to quickly check the date of the latest version look at http://www.tldp.org/HOWTO/Serial−HOWTO.html and compare it to this version: v2.22 December 2003 .

1.3 New in Recent Versions

For a full revision history going back to the time I started maintaining this HOWTO, see the source file (in linuxdoc format) at
http://www.ibiblio.org/pub/linux/docs/HOWTO/other−formats/sgml/Serial−HOWTO.sgml.gz.

- v2.22 Dec. 2003: revised Complex Flow Control Example, more on devfs
- v2.21 Nov. 2003 Kernel compile USB options for serial ports, revised setserial
- v2.20 Oct. 2003: MAKEDEV is often only in /sbin and not in /dev.
- v2.19 September 2003: linux−serial email now at kernel.org, new section: Servers, pinout diagram
- v2.18 May 2003: EIA−485 features not supported by Linux, Flow control "typos" fixed
- v2.17 Feb 2003: url signum−>cendio, Mac port names, clarity when stopping data flow when printing, ide2 address conflict

1.4 Related HOWTO's re the Serial Port

Modems, Text−Terminals, some printers, and other peripherals often use the serial port. Get these HOWTOs from the nearest mirror site as explained above.

- `Modem-HOWTO` is about installing and configuring modems
- `Printing-HOWTO` has info for serial printers using old lpr command
- `LPRng-HOWTO` (not a LDP HOWTO, may come with software) has info for serial printing for "Next Generation" lpr
- `Serial-Programming-HOWTO` helps you write C programs that read and write to the serial port and/or check/set its state. A version written by Vern Hoxie but not submitted is at Internet.
- `Text-Terminal-HOWTO` is about how they work, how to install configure, and repair them. It includes a section on "Make a Terminal the Console" which is useful for using a remote terminal to control a server (via the serial port).
- `Remote-Serial-Console-HOWTO` is about making a text−terminal be the console so it can display boot−time messages, etc.

1.5 Feedback

Please send me any suggestions, or additional material. Tell me what you don't understand, or what could be clearer. You can reach me via email at `mailto:dave@lafn.org`.

1.6 What is a Serial Port?

The conventional serial port (not the newer USB port, or HSSI port) is a very old I/O port. Almost all PC's have them. Macs (Apple Computer) after mid−1998 only have the USB port. However, it's possible, to put a conventional serial port device on the USB.

Each serial port has a "file" associated with it in the /dev directory. It isn't really a file but it seems like one. For newer versions of Linux which use the Device File System (devfs) an ordinary serial port is for example: /dev/tts/0. Formerly (before the devfs) this port was called /dev/ttyS0. Other serial ports are /dev/tts/1, /dev/tts/2, etc. But ports on the USB bus, multiport cards, etc. have different names.

The common specification for the conventional serial port is RS−232 (or EIA−232). The connector for the serial port is often seen as one or two 9−pin connectors (in some cases 25−pin) on the back of a PC. But the serial port is more than just that. It includes the associated electronics which must produce signals conforming to the EIA−232 specification. See Voltage Waveshapes. One pin is used to send out data bytes and another to receive data bytes. Another pin is a common signal ground. The other "useful" pins are used mainly for signalling purposes with a steady negative voltage meaning "off" and a steady positive voltage meaning "on".

The UART (Universal Asynchronous Receiver−Transmitter) chip does most of the work. Today, the functionality of this chip is usually built into another chip. See What Are UARTs? These have improved over time and old models (several years old) are now obsolete.

The serial port was originally designed for connecting modems but it's used to connect many other devices also such as mice, text−terminals, some printers, etc. to a computer. You just plug these devices into the serial port using the correct cable. Many internal modem cards have a built−in serial port so when you install one inside your PC it's as if you just installed another serial port in your PC.

2. How the Hardware Transfers Bytes

Below is an introduction to the topic, but for a more advanced treatment of it see FIFOs.

2.1 Transmitting

Transmitting is sending bytes out of the serial port away from the computer. Once you understand transmitting, receiving is easy to understand since it's similar. The first explanation given here will be grossly oversimplified. Then more detail will be added in later explanations. When the computer wants to send a byte out the serial port (to the external cable) the CPU sends the byte on the bus inside the computer to the I/O address of the serial port. The serial port takes the byte, and sends it out one bit at a time (a serial bit−stream) on the transmit pin of the serial cable connector. For what a bit (and byte) look like electrically see Voltage Waveshapes.

Here's a replay of the above in a little more detail (but still very incomplete). Most of the work at the serial port is done by the UART chip (or the like). To transmit a byte, the serial device driver program (running on the CPU) sends a byte to the serial port"s I/O address. This byte gets into a 1−byte "transmit shift register" in the serial port. From this shift register bits are taken from the byte one−by−one and sent out bit−by−bit on the serial line. Then when the last bit has been sent and the shift register needs another byte to send it could just ask the CPU to send it another byte. Thus would be simple but it would likely introduce delays since the CPU might not be able to get the byte immediately. After all, the CPU is usually doing other things besides just handling the serial port.

A way to eliminate such delays is to arrange things so that the CPU gets the byte before the shift register needs it and stores it in a serial port buffer (in hardware). Then when the shift register has sent out its byte and needs a new byte immediately, the serial port hardware just transfers the next byte from its own buffer to the shift register. No need to call the CPU to fetch a new byte.

The size of this serial port buffer was originally only one byte, but today it is usually 16 bytes (more in higher priced serial ports). Now there is still the problem of keeping this buffer sufficiently supplied with bytes so that when the shift register needs a byte to transmit it will always find one there (unless there are no more bytes to send). This is done by contacting the CPU using an interrupt.

First we'll explain the case of the old fashioned one−byte buffer, since 16−byte buffers work similarly (but are more complex). When the shift register grabs the byte out of the buffer and the buffer needs another byte, it sends an interrupt to the CPU by putting a voltage on a dedicated wire on the computer bus. Unless the CPU is doing something very important, the interrupt forces it to stop what it was doing and start running a program which will supply another byte to the port's buffer. The purpose of this buffer is to keep an extra byte (waiting to be sent) queued in hardware so that there will be no gaps in the transmission of bytes out the serial port cable.

Once the CPU gets the interrupt, it will know who sent the interrupt since there is a dedicated interrupt wire for each serial port (unless interrupts are shared). Then the CPU will start running the serial device driver which checks registers at I/0 addresses to find out what has happened. It finds out that the serial's transmit buffer is empty and waiting for another byte. So if there are more bytes to send, it sends the next byte to the serial port's I/0 address. This next byte should arrive when the previous byte is still in the transmit shift register and is still being transmitted bit−by−bit.

In review, when a byte has been fully transmitted out the transmit wire of the serial port and the shift register is now empty the following 3 things happen almost simultaneously:

1. The next byte is moved from the transmit buffer into the transmit shift register
2. The transmission of this new byte (bit−by−bit) begins
3. Another interrupt is issued to tell the device driver to send yet another byte to the now empty transmit buffer

Thus we say that the serial port is interrupt driven. Each time the serial port issues an interrupt, the CPU sends it another byte. Once a byte has been sent to the transmit buffer by the CPU, then the CPU is free to pursue some other activity until it gets the next interrupt. The serial port transmits bits at a fixed rate which is selected by the user (or an application program). It's sometimes called the baud rate. The serial port also adds extra bits to each byte (start, stop and perhaps parity bits) so there are often 10 bits sent per byte. At a rate (also called speed) of 19,200 bits per second (bps), there are thus 1,920 bytes/sec (and also 1,920 interrupts/sec).

Doing all this is a lot of work for the CPU. This is true for many reasons. First, just sending one 8−bit byte at a time over a 32−bit data bus (or even 64−bit) is not a very efficient use of bus width. Also, there is a lot of overhead in handing each interrupt. When the interrupt is received, the device driver only knows that something caused an interrupt at the serial port but doesn't know that it's because a character has been sent. The device driver has to make various checks to find out what happened. The same interrupt could mean that a character was received, one of the control lines changed state, etc.

A major improvement has been the enlargement of the buffer size of the serial port from 1−byte to 16−bytes. This means that when the CPU gets an interrupt it gives the serial port up to 16 new bytes to transmit. This is fewer interrupts to service but data must still be transferred one byte at a time over a wide bus. The 16−byte buffer is actually a FIFO (First In First Out) queue and is often called a FIFO. See FIFOs for details about the FIFO along with a repeat of some of the above info.

2.2 Receiving

Receiving bytes by a serial port is similar to sending them only it's in the opposite direction. It's also interrupt driven. For the obsolete type of serial port with 1−byte buffers, when a byte is fully received from the external cable it goes into the 1−byte receive buffer. Then the port gives the CPU an interrupt to tell it to pick up that byte so that the serial port will have room for storing the next byte which is currently being received. For newer serial ports with 16−byte buffers, this interrupt (to fetch the bytes) may be sent after 14 bytes are in the receive buffer. The CPU then stops what it was doing, runs the interrupt service routine, and picks up 14 to 16 bytes from the port. For an interrupt sent when the 14th byte has been received, there could be 16 bytes to get if 2 more bytes have arrived since the interrupt. But if 3 more bytes should arrive (instead of 2), then the 16−byte buffer will overrun. It also may pick up less than 14 bytes by setting it that way or due to timeouts. See FIFOs for more details.

2.3 The Large Serial Buffers

We've talked about small 16−byte serial port hardware buffers but there are also much larger buffers in main memory. When the CPU takes some bytes out of the receive buffer of the hardware, it puts them into a much larger (say 8k−byte) receive buffer in main memory. Then a program that is getting bytes from the serial port takes the bytes it's receiving out of that large buffer (using a "read" statement in the program). A similar situation exists for bytes that are to be transmitted. When the CPU needs to fetch some bytes to be transmitted it takes them out of a large (8k−byte) transmit buffer in main memory and puts them into the small 16−byte transmit buffer in the hardware.

3. Serial Port Basics

You don't have to understand the basics to use the serial port But understanding it may help to determine what is wrong if you run into problems. This section not only presents new topics but also repeats some of what was said in the previous section How the Hardware Transfers Bytes but in greater detail.

3.1 What is a Serial Port ?

Intro to Serial

The UART serial port (or just "serial port for short" is an I/O (Input/Output) device.

An I/O device is just a way to get data into and out of a computer. There are many types of I/O devices such as serial ports, parallel ports, disk drive controllers, ethernet boards, universal serial buses, etc. 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 part of a chip) known as a UART. For more details on UARTs see the section

What Are UARTS? 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.

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 section Pinout and Signals for more details.

RS−232 or EIA−232, etc.

The serial port (not the USB) is usually a RS−232−C, EIA−232−D, or EIA−232−E. These three are almost the same thing. The original RS (Recommended Standard) prefix became EIA (Electronics Industries Association) and later EIA/TIA after EIA merged with TIA (Telecommunications Industries Association). The EIA−232 spec provides also for synchronous (sync) communication but the hardware to support sync is almost always missing on PC's. The RS designation is obsolete but is still widely used. EIA will be used in this howto. Some documents use the full EIA/TIA designation. For info on other (non−EIA−232) serial ports see the section Other Serial Devices (not async EIA−232)

3.2 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 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.

3.3 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 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 which 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 Serial Port Devices /dev/ttyS2, etc. for more details>

3.4 Interrupts

When the serial port receives a number of bytes (may be set to 1, 4, 8, or 14) into its FIFO buffer, it signals the CPU to fetch them by sending an electrical signal known as an interrupt on a certain wire normally used only by that port. Thus the FIFO waits until it has received a number of bytes and then issues an interrupt.

However, 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 from someone typing on a terminal keyboard) there may be an interrupt issued for every byte received. For some UART chips the rule is like this: If 4 bytes in a row could have been received in an interval of time, but none of these 4 show up, then the port gives up waiting for more bytes and issues an interrupt to fetch the bytes currently in the FIFO. Of course, if the FIFO is empty, no interrupt will be issued.

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 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.

There is no Flow Control to prevent this.

Interrupts are also issued when the serial port has just sent out all of its bytes from its small transmit FIFO 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 then 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.

3.5 Data Flow (Speeds)

Data (bytes representing letters, pictures, etc.) flows into and out of your serial port. Flow rates (such as 56k (56000) bits/sec) are (incorrectly) called "speed". But almost everyone says "speed" instead of "flow rate".

It's important to understand that the average speed is often less than the specified speed. Waits (or idle time) result in a lower average speed. These waits may include long waits of perhaps a second due to Flow Control. At the other extreme there may be very short waits (idle time) of several micro−seconds between bytes. If the device on the serial port (such as a modem) can't accept the full serial port speed, then the average speed must be reduced.

3.6 Flow Control

Flow control means the ability to slow down the flow of bytes in a wire. For serial ports this means the ability to stop and then restart the flow without any loss of bytes. Flow control is needed for modems and other hardware to allow a jump in instantaneous flow rates.

Example of Flow Control

For example, consider the case where you connect a 33.6k external modem via a short cable to your serial port. The modem sends and receives bytes over the phone line at 33.6k bits per second (bps). Assume 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 over the short cable is at 115.2k bps. However the flow from your modem out the phone line is 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 soon overrun (run out of free storage space) unless the high 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 buffer. Since nothing is coming into this buffer, the number 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 the short cable (in this case 33.6k) which is significantly less than the "on" flow rate of 115.2k bps. This is "start−stop" flow control.

In the above simple example it was assumed that the modem did no data compression. This could happen when the modem is sending a file which is already compressed and can't be compressed further. Now let's consider the opposite extreme where the modem is compressing the data with a high compression ratio. In such a case the modem might need an input flow rate of say 115.2k bps to provide an output (to the phone line) of 33.6k bps (compressed data). This compression ratio is 3.43 (115.2/33.6). In this case the modem is able to compress the 115.2 bps PC−to−modem flow and send the same data (in compressed form) out the phone line at 33.6bps. There's no need for flow control here so long as the compression ratio remains higher that 3.43. But the compression ratio varies from second to second and if it should drop below 3.43, flow control will be needed

In the above example, the modem was an external modem. But the same situation exists (as of early 2003) for most internal modems. There is still a speed limit on the PC−to−modem speed even though this flow doesn't take place over an external cable. This makes the internal modems compatible with the external modems.

In the above example of flow control, the flow was 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 involves 3 buffers: 1. in the modem 2. in the UART chip (called FIFOs) and 3. in main memory managed by the serial driver. Flow control protects all buffers (except the FIFOs) from overflowing.

Under Linux, the small UART FIFO buffers are not protected by flow control but instead rely on a fast response to the interrupts they issue. Some UART chips can be set to do hardware flow control to protect their FIFOs but Linux (as of early 2003) doesn't seem to support it. FIFO stand for "First In, First Out" which is the way it handles bytes in a queue. All the 3 buffers use the FIFO rule but only the one in the UART is named "FIFO". This is the essence of flow control but there are still some more details.

Symptoms of No Flow Control

Understanding flow−control theory can be of practical use. The symptom of no flow control is that chunks of data missing from files sent without the benefit of flow control. When overflow happens, often hundreds or even thousands of bytes get lost, and all in contiguous chunks.

Hardware vs. Software Flow Control

If feasible, it's best to use "hardware" flow control that uses two dedicated "modem control" wires to send the "stop" and "start" signals. Hardware flow control at the serial port works like this: The two pins, RTS (Request to send) and CTS (Clear to send) are used. When the computer is ready to receive date it asserts RTS by putting a positive voltage on the RTS pin (meaning "Request To Send to me"). When the computer is not able to receive any more bytes, it negates RTS by putting a negative voltage on the pin saying: "stop sending to me". The RTS pin is connected by the serial cable to another pin on the modem, printer, terminal, etc. This other pin's only function is to receive the signal.

For the case of a modem this "other" pin will be the modem's RTS pin. But for a printer, another PC, or a non−modem device, it's usually a CTS pin so a "crossover" or "null modem" cable is required. This cable connects the CTS pin at one end with the RTS pin at the other end (two wires since each end of the cable has a CTS pin). For a modem, a straight−thru cable is used.

For the opposite direction of flow a similar scheme is used. For a modem, the CTS pin is used to send the flow control signal to the CTS pin on the PC. For a non−modem, the RTS pin sends the signal. Thus modems and non−modems have the roles of their RTS and CTS pins interchanged. Some non−modems such as dumb terminals may use other pins for flow control such as the DTR pin instead of RTS.

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.

3.7 Data Flow Path; Buffers

It's been mention that there are 3 buffers for each direction of flow (3 pairs altogether): 16−byte FIFO buffers (in the UART), a pair of larger buffers inside a device connected to the serial port (such as a modem), and a pair of buffers (say 8k) in main memory. When an application program sends bytes to the serial port they first get stashed in the transmit serial port buffer in main memory. The other member of the pair consists of a receive buffer for the opposite direction of byte−flow. Here's an example diagram for the case of browsing the Internet with a browser. Transmit data flow is left to right while receive flow is right to left. There is a separate buffer for each direction of flow.

```
application     8k−byte         16−byte        1k−byte        tele−
BROWSER ------- MEMORY -------- FIFO --------- MODEM -------- phone
program         buffer          buffer         buffer         line
```

For the transmit case, the serial device driver takes out say 16 bytes from this transmit buffer (in main memory), one byte at a time and puts them into the 16−byte transmit buffer in the serial UART for transmission. Once in that transmit buffer, there is no way to stop them from being transmitted. They are then transmitted to the modem or (other device connected to the serial port) 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. At the same time, the bytes stored in the FIFO and MODEM continue to be sent out until these buffers empty.

When the memory buffer 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 since it may switch to running other processes while it's waiting at a flow control stop.

The above was a little oversimplified in three ways. First, some UARTs can do automatic hardware flow control which can stop the transmission out of the FIFO buffers if needed (not yet supported by Linux). Second, while an application process is waiting to write to the transmit buffer, it could possibly perform other tasks. Third, the serial driver (located between the memory buffer and the FIFO) has it's own buffer (in main memory) used to process characters.

3.8 Complex Flow Control Example

For many situations, there is a transmit path involving several links, each with its own flow control. For example, I type at a text−terminal connected to a PC and the PC (under my control) dials out to another computer using a modem. Today, a "text−terminal" is likely to be just another PC emulating a text−terminal. The main (server) PC, in addition to serving my text−terminal, could also have someone else sitting at it doing something else. Note that calling this PC a "server" is not technically correct but it does serve the terminal.

The text−terminal uses a command−line interface with no graphical display. Every letter I type at the text−terminal goes over the serial cable to my main PC and then over the phone line to the computer that I've dialed out to. To dial out, I've used the communication software: "minicom" which runs on my PC.

This sounds like a simple data path. I hit a key and the byte that key generates flows over just two cables (besides the keyboard cable): 1. the cable from my text−terminal to my PC and 2. the telephone line cable to some other computer. Of course, the telephone cable is actually a number of telephone system cables and includes switches and electronics so that a single physical cable can transmit many phone calls. But I can think of it like one cable (or one link).

Now, let's count the number and type of electronic devices each keystroke−byte has to pass thru. The terminal−to−PC cable has a serial port at each end. The telephone cable has both a serial port and a modem at each end. This adds up to 4 serial ports and 2 modems. Since each serial port has 2 buffers, and each modem one buffer, that adds up to 10 buffers. And that's just for one direction of flow. Each byte also must pass thru the minicom software as well.

While there's just 2 cables in the above scenario, if external modems were used there would be an additional cable between each modem and it's serial port. This makes 4 cables in all. Even with internal modems it's like there is a "virtual cable" between the modem and its serial port. On all these 4 links (or cables), flow control takes place.

Now lets consider an example of the operation of flow control. Consider the flow of bytes from the remote computer at the other end of the phone line to the screen on the text−terminal that I'm sitting at. A real text−terminal has a limit to the speed at which bytes can be displayed on its screen and issues a flow control "stop" from time to time to slow down the flow. This "stop" propagates in a direction opposite to the flow of bytes it controls. What happens when such a "stop" is issued? Let's consider a case where the "stop" waits long enough before canceling it with a "start", so that it gets thru to the remote computer at the other end of the phone line. When it gets there it will stop the program at the remote computer which is sending out the bytes.

Let's trace out the flow of this "stop" (which may be "hardware" on some links and "software" on others). First, suppose I'm "capturing" a long file from the remote computer which is being sent simultaneously to both my text−terminal and a to file on my hard−disk. The bytes are coming in faster than the terminal can handle them so it sends a "stop" out its serial port to a serial port on my PC. The device driver detects it and stops sending bytes from the 8k PC serial buffer (in main memory) to the terminal. But minicom still keeps sending out bytes for the terminal into this 8k buffer.

When this 8k transmit buffer (on the first serial port) is full, minicom must stop writing to it. Minicom stops and waits. But this also causes minicom to stop reading from the 8k receive buffer on the 2nd serial port connected to the modem. Flow from the modem continues until this 8k buffer too fills up and sends a different "stop" to the modem. Now the modem's buffer ceases to send to the serial port and also fills up. The modem (assuming error correction is enabled) sends a "stop signal" to the other modem at the remote computer. This modem stops sending bytes out of its buffer and when its buffer gets fill, another stop signal is sent to the serial port of the remote computer. At the remote computer, the 8−k (or whatever) buffer fills up and the program at the remote computer can't write to it anymore and thus temporarily halts.

Thus a stop signal from a text terminal has halted a program on a remote computer computer. What a long sequence of events! Note that the stop signal passed thru 4 serial ports, 2 modems, and one application program (minicom). Each serial port has 2 buffers (in one direction of flow): the 8k one and the hardware 16−byte one. The application program may have a buffer in its C_code. This adds up to 11 different buffers the data is passing thru. Note that the small serial hardware buffers do not participate directly in flow control. Also note that the two buffers associated with the text−terminal's serial port are going to be dumping their contents to the screen during this flow control halt. This leaves 9 other buffers that may be getting filled up during the flow control halt.

If the terminal speed limitation is the bottleneck in the flow from the remote computer to the terminal, then its flow control "stop" is actually stopping the program that is sending from the remote computer as explained above. But you may ask: How can a "stop" last so long that 9 buffers (some of them large) all get filled up? It seldom happens, but it can actually happen this way if all the buffers were near their upper limits when the terminal sent out the "stop".

But if you were to run a simulation on this you would discover that it's usually more complicated than this. At an instant of time some links are flowing and others are stopped (due to flow control). A "stop" from the terminal seldom propagates back to the remote computer neatly as described above. It may take a few "stops" from the terminal to result in one "stop" at the remote computer, etc. To understand what is going on you really need to observe a simulation which can be done for a simple case with coins on a table. Use only a few buffers and set the upper level for each buffer at only a few coins.

Does one really need to understand all this? Well, understanding this explained to me why capturing text from a remote computer was loosing text. The situation was exactly the above example but modem−to−modem flow control was disabled. Chunks of captured text that were supposed to also get to my hard−disk never got there because of an overflow at my modem buffer due to flow control "stops" from the terminal. Even though the remote computer had a flow path to the hard−disk without bottlenecks, the same flow also went to a terminal which issued flow control "stops" with disastrous results for the branch of the flow going to a file on the hard−disk. The flow to the hard−disk passed thru my modem and since the overflow happened at the modem, bytes intended for the hard−disk were lost.

3.9 Serial 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. >From kernel 2.2 on, this module will normally get loaded automatically if it's needed. 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 (and sometimes still is). If it's built−in don't let the serial module load or else you will have two serial drivers running at the same time. With 2 drivers there are all sorts of errors including a possible "I/O error" when attempting to open a serial port. Use "lsmod" to see if the module is loaded.

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 similar to the first but with the correct IRQ, etc. See Serial Module See What is Setserial for more info on `setserial`.

---

4. Is the Serial Port Obsolete?

4.1 Introduction

The answer is yes, but ... The serial port is somewhat obsolete but it's still needed, especially for Linux. The serial port has many shortcomings but almost all new PC's seem to come with them. Linux supports ordinary telephone modems only if they work thru a serial port (although the port may be built into the modem).

The serial port must pass data between the computer and the external cable. Thus it has two interfaces: the serial−port−to cable and the serial−port−to−computer−bus. Both of these interfaces are slow. First we'll consider the interface via external cable to the outside world.

4.2 EIA−232 Cable Is Low Speed & Short Distance

The conventional EIA−232 serial port is inherently low speed and is severely limited in distance. Ads often read "high speed" but it can only work at "high speed" over very short distances such as to a modem located right next to the computer. Compared to a network card, even this "high speed" is actually low speed. All of the EIA−232 serial cable wires use a common ground return wire so that twisted−pair technology (needed for high speeds) can't be used without additional hardware. More modern interfaces for serial ports exist but they are not standard on PC's like the EIA−232 is. See Successors to EIA−232. Some multiport serial cards support them.

It is somewhat tragic that the RS−232 standard from 1969 did not use twisted pair technology which could operate about a hundred times faster. Twisted pairs have been used in telephone cables since the late 1800's. In 1888 (over 110 years ago) the "Cable Conference" reported its support of twisted−pair (for telephone systems) and pointed out its advantages. But over 80 years after this approval by the "Cable Conference", RS−232 failed to utilize it. Since RS−232 was originally designed for connecting a terminal to a low speed modem located nearby, the need for high speed and longer distance transmission was apparently not recognized.

4.3 Inefficient Interface to the Computer

To communicate with the computer, any I/O device needs to have an address so that the computer can write to it and read from it. For this purpose many I/O devices (such as serial ports) use a special type of address

known as an I/O addresses (sometimes called an I/O port). The I/O address of a certain device (such as ttys/2) will be a range of addresses. The lower address in this range is the base address. "address" usually means just the "base address".

Instead of using I/O addresses, some I/O devices read and write directly from/to main memory. This provides more bandwidth since the conventional serial I/O system only moves a byte at a time. There are various ways to read/write directly to main memory. One way is called shared memory I/O (where the shared memory is usually on the same card as the I/O device). Other methods are DMA (direct memory access) on the PCI or ISA bus. Strictly speaking, for the PCI bus this is "bus mastering" and not DMA but it's like DMA and a lot faster than true DMA on the old ISA bus.

DMA is a lot faster than the serial−port−to computer−bus interface. Also, the computer bus is capable of 4 (or 8) byte transfers but the serial port only transfers a byte at a time. Interrupts to initiate a series of such single−byte transfers often happen only every 14 bytes.

---

5. Multiport Serial Boards/Cards/Adapters

5.1 Intro to Multiport Serial

Multiport serial cards install in slots in a PC on the ISA or PCI bus. Instead of being called "... cards" they are also called "... adapters" or "... boards". Each such card provides you with many serial ports. Today they are commonly used for the control of external devices (including automation for both industry and the home). They can connect to computer servers for the purpose of monitoring/controlling the server from a remote location. They were once mainly used for connecting up many dumb terminals and/or modems to serial ports. Today, use of dumb terminals has declined, and several modems (or digital modems) can now be built into an internal card. So multiport serial cards are not as significant as they once were.

Each multiport card has a number of external connecters (DB−25 or RJ45) so that one may connect up a number of devices (modems, terminals, etc.). Each such physical device would then be connected to its own serial port. Since the space on the external−facing part of the card is limited there is often not enough room for all the serial port connectors. To solve this problem, the connectors may be on the ends of cables which come out (externally) from the card (octopus cable). Or they may be on an external box (possibly rack mountable) which is connected by a cable to a multiport card.

5.2 Modem Limitations

For a modem to transmit at nearly 56k requires that it be a special digital modem and have a digital connection to a digital phone line (such as a T1 line). Modem banks that connect to multiport cards do exist, and some have a card that can access multiplexed digital phone lines. Thus one can use a multiport card with a few 56k digital modems.

For both analog and digital modem there is one modem on each serial port so there needs to be an external cable (modem bank to multiport) for each modem. This can lead to a large number of cables. So it's less clutter (and cheaper) to use internal modems without a multiport card. It's somewhat analogous to the lower cost of an internal modem for a desktop PC as compared to the higher cost (and more cabling) for an external modem. See Modem−HOWTO: Modem Pools, Digital Modems.

5.3 Dumb vs. Smart Cards

Dumb multiport cards are not too much different than ordinary serial ports. They are interrupt driven and the CPU of the computer does most all the work servicing them. They usually have a system of sharing a single interrupt for all the ports. This doesn't decrease the load on the CPU since the single interrupt will be sent to the CPU each time any one port needs servicing. Such devices usually require special drivers that you must either compile into the kernel or use as a module. In rare cases thy activate by putting a #define into the source code (or the like).

Smart boards may use ordinary UARTs but handle most interrupts from the UARTs internally within the board. This frees the CPU from the burden of handling all these interrupts. The board may save up bytes in its large internal FIFOs and transfer perhaps 1k bytes at a time to the serial buffer in main memory. It may use the full bus width of 32 bits for making data transfers to main memory (instead of transferring only 8−bit bytes like dumb serial cards do). Not all "smart" boards are equally efficient. Many boards today are Plug−and−Play.

5.4 Getting/Enabling a Driver

Introduction

For a multiport board to work, a special driver for it must be used. This driver may either be built into the kernel source code or supplied as a module. Support for dumb boards is likely to the built into the kernel while smart boards usually need a module.

If you need driver built into the kernel (mostly dumb boards)

A pre−compiled kernel may not have multiport support built in. So you may need to compile it yourself. In kernel 2.4 you should select "CONFIG_SERIAL_EXTENDED" when configuring the kernel (just before you compile). If you select this there will be still more choices presented to you. Even after you do this you may need to edit the resulting source code a little (depending on the card).

If you use a module (mostly for smart boards)

A pre−compiled kernel may come with a pre−compiled module for the board so that you don't need to recompile the kernel. This module must be loaded in order to use it, but the kernel may automatically do this for you if a program is trying to use a device on the smart board (provided there exists a table showing which module to load for the device). This table may be in /etc/modules.conf and/or be internal to the kernel. Also certain parameters may need to be passed to the driver (via lilo's append command or via /etc/modules.conf). For kernel 2.4 the modules should be found in `/lib/modules/.../kernel/drivers/char`.

Getting info on multiport boards

The board's manufacturer should have info on their website. Unfortunately, info for old boards is sometimes not there but might be found somewhere else on the Internet (including discussion groups). You might also want to look at the kernel documentation in /usr/share/kernel−doc... For configuring the kernel or modules prior to compiling see: Configure.help and search for "serial", etc. There are also kernel documentation files for certain boards including computone, hayes−esp, moxa−smartio, riscom8, specialix, stallion, and sx (specialix).

5.5 Multiport Devices in the /dev Directory,

The serial ports your multiport board uses depends on what kind of board you have. Some have their own device names like /dev/ttyE27 (Stallion) or /dev/ttyD2 (Digiboard), etc. For various other brands, see see devices.txt in the kernel documentation. Some use the standard names like /dev/ttyS14 (/dev/tts/14) and may be found in configuration files that used as arguments to `setserial`. Such files may be included in a setserial or serial package.

For the device file system (devfs), for example, /dev/ttyF9 becomes /dev/ttf/9, or in a later version /dev/tts/F9. Substitute for F (or f) whatever letter(s) you multiport board uses for this purpose. Your multiport driver is supposed to create a devfs name similar to the above and put it into the /dev directory

5.6 Making Legacy Multiport Devices in the /dev Directory

If you're using the device file system (devfs), then the device driver should create the device name and put it in the /dev directory. Otherwise for a legacy (non−devfs), an installation script may do this for you. But if not, here's some examples of how to create a device name in the /dev directory.

For the legacy names and numbers of other types of serial ports other than ttyS.. See devices.txt in the kernel documentation. Either use the `mknod` command, or the `MAKEDEV` script. Typing "man makedev" may show instructions on using it.

Using the `MAKEDEV` script, you would first become the superuser (root) and type (for example) either:

```
linux# MAKEDEV ttyS17
```

Or if the above doesn't work cd to /dev before giving the above command>. Substitute whatever your port is for ttyS17.

Using `mknod` is a more complicated option since you need to know the major and minor device numbers. These numbers are in the "devices" file in the kernel documentation. For ttyS serial ports the minor number is: 64 + port number (=81 for the example below). Note the "major" number is always 4 for ttyS devices (and 5 for the obsolete cua devices). So, if you wanted to create a device for `ttyS17` using `mknod`, you would type:

```
linux# mknod −m 666 /dev/ttyS17 c 4 81
```

5.7 Standard PC Serial Cards

In olden days, PCs came with a serial card installed. Later on, the serial function was put on the hard−drive interface card. Today, one or two serial ports are usually built into the motherboard. Most of them (as of 2002) use a 16550 but some use 16650 (32−byte FIFOs). But one may still buy the individual PC serial cards if they need 1−4 more serial ports. These are for tts/0−tts/3 (COM1 − COM4). They can be used to connect external serial devices (modems, serial mice, etc...). Only a tiny percentage of retail computer stores carry such cards. But one can purchase them on the Internet. Before getting a PCI one, make sure Linux supports it.

Here's a list of a few popular brands:

- Byte Runner (may order directly, shows prices) http://www.byterunner.com
- SIIG http://www.siig.com/products/io/

- Dolphin http://www.dolphinfast.com/sersol.html

Note: due to address conflicts, you may not be able to use COM4 and IBM8514 video card (or some others) simultaneously. See Avoiding IO Address Conflicts with Certain Video Boards

5.8 Dumb Multiport Serial Boards (with standard UART chips)

They are also called "serial adapters". Each port has its own address. They often have a special method of sharing interrupts which requires that you compile support for them into the kernel.

* => The file that ran setserial in Debian shows some details of configuring
# => See note below for this board

- AST FourPort and clones (4 ports) * #
- Accent Async−4 (4 ports) *
- Arnet Multiport−8 (8 ports)
- Bell Technologies HUB6 (6 ports)
- Boca BB−1004 (4 ports), BB−1008 (8 ports), BB−2016 (16 ports; See the Boca mini−howto revised in 2001) * #
- Boca IOAT66 or? ATIO66 (6 ports, Linux doesn't support its IRQ sharing ?? Uses odd−ball 10−cond RJ45−like connectors)
- Boca 2by4 (4 serial ports, 2 parallel ports)
- Byte Runner http://www.byterunner.com
- Computone ValuePort V4−ISA (AST FourPort compatible) *
- Digi PC/8 (8 ports) #
- Dolphin http://www.dolphinfast.com/sersol/
- Globetek http://www.globetek.com/
- GTEK BBS−550 (8 ports; See the mini−howto)
- Hayes ESP (after kernel 2.1.15)
- HUB−6 See Bell Technologies.
- Longshine LCS−8880, Longshine LCS−8880+ (AST FourPort compatible) *
- Moxa C104, Moxa C104+ (AST FourPort compatible) *
- NI−SERIAL by National Instruments
- NetBus (2 ports) http://www.netbus.com using patch from http://lists.insecure.org/linux−kernel/2001/Feb/2809.html
- PC−COMM (4 ports)
- Sealevel Systems COMM−2 (2 ports), COMM−4 (4 ports) and COMM−8 (8 ports)
- SIIG I/O Expander 2S IO1812 (4 ports) #
- STB−4COM (4 ports)
- Twincom ACI/550
- Usenet Serial Board II (4 ports) *
- VScom (uses same driver as ByteRunner)

In general, Linux will support any serial board which uses a 8250, 16450, 16550, 16550A, 16650, etc. UART. See the latest man page for "setserial" for a more complete list.

Notes:

AST Fourport: You might need to specify `skip_test` in `rc.serial`.

BB−1004 and BB−1008 do not support DCD and RI lines, and thus are not usable for dialin modems. They will work fine for all other purposes.

Digi PC/8 Interrupt Status Register is at 0x140.

SIIG IO1812 manual for the listing for COM5−COM8 is wrong. They should be COM5=0x250, COM6=0x258, COM7=0x260, and COM8=0x268.

5.9 Intelligent Multiport Serial Boards

Make sure that a Linux−compatible driver is available and read the information that comes with it. These boards use special devices (in the /dev directory), and not the standard tts ones. This information varies depending on your hardware. If you have updated info which should be shown here please email it to me.

Names of Linux driver modules are *.o but these may not work for all models shown. See Modules (mostly for smart boards) The needed module may have been supplied with your Linux distribution. Also, parameters (such as the io and irq often need to be given to the module so you need to find instructions on this (possibly in the source code tree).

There are many different brands, each of which often offers many different cards. No attempt is currently being made to list all the cards here (and many listed are obsolete). But all major brands and websites should be shown here so it something is missing let me know. Go the the webpage shown for more information. These websites often also have info (ads) on related hardware such as modem pools, remote access servers (RASs), and terminal servers. Where there is no webpage, the cards are likely obsolete. If you would like to put together a better list, let me know.

- Chase Research now Perle Systems Ltd (UK based, ISA/PCI cards)
  webpage: `http://www.perle.com`
  driver status: included in kernel 2.4+ for PCI only; otherwise supported by Perle
  driver location: http://www.perle.com/downloads/multi_port.html
- Comtrol RocketPort (36MHz ASIC; 4, 8, 16, 32, up to 128 ports)
  webpage: `http://www.comtrol.com`
  driver status: supported by Comtrol. rocket.o
  driver location: `ftp://tsx−11.mit.edu/pub/linux/packages/comtrol`
- Computone IntelliPort II (ISA, PCI and EISA busses up to 64 ports)
  webpage: http://www.computone.com
  driver location: old patch at http://www.wittsend.com/computone/linux−2.2.10−ctone.patch.gz
  mailing list: mailto:majordomo@lazuli.wittsend.com with "subscribe linux−computone" in body
  note: Old ATvantage and Intelliport cards are not supported by Computone
- Connecttech
  website: `http://www.connecttech.com/`
  driver location: ftp://ftp.connecttech.com/pub/linux/
- Cyclades
  Cyclom−Y (Cirrus Logic CD1400 UARTs; 8 − 32 ports),
  Cyclom−Z (MIPS R3000; 8 − 64 ports)
  website: `http://www.cyclades.com/products/svrbas/zseries.php`
  driver status: supported by Cyclades
  driver location: `ftp://ftp.cyclades.com/pub/cyclades` and included in Linux kernel since version 1.1.75: cyclades.o

- Decision PCCOM (2–8 ports; ISA and PCI; aka PC COM)
  ISA:
  contact: mailto:info@cendio.se
  driver location: (dead link) ftp://ftp.cendio.se/pub/pccom8
  PCI:
  drivers: http://www.decision.com.tw
  driver status: Support in serial driver 5.03. For an earlier driver, there exists a patch for kernel 2.2.16 at http://www.qualica.com/serial/ and for kernels 2.2.14–2.2.17 at http://www.pccompci.com/mains/installing_pci_linux1.html
- Digi PC/Xi (12.5MHz 80186; 4, 8, or 16 ports),
  PC/Xe (12.5/16MHz 80186; 2, 4, or 8 ports),
  PC/Xr (16MHz IDT3041; 4 or 8 ports),
  PC/Xem (20MHz IDT3051; 8 – 64 ports)
  website: http://www.dgii.com
  driver status: supported by Digi
  driver location: ftp://ftp.dgii.com/drivers/linux and included in Linux kernel since version 2.0. epca.o
- Digi COM/Xi (10MHz 80188; 4 or 8 ports)
  contact: Simon Park, si@wimpol.demon.co.uk
  driver status: ?
  note: Simon is often away from email for months at a time due to his job. Mark Hatle, mailto:fray@krypton.mankato.msus.edu has graciously volunteered to make the driver available if you need it. Mark is not maintaining or supporting the driver.
- Equinox SuperSerial Technology (30MHz ASIC; 2 – 128 ports)
  website: http://www.equinox.com
  driver status: supported by Equinox
  driver location: ftp://ftp.equinox.com/library/sst
- Globetek
  website: http://www.globetek.com/products.shtml
  driver location: http://www.globetek.com/media/files/linux.tar.gz
- GTEK Cyclone (16C654 UARTs; 6, 16 and 32 ports),
  SmartCard (24MHz Dallas DS80C320; 8 ports),
  BlackBoard–8A (16C654 UARTs; 8 ports),
  PCSS (15/24MHz 8032; 8 ports)
  website: http://www.gtek.com
  driver status: supported by GTEK
  driver location: ftp://ftp.gtek.com/pub
- Hayes ESP (COM–bic; 1 – 8 ports)
  website: http://www.nyx.net/~arobinso
  driver status: Supported by Linux kernel (1998) since v. 2.1.15. esp.o. Setserial 2.15+ supports. Also supported by author
  driver location: http://www.nyx.net/~arobinso
- Intelligent Serial Interface by Multi–Tech Systems
  PCI: 4 or 8 port. ISA 8 port. DTE speed 460.8k
  webpage: http://www.multitech.com/products/
- Maxpeed SS (Toshiba; 4, 8 and 16 ports)
  website: http://www.maxpeed.com
  driver status: supported by Maxpeed
  driver location: ftp://maxpeed.com/pub/ss
- Microgate SyncLink ISA and PCI high speed multiprotocol serial. Intended for synchronous HDLC.
  website: http://ww/microgate.com/products/sllinux/hdlcapi.htm

driver status: supported by Microgate: synclink.o

- Moxa C218 (12MHz 80286; 8 ports),
  Moxa C320 (40MHz TMS320; 8 – 32 ports)
  website: `http://www.moxa.com`
  driver status: supported by Moxa
  driver locations: `http://www.moxa.com/support/download/download.php3>`
  `ftp://ftp.moxa.com/drivers/linux` (also from Taiwan at www.moxa.com.tw/...) where ... is the same as above)
- SDL RISCom/8 (Cirrus Logic CD180; 8 ports)
  website: `http://www.sdlcomm.com`
  driver status: supported by SDL
  driver location: `ftp://ftp.sdlcomm.com/pub/drivers`
- Specialix SX (25MHz T225; 8? – 32 ports),
  SIO/XIO (20 MHz Zilog Z280; 4 – 32 ports)
  webpage: Old link is broken. Out of business?
  driver status: Was supported by Specialix
  driver location: http://www.BitWizard.nl/specialix/
  old driver location: ftp://metalab.unc.edu/pub/Linux/kernel/patches/serial
- Stallion EasyIO–4 (4 ports), EasyIO–8 (8 ports), and
  EasyConnection (8 – 32 ports) – each with Cirrus Logic CD1400 UARTs,
  Stallion (8MHz 80186 CPU; 8 or 16 ports),
  Brumby (10/12 MHz 80186 CPU; 4, 8 or 16 ports),
  ONboard (16MHz 80186 CPU; 4, 8, 12, 16 or 32 ports),
  EasyConnection 8/64 (25MHz 80186 CPU; 8 – 64 ports)
  contact: `sales@stallion.com` or `http://www.stallion.com`
  driver status: supported by Stallion
  driver location: `ftp://ftp.stallion.com/drivers/ata5/Linux` and included in linux kernel since 1.3.27 moved; it's now at
- System Base website: http://www.sysbas.com/

A review of Comtrol, Cyclades, Digi, and Stallion products was printed in the June 1995 issue of the *Linux Journal*. The article is available at `http://www.linuxjournal.com/article.php?sid=1097 name="http://www.ssc.com/lj/issue14">`.

Besides the listing of various brands of multiports found above in this HOWTO there is Gary's Encyclopedia – Serial Cards. It's not as complete, but may have some different links.

5.10 Unsupported Multiport Boards

The following boards don't mention any Linux support as of 1 Jan. 2000. Let me know if this changes.

- Aurora (PCI only) www.auroratech.com

6. Servers for Serial Ports

A computer that has many serial ports (with many serial cables connected to it) is often called a server. Of course, most servers serve other functions besides just serving serial ports, and many do not serve serial ports at all (although they likely have a serial port on them). For example, a "serial server" may have serial cables, each of which runs to a different (non–serial) server. The serial server (perhaps called a "console server")

controls, via a console, all the other servers. The console may be physically located remote from the serial server, communicating with the server over a network.

There are two basic types of serial servers. One type is just an ordinary computer (perhaps rack mounted) that uses multiport cards on a PCI bus (or the like). The other type is a proprietary server that is a dedicated computer that serves a special purpose. Servers of both types may be called: serial servers, console servers, print servers, or terminal servers. They are not the same.

The terminal server was originally designed to provide many serial ports, each connected to a dumb text−terminal. Today, a terminal server often connects to graphic terminals over a fast network and doesn't use serial ports since they are too slow. One network cable takes the place of many serial cables and each graphic terminal uses far more bandwidth than the text−terminals did. However, graphic terminals may be run in text mode to reduce the bandwidth required. A more detailed discussion of terminal servers (serial port) is in Text−Terminal−HOWTO. For networked terminal servers (not serial port) see Linux Terminal Server Project (LTSP)

(To−do: Discuss other types of serial servers, but the author knows little about them.)

7. Configuring Overview

Configuring of the serial port should be done automatically, both the serial driver software and by your application software. But sometimes it isn't and you thus need to do it yourself. Or perhaps you need to configure it in a special way, etc. This HOWTO only covers configuration of the serial port itself and not the configuring of any devices attached to the port (such as a modem).

The first part (locating the hardware or low−level configuring) is assigning each port 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" program is sometimes used to tell the driver. PnP methods, jumpers, etc, are used to set the I0 and IRQ in 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. See Locating the Serial Port: IO address IRQs What is Setserial

The second part (high−level configuring) is assigning it a speed (such as 115.2k 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 your computer). However you will need to tell these programs what speed you want, etc. by using a menu or a configuration file. This high−level configuring may also be done manually with the `stty` program. `stty` is also useful to view the current status if you're having problems. See the section Stty

8. Locating the Serial Port: IO address, IRQs

8.1 IO & IRQ Overview

For a serial port to work properly, it must have both an IRQ and an IO address. Without an IO address, it can't be located and will not work at all. Without an IRQ it will need to use inefficient polling methods for which one must set the IRQ to 0 in the serial driver. So every serial port needs an IO address and IRQ. In olden days this was set by jumpers on a serial port card. Today it's set by digital signals sent to the hardware and this is part of "Plug−and−Play (PnP). It all should get configured automatically so that you only need to read this if you're having problems or if you want to understand how it works.

The driver must also know both the IO address and IRQ so that it can talk to the port chip. Modern serial port drivers (kernel 2.4) try to determine this by PnP methods so one doesn't normally need to tell the driver (by using "setserial"). The modern driver might also enable the serial port and set an IO address or IRQ in the hardware. But unfortunately, there is some PCI serial port hardware that the driver doesn't recognize so you might need to enable the port yourself. See PCI: Enabling a disabled port

The driver also probes likely ISA serial port addresses to see if there are any serial ports there. This works for the case of jumpers and sometimes works for a ISA PnP port when the driver doesn't do ISA PnP (prior to kernel 2.4).

Locating the serial port by giving it an IRQ and IO address is low−level configuring. It's often automatically done by the serial driver but sometimes you have to do it yourself. What follows repeats what was said above but in more detail.

The low−level configuring consists of assigning an IO address, IRQ, and name (such as ttyS2 = tts/2). This IO−IRQ pair must be set in both the hardware and told to the serial driver. And the driver needs to call this pair a name (such as ttyS2). We could call this "io−irq" configuring for short. The "setserial" program is one way to tell the driver. The other way is for the driver to use PnP methods to determine/set the IO/IRQ and then remember what it did. For jumpers there is no PnP but the driver might detect the port if the jumpers are set to the usual I0/IRQ. If you need to configure but don't understand certain details it's easy to get into trouble.

When Linux starts, some effort is made to detect and configure (low−level) serial ports. Exactly what happens depends on your BIOS, hardware, Linux distribution, kernel version, etc. If the serial ports work OK, there may be no need for you to do any more low−level configuring.

If you're having problems with the serial ports, then you may need to do low−level configuring. If you have kernel 2.2 or lower, then you need to do it if you:

- Plan to use more than 2 ISA serial ports
- Are installing a new serial port (such as an internal modem)

Starting with kernel 2.2 you may be able to use more that 2 serial ports without doing any low−level configuring by sharing interrupts. All PCI ports should support this but for ISA it only works for some hardware. It may be just as easy to give each port a unique interrupt if they is available. See Interrupt sharing and Kernels 2.2+

The low−level configuring (setting the IRQ and IO address) seems to cause people more trouble than the high−level stuff, although for many it's fully automatic and there is no configuring to be done. Until the serial driver knows the correct IRQ and IO address, the port will not usually not work at all. Also, PnP ports can be disabled so that they don't have any IO address and thus can't be used. However PnP tools such as lspci should be able to find them. Applications, and utilities such as "setserial" and "scanport" (Debian only ??) only probe I0 addresses, don't use PnP tools, and thus can't detect disabled ports.

Even if an ISA port can be found by the probing done by the serial driver it may work extremely slow if the IRQ is wrong. See Extremely Slow: Text appears on the screen slowly after long delays. PCI ports are less likely to get the IRQ wrong.

In the Wintel world, the IO address and IRQ are called "resources" and we are thus configuring certain resources. But there are many other types of "resources" so the term has many other meanings. In summary, the low−level configuring consists of enabling the device, giving it a name (ttyS2 for example) and putting two values (an IRQ number and IO address) into two places:

1. the device driver (done by PnP or "`setserial`")
2. configuration registers of the serial port hardware itself (done by PnP or jumpers)

You may watch the start−up (= boot−time) messages. They are usually correct. But if you're having problems, your serial port may not show up at all or if you do see a message from "setserial" it may not show the true configuration of the hardware (and it is not necessarily supposed to). See I/O Address & IRQ: Boot−time messages.

8.2 PCI Bus Support

Introduction

If you have kernel 2.4, then there should be support for PnP (either built−in or by modules). Some PCI serial ports can be automatically detected and low−level configured by the serial driver. Others may not be.

While kernel 2.2 supported PCI in general, it had no support for PCI serial ports (although some people got them working anyway). The 2.4 serial driver will read the id number digitally stored in the serial hardware to determine how to support it (if it knows how). It should assign an I/O address to it, determine it's IRQ, etc. So you don't need to use "setserial" for it.

There is a possible problem if you don't use the device filesystem. The driver may assign the port to say "`ttyS04" per a boot−time message (use dmesg to see it). But if you don't have a "file"` dev/ttyS4 then the port will not work. So you will then need to create it, using `cd /dev` and then `./MAKEDEV ttyS4`
For the device filesystem, the driver should create the device `tts/1`

More info on PCI

PCI ports are not well standardized. Some use main memory for communication with the PC. Some require special enabling of the IRQ. The output of "lspci −vv" can help determine if one can be supported. If you see a 4−digit IO port, the port might work by just telling "setserial" the IO port and the IRQ. For example, if lspci shows IRQ 10, I/O at 0xecb8 and you decide to name it ttyS2 then the command is:

setserial /dev/ttyS2 irq 10 port 0xecb8 autoconfig

Note that the boot−time message "Probing PCI hardware" means reading the PnP configuration registers in the PCI cards which reveals the IO addresses and IRQs. This is different that the probing of IO addresses by the serial driver which means reading certain IO addresses to see if what's read looks like there's a serial port at that address.

8.3 Common mistakes made re low−level configuring

Here are some common mistakes people make:

- setserial command: They run it (without the "autoconfig" and auto_irq options) and think it has checked the hardware to see if what it shows is correct (it hasn't).
- setserial messages: They see them displayed on the screen at boot−time (or by giving the setserial command) and erroneously think that the result always 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 and /proc/tty/driver/serial: People think this shows the actual hardware configuration when it only shows about the same info (possibly erroneous) as setserial.

8.4 IRQ & IO Address Must be Correct

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 usually sets and shows.). 2. What is actually set in the hardware. Both 1. and 2. above should be the same. If they're not it spells trouble since the driver has incorrect info on the physical serial port. In some cases the hardware is disabled so it has no IO address or IRQ.

If the driver has the wrong IO address it will try to send data to a non−existing serial port −−or even worse, to some other device. If it has the wrong IRQ the driver will not get interrupt service requests from the serial port, resulting in a very slow or no response. See Extremely 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 how they are set in both the driver and the hardware.

8.5 What is the IO Address and IRQ per the driver ?

Introduction

What the driver thinks is not necessarily how the hardware is actually set. If everything works OK then what the driver thinks is likely correct (set in the hardware) and you don't need to investigate (unless you're curious or want to become a guru). Ways to determine what the driver thinks include: boot−time messages I/O Address & IRQ: Boot−time messages, the /proc directory "files" The /proc directory and setserial, and the "setserial" command.

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). These BIOS messages may be frozen by pressing the Pause key. Use Shift−PageUp to scroll back to the messages after they have flashed by. Shift−PageDown will scroll in the opposite direction. The `dmesg` command (or looking at logs in /var/log) will show some of the messages but they seem to miss important ones from setserial. 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 setserial shows you  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. dmesg may not display the second message. In most cases the second 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 could be 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 by default 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 port addresses but it doesn't probe for IRQs. 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` and `auto_irq` options but this isn't guaranteed to be correct either.

The data shown by the BIOS messages (which you see at first before Linux is booted) is what is initially set in the hardware. If your serial port is Plug−and−Play (PnP) then it's possible that "isapnp" or "setpci" 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 or setpci). 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, they can only be found by PnP software unless the IRQ and IO has been set inside the hardware by Plug−and−Play software. Prior to kernel 2.4 this was a common reason why the start−up messages did not show a serial port that physically exists. A PnP BIOS may automatically low−level configure them. PnP configuring will be explained later.

The /proc directory and setserial

Type "setserial −g /dev/ttyS*". There are some other ways to find this info by looking at "files" in the /proc directory. Be warned that there is no guarantee that the same is set in the hardware.

`/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). It shows how many interrupts have actually be issued. `/proc/tty/driver/serial` shows much of the above, plus the number of bytes that have been received and sent (even if the device is not now open).

Note that for the IO addresses and IRQ assignments, you are only seeing what the driver thinks and not necessarily what is actually set in the hardware. The data on the actual number of interrupts issued and bytes processed is real however. If you see a large number of interrupts and/or bytes then it probably means that the device is (or was) working. But the interrupts might be from another device. If there are no bytes received (rx:0) but bytes were transmitted (tx:3749 for example), then only one direction of flow is working (or being utilized).

Sometimes a showing of just a few interrupts doesn't mean that the 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. 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.

8.6 What is the IO Address & IRQ of my Serial Port Hardware?

Introduction

If it's PCI or ISA PnP then what's set in the hardware has been done by PnP methods. Even if nothing has been set or the port disabled, PnP ports may still be found by using "lspci −v" or "isapnp −−dumpregs". Ports disabled by jumpers (or hardware failures) are completely lost. See ISA PnP ports, PCI: What IOs and IRQs have been set?, PCI: Enabling a disabled port

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. That's why a PnP port is more likely to be found in a disabled state than an old non−PnP one.

PCI: What IOs and IRQs have been set?

For PCI, the BIOS almost always sets the IRQ and may set the IO address as well. To see how it's set use "lspci −vv" (best) or look in /proc/bus/pci (or for kernels <2.2 /proc/pci). The modem's serial port is often called a "Communication controller". Look for this. If lspci shows "I/O ports at ... [disabled]" then the serial port is disabled and the hardware has no IO address so it's lost and can't be used. See PCI: Enabling a disabled port for how to enable it.

If more than one IO address is shown, the first one is more likely to be it. You can't change the IRQ (at least not with "setpci") This is because if one writes an IRQ it it's hardware register no action is taken on it. It's the BIOS that should actually set up the IRQs and then write the correct value to this register for lspci to view. If you must, change the IO address with "setpci" by changing the BASE_ADDRESS_0 or the like.

PCI: Enabling a disabled port

If the port communicates via an IO address then "lspci −vv" should show "Control: I/O+ ..." with + meaning that the IO address is enabled. If it shows "I/O−" (and "I/O ports at ... [disabled]") then you may need to use the setpci command to enable it. For example "setpci −d 151f:000 command=101". 151f is the vendor id, and 000 is the device id both obtained from "lspci −n −v" or from /proc/bus/pci or from "scanpci −v". The "command=101" means that 101 is put into the command register which is the same as the "Control" register displayed by "lspci". The 101h sets two bits: the 1 sets I/O to + and the 100 part keeps SERR# set to +. In this case only the SERR# bit of the Control register was initially observed to be + when the lspci command was run. So we kept it enabled to + by setting bit 8 (where bit 0 is I/O) to 1 by the first 1 in 101. Some serial cards don't use SERR# so if you see SERR#− then there's no need to enable it so then use: command=1. Then you'll need to set up "setserial" to tell the driver the IO and IRQ.

Bit 8 is actually the 9th bit since we started counting bits from 0. Don't be alarmed that lspci shows a lot of − signs showing that the card doesn't have many features available (or enabled). Serial ports are relatively slow and don't need these features.

Another way to enable it is to let the BIOS do it by telling the BIOS that you don't have a plug−and−play operating system. Then the BIOS should enable it when you start your PC. If you have MS Windows9x on the same PC then doing this might cause problems with Windows (see Plug−and−Play−HOWTO).

ISA PnP ports

For an ISA Plug−and−Play (PnP) port 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. It should also find an ISA PnP port that is disabled. The address it "trys" is not the device's IO address, but a special address used for communicating with PnP cards.

Finding a port that is not disabled (ISA, PCI, PnP, non−PnP)

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 setpci can. Starting with kernel 2.4, the serial driver can make such changes for many (but not all) serial ports.

Using "scanport" (Debian only ??) will probe all I/O ports and will indicate what it thinks may be serial port. After this you could try probing with setserial using the "autoconfig" option. You'll need to guess the addresses to probe at (using clues from "scanport"). See What is Setserial.

For a port set with jumpers, the IO ports and IRQs are set per the jumpers. 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.

Exploring via MS Windows (a last resort)

For PnP ports, checking on how it's configured under DOS/Windows may (or may not) imply how it's under Linux. MS Windows stores its configuration info in its Registry which is not used by Linux so they are not necessarily configured the same. 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 starting Linux) then Linux should use whatever configuration is in the BIOS's non−volatile memory. Windows also makes use of the same non−volatile memory but doesn't necessarily configure it that way.

8.7 Choosing Serial IRQs

If you have Plug−and−Play ports then either a PnP BIOS or a serial driver may configure all your devices for you so then you may not need to choose any IRQs. PnP software determines what it thinks is best and assigns them (but it's not always best). But if you directly use isapnp (ISA bus) or jumpers 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).

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. Such polling puts more load on the CPU but can be tried if there is an interrupt conflict or mis−set interrupt. The advantage of assigning IRQ 0 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.

Interrupt sharing, Kernels 2.2+

Sharing of IRQs is where two devices use the same IRQ. As a general rule, this wasn't allowed for the ISA bus. The PCI bus may share IRQs but one can't share the same IRQ between the ISA and the PCI bus. Most multi−port boards may share IRQs. Sharing is not 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 nicer to allocate every device its own interrupt.

Prior to kernel 2.2, serial IRQs could not be shared with each other except for most multiport boards. Starting with kernel 2.2 serial IRQs may be sometimes shared between 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). Since the PCI bus specs permit sharing, any PCI card should allow sharing.

What IRQs to choose?

The serial hardware often has only a limited number of IRQs. 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 or 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. Try to find one that 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 you might want to double−check `/proc/interrupts` when programs that use interrupts are being run and make sure there are no conflicts.

8.8 Choosing Addresses −−Video card conflict with ttyS3

Here's a problem with some old serial cards. 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 with the IO address of `ttyS3` at 0x02e8. Your may think that this shouldn't happen since the addresses are different in the high order digit (the leading 0 in 02e8). You're right, but a poorly designed serial port may ignore the high order digit and respond to any address that ends in 2e8. That is bad news if you try to use `ttyS3` (ISA bus) at this IO address.

For the ISA bus you should try to use the default addresses shown below. PCI cards use different addresses so as not to conflict with ISA addresses. The addresses shown below 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 commonly used serial ports on the ISA bus:

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

Suppose there is an address conflict (as reported by `setserial -g /dev/ttyS*`) between a real serial port and another port which does not physically exist (and shows UART: unknown). Such a conflict shouldn't cause problems but it sometimes does in older kernels. To avoid this problem don't permit such address conflicts or delete /dev/ttySx if it doesn't physically exist.

8.9 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 devices on ttyS0 (Com1) and ttyS1 (Com2))
- Letting a PnP BIOS automatically configure a PnP serial port See Using a PnP BIOS to I0−IRQ Configure
- Doing nothing if the serial driver recognized your card OK
- Using `isapnp` for a PnP serial port non−PCI)
- Using setpci (pciutils or 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 (PnP) was designed to automate this io−irq configuring, but for Linux it initially made life much more complicated. In modern Linux (2.4 kernels −−partially in 2.2 kernels), each device driver has to do it's own PnP (using supplied software which it may utilize). There is unfortunately no centralized planning for assigning IO addresses and IRQs as there is in MS Windows. But it usually works out OK in Linux anyway.

Using a PnP BIOS to I0−IRQ Configure

While the explanation of how to use setpci 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. This menu may be hard to find. For an "Award" BIOS it was found under "chipset features setup" There is often little to choose from. For ISA serial ports, the first two ports normally get set at the standard IO addresses and IRQs. See More on Serial Port Names

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 Linux is in some sense) or if thinks you don't have a PnP OS it may fully configure all the PnP devices but not configure the device drivers.

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. An indirect way to control what the BIOS does (if you have Windows 9x on the same PC) is to "force" a configuration under Windows. See Plug−and−Play−HOWTO and search for "forced". It's easier to use the CMOS BIOS menu which may override what you "forced" under Windows. There could be a BIOS option that can set or disable this "override" capability.

If you add a new PnP device, the BIOS should PnP configure 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 its setup menus of via messages on the screen when you turn on your computer. See What is set in my serial port hardware?. Other ways of finding out is to use lspci for the PCI bus or isapnp −−dumpregs for the ISA bus. The cryptic results it shows you may not be clear to a novice.

8.10 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 Boot−time Configuration

---

9. Configuring the Serial Driver (high−level) "stty"

9.1 Overview

See the section Stty. The "stty" command sets many things such as flow control, speed, and parity. The only one discussed in this section is flow control.

9.2 Flow Control

Configuring Flow Control: Hardware Flow Control is Usually Best See Flow Control for an explanation of it. It's usually better to use hardware flow control rather than software flow control using Xon/Xoff. To use full hardware flow control you must normally have two wires for it in the cable between the serial port and the device. If the device is on a card or the motherboard, then it should always be possible to use hardware flow control.

Many applications (and the getty program) give you an option regarding flow control and will set it for you. It might even set hardware flow control by default. It must be set both in the serial driver and in the hardware connected to the serial port. How it's set into the hardware is hardware dependent. Sometimes there is a certain "init string" you send to the hardware device via the serial port from your PC. For a modem, the communication program should set it in both places.

If a program you use doesn't set flow control in the serial driver, then you may do it yourself using the `stty` command. Since the driver doesn't remember the setting after you stop Linux, you could put the stty command in a file that runs at start−up or when you login (such as /etc/profile for the bash shell). Here's what you would add for hardware flow control for port ttyS2:

```
stty crtscts < /dev/ttyS2
or for stty version >= 1.17:
stty -F /dev/ttyS2 crtscts
```

`crtscts` stands for a Control setting to use the RTS and CTS pins of the serial port for hardware flow control. Note that RTS+CTS almost spells: `crtscts`.

---

10. Serial Port Devices /dev/tts/2 = /dev/ttyS2, etc.

10.1 Serial Port Names: ttyS2, tts/1, etc.

Once upon a time the names of the serial ports were simple. Except for some multiport serial cards they were named /dev/ttyS0, /dev/ttyS1, etc. Then around the year 2000 came the USB bus with names like /dev/ttyUSB0 and /dev/ttyACM1 (for the ACM modem on the USB bus). A little later with kernel 2.4 came the "device file system" (devfs) with a whole new set of names for everything. ttyS1 is now tts/1, ttyUSB1 is now /usb/tts/1, and ttyACM1 is now /usb/acm/1. Note that the the number 1 above is just an example. It could be replaced by 0, 2, 3, etc. The use of the device file system is optional and many are still using the legacy system. Others use devfs but have the old legacy names linked (via symlinks) to the new names. So they use the new system with the old names but may also use some of the new names for some devices. It's even possible ?? to use the new names for the old (non−devfs) system.

10.2 Devfs (The Device File System)

Starting with kernel 2.4, a new system of device naming was created. It makes it easy to deal with a huge number of devices. But there's an option to continue using the old names. However, a new device may not have an old−style name so then one must use the new devfs. For a detailed description of it see: http://www.atnf.csiro.au/~rgooch/linux/docs/devfs.html Also see the kernel documentation tree: filesystems/devfs.

Some more examples of name changes are: ttyS2 becomes tts/2 (Serial port), tty3 becomes vc/3 (Virtual Console), ptyp1 becomes pty/m1 (PTY master), ttyp2 becomes pty/s2 (PTY slave). "tts" looks like a directory which contains devices "files": 0, 1, 2, etc. All of these new names should still be in the /dev directory although optionally one may put them elsewhere.

Device names in the /dev directory are created automatically by the corresponding driver. Thus, if serial support comes from a module and that module isn't loaded yet, there will not be any serial devices in the /dev directory. This can be confusing: you physically have serial ports but don't see them in the /dev directory. However, if a device name is told to a communication program and the serial module isn't loaded, the kernel is supposed to try to find a driver for it and create a name for it in the /dev directory.

This is works OK if it finds a driver. But suppose there is no driver found for it. For example, if you try to use "setserial" to configure a port that the driver failed to detect, it claims there is no such port. How does one create a devfs port in this case?

10.3 Legacy Serial Port Device Names & Numbers

Before the device file system, devices in Linux had major and minor numbers. The serial port ttySx (x=0,1,2, etc.) was major number 4. You could see this (and the minor numbers too) by typing: "ls −l ttyS*" in the /dev directory. To find the old device names for various devices, see the "devices" file in the kernel documentation.

There formerly was a "cua" name for each serial port and it behaved just a little differently. For example, ttyS2 would correspond to cua2. It was mainly used for modems. The cua major number was 5 and minor numbers started at 64. You may still have the cua devices in your /dev directory but they are now deprecated. For details see Modem−HOWTO, section: cua Device Obsolete.

For creating the old devices in the device directory see:

Creating Devices In the /dev directory

10.4 More on Serial Port Names

Dos/Windows use the COM name while the messages from the serial driver use ttyS00, ttyS01, etc. Older serial drivers (2001 ?) used just tty00, tty01, etc.

The tables below shows some examples of serial device names. The IO addresses are the default addresses for the old ISA bus (not for the newer PCI and USB buses). The major/minor numbers aren't needed for the devfs, but they often exist anyway just in case the devfs method of locating drivers can't be used.

```
                              old   old  ISA IO
 dos     devfs     old name major minor address
COM1  /dev/tts/0  /dev/ttyS0  4,  64;   3F8
```

```
COM2  /dev/tts/1  /dev/ttyS1  4,  65;   2F8
COM3  /dev/tts/2  /dev/ttyS2  4,  66;   3E8
COM4  /dev/tts/3  /dev/ttyS3  4,  67;   2E8

     DEVICES-ON-THE-USB-BUS (acm is a certain type of modem)
     devfs       legacy name           devfs      legacy name
/dev/usb/tts/0  /dev/ttyUSB0  |  /dev/usb/acm/0  /dev/ttyACM0
/dev/usb/tts/1  /dev/ttyUSB1  |  /dev/usb/acm/1  /dev/ttyACM1
/dev/usb/tts/2  /dev/ttyUSB2  |  /dev/usb/acm/2  /dev/ttyACM2
/dev/usb/tts/3  /dev/ttyUSB3  |  /dev/usb/acm/3  /dev/ttyACM3
```

10.5 USB (Universal Serial Bus) Serial Ports

For more info see the usb subdirectory in the kernel documentation directory for files: usb−serial, acm, etc.

10.6 Link ttySN to /dev/modem

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

Historical note: Formerly (in the 1990s) the use of `/dev/modem` was discouraged since lock files might not realize that it was really say `/dev/ttyS2`. The newer lock file system doesn't fall into this trap so it's now OK to use such links.

10.7 Which Connector on the Back of my PC is ttyS1, etc?

Inspect the connectors

Inspecting the connectors may give some clues but is often not definitive. The serial connectors on the back side of a PC are usually DB connectors with male pins. 9−pin is the most common but some are 25−pin (especially older PCs like 486s). There may be one 9−pin (perhaps ttyS0 ??) and one 25−pin (perhaps ttyS1 ??). For two 9−pin ones the top one might be ttyS0.

If you only have one serial port connector on the back of your PC, this may be easy. If you also have an internal modem, a program like wvdial may be able to tell you what port it's on (unless it's a PnP that hasn't been enabled yet). A report from setserial (at boot−time or run by you from the command line) should help you identify the non−modem port.

If you have two serial connectors it may be more difficult. First check manuals (if any) for your computer. Look at the connectors for meaningful labels. You might even want to take off the PC's cover and see if there are any meaningful labels on the card where the internal ribbon cables plug in. Labels (if any) are likely to say something like "serial 1", "serial 2" or A, B. Which com port it actually is will depend on jumper or PnP settings (sometimes shown in a CMOS setup menu). But 1 or A are more likely to be ttyS0 with 2 or B ttyS1.

Send bytes to the port

Labels are not apt to be definitive so here's another method. If the serial ports have been configured correctly per setserial, then you may send some bytes out a port and try to detect which connector (if any) it's coming

out of. One way to send such a signal is to copy a long text file to the port using a command like: cp my_file_name /dev/ttyS1. A voltmeter connected to the DTR pin (see Serial−HOWTO for Pinout) will display a positive voltage as soon as you give the copy command.

The transmit pin should go from several volts negative to a voltage fluctuating around zero after you start sending the bytes. If it doesn't (but the DTR went positive) then you've got the right port but it's blocked from sending. This may be due to a wrong IRQ, −clocal being set, etc. The command "`stty -F /dev/ttyS1 -a`" should show clocal (and not −clocal). If not, change it to clocal.

Another test is to jumper the transmit and receive pins (pins 2 and 3 of either the 25−pin or 9−pin connector) of a test serial port. Then send something to each port (from the PCs keyboard) and see if it gets sent back. If it does it's likely the port with the jumper on it. Then remove the jumper and verify that nothing gets sent back. Note that if "echo" is set (per stty) then a jumper creates an infinite loop. Bytes that pass thru the jumper go into the port and come right back out of the other pin back to the jumper. Then they go back in and out again and again. Whatever you send to the port repeats itself forever (until you interrupt it by removing the jumper, etc.). This may be a good way to test it as the repeating test messages halt when the jumper is removed.

As a jumper you could use a mini (or micro) jumper cable (sold in some electronic parts stores) with mini alligator clips. A scrap of paper may be used to prevent the mini clips from accidentally touching the metal of the connector. Metal paper clips can sometimes be bent to use as jumpers. Whatever you use as a jumper take care not to bend or excessively scratch the pins. To receive something from a port, you can go to a virtual terminal (Alt−F2 for example) and type something like "cp /dev/ttyS2 /dev/tty". Then at another virtual terminal you may send something to ttyS2 (or whatever) by "echo test_message > /dev/ttyS2". Then go back to the receive virtual terminal and look for the test_message. See Serial Electrical Test Equipment for more info.

Connect a device to the connector

Another way to try to identify a serial port is to connect some physical serial device to it and see if it works. But a problem here is that it might not work because it's not configured right. A serial mouse might get detected at boot−time if connected.

Missing connectors

If the software shows that you have more serial ports than you have connectors for (including an internal modem which counts as a serial port) then you may have a serial port that has no connector. Some motherboards come with a serial port with no cable or serial DB connector. Someone may build a PC from this and omit the connector. There may be a "serial" connector and label on the motherboard but no ribbon cable connects to its pins. To use this port you must get a ribbon cable and connector. I've seen different wiring arrangements for such ribbon cables so beware.

10.8 Creating Devices In the /dev directory

If you don't use devfs (which automatically creates such devices) and don't have a device "file" that you need, you will have to create it. Use the `mknod` command or with the MAKEDEV shell script. Example, suppose you needed to create `ttyS0`:

```
linux# mknod -m 666 /dev/ttyS0 c 4 64
```

The MAKEDEV script is easier to use. See the man page for it. For example, if you needed to make the device for `ttyS0` you would just type:

```
linux# MAKEDEV ttyS0
```

If the above command doesn't work (and your are the root user), look for the MAKEDEV script in the the /dev directory and run it.

This handles the devices creation and should set the correct permissions. For making multiport devices see Making multiport devices in the /dev directory.

---

11. Interesting Programs You Should Know About

Most info on getty has been moved to Modem−HOWTO with a little info on the use of getty with directly connected terminals now found in Text−Terminal−HOWTO.

11.1 Serial Monitoring/Diagnostics Programs

A few Linux programs (and one "file") will monitor various modem control lines and indicate if they are positive (1 or green) or negative (0 or red).

- The "file": /proc/tty/driver/serial lists those that are asserted (positive voltage)
- modemstat (Only works correctly on Linux PC consoles. Status monitored in a tiny window. Color−coded and compact. Must kill it (a process) to quit.
- statserial (Info displayed on entire screen)
- serialmon (Doesn't monitor RTS, CTS, DSR but logs other functions)

As of June 1998, I know of no diagnostic program in Linux for the serial port.

11.2 Changing Interrupt Priority

- `irqtune` will give serial port interrupts higher priority to improve performance.
- `hdparm` for hard−disk tuning may help some more.

11.3 What is Setserial ?

This part is in 3 HOWTOs: Modem, Serial, and Text−Terminal. There

are some minor differences, depending on which HOWTO it appears in.

Important information

If you have a Laptop (PCMCIA) don't use `setserial` until you read Laptops: PCMCIA.

The term PnP means Plug−and−Play. All serial ports on the PCI bus are PnP and many on the ISA bus are also PnP. Linux software with the term "pnp" in it is usually only for the ISA bus. The PCI bus is inherently PnP so sofware for dealing with PnP on the PCI bus usually contains the term "pci" but not "pnp".

Introduction

`setserial` is a program which allows you (or a shell script) to talk to the serial device driver software. You may use it to tell the device driver how the hardware is phsically set (I/O addrees, IRQ, etc.). But normally, the device drivers finds out this infomation itself so that you don't need to use `setserial`. But if the device driver can't find a serial port (or perhaps gets the IRQ wrong) and you can find out this infomation while the driver can't, then `setserial` is essential. How to find serial ports is covered elsewhere, and in a minority of cases, `setserial` will be useful to help find them.

`setserial` permits you (or a script) to configure the serial port by telling the device driver 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. The name `setserial` is somewhat of a misnomer since it doesn't set the I/O address nor IRQ in the hardware, it just "sets" them in the driver software. And the driver naively believes that what `setserial` tells it has already been set in the hardware. Since the device driver is considered to be part of the kernel, the word "kernel" is often used in other documentation with no mention made of any "serial driver".

Some distributions (and versions) set things up so that `setserial` is run at boot−time by an initialization shell script (usually in the /etc directory tree). In other cases, you have to take some action to run `setserial` at boot−time. `setserial` will not work without either serial support built into the kernel or loaded as a module. The module may get loaded automatically if you (or a script) attempt to use setserial.

In addition to using `setserial` to configure, `setserial` can show you how the driver is currently configured (set). Hopefully, the hardware is set the same. In addition, it can be made to probe the hardware I0 port addresses to try to determine the UART type and IRQ, but this has severe limitations. See Probing. Note that it can't set the IRQ or the port address in the hardware of PnP serial ports (but the plug−and−play features of the serial driver may do this). It also can't directly read the PnP data stored in configuration registers in the hardware. But since the device driver can read these registers, `setserial` could be telling you what's in them, or it could be telling you what `setserial` had previously (and perhaps erroneously) told the driver. There's no way to know for sure without doing some other checks.

The serial driver (for Linux 2.4+) looks for a few "standard" legacy serial ports, for PnP ports on the ISA bus, and for all supported port hardware on the PCI bus. If it finds these OK then there's no need to use `setserial`. The driver doesn't probe for legacy IRQs and may get these wrong.

Besides the man page for `setserial`, check out info in `/usr/doc/setserial.../` or `/usr/share/doc/setserial`. It should tell you how setserial is handled in your distribution of Linux. While `setserial` behaves the same in all distributions, the scripts for running it, how to configure such scripts (including automatic configuration), and the names and locations of the script files, etc., are all distribution−dependent.

Serial module unload

If a serial module gets unloaded, the changes previously made by `setserial` will be forgotten by the driver. But while the driver forgets it, a script provided by the distribution may save it in a file somewhere so that it can the restored if the module is reloaded.

Slow baud rates of 1200 or less

There once was a problem with slow serial printers (especially the old ones of the 1980s). The printing program would close the serial port at the "end" of printing well before all the characters from the large serial

buffer (in main memory) were sent to the printer. The result was a truncated print job that didn't print the last paragraph or last page, etc.

But the newer lprng print program (and possibly other printing programs) keeps the port open until printing is finished so "problem solved", even if you're using an antique printer. Setserial can modify the time that the port will keep operating after it's closed (in order to output any characters still in its buffer in main RAM). This is done by the "closing_wait" option per the `setserial` man page. For "bad" software that closes the port too soon, it might also be needed at speeds above 1200 if there are a lot of "flow control" waits.

Giving the `setserial` command

Remember, that `setserial` can't set any I/O addresses or IRQs in the hardware. That's done either by plug−and−play software (run by the driver) or by jumpers for legacy serial ports. Even if you give an I/O address or IRQ to the driver via `setserial` it will not set such values and assumes that they have already been set. If you give it wrong values, the serial port will not work right (if at all).

For legacy ports, 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 by just typing `setserial` with no arguments. This fails to show you the one−letter options such as −v for verbose which you should normally use when troubleshooting. Note that setserial calls an IO address a "port". If you type:

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

you'll see some info about how the device driver is configured for your ports. Note that where it says `"UART: unknown"` it probably means that no uart exists. In other words, you probably have no such serial port and the other info shown about the port is meaningless and should be ignored. If you really do have such a serial port, setserial doesn't recognize it and that needs to be fixed.

If you add −a to the option −g you will see more info although few people need to deal with (or understand) this additional info since the default settings you see usually work fine. In normal cases the hardware is set up the same way as "setserial" reports. But if you are having problems there is a good chance that `setserial` has it wrong. In fact, you can run "setserial" and assign a purely fictitious I/O port 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 you've supplied to the driver. They will also be officially registered with the kernel as displayed (at the top of the screen) by the "scanport" command (Debian). Of course the serial port driver will not work correctly (if at all) if you attempt to use such a port. Thus, when giving parameters to `setserial`, "anything goes". Well almost. If you assign one port a base address that is already assigned (such as 3e8) it may not accept it. But if you use 3e9 it will accept it. Unfortunately 3e9 is actually assigned since it is within the range starting at base address 3e8. Thus the moral of the story is to make sure your data is correct before assigning resources with setserial.

Configuration file

While assignments made by setserial are lost when the PC is powered off, a configuration file usually restores them when the PC is started up again. In newer versions, what you change by setserial might get automatically saved to a configuration file. When `setserial` runs it uses the info from the the configuration file. In Debian there are 4 options for use of this configuration file:

1. Don't use this file at all. At each boot, the serial driver alone detects the ports and setserial doesn't ever run. ("kernel" option)
2. Save what `setserial` reports when the system is first shutdown and put it in the configuration file. After that, don't ever make any changes to the configuration file, even if someone has made changes by running the `setserial` command on the command line and then shuts down the system. ("autosave−once" option)
3. At every shutdown, save whatever `setserial` detects to the configuration file. ("autosave" option)
4. Manually edit the configuration file to set the configuration. Don't ever do any automatic saves to it. ("manual" option)

In old versions (perhaps before 2000), there wasn't any configuration file and the configuration was manually set (hard coded) inside the shell script that ran `setserial`. See Edit a script (prior to version 2.15).

Probing

You probe for a port only when you suspect that it has been enabled (by PnP methods ,the BIOS, jumbers, etc.), otherwise setserial probing will never find it since its address doesn't exist. Probling is where the software looks for a port at specified I/O addresses. Prior to probing with "setserial", one may run the "scanport" (Debian) command to check all possible ports in one scan. It makes crude guesses as to what is on some ports but doesn't determine the IRQ. It's a fast first start. It may hang your PC but so far it's worked fine for me. Note that non−Debian distributions don't seem to supply "scanport". Is there another scan program?

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 (2F8). If you tell setserial that ttyS2 is at a different address, then it will probe at that address, etc. See Probing

The purpose of such probing is to see if there is a uart there, and if so, what its IRQ is. Use `setserial` mainly as a last resort as there are faster ways to attempt it such as wvdialconf to detect modems, looking at very early boot−time messages, or using `pnpdump --dumpregs`. To try to detect the physical hardware use for example :

```
setserial /dev/ttyS2 -v autoconfig
```

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 supposedly 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 a serial port there.

Besides auto−probing for a uart type, setserial can auto−probe for IRQ's but this doesn't always work right either. In one case it first gave the wrong irq but when the command was repeated it found the correct irq. In versions of setserial >= 2.15, the results of your last probe test could be automatically saved and put into a configuration file such as `/etc/serial.conf` or `/var/lib/setserial/autoserial.conf`. This will be used next time you start Linux.

It may be that two serial ports both have the same IO address set in the hardware. Of course this is not normally permitted for the ISA bus but it sometimes happens anyway. Probing detects one serial port when actually there are two. However if they have different IRQs, then the probe for IRQs may show IRQ = 0. For me, it only did this if I first used `setserial` to give the IRQ a fictitious value.

Boot−time Configuration

While `setserial` may run via an initialization script, something akin to `setserial` also runs earlier when the serial module is loaded (or when the kernel starts the built−in serial driver if it was compiled into the kernel). Thus when you watch the start−up messages on the screen it may look like it ran twice, and in fact it has.

But the IRQs shown may be wrong since it doesn't probe for IRQs. The second report of the same is the result of a script such as /etc/init.d/setserial. It usually does no probing and thus could be wrong about how the hardware is actually set. It only shows configuration data that got saved in a configuration files. The old method, prior to setserial 2.15, was to manually write such data directly into the script.

When the kernel loads the serial module (or if the "module equivalent" is built into the kernel) then all supported PnP ports are detected. For legacy (non−PnP) ports, only `ttyS{0−3}` are auto−detected and the driver is set to use only IRQs 4 and 3 (regardless of what IRQs are actually set in the hardware). No probing is done for IRQs but it's possible to do this manually. You see this as a boot−time message just as if `setserial` had been run.

For legacy ports, to correct possible errors in IRQs (or for other reasons) there is likely a script file somewhere that runs `setserial` again. Unfortunately, if this file has some IRQs wrong, the kernel will still have incorrect info about the IRQs. This file is usually part of the initialization done at boot−time and whether it runs or not depends on how you (and/or your distribution) have set up the running to these initialization scripts. It also depends on the runlevel.

Before modifying a configuration file, you can test out a "proposed" `setserial` command by just typing it on the command line. In some cases the results of this use of `setserial` will automatically get saved in /etc/serial.conf when you shutdown. So if it worked OK (and solved your problem) then there's no need to modify any configuration file. See Configuration method using /etc/serial.conf, etc..

Edit a script (required prior to version 2.15)

This is how it was done prior to `setserial` 2.15 (1999) The objective was to modify (or create) a script file in the /etc tree that runs setserial at boot−time. Most distributions provided such a file (but it may not have initially resided in the /etc tree).

So prior to version 2.15 (1999) it was simpler. All you did was edit a script. There was no /etc/serial.conf file (or the like) to configure setserial. Thus you needed to find the file that runs "setserial" at boot time and edit it. If it didn't exist, you needed to create one (or place the commands in a file that ran early at boot−time). If such a file was currently being used it's likely was somewhere in the /etc directory−tree. But Redhat <6.0 has supplied it in /usr/doc/setserial/ but you need to move it to the /etc tree before using it.

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 was `/etc/rc.d/rc.local` but it's may not have run early enough. It's was reported that other processes may try to open the serial port before rc.local ran resulting in serial communication failure. Later on it's most likely was found in /etc/init.d/ but wasn't normally intended to be edited.

If such a file was supplied, it likely contained a number of commented−out examples. By uncommenting some of these and/or modifying them, you could set things up correctly. It was important use 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 was a lot faster than doing repeated reboots to get it right.

For versions >= 2.15 (provided your distribution implemented the change, Redhat didn't) it may be more tricky 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. See Configuration method using /etc/serial.conf, etc..

An example line in such a script was"

```
/sbin/setserial /dev/ttyS3 irq 5 uart 16550A  skip_test
```

or, if you wanted setserial to automatically determine the uart and the IRQ for ttyS3 you would have used something like this:

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

This was done for every serial port you wanted to auto configure, using a device name that really does exist on your machine. In some cases it didn't work right due to the hardware.

Configuration method using /etc/serial.conf, etc.

Prior to setserial version 2.15 (1999), the way to configure setserial was to manually edit the shell−script that ran setserial at boot−time. See Edit a script (before version 2.15). Today the script and configuration file are two different files instead of one. This shell−script is not edited but gets its data from a configuration file such as `/etc/serial.conf`.

Furthermore you may not even need to edit serial.conf because using the "setserial" command on the command line may automatically cause serial.conf to be edited appropriately. This was done so that you don't need to edit any file in order to set up (or change) what setserial does each time that Linux is booted.

What often happens is this: When you shut down your PC the script that ran "setserial" at boot−time is run again, but this time it only does what the part for the "stop" case says to do: It uses "setserial" to find out what the current state of "setserial" is, and it puts that info into the serial configuration file such as `serial.conf`. Thus when you run "setserial" to change the serial.conf file, it doesn't get changed immediately but only when and if you shut down normally.

Now you can perhaps guess what problems might occur. Suppose you don't shut down normally (someone turns the power off, etc.) and the changes don't get saved. Suppose you experiment with "setserial" and forget to run it a final time to restore the original state (or make a mistake in restoring the original state). Then your "experimental" settings are saved. There's an option to avoid this in Debian known as "AUTOSAVE−ONCE" which will be discussed later on.

If you manually edit serial.conf, then your editing is destroyed when you shut down because it gets changed back to the state of setserial at shutdown. There is a way to disable the changing of serial.conf at shutdown and that is to remove "###AUTOSAVE###" or the like from first line of serial.conf. In the Debian distribution, the removal of "###AUTOSAVE###" from the first line was once automatically done after the first time you shutdown just after installation. To retain this effect the "AUTOSAVE−ONCE" option was created which only saves the first time the system is shut down.

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 should not normally be edited. For

2.15, Redhat 6.0 just had a file /usr/doc/setserial−2.15/rc.serial which you have to move to /etc/init.d/ if you want setserial to run at boot−time.

To disable a port, use `setserial` to set it to "uart none". This will not be saved. 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.

BUG: As of July 1999 there is a bug/problem in Debian since with ###AUTOSAVE### only the setserial parameters displayed by "setserial −Gg /dev/ttyS*" get saved but the other parameters don't get saved. Use the −a flag to "setserial" to see all parameters. This will only affect a small minority of users since the defaults for the parameters not saved are usually OK for most situations. 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.

In some cases you may wind up with both the old and new configuration methods installed but hopefully only one of them runs at boot−time. Debian labeled obsolete files with "...pre−2.15".

IRQs

By default, both ttyS0 and ttyS2 will share IRQ 4, while ttyS1 and ttyS3 share IRQ 3. But actually sharing serial interrupts (using them in running programs) 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

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 a legacy internal modem (without plug−and−play) 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".

Laptops: PCMCIA

If you have a Laptop, read PCMCIA−HOWTO for info on the serial configuration. For serial ports on the motherboard, setserial is used just like it is for a desktop. But for PCMCIA cards (such as a modem) it's a different story. The configuring of the PCMCIA system should automatically run setserial so you shouldn't need to run it. If you do run it (by a script file or by /etc/serial.conf) it might be different and cause trouble. The autosave feature for serial.conf shouldn't save anything for PCMCIA cards (but Debian did until 2.15−7). Of course, it's always OK to use setserial to find out how the driver is configured for PCMCIA cards.

11.4 Stty

Introduction

`stty` does much of the configuration of the serial port but since application programs (and the getty program) often handle it, you may not need to use it much. It's handy if you're having problems or want to see how the port is set up. Try typing ``stty −a'' at your terminal/console to see how it's now set. Also try typing it without the −a (all) for a short listing which shows how it's set different than normal. Don't try to learn all the setting unless you want to become a serial guru. Most of the defaults should work OK and some of the settings are needed only for certain obsolete dumb terminals made in the 1970's.

`stty` is documented in the man pages with a more detailed account in the info pages. Type `"man stty"` or `"info stty"`.

Whereas `setserial` only deals with actual serial ports, stty is used both for serial ports and for virtual terminals such as the standard Linux text interface at a PC monitor. For the PC monitor, many of the stty settings are meaningless. Changing the baud rate, etc. doesn't appear to actually do anything.

Here are some of the items stty configures: speed (bits/sec), parity, bits/byte, # of stop bits, strip 8th bit?, modem control signals, flow control, break signal, end−of−line markers, change case, padding, beep if buffer overrun?, echo what you type to the screen, allow background tasks to write to terminal?, define special (control) characters (such as what key to press for interrupt). See the `stty` man or info page for more details. Also see the man page: `termios` which covers the same options set by stty but (as of mid 1999) covers features which the stty man page fails to mention.

With some implementations of getty (getty_ps package), the commands that one would normally give to stty are typed into a getty configuration file: /etc/gettydefs. Even without this configuration file, the getty command line may be sufficient to set things up so that you don't need stty.

One may write C programs which change the stty configuration, etc. Looking at some of the documentation for this may help one better understand the use of the stty command (and its many possible arguments). Serial−Programming−HOWTO is useful. The manual page: termios contains a description of the C−language structure (of type termios) which stores the stty configuration in computer memory. Many of the flag names in this C−structure are almost the same (and do the same thing) as the arguments to the stty command.

Flow control options

To set hardware flow control use "crtscts". For software flow control there are 3 settings: ixon, ixoff, and ixany.

ixany: Mainly for terminals. Hitting any key will restarts the flow after a flow−control stop. If you stop scrolling with the "stop scroll" key (or the like) then hitting any key will resume scrolling. It's seldom needed since hitting the "scroll lock" key again will do the same thing.

ixon: Enables the port to listen for Xoff and to stop transmitting when it gets an Xoff. Likewise, it will resume transmitting if it gets an Xon.

ixoff: enables the port to send the Xoff signal out the transmit line when its buffers in main memory are nearly full. It protects the device where the port is located from being overrun.

For a slow dumb terminal (or other slow device) connected to a fast PC, it's unlikely the the PC's port will be overrun. So you seldom actually need to enable ixoff. But it's often enabled "just in case".

Using stty at a "foreign" terminal

Using `stty` to configure the terminal that you are currently using is easy. Doing it for a different (foreign) terminal or serial port may be impossible. For example, let's say you are at the PC monitor (tty1) and want to use `stty` to deal with the serial port ttyS2. Prior to about 2000 you needed to use the redirection operator "<". After 2000 (provided your version of setserial is >= 1.17 and stty >= 2.0) there is a better method using the −F option. This will work when the old redirection method fails. Even with the latest versions be warned that if there is a terminal on ttyS2 and a shell is running on that terminal, then what you see will likely be deceptive and trying to set it will not work. See Two interfaces at a terminal to understand it.

The new method is ``stty −F /dev/ttyS2 ..." (or −−file instead of F). If ... is −a it displays all the stty settings. The old redirection method (which still works in later versions) is to type ``stty ... </dev/ttyS2". If the new method works but the old one hangs, it implies that the port is hung due to a modem control line not being asserted. Thus the old method is still useful for troubleshooting. See the following subsection for details.

Old redirection method

Here's a problem with the old redirection operator (which doesn't happen if you use the newer −F option instead). Sometimes when trying to use stty, the command hangs and nothing happens (you don't get a prompt for a next command even after hitting <return>). This is likely due to the port being stuck because it's waiting for one of the modem control lines to be asserted. For example, unless you've set "clocal" to ignore modem control lines, then if no CD signal is asserted the port will not open and stty will not work for it (unless you use the newer −F option). A similar situation seems to exist for hardware flow control. If the cable for the port doesn't even have a conductor for the pin that needs to be asserted then there is no easy way to stop the hang.

One way to try to get out of the above hang is to use the newer −F option and set "clocal" and/or "crtscts" as needed. If you don't have the −F option then you may try to run some program (such as minicom) on the port that will force it to operate even if the control lines say not to. Then hopefully this program might set the port so it doesn't need the control signal in the future in order to open: clocal or −crtscts. To use "minicom" to do this you likely will have to reconfigure minicom and then exit it and restart it. Instead of all this bother, it may be simpler to just reboot the PC.

The old redirection method makes ttyS2 the standard input to stty. This gives the stty program a link to the "file" ttyS2 so that it may "read" it. But instead of reading the bytes sent to ttyS2 as one might expect, it uses the link to find the configuration settings of the port so that it may read or change them. Some people tried to use ``stty ... > /dev/ttyS2" to set the terminal. This will not do it. Instead, it takes the message normal displayed by the stty command for the terminal you are on (say tty1) and sends this message to ttyS2. But it doesn't change any settings for ttyS2.

Two interfaces at a terminal

When using a shell (such as bash) with command−line−editing enabled there are two different terminal interfaces (what you see when you type stty −a). When you type in modern shells at the command line you have a temporary "raw" interface (or raw mode) where each character is read by the command−line−editor as you type it. Once you hit the <return> key, the command−line−editor is exited and the terminal interface is changed to the nominal "cooked" interface (cooked mode) for the terminal. This cooked mode lasts until the next prompt is sent to the terminal (which is only a small fraction of a second). Note that one never gets to

type anything to this cooked mode but what was typed in raw mode gets executed while in cooked mode.

When a prompt is sent to the terminal, the terminal goes from "cooked" to "raw" mode (just like it does when you start an editor since you are starting the command−line editor). The settings for the "raw" mode are based only on the basic settings taken from the "cooked" mode. Raw mode keeps these setting but changes several other settings in order to change the mode to "raw". It is not at all based on the settings used in the previous "raw" mode. Thus if one uses stty to change settings for the raw mode, such settings will be permanently lost as soon as one hits the <return> key at the terminal that has supposedly been "set".

Now when one types stty to look at the terminal interface, one may either get a view of the cooked mode or the raw mode. You need to figure out which one you're looking at. It you use stty from another (foreign) terminal then you will see the raw mode settings. Any changes made will only be made to the raw mode and will be lost when someone presses <return> at the terminal you tried to "set". But if you type a stty command at your terminal (without the −F option or redirection) and then hit <return> it's a different story. The <return> puts the terminal in cooked mode. Your changes are saved and will still be there when the terminal goes back into raw mode (unless of course it's a setting not allowed in raw mode).

This situation can create problems. For example, suppose you corrupt your terminal interface. To restore it you go to another terminal and "stty −F dev/ttyS1 sane" (or the like). It will not work! Of course you can try to type "stty sane ..." at the terminal that is corrupted but you can't see what you typed. All the above not only applies to dumb terminals but to virtual terminals used on a PC Monitor as well as to the terminal windows in X. In other words, it applies to almost everyone who uses Linux.

Luckily, when you start up Linux, any file that runs stty at boot−time will likely deal with a terminal (or serial port with no terminal) that has no shell running on it so there's no problem for this special case.

Where to put the stty command ?

Should you need to have `stty` set up the serial interface each time the computer starts up then you need to put the `stty` command in a file that will be executed each time the computer is started up (Linux boots). It should be run before the serial port is used (including running getty on the port). There are many possible places to put it. If it gets put in more than one place and you only know about (or remember) one of those places, then a conflict is likely. So make sure to document what you do.

One place to put it would be in the same file that runs setserial when the system is booted. The location is distribution and version dependent. It would seem best to put it after the setserial command so that the low level stuff is done first. If you have directories in the /etc tree where every file in them is executed at boot−time (System V Init) then you could create a file named "stty" for this purpose.

11.5 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. 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.

11.6 What is slattach?

It's "serial line attach". It puts the serial line into a networking mode. You can thus network two computers together via a serial line using, for example, the slip protocol. But for the ppp protocol, you need to start pppd on the serial line.

---

12. Speed (Flow Rate)

By "speed" we really mean the "data flow rate" but almost everybody incorrectly calls it speed. The speed is measured in bits/sec (or baud). Speed is set using the "stty" command or by a program which uses the serial port. See Stty

12.1 Very High Speeds

Speeds over 115.2k

The top speed of 115.2k has been standard since the mid 1990's. But by the year 2000, most new serial ports supported higher speeds of 230.4k and 460.8k. Some also support 921.6k. Unfortunately Linux seldom uses these speeds due to lack of drivers. Thus such ports behave just like 115.2k ports unless the higher speeds are enabled by special software. To get these speeds you need to compile the kernel with special patches or use modules until support is built into the kernel's serial driver.

Unfortunately serial port manufacturers never got together on a standard way to support high speeds, so the serial driver needs to support a variety of hardware. Once high speed is enabled, a standard way to choose it is to set baud_base to the highest speed with setserial (unless the serial driver does this for you). The software will then use a divisor of 1 to set the highest speed. All this will hopefully be supported by the Linux kernel sometime in 2003.

A driver for the w83627hf chip (used on many motherboards such as the Tyan S2460) is at https://www.muru.com/linux/w83627hf/

A non−standard way that some manufacturers have implemented high speed is to use a very large number for the divisor to get the high speed. This number isn't really a divisor at all since it doesn't divide anything. It's just serves as a code number to tell the hardware what speed to use. In such cases you need to compile the kernel with special patches.

One patch to support this second type of high−speed hardware is called shsmod (Super High Speed Mode). There are both Windows and Linux versions of this patch. See http://www.devdrv.com/shsmod/. There is also a module for the VIA VT82C686 chip http://www.kati.fi/viahss/. Using it may result in buffer overflow.

For internal modems, only a minority of them advertise that they support speeds of over 115.2k for their built−in serial ports. Does shsmod support these ??

How speed is set in hardware: the divisor and baud_base

Speed is set by having the serial port's clock change frequency. But this change happens not by actually changing the frequency of the oscillator driving the clock but by "dividing" the clock's frequency. For example, to divide by two, just ignore every other clock tick. This cuts the speed in half. Dividing by 3 makes

the clock run at 1/3 frequency, etc. So to slow the clock down (meaning set speed), we just send the clock a divisor. It's sent by the serial driver to a register in the port. Thus speed is set by a divisor.

If the clock runs at a top speed of 115,000 bps (common), then here are the divisors for various 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).

There are exceptions to the above since for certain serial port hardware, speeds above 115.2k are set by using a very high divisor. Keep that exception in mind as you read the rest of this section. 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 sets the highest speed.

Besides using a very high divisor to set high speed, the conventional way to do it is as follows: If you happen to have hardware with a maximum speed of say 230.4k (and the 230.4k speed has been enabled in the hardware), then specifying 115.2k will result in divisor 1. For some hardware this 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. All the above assumes that you don't use "setserial" to modify things.

Setting the divisor, speed accounting

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.

If you have very old software which will not allow you to tell it such a high speed (but your hardware has it enabled) then you might want to look into using the "spd_cust" parameter. This allows you to tell the application that the speed is 38,400 but the actual speed for this case is determined by the value of "divisor" which has also been set in setserial. I think it best to try to avoid using this kludge.

There are some brands of UARTs that uses a very high divisor to set high speeds. There isn't any satisfactory way to use "setserial" (say set "divisor 32770") to get such a speed since then setserial would then think that the speed is very low and disable the FIFO in the UART.

Crystal frequency is higher than baud_base

Note that the baud_base setting is usually much lower than the frequency of the crystal oscillator since the crystal frequency of say 1.8432 MHz is divided by 16 in the hardware to get the actual top speed of 115.2k. The reason the crystal frequency needs to be higher is so that this high crystal speed can generate clock ticks to take a number of samples of each bit to determine if it's a 1 or a 0.

Actually, the 1.8432 MHz "crystal frequency" may be obtained from a 18.432 MHz crystal oscillator by dividing by 10 before being fed to the UART. Other schemes are also possible as long as the UART performs properly.

12.2 Higher Serial Throughput

If you are seeing slow throughput and serial port overruns on a system with (E)IDE disk drives, you can get `hdparm`. This is a utility that can modify (E)IDE parameters, including unmasking other IRQs during a disk IRQ. This will improve responsiveness and will help eliminate overruns. Be sure to read the man page very carefully, since some drive/controller combinations don't like this and may corrupt the filesystem.

Also have a look at a utility called `irqtune` that will change the IRQ priority of a device, for example the serial port that your modem is on. This may improve the serial throughput on your system. The `irqtune` FAQ is at http://www.best.com/~cae/irqtune

---

13. Locking Out Others

13.1 Introduction

When you are using a serial port, you may want to prevent others from using it at the same time. However there may be cases where you do want others to use it, such as sending you an important message if you are using a text−terminal.

There are various ways of preventing others (or other processes) from using your serial port when you are using it (locking). This should all happen automatically but it's important to know about this if it gives you trouble. If a program is abnormally exited or the PC is abruptly turned off (by pulling the plug, etc.) your serial port might wind up locked. Even if the lock remains, it's usually automatically removed when you want to use the serial port again. But in rare cases it isn't. That's when you need to understand what happened.

One way to implement locking is to design the kernel to handle it but Linux thus far has shunned this solution (with an exception involving the cua device which is now obsolete). Two solutions used by Linux is to:

1. create lock−files
2. modify the permissions and/or owners of devices such as /dev/ttyS2

13.2 Lock−Files

If you use the new device−filesystem (devfs) then see the next section. A lock−file is simply a file created to mean that a particular device is in use. They are kept in `/var/lock`. Formerly they were in `/usr/spool/uucp`. Linux lock−files are usually named `LCK..`*name*, where *name* may be a device name, a process id number, a device's major and minor numbers, or a UUCP site name. Most processes (an exception is getty) create these locks so that they can have exclusive access to devices. For instance if you dial out on your modem, some lockfiles will appear to tell other processes that someone else is using the modem. In older versions (in the 1990s) there was usually only one lockfile per process. Lock files contain the PID of the process that has locked the device. Note that if a process insists on using a device that is locked, it may ignore the lockfile and use the device anyway. This is useful in sending a message to a text−terminal, etc.

When a program wants to use a serial port but finds it locked with lock−files it should check to see if the lock−file's PID is still in use. If it's not it means that the lock is stale and it's OK to go ahead and use the port anyway (after removing the stale lock−files). Unfortunately, there may be some programs that don't do this and give up by telling you that a device is already in use when it really isn't.

When there were only lockfiles with device names, the following problem could arise: If the same device has two different names then two different processes could each use a different name for the same device. This results in lockfiles with different names that actually are the same device. Formerly each physical serial port was known by two different device names: ttyS0 and cua0. To solve this lockfile alias problem, 3 methods have been used. It may be overkill since any one of these methods would have fixed the problem.

1. The lock checking software was made aware of ttyS vs. cua.
2. The device cua was deprecated
3. Additional locks were created which use unique device numbers instead of names.

Using alternate names such as /dev/modem for /dev/ttyS2 may cause problems with older versions. For dumb terminals, lockfiles are not used since this would not permit someone else to send a message to your terminal using the write or talk program.

13.3 Lock–Files and devfs Problems

The new device–filesystem (devfs) has the /dev directory with subdirectories. As of late 2001, there were problems with lockfiles. For example, the lockfile mechanism considered dev/usb/tts/0 and /dev/tts/0 to be the same device with name "0". Ditto for all other devices that had the same "leaf" name.

Also, if some applications use the old name for a device and other applications use the new name (devfs) for the same device, then the lockfiles will have different names. But the serial driver should know they are the same.

13.4 Change Owners, Groups, and/or Permissions of Device Files

In order to use a device, you (or the program you run if you have "set user id") needs to have permission to read and write the device "file" in the /dev directory. So a logical way to prevent others from using a device is to make yourself the temporary owner of the device and set permissions so that no one else can use it. A program may do this for you. A similar method can be used with the group of the device file.

While lock files prevent other process from using the device, changing device file owners/permissions restricts other users (or the group) from using it. One case is where the group is permitted to write to the port, but not to read from it. Writing to the port might just mean a message sent to a text–terminal while reading means destructive reading. The original process that needs to read the data may find data missing if another process has already read that data. Thus a read can do more harm that a write since a read causes loss of data while a write only adds extra data. That's a reason to allow writes but not reads. This is exactly the opposite of the case for ordinary files where you allow others to read the file but not write (modify) it. Use of a port normally requires both read and write permissions.

A program that changes the device file attributes should undo these changes when it exits. But if the exit is abnormal, then a device file may be left in such a condition that it gives the error "permission denied" when one attempts to use it again.

14. Communications Programs And Utilities

14.1 List of Software

Here is a list of some communication software you can choose from, available via FTP, if they didn't come with your distribution.

- `ecu` – a communications program
- C–Kermit – portable, scriptable, serial and TCP/IP communications including file transfer, character–set translation, and zmodem support
- `gkermit` Tiny GPLed kermit run only from the command line. Can't connect to another computer
- `minicom` – telix–like communications program
- `pppd` – establishes a ppp connection on the serial line
- `seyon` – X based communication program
- `xc` – xcomm communication package
- `term` and `SLiRP` offer TCP/IP functionality using a shell account.
- `screen` is another multi–session program. This one behaves like the virtual consoles.
- `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).
- `mgetty+fax` handles FAX stuff, and provides an alternate `ps_getty`.
- `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 (if not in your Linux distribution) can be found at `ftp://metalab.unc.edu/pub/Linux/system/network/serial`.

14.2 kermit and zmodem

For use of kermit with modems see the Modem–HOWTO. One can run zmodem within the kermit program. To do this (for ttyS3), add the following to your `.kermrc` file:

```
define rz !rz < /dev/ttyS3 > /dev/ttyS3
define sz !sz \%0 > /dev/ttyS3 < /dev/ttyS3
```

Be sure to put in the correct port your modem is on. Then, to use it, just type `rz` or `sz <filename>` at the `kermit` prompt.

---

15. Serial Tips And Miscellany

15.1 Serial Module

Often the serial driver is provided as a module. Parameters may be supplied to certain modules in /etc/modules.conf. Since kernel 2.2 you don't edit this file but use the program update–modules to change it. The info that is used to update modules.conf is put in /etc/modutils/. The Debian/GNU Linux has a file here named /etc/modutils/setserial which runs the serial script in /etc/init.d/ every time the serial module is loaded or unloaded. When the serial module is unloaded this script will save the state of the module in /var/run/setserial.conf. Then if the module loads again this saved state is restored. When the serial module first loads at boot–time, there's nothing in /var/run/setserial.conf so the state is obtained from /etc/serial.conf. So there are two files that save the state. Other distributions may do something similar.

One may modify the serial driver by editing the source code. Much of the serial driver is found in the file serial.c. For info regarding writing of programs for the serial port see Serial−Programming−HOWTO. It was revised in 1999 by Vern Hoxie but it's not at LDP. Get it from scicom.alphacdc.com/pub/linux

15.2 Serial Console (console on the serial port)

See the kernel documentation in: Documentation/serial−console.txt. Kernel 2.4+ has better documentation. See also "Serial Console" in Text−Terminal−HOWTO.

15.3 Line Drivers

For a text terminal, the EIA−232 speeds are fast enough but the usable cable length is often too short. Balanced technology could fix this. The common method of obtaining balanced communication with a text terminal is to install 2@ line drivers in the serial line to convert unbalanced to balanced (and conversely). They are a specialty item and are expensive if purchased new.

15.4 Stopping the Data Flow when Printing, etc.

Normally flow control and/or application programs stop the flow of bytes when its needed. But sometimes they don't. The problem is that output to the serial port first passes thru the large serial buffer in the PC's main memory. So if you want to abort printing, whatever is in this buffer should be removed. When you tell an application program to stop printing, it may not empty this buffer so printing continues until it's empty. In addition, your printer has it's own buffer which needs to be cleared. So telling the PC to stop printing may not work due to these two buffers that continue to supply bytes for the printer. It's a problem with printer software not knowing about the serial port and that modem control lines need to be dropped to stop the printer.

One way to insure that printing stops is to just turn off the printer. With newer serial drivers, this works OK. The buffers are cleared and printing doesn't resume. With older serial drivers, the PC's serial buffer didn't clear and it would sometimes continue to print when the printer was turned back on. To avoid this, you must wait a time specified by setserial's closing_wait before turning the printer back on again. You may also need to remove the print job from the print queue so it won't try to resume.

15.5 Known IO Address Conflicts

Avoiding IO Address Conflicts with Certain Video Boards

The IO address of the IBM 8514 video board (and others) is allegedly 0x?2e8 where ? is 2, 4, 8, or 9. This may conflict (but shouldn't if the serial port is well designed) with the IO address of `ttyS3` at 0x02e8 if the serial port ignores the leading 0 hex digit when it decodes the address (many do). That is bad news if you try to use `ttyS3` at this IO address. Another story is that Linux will not detect your internal modem on `ttyS3` but that you can use `setserial` to put `ttyS3` at this address and the modem will work fine.

IO address conflict with ide2 hard drive

The address of ttyS2 is 3e8−3ef while hard drive ide2 uses 3ee which is in this range. So when booting Linux you may see a report of this conflict. Most people don't use ide2 (the 3rd hard drive cable) and may ignore this conflict message. You may have 2 hard drives on ide0 and two more on ide1 so most people don't need ide2.

15.6 Known Defective Hardware

Problem with AMD Elan SC400 CPU (PC−on−a−chip)

This has a race condition between an interrupt and a status register of the UART. An interrupt is issued when the UART transmitter finishes the transmission of a byte and the UART transmit buffer becomes empty (waiting for the next byte). But a status register of the UART doesn't get updated fast enough to reflect this. As a result, the interrupt service routine rapidly checks and determines (erroneously) that nothing has happened. Thus no byte is sent to the port to be transmitted and the UART transmitter waits in vain for a byte that never arrives. If the interrupt service routine had waited just a bit longer before checking the status register, then it would have been updated to reflect the true state and all would be OK.

There is a proposal to fix this by patching the serial driver. But Should linux be patched to accommodate defective hardware, especially if this patch may impair performance of good hardware?

---

16. Troubleshooting

See Modem−HOWTO for troubleshooting related to modems or getty for modems. For a Text−Terminal much of the info here will be of value as well as the troubleshooting info in Text−Terminal−HOWTO.

16.1 Serial Electrical Test Equipment

Breakout Gadgets, etc.

While a multimeter (used as a voltmeter) may be all that you need for just a few serial ports, simple special test equipment has been made for testing serial port lines. Some are called "breakout ... " where breakout means to break out conductors from a cable. These gadgets have a couple of connectors which connect to serial port connectors (either at the ends of serial cables or at the back of a PC). Some have test points for connecting a voltmeter. Others have LED lamps which light when certain modem control lines are asserted (turned on). The color of the light may indicate the polarity of the signal (positive or negative voltage). Still others have jumpers so that you can connect any wire to any wire. Some have switches.

Radio Shack sells (in 2002) a "RS−232 Troubleshooter" (formerly called "RS−232 Line Tester") Cat. #276−1401. It checks TD, RD, CD, RTS, CTS, DTR, and DSR. A green light means on (+12 v) while red means off (−12 v). They also sell a "RS−232 Serial Jumper Box" Cat. #276−1403. This permits connecting the pins anyway you choose. Both these items are under the heading of "Peripheral hookup helpers". Unfortunately, they are not listed in the index to the printed catalog. They are on the same page as the D type connecters so look in the index under "Connectors, Computer, D−Sub". A store chain named "Active Components" may have them.

Measuring voltages

Any voltmeter or multimeter, even the cheapest that sells for about $10, should work fine. Trying to use other methods for checking voltage is tricky. Don't use a LED unless it has a series resistor to reduce the voltage across the LED. A 470 ohm resistor is used for a 20 ma LED (but not all LED's are 20 ma). The LED will only light for a certain polarity so you may test for + or − voltages. Does anyone make such a gadget for automotive circuit testing?? Logic probes may be damaged if you try to use them since the TTL voltages for which they are designed are only 5 volts. Trying to use a 12 V incandescent light bulb is not a good idea. It

won't show polarity and due to limited output current of the UART it probably will not even light up.

To measure voltage on a female connector you may plug in a bent paper clip into the desired opening. The paper clip's diameter should be no larger than the pins so that it doesn't damage the contact. Clip an alligator clip (or the like) to the paper clip to connect up. Take care not to touch two pins at the same time with any metal object.

Taste voltage

As a last resort, if you have no test equipment and are willing to risk getting shocked (or even electrocuted) you can always taste the voltage. Before touching one of the test leads with your tongue, test them to make sure that there is no high voltage on them. Touch both leads (at the same time) to one hand to see if they shock you. Then if no shock, wet the skin contact points by licking and repeat. If this test gives you a shock, you certainly don't want to use your tongue.

For the test for 12 V, Lick a finger and hold one test lead in it. Put the other test lead on your tongue. If the lead on your tongue is positive, there will be a noticeable taste. You might try this with flashlight batteries first so you will know what taste to expect.

16.2 Serial Monitoring/Diagnostics

A few Linux programs will monitor the modem control lines and indicate if they are positive (1) or negative (0). See section Serial Monitoring/Diagnostics

16.3 (The following subsections are in both the Serial and Modem HOWTOs)

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

If a physical device (such as a modem) doesn't work at all it's often because it's disabled and has no address (PnP hasn't enabled it) or that it is enabled but is not at the I/O address that setserial thinks it's at. Thus it can't be found.

First check BIOS messages at boot−time (and possibly the BIOS menu for the serial port). Then for the PCI bus use lspci or scanpci. If it's an ISA bus PnP serial port, try "pnpdump −−dumpregs" and/or see Plug−and−Play−HOWTO. If the port happens to be enabled then the following two paragraphs may help find it:

Using "scanport" (Debian only ??) will scan all enabled bus ports and may discover an unknown port that could be a serial port (but it doesn't probe the port). It could hang your PC. You may try probing with setserial. See Probing.

If nothing seems to get thru the port it may be accessible but have a bad interrupt. See Extremely Slow: Text appears on the screen slowly after long delays. Use `setserial -g` to see what the serial driver thinks and check for IRQ and I0 address conflicts. Even if you see no conflicts the driver may have incorrect information (view it by "setserial") and conflicts may still exist.

If two ports have the same IO address then probing it will erroneously indicate only one port. Plug−and−play

detection will find both ports so this should only be a problem if at least one port is not plug−and−play. All sorts of errors may be reported/observed for devices illegally "sharing" a port but the fact that there are two devices on the same a port doesn't seem to get detected (except hopefully by you). In the above case, if the IRQs are different then probing for IRQs with setserial might "detect" this situation by failing to detect any IRQ. See Probing.

16.5 Extremely 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 serial 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. It may be just an invisible <return> character so all you notice is that 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

Interrupt Problem Details and/or Interrupt Conflicts and/or Mis−set Interrupts. If it involves Plug−and−Play devices, see also Plug−and−Play−HOWTO.

As a quick check to see if it really is an interrupt problem, set the IRQ to 0 with "setserial". This will tell the driver to use polling instead of interrupts. If this seems to fix the "slow" problem then you had an interrupt problem. You should still try to solve the problem since polling uses excessive computer resources.

Checking to find the interrupt conflict may not be easy since Linux supposedly doesn't permit any interrupt conflicts and will send you a /dev/ttyS?: Device or resource busy error message if it thinks you are attempting to create a conflict. But a real conflict can be created if "setserial" has told the kernel incorrect info. The kernel has been lied to and thus doesn't think there is any conflict. Thus using "setserial" will not reveal the conflict (nor will looking at /proc/interrupts which bases its info on "setserial"). You still need to know what "setserial" thinks so that you can pinpoint where it's wrong and change it when you determine what's really set in the hardware.

What you need to do is to check how the hardware is set by checking jumpers or using PnP software to check how the hardware is actually set. For PnP run either "pnpdump −−dumpregs" (if ISA bus) or run "lspci" (if PCI bus). Compare this to how Linux (e.g. "setserial") thinks the hardware is set.

16.6 Somewhat Slow: I expected it to be a few times faster

An obvious reason is that the baud rate is actually set too slow. It's claimed that this happened by trying to set the baud rate to a speed higher than the hardware can support (such as 230400).

Another reason may be that whatever is on the serial port (such as a modem, terminal, printer) doesn't work as fast as you thought it did.

Another possible reason is that you have an obsolete serial port: UART 8250, 16450 or early 16550 (or the serial driver thinks you do). See

What Are UARTS? Use "setserial –g /dev/ttyS*". If it shows anything less than a 16550A, this may be your problem. If you think that "setserial" has it wrong check it out. See What is Setserial for more info. If you really do have an obsolete serial port, lying about it to setserial will only make things worse.

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

For non–PnP 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" which is the same as ttyS2). You may need to use `setserial` to tell Linux the IRQ you are using.

16.8 "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 will work for you if it shows 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 (and secure) ways to get access like belonging to a "group" that has group permission. Some programs change the permissions when they run but restore them when the program exists normally. But if someone pulls the plug on your PC it's an abnormal exit and correct permissions may not be restored.

16.9 "Operation not supported by device" for ttyS?

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 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.

16.10 "Cannot create lockfile. Sorry"

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. Giving rwx permissions for the root owner and the group should work, provided that the users that need to

dialout belong to that group. Others should have r–x permission. Even with this scheme, there may be a security risk. Use "chmod" to change permissions and "chgrp" to change groups. Of course, if there is no "lock" directory no lockfile can be created there. For more info on lockfiles see What Are Lock Files

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

This means that someone else (or some other process) is supposedly using the serial port. There are various ways to try to find out what process is "using" it. One way is to look at the contents of the lockfile (/var/lock/LCK...). It should be the process id. If the process id is say 100 type "ps 100" to find out what it is. Then if the process is no longer needed, it may be gracefully killed by "kill 100". If it refuses to be killed use "kill –9 100" to force it to be killed, but then the lockfile will not be removed and you'll need to delete it manually. Of course if there is no such process as 100 then you may just remove the lockfile but in most cases the lockfile should have been automatically removed if it contained a stale process id (such as 100).

16.12 "/dev/tty? Device or resource busy"

This means that the device you are trying to access (or use) is supposedly busy (in use) or that a resource it needs (such as an IRQ) is supposedly being used by another device and can't be shared. This message is easy to understand if it only means that the device is busy (in use). But it sometimes means that a needed resource is already in use (busy). What makes it even more confusing is that in some cases neither the device nor the resources that it needs are actually "busy".

In olden days, if a PC was shutdown by just turning off the power, a bogus lockfile might remain and then later on one would get this bogus message and not be able to use the serial port. Software today is supposed to automatically remove such bogus lockfiles, but as of 2003 there is still a problem with the "wvdial" dialer program related to lockfiles. If wvdial can't create a lockfile because it doesn't have write permission in the /var/lock/ directory, you will see this erroneous message.

The following example is where interrupts can't be shared (at least one of the interrupts is on the ISA bus). The ``resource busy'' part often means (example for `ttyS2`) ``You can't use `ttyS2` since another device is using ttyS2's interrupt.'' The potential interrupt conflict is inferred from what "setserial" thinks. A more accurate error message would be ``Can't use `ttyS2` since the setserial data (and kernel data) indicates that another device is using `ttyS2`'s interrupt''. If two devices use the same IRQ and you start up only one of the devices, everything is OK because there is no conflict yet. But when you next try to start the second device (without quitting the first device) you get a "... busy" error message. This is because the kernel only keeps track of what IRQs are actually in use and actual conflicts don't happen unless the devices are in use (open). The situation for I/O address (such as 0x3f8) conflict is similar.

This error is sometimes due to having two serial drivers: one a module and the other compiled into the kernel. Both drivers try to grab the same resources and one driver finds them "busy".

There are two possible cases when you see this message:

1. There may be a real resource conflict that is being avoided.
2. Setserial has it wrong and the only reason `ttyS2` can't be used is that setserial erroneously predicts a conflict.

What you need to do is to find the interrupt setserial thinks `ttyS2` is using. Look at /proc/tty/driver/serial. You should also be able to find it with the "setserial" command for `ttyS2`.

Bug in old versions: Prior to 2001 there was a bug which wouldn't let you see it with "setserial". Trying to see it would give the same "... busy" error message.

To try to resolve this problem reboot or: exit or gracefully kill all likely conflicting processes. If you reboot: 1. Watch the boot−time messages for the serial ports. 2. Hope that the file that runs "setserial" at boot−time doesn't (by itself) create the same conflict again.

If you think you know what IRQ say `ttyS2` is using then you may look at /proc/interrupts to find what else (besides another serial port) is currently using this IRQ. You might also want to double check that any suspicious IRQs shown here (and by "setserial") are correct (the same as set in the hardware). A way to test whether or not it's a potential interrupt conflict is to set the IRQ to 0 (polling) using "setserial". Then if the busy message goes away, it was likely a potential interrupt conflict. It's not a good idea to leave it permanently set at 0 since it will put more load on the CPU.

16.13 "Input/output error" from setserial, stty, pppd, etc.

This means that communication with the serial port isn't working right. It could mean that there isn't any serial port at the IO address that setserial thinks your port is at. It could also be an interrupt conflict (or an IO address conflict). It also may mean that the serial port is in use (busy or opened) and thus the attempt to get/set parameters by setserial or stty failed. It will also happen if you make a typo in the serial port name such as typing "ttys" instead of "ttyS".

16.14 "LSR safety check engaged"

LSR is the name of a hardware register. It's claimed that this means there is no serial port at the specified address.

16.15 Overrun errors on serial port

This is an overrun of the hardware FIFO buffer and you can't increase its size. Bug note (reported in 2002): Due to a bug in some kernel 2.4 versions, the port number may be missing and you will only see "ttyS" (no port number). But if devfs notation such as "tts/2" is being used, there is no bug. See

16.16 Port gets characters only sporadically

There could be some other program running on the port. Use "top" (provided you've set it to display the port number) or type "ps −alxw". Look at the results to see if the port is being used by another program. Be on the lookout for the gpm mouse program which often runs on a serial port.

16.17 Troubleshooting Tools

These are some of the programs you might want to use in troubleshooting:

- "lsof /dev/ttyS*" will list serial ports which are open.
- "setserial" shows and sets the low−level hardware configuration of a port (what the driver thinks it is). See What is Setserial
- "stty" shows and sets the configuration of a port (except for that handled by "setserial"). See the section Stty

- "modemstat" or "statserial" will 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.
- "lspci" shows the actual IRQs, etc. of hardware on the PCI bus.
- "pnpdump −−dumpregs" shows the actual IRQs, etc. of hardware for PnP devices on the ISA bus.
- Some "files" in the /proc tree (such as ioports, interrupts, and tty/driver/serial).

16.18 Almost all characters are wrong; Many missing or many extras

Perhaps a baud mismatch. If one port sends at twice the speed that the other port is set to receive, then every two characters sent will be received as one character. Each bit of this received character will be a sample of two bits of what is sent and will be wrong. Also, only half the characters sent seem to get received. For flow in the reverse direction, it's just the opposite. Twice as many characters get received than were sent. A worse mismatch will produce even worse results.

A speed mismatch is not likely to happen with a modem since the modem autodetects the speed. One cause of a mismatch may be due to serial port hardware that has been set to run at very fast speeds. It may actually operate at a speed say 8 times that of which you (or an application) set it via software. See Very High Speeds

17. Interrupt Problem Details

While the section Troubleshooting lists problems by symptom, this section explains what will happen if interrupts are set incorrectly. This section helps you understand what caused the symptom, what other symptoms might be due to the same problem, and what to do about it.

17.1 Types of interrupt problems

The "setserial" program will show you how serial driver thinks the interrupts are set. If the serial driver (and setserial) has it right then everything regarding interrupts should be OK. Of course a /dev/ttyS must exist for the device and Plug−and−Play (or jumpers) must have set an address and IRQ in the hardware. Linux will not knowingly permit an interrupt conflict and you will get a "Device or resource busy" error message if you attempt to do something that would create a conflict.

Since the kernel tries to avoid interrupt conflicts and gives you the "resource busy" message if you try to create a conflict, how can interrupt conflicts happen? Easy. "setserial" may have it wrong and erroneously predicts no conflict when there will actually be a real conflict based on what is set in the hardware. When this happens there will be no "... busy" message but a conflict will physically happen. Performance is likely to be extremely slow. Both devices will send identical interrupt signals on the same wire and the CPU will erroneously think that the interrupts only come from one device. This will be explained in detail in the following sections.

Linux doesn't complain when you assign two devices the same IRQ provided that neither device is in use. As each device starts up (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 in use, you'll see this "... busy" error message. Thus if two devices use the same IRQ and you start up only one of the devices, everything is OK. But when you next try to start the second device (without quitting the first device) you get "... busy" error

message.

17.2 Symptoms of Mis−set or Conflicting Interrupts

The symptoms depend on whether or not you have a modern serial port with FIFO buffers or an obsolete serial port without FIFO buffers. It's important to understand the symptoms for the obsolete ones also since sometimes modern ports seem to behave that way.

For the obsolete serial ports, only one character gets thru every several seconds. This is so slow that it seems almost like nothing is working (especially if the character that gets thru is invisible (such a space or newline). For the modern ports with FIFO buffers you will likely see bursts of up to 16 characters every several seconds.

If you have a modem on the port and dial a number, it seemingly may not connect since the CONNECT message may not make it thru. But after a long wait it may finally connect and you may see part of a login message (or the like). The response from your side of the connection may be so delayed that the other side gives up and disconnects you, resulting in a NO CARRIER message.

If you use minicom, a common test to see if things are working is to type the simplest "AT" command and see if the modem responds. Typing just at<enter> should normally (if interrupts are OK) result in an immediate "OK" response from the modem. With bad interrupts you type at<enter> and may see nothing. But then after 10 seconds or so you see the cursor drop down one line. What is going on is that the FIFO is behaving like it can only hold one byte. The "at" you typed caused it to overrun and both letters were lost. But the final <enter> eventually got thru and you "see" this invisible character by noticing that the cursor jumped down one line. If you were to type a single letter and then wait about 10 seconds, you should see it echo back to the screen. This is fine if your typing speed is less that one word per minute :−)

17.3 Mis−set Interrupts

If you don't understand what an interrupt does see Interrupts. If a serial port has one IRQ set in the hardware but a different one set in the device driver, the device driver will not catch any interrupts sent by the serial port. Since the serial port uses interrupts to call its driver to service the port (fetching bytes from its 16−byte receive buffer or putting another 16−bytes in its transmit buffer) one might expect that the serial port would not work at all.

But it still may work anyway −−sort of. Why? Well, besides the interrupt method of servicing the port there's a slow polling method that doesn't need interrupts. The way it works is that every so often the device driver checks the serial port to see if it needs anything such as if it has some bytes that need fetching from its receive buffer. If interrupts don't work, the serial driver falls back to this polling method. But this polling method was not intended to be used a substitute for interrupts. It's so slow that it's not practical to use and may cause buffer overruns. Its purpose may have been to get things going again if just one interrupt is lost or fails to do the right thing. It's also useful in showing you that interrupts have failed. Don't confuse this slow polling method with the fast polling method that operates on ports that have their IRQs set to 0.

For the 16−byte transmit buffer, 16 bytes will be transmitted and then it will wait until the next polling takes place (several seconds later) before the next 16 bytes are sent out. Thus transmission is very slow and in small chunks. Receiving is slow too since bytes that are received by the receive buffer are likely to remain there for several seconds until it is polled.

This explains why it takes so long before you see what you typed. When you type say AT to a modem, the AT goes out the serial port to the modem. The modem then echos the AT back thru the serial port to the screen. Thus the AT characters have to pass twice thru the serial port. Normally this happens so fast that AT seems to appear on the screen at the same time you hit the keys on the keyboard. With slow polling delays at the serial port, you don't see what you typed until many seconds later.

What about overruns of the 16−byte receive buffer? This will happen with an external modem since the modem just sends to the serial port at high speed which is likely to overrun the 16−byte buffer. But for an internal modem, the serial port is on the same card and it's likely to check that this receive buffer has room for more bytes before putting received bytes into it. In this case there will be no overrun of this receive buffer, but text will just appear on your screen in 16−byte chunks spaced at intervals of several seconds.

Even with an external modem you might not get overruns. If just a few characters (under 16) are sent you don't get overruns since the buffer likely has room for them. But attempts to send a larger number of bytes from your modem to your screen may result in overruns. However, more than 16 (with no gaps) can get thru without overruns if the timing is right. For example, suppose a burst of 32 bytes is sent into the port from the external cable. The polling might just happen after the first 16 bytes came in so it would pick up these 16 bytes OK. Then there would be space for the next 16 bytes so that entire 32 bytes gets thru OK. While this scenario is not very likely, similar cases where 17 to 31 bytes make thru are more likely. But it's even more likely that only an occasional 16−byte chunk will get thru with possible loss of data.

If you have an obsolete serial port with only a 1−byte buffer (or it's been incorrectly set to work like a 1−byte buffer) then the situation will be much worse than described above and only one character will occasionally make it thru the port. Every character received causes an overrun (and is lost) except for the last character received. This character is likely to be just a line−feed since this is often the last character to be transmitted in a burst of characters sent to your screen. Thus you may type AT<return> to the modem but never see AT on the screen. All you see several seconds later is that the cursor drops down one line (a line feed). This has happened to me with a 16−byte FIFO buffer that was behaving like a 1−byte buffer.

When a communication program starts up, it expects interrupts to be working. It's not geared to using this slow polling−like mode of operation. Thus all sorts of mistakes may be made such as setting up the serial port and/or modem incorrectly. It may fail to realize when a connection has been made. If a script is being used for login, it may fail (caused by timeout) due to the polling delays.

17.4 Interrupt Conflicts

When two devices have the same IRQ number it's called sharing interrupts. Under some conditions this sharing works out OK. Starting with kernel version 2.2, ISA serial ports may, if the hardware is designed for this, share interrupts with other serial ports. Devices on the PCI bus may share the same IRQ interrupt with other devices on the PCI bus (provided the software supports this). In other cases where there is potential for conflict, there should be no problem if no two devices with the same IRQ are ever "in use" at the same time. More precisely, "in use" really means "open" (in programmer jargon). In cases other than the exceptions mentioned above (unless special software and hardware permit sharing), sharing is not allowed and conflicts arise if sharing is attempted.

Even if two processes with conflicting IRQs run at the same time, one of the devices will likely have its interrupts caught by its device driver and may work OK. The other device will not have its interrupts caught by the correct driver and will likely behave just like a process with mis−set interrupts. See Mis−set Interrupts for more details.

17.5 Resolving Interrupt Problems

If you are getting a very slow response as described above, then one test is to change the IRQ to 0 (uses fast polling instead of interrupts) and see if the problem goes away. Note that the polling due to IRQ=0 is orders of magnitude faster than the slow "polling" due to bad interrupts. If IRQ=0 seems to fix the problem, then there was likely something wrong with the interrupts. Using IRQ=0 is very resource intensive and is only a temporary fix. You should try to find the cause of the interrupt problem and not permanently use IRQ=0.

Check /proc/interrupts to see if the IRQ is currently in use by another process. If it's in use by another serial port you could try "top" (type f and then enable the TTY display) or "ps −e" to find out which serial ports are in use. If you suspect that setserial has a wrong IRQ then see What is the current IO address and IRQ of my Serial Port ?

---

18. What Are UARTs? How Do They Affect Performance?

18.1 Introduction to UARTS

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 its 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 its 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.

18.2 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, and higher numbers)

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.

18.3 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 obsolete 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 obsolete 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".

FIFO UARTs help solve this problem. 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 to 16) bytes at a time. This is a significant advantage over the obsolete UARTs, which only had 1 byte buffers. The CPU receives less interrupts, and is free to do other things. Data is rarely lost. Note that the interrupt threshold of FIFO buffers (trigger level) may be set at less than 14. 1, 4 and 8 are other possible choices. As of late 2000 there was no way the Linux user could set these directly (setserial can't do it). While many 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 gets full (at say a "trigger level" of 14 bytes for a 16−byte buffer). This allows room for a couple more bytes to be received before the interrupt service routine is able to actually fetch all these bytes. The trigger level may be set to various permitted values by kernel software. A trigger level of 1 will be almost like an obsolete UART (except that it still has room for 15 more bytes after it issues the interrupt).

Now consider the case where you're on the Internet. It's just sent you a short webpage of text. All of this came in thru the serial port. If you had a 16−byte buffer on the serial port which held back characters until it had 14 of them, some of the last several characters on the screen might be missing as the FIFO buffer waited to get the 14th character. But the 14th character doesn't arrive since you've been sent the entire page (over the phone line) and there are no more characters to send to you. It could be that these last characters are part of the HTML formatting, etc. and are not characters to display on the screen but you don't want to lose format either.

There is a "timeout" to prevent the above problem. 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 anyway. This results in fetching all of the characters in the FIFO buffer, even if only a few (or only one) are present. There is also "timeout" for the transmit buffer as well.

18.4 Why FIFO Buffers are Small

You may wonder why the FIFO buffers are not larger. After all, memory is cheap and it wouldn't cost much more to use buffers in the kilo−byte range. The reason is flow control. Flow control stops the flow of data (bytes) on serial line when necessary. If a stop signal is sent to serial port, then the stop request is handled by software (even if the flow control is "hardware"). The serial port hardware knows nothing about flow control.

If the serial port buffer contains 64 bytes ready to send when it receives a flow control signal to stop sending, it will send out the 64 bytes anyway in violation of the stop request. There is no stopping it since it doesn't know about flow control. If the buffer was large, then many more bytes would be sent in violation of flow control's request to stop.

18.5 UART Model Numbers

Here's a list of some 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; 115.2 kbps standard, many support 230.4 or 460.8 kbps
- 16650: 32−byte buffers. 460.8 kbps
- 16750: 64−byte buffer for send, 56−byte for receive. 921.6 kbps
- 16850, 16C850: 128−byte buffers. 460.8 kbps or 1.5 mbps
- 16950
- Hayes ESP: 1k−byte buffers.

For V.90 56k modems, it may be a several percent faster with a 16650 (especially if you are downloading large 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.

Many 486 PCs (old) and all Pentiums (or the like) should have 16550As (usually called just 16550's) with FIFOs. Some better motherboards today (2000) even have 16650s. For replacing obsolete UARTs with newer ones in pre 1990 hardware see the Appendix: Obsolete ...

---

19. Pinout and Signals

19.1 Pinout of 9−pin and 25−pin serial connectors

The pin numbers are often engraved in the plastic of the connector but you may need a magnifying glass to read them. Note DCD is sometimes labeled CD. The numbering of the pins on a female connector is read from right to left, starting with 1 in the upper right corner (instead of 1 in the upper left corner for the male connector as shown below). −−> direction is out of PC.

```
 _________                                  _______________________________
\1 2 3 4 5/  Looking at pins   \1  2  3  4  5  6  7  8  9  10 11 12 13/
```

```
   \6 7 8 9/  on male connector   \14 15 16 17 18 19 20 21 22 23 24 25/
    ------                        -----------------------------------
Pin #   Pin #   Acronym  Full-Name   Direction  What-it-May-Do/Mean
9-pin   25-pin
 3       2      TxD     Transmit Data     -->   Transmits bytes out of PC
 2       3      RxD     Receive Data      <--        Receives bytes into PC
 7       4      RTS     Request To Send   -->   RTS/CTS flow control
 8       5      CTS     Clear To Send     <--        RTS/CTS flow control
 6       6      DSR     Data Set Ready    <--        I'm ready to communicate
 4      20      DTR     Data Terminal Ready-->  I'm ready to communicate
 1       8      DCD     Data Carrier Detect<--       Modem connected to another
 9      22      RI      Ring Indicator    <--        Telephone line ringing
 5       7      SG      Signal Ground

  9-Pin DB9 Connector                   25-Pin DB-25 Connector
1     DCD   Carrier Detect             1           Chassis Ground
2     RxD   Receive Data               2     TxD   Transmit Data
3     TxD   Transmit Data              3     RxD   Receive Data
4     DTR   Data Terminal Ready        4     RTS   Request To Send
5     SG    Signal Ground              5     CTS   Clear To Send
6     DSR   Data Set Ready            6     DSR   Data Set Ready
7     RTS   Request To Send            7     SG    Signal Ground
8     CTS   Clear To Send              8     DCD   Carrier Detect
9     RI    Ring Indicator            20     DTR   Data Terminal Ready
                                      22     RI    Ring Indicator
```

19.2 Signals May Have No Fixed Meaning

Only 3 of the 9 pins have a fixed assignment: transmit, receive and signal ground. This is fixed by the hardware and you can't change it. But the other signal lines are controlled by software and may do (and mean) almost anything at all. However they can only be in one of two states: asserted (+12 volts) or negated (−12 volts). Asserted is "on" and negated is "off". For example, Linux software may command that DTR be negated and the hardware only carries out this command and puts −12 volts on the DTR pin. A modem (or other device) that receives this DTR signal may do various things. If a modem has been configured a certain way it will hang up the telephone line when DTR is negated. In other cases it may ignore this signal or do something else when DTR is negated (turned off).

It's like this for all the 6 signal lines. The hardware only sends and receives the signals, but what action (if any) they perform is up to the Linux software and the configuration/design of devices that you connect to the serial port. However, most pins have certain functions which they normally perform but this may vary with the operating system and the device driver configuration. Under Linux, one may modify the source code to make these signal lines behave differently (some people have).

19.3 Cabling Between Serial Ports

A cable from a serial port always connects to another serial port. An external modem or other device that connects to the serial port has a serial port built into it. For modems, the cable is always straight thru: pin 2 goes to pin 2, etc. The modem is said to be DCE (Data Communications Equipment) and the computer is said to be DTE (Data Terminal Equipment). Thus for connecting DTE−to−DCE you use straight−thru cable. For connecting DTE−to−DTE you must use a null−modem cable (also called a crossover cable). There are many ways to wire such cable (see examples in Text−Terminal−HOWTO subsection: "Direct Cable Connection")

There are good reasons why it works this way. One reason is that the signals are unidirectional. If pin 2 sends a signal out of it (but is unable to receive any signal) then obviously you can't connect it to pin 2 of the same

type of device. If you did, they would both send out signals on the same wire to each other but neither would be able to receive any signal. There are two ways to deal with this situation. One way is to have a two different types of equipment where pin 2 of the first type sends the signal to pin 2 of the second type (which receives the signal). That's the way it's done when you connect a PC (DTE) to a modem (DCE). There's a second way to do this without having two different types of equipment: Connect pin sending pin 2 to a receiving pin 3 on same type of equipment. That's the way it's done when you connect 2 PCs together or a PC to a terminal (DTE−to−DTE). The cable used for this is called a null−modem cable since it connects two PCs without use of a modem. A null−modem cable may also be called a cross−over cable since the wires between pins 2 and 3 cross over each other (if you draw them on a sheet of paper). The above example is for a 25 pin connector but for a 9−pin connector the pin numbers are just the opposite.

The serial pin designations were originally intended for connecting a dumb terminal to a modem. The terminal was DTE (Data Terminal Equipment) and the modem was DCE (Data Communication Equipment). Today the PC is usually used as DTE instead of a terminal (but real terminals may still be used this way). The names of the pins are the same on both DTE and DCE. The words: "receive" and "transmit" are from the "point of view" of the PC (DTE). The transmit pin from the PC transmits to the "transmit" pin of the modem (but actually the modem is receiving the data from this pin so from the point of view of the modem it would be a receive pin).

The serial port was originally intended to be used for connecting DTE to DCE which makes cabling simple: just use a straight−thru cable. Thus when one connects a modem one seldom needs to worry about which pin is which. But people wanted to connect DTE to DTE (for example a computer to a terminal) and various ways were found to do this by fabricating various types of special null−modem cables. In this case what pin connects to what pin becomes significant.

19.4 RTS/CTS and DTR/DSR Flow Control

This is "hardware" flow control. Flow control was previously explained in the Flow Control subsection but the pins and voltage signals were not. Linux only supports RTS/CTS flow control at present (but a special driver may exist for a specific application which supports DTR/DSR flow control). Only RTS/CTS flow control will be discussed since DTR/DSR flow control works the same way. To get RTS/CTS flow control one needs to either select hardware flow control in an application program or use the command: stty −F /dev/ttyS2 crtscts (or the like). This enables RTS/CTS hardware flow control in the Linux device driver.

Then when a DTE (such as a PC) wants to stop the flow into it, it negates RTS. Negated "Request To Send" (−12 volts) means "request NOT to send to me" (stop sending). When the PC is ready for more bytes it asserts RTS (+12 volts) and the flow of bytes to it resumes. Flow control signals are always sent in a direction opposite to the flow of bytes that is being controlled. DCE equipment (modems) works the same way but sends the stop signal out the CTS pin. Thus it's RTS/CTS flow control using 2 lines.

On what pins is this stop signal received? That depends on whether we have a DCE−DTE connection or a DTE−DTE connection. For DCE−DTE it's a straight−thru connection so obviously the signal is received on a pin with the same name as the pin it's sent out from. It's RTS−−>RTS (PC to modem) and CTS<−−CTS (modem to PC). For DTE−to−DTE the connection is also easy to figure out. The RTS pin always sends and the CTS pin always receives. Assume that we connect two PCs (PC1 and PC2) together via their serial ports. Then it's RTS(PC1)−−>CTS(PC2) and CTS(PC1)<−−RTS(PC2). In other words RTS and CTS cross over. Such a cable (with other signals crossed over as well) is called a "null modem" cable. See Cabling Between Serial Ports

What is sometimes confusing is that there is the original use of RTS where it means about the opposite of the previous explanation above. This original meaning is: I Request To Send to you. This request was intended to be sent from a terminal (or computer) to a modem which, if it decided to grant the request, would send back an asserted CTS from its CTS pin to the CTS pin of the computer: You are Cleared To Send to me. Note that in contrast to the modern RTS/CTS bi−directional flow control, this only protects the flow in one direction: from the computer (or terminal) to the modem. This original use appears to be little used today on modern equipment (including modems).

The DTR and DSR Pins

Just like RTS and CTS, these pins are paired. For DTE−to−DTE connections they are likely to cross over. There are two ways to use these pins. One way is to use them as a substitute for RTS/CTS flow control. The DTR pin is just like the RTS pin while the DSR pin behaves like the CTS pin. Although Linux doesn't support DTR/DSR flow control, it can be obtained by connecting the RTS/CTS pins at the PC to the DSR/DTR pins at the device that uses DTR/DSR flow control. DTR flow control is the same as DTR/DSR flow control but it's only one−way and only uses the DTR pin at the device. Many text terminals and some printers use DTR/DSR (or just DTR) flow control. In the future, Linux may support DTR/DSR flow control. The software has already been written but it's not clear when (or if) it will incorporated into the serial driver.

The normal use of DTR and DSR (not for flow control) is as follows: A device asserting DTR says that its powered on and ready to operate. For a modem, the meaning of a DTR signal from the PC depends on how the modem is configured. Negating DTR is sometimes called "hanging up" but it doesn't always do this. One way to "hang up" (negate DTR) is to set the baud rate to 0 using the command "stty 0". Trying to do this from a "foreign" terminal may not work due to the two−interface problem. See Two interfaces at a terminal. For internal modem−serial_ports it worked OK with a port using minicom but didn't work if the port was using wvdial. Why?

19.5 Preventing a Port From Opening

If "stty −clocal" (or getty is used with the "local" flag negated) then a serial port can't open until DCD gets an assert (+12 volts) signal.

---

20. Voltage Waveshapes

20.1 Voltage for a Bit

At the EIA−232 serial port, voltages are bipolar (positive or negative with respect to ground) and should be about 12 volts in magnitude (some are 5 or 10 volts). For the transmit and receive pins +12 volts is a 0−bit (sometimes called "space") and −12 volts is a 1−bit (sometimes called "mark"). This is known as inverted logic since normally a 0−bit is both false and negative while a one is normally both true and positive. Although the receive and transmit pins are inverted logic, other pins (modem control lines) are normal logic with a positive voltage being true (or "on" or "asserted") and a negative voltage being false (or "off" or "negated"). Zero voltage has no meaning (except it usually means that the unit is powered off).

A range of voltages is allowed. The specs say the magnitude of a transmitted signal should be between 5 and 15 volts but must never exceed 25 V. Any voltage received under 3 V is undefined (but some devices will accept a lower voltage as valid). One sometimes sees erroneous claims that the voltage is commonly 5 volts (or even 3 volts) but it's usually 11−12 volts. If you are using a EIA−422 port on a Mac computer as an EIA−232 (requires a special cable) or EIA−423 then the voltage will actually be only 5 V. The discussion here

assumes 12 V.

Note that normal computer logic normally is just a few volts (5 volts was once the standard) so that if you try to use test equipment designed for testing 3−5 volt computer logic (TTL) on the 12 volts of a serial port, it may damage the test equipment.

20.2 Voltage Sequence for a Byte

The transmit pin (TxD) is held at −12 V (mark) at idle when nothing is being sent. To start a byte it jumps to +12 V (space) for the start bit and remains at +12 V for the duration (period) of the start bit. Next comes the low−order bit of the data byte. If it's a 0−bit nothing changes and the line remains at +12 V for another bit−period. If it's a 1−bit the voltage jumps from +12 to −12 V. After that comes the next bit (−12 V if a 1 or +12 V if a 0), etc., etc. After the last data bit a parity bit may be sent and then a −12 V (mark) stop bit. Then the line remains at −12 V (idle) until the next start bit. Note that there is no return to 0 volts and thus there is no simple way (except by a synchronizing signal) to tell where one bit ends and the next one begins for the case where 2 consecutive bits are the same polarity (both zero or both one).

A 2nd stop bit would also be −12 V, just the same as the first stop bit. Since there is no signal to mark the boundaries between these bits, the only effect of the 2nd stop bit is that the line must remain at −12 V idle twice as long. The receiver has no way of detecting the difference between a 2nd stop bit and a longer idle time between bytes. Thus communications works OK if one end uses one stop bit and the other end uses 2 stop bits, but using only one stop bit is obviously faster. In rare cases 1 1/2 stop bits are used. This means that the line is kept at −12 V for 1 1/2 time periods (like a stop bit 50% wider than normal).

20.3 Parity Explained

Characters are normally transmitted with either 7 or 8 bits of data. An additional parity bit may (or may not) be appended to this resulting in a byte length of 7, 8 or 9 bits. Some terminal emulators and older terminals do not allow 9 bits. Some prohibit 9 bits if 2 stop bits are used (since this would make the total number of bits too large: 12 bits total after adding the start bit).

The parity may be set to odd, even or none (mark and space parity may be options on some terminals or other serial devices). With odd parity, the parity bit is selected so that the number of 1−bits in a byte, including the parity bit, is odd. If a such a byte gets corrupted by a bit being flipped, the result is an illegal byte of even parity. This error will be detected and if it's an incoming byte to the terminal an error−character symbol will appear on the screen. Even parity works in a similar manner with all legal bytes (including the parity bit) having an even number of 1−bits. During set−up, the number of bits per character usually means only the number of data bits per byte (7 for true ASCII and 8 for various ISO character sets).

A "mark" is a 1−bit (or logic 1) and a "space" is a 0−bit (or logic 0). For mark parity, the parity bit is always a one−bit. For space parity it's always a zero−bit. Mark or space parity (also known as "sticky parity") only wastes bandwidth and should be avoided if feasible. The `stty` command can't set sticky parity but it's supported by serial hardware and can be dealt with by programming in C. "No parity" means that no parity bit is added. For terminals that don't permit 9 bit bytes, "no parity" must be selected when using 8 bit character sets since there is no room for a parity bit.

20.4 Forming a Byte (Framing)

In serial transmission of bytes via EIA−232 ports, the low−order bit is always sent first. Serial ports on PC's use asynchronous communication where there is a start bit and a stop bit to mark the beginning and end of a byte. This is called framing and the framed byte is sometimes called a frame. As a result a total of 9, 10, or 11 bits are sent per byte with 10 being the most common. 8−N−1 means 8 data bits, No parity, 1 stop bit. This adds up to 10 bits total when one counts the start bit. One stop bit is almost universally used. At 110 bits/sec (and sometimes at 300 bits/sec) 2 stop bits were once used but today the 2nd stop bit is used only in very unusual situations (or by mistake since it still works OK that way but wastes bandwidth).

Don't confuse this type of framing with the framing used for a packet of bytes on a network. The serial port just frames every byte. For a network many bytes are framed into a packet (sometimes called a frame). For a network frame, instead of a start bit, there is a sequence of bytes called a header. On a network that uses serial ports (with modems), a report of a frame error usually refers to a multi−byte frame and not the serial port frame of a single byte.

20.5 How "Asynchronous" is Synchronized

The EIA−232 serial port as implemented on PC is asynchronous which in effect means that there is no "clock" signal sent with "ticks" to mark when each bit is sent.. There are only two states of the transmit (or receive) wire: mark (−12 V) or space (+12 V). There is no state of 0 V. Thus a sequence of 1−bits is transmitted by just a steady −12 V with no markers of any kind between bits. For the receiver to detect individual bits it must always have a clock signal which is in synchronization with the transmitter clock. Such a clock would generate a "tick" in synchronization with each transmitted (or received) bit.

For asynchronous transmission, synchronization is achieved by framing each byte with a start bit and a stop bit (done by hardware). The receiver listens on the negative line for a positive start bit and when it detects one it starts its clock ticking. It uses this clock tick to time the reading of the next 7, 8 or 9 bits. (It actually is a little more complex than this since several samples of a bit are normally taken and this requires additional timing ticks.) Then the stop bit is read, the clock stops and the receiver waits for the next start bit. Thus async is actually synchronized during the reception of a single byte but there is no synchronization between one byte and the next byte.

---

21. Other Serial Devices (not async EIA−232)

21.1 Successors to EIA−232

A number of EIA standards have been established for higher speeds and longer distances using twisted−pair (balanced) technology. Balanced transmission make possible higher speeds, and can be a hundred times faster than unbalanced EIA−232. For a given speed, the distance (maximum cable length) may be many times longer with twisted pair. But PC−s keep being made with the "obsolete" EIA−232 since it works OK with modems and mice since the cable length is short. If this appears in the latest version of this HOWTO, please let me know if any of the non−EIA−232 listed below are supported by Linux.

High speed serial ports (over 460.8 kbps) will often support both EIA−232 and EIA−485/EIA−422 modes. At such high speeds EIA−232 is not of much use (except for a very short cable).

21.2 EIA−422−A (balanced) and EIA−423−A (unbalanced)

EIA−423 is just like the unbalanced EIA−232 except that the voltage is only 5 volts. Since this falls within

EIA−232 specs it can be connected to a EIA−232 port. Its specs call for somewhat higher speeds than the EIA−232 (but this may be of little help on a long run where it's the unbalance that causes interference). Since EIA−423 is not much of an improvement over EIA−232, it is seldom used except on old Mac computers.

EIA−422 is twisted pair (known as "balanced" or "differential) and is (per specs) exactly 100 times as fast as EIA−423 (which in turn is somewhat faster than EIA−232). Apple's Mac computer used it prior to mid−1998 with its EIA−232/EIA−422 port. The Mac used a small round "mini−DIN−8" connector and named these serial ports as "modem port", "printer port", and/or "GeoPort".

Mac also provided conventional EIA−232 but at only at 5 volts (which is still legal EIA−232). To make it work like at EIA−232 one must use a special cable which (signal) grounds RxD+ (one side of a balanced pair) and use RxD− as the receive pin. While TxD− is used as the transmit pin, for some reason TxD+ should not be grounded. See Macintosh Communications FAQ. However, due to the fact that Macs (and upgrades for them) cost more than PC's, they are not widely as host computers for Linux.

21.3 EIA−485

This is like EIA−422 (balanced = differential). It is half−duplex. It's not just point−to−point but is like ethernet or the USB since all devices (nodes) on it share the same "bus". It may be used for a multidrop LAN (up to 32 nodes or more). Unfortunately, Linux currently doesn't support this and you can only use it under Linux only for point−to−point where it behaves like EIA−232. So read further only if you are curious about how its features would work if only Linux supported them.

Since many nodes share the same twisted pair, there's a need to use the electrical tri−state mode. Thus, besides the 0 and 1 binary states, there is also an open circuit state to permit other nodes to use the twisted pair line. Instead of a transmitter keeping a 1−state voltage on the line during line idle, the line is open circuited and all nodes just listen (receive mode).

The most common architecture is master/slave. The master polls the slaves to see if they have anything to send. A slave can only transmit just after it's been polled. But EIA−485 is just an electrical specification and doesn't specify any protocol for the master/slave interaction. In fact, it doesn't even specify that there must be a master and slaves. So various protocols have been used. Based on a discussion of 485 on the linux−serial mailing list in March 2003, it seems likely that none of these master/slave protocols are currently supported by Linux.

There is an alternative implementation where two pair of wires are used for sending data. One pair is only for the Master to send to the Slaves. Since no one transmits on this line except the master, there is no need for it to be tri−state. Thus the Master may just be EIA−232 but the slaves must still be EIA−485. See http://www.hw.cz/english/docs/rs485/rs485.html for more details.

21.4 EIA−530

EIA−530−A (balanced but can also be used unbalanced) at 2Mbits/s (balanced) was intended to be a replacement for EIA−232 but few have been installed. It uses the same 25−pin connector as EIA−232.

21.5 EIA−612/613

The High Speed Serial Interface (HSSI = EIA−612/613) uses a 50−pin connector and goes up to about 50 Mbits/s but the distance is limited to only several meters. For Linux there are PCI cards supporting HSSI. The

companies that sell the cards often provide (or point you to) a Linux driver. A howto or the like is needed for this topic.

21.6 The Universal Serial Bus (USB)

The Universal Serial Bus (USB) is being built into PCI chips. Newer PC's have them. It is 12 Mbps (with 200 Mbps planned) over a twisted pair with a 4−pin connector (2 wires are power supply). It also is limited to short distances of at most 5 meters (depends on configuration). Linux supports the bus, although not all devices that can plug into the bus are supported.

It is synchronous and transmits in special packets like a network. Just like a network, it can have several devices physically attached to it, including serial ports. Each device on it gets a time−slice of exclusive use for a short time. A device can also be guaranteed the use of the bus at fixed intervals. One device can monopolize it if no other device wants to use it. It's not simple to describe in detail.

For serial ports on the USB bus, there are numerous configuration options to use when compiling the kernel. They all start with: CONFIG_USB_SERIAL. Each one is usually for a certain brand/model of serial port, although generic is also an option. See the Configuration Help file in the kernel documentation.

For documentation, see the USB directory in /usr/share/doc/kernel ... and look at the file: usb−serial.txt. The modules that support usb serial devices are found in the modules tree: kernel/drivers/usb/serial. It would be nice to have a HOWTO on the USB. See also http://www.linux−usb.org and/or http://.www.qbik.ch/usb/.

21.7 Firewire

Firewire (IEEE 1394) is something like the USB only faster (800 Mbps is planned). The protocol on the bus is claimed to be more efficient than USB's. It uses two twisted pair for data plus two power conductors (6 conductors in all). A variants uses only 4 conductors. You may compile firewire support into the Linux kernel. Like USB, it's also limited to short distances.

21.8 MIDI

Sound cards often have a 15−pin game port connector used for MIDI. They are for connecting a musical keyboard to a PC so that you can create musical recordings. You could also connect a MIDI sound system. The MIDI standard uses 31250 baud (1M/32) which is not available on an ordinary serial port. Some MIDI devices are designed so that they can be connected directly to an ordinary serial port.

Besides the 15−pin connector, a 5−pin DIN connector is also a MIDI standard but the flow of sound is only one way thru it so for bidirectional sound you need 2 of them. Breakout cables often have a 15−pin connector on one end and 2 or more 5−pin connectors on the other end. The /dev/midi00 is for MIDI.

21.9 Synchronization & Synchronous

Beside the asynchronous EIA−232 (and others) there are a number of synchronous serial port standards. In fact EIA−232 includes synchronous specifications but they aren't normally implemented for serial ports on PC's. But first we'll explain what a synchronous means.

Defining Asynchronous vs Synchronous

Asynchronous (async) means "not synchronous". In practice, an async signal is what the async serial port sends and receives which is a stream of bytes with each byte framed by a start and stop bit. Synchronous (sync) is most everything else. But this doesn't explain the basic concepts.

In theory, synchronous means that bytes are sent out at a constant rate one after another in step with a clock signal tick. There is often a separate wire or channel for sending the clock signal. The clock signal might also be embedded in the transmitted bytes. Asynchronous bytes may be sent out erratically with various time intervals between bytes (like someone typing characters at a keyboard).

When a file is being sent thru the async serial port, the flow of bytes will likely be at the speed of the port (say 115.2k) which is a constant rate. This flow may frequently start and stop due to flow control. Is this sync or async? Ignoring the flow control stops, it might seem like sync since it's a steady flow. But it's not because there is no clock signal and the bytes could have been sent erratically since they are framed by start/stop bits.

Another case is where data bytes (without any start−stop bits) are put into packets with possible erratic spacing between one packet and the next. This is called sync since the bytes within each packet are transmitted synchronously.

Synchronous Communication

Did you ever wonder what all the unused pins are for on a 25−pin connector for the serial port? Most of them are for use in synchronous communication which is seldom implemented in chips for PC's. There are pins for sync timing signals as well as for a sync reverse channel. The EIA−232 spec provides for both sync and async but PC's use a UART (Universal Asynchronous Receiver/Transmitter) chip such as a 16450, 16550A, or 16650 and can't deal with sync. For sync one needs a USRT chip or the equivalent where the "S" stands for Synchronous. A USART chip supports both synchronous and asynchronous. Since sync is a niche market, a sync serial port is likely to be quite expensive.

SCC stands for "Serial Communication Controller" or "Serial Controller Chip". It's likely old terminology and since it doesn't say "sync" or "async" it might support both.

Besides the sync part of the EIA−232, there are various other EIA synchronous standards. For EIA−232, 3 pins of the connector are reserved for clock (or timing) signals. Sometimes it's a modem's task to generate some timing signals making it impossible to use synchronous communications without a synchronous modem (or without a device called a "synchronous modem eliminator" which provides the timing signals).

Although few serial ports are sync, synchronous communication does often take place over telephone lines using modems which use V.42 error correction. This strips off the start/stop bits and puts the data bytes in packets resulting in synchronous operation over the phone line.

---

22. Other Sources of Information

22.1 Books

1. Axleson, Jan: Serial Port Complete, Lakeview Research, Madison, WI, 1998.
2. Black, Uyless D.: Physical Layer Interfaces & Protocols, IEEE Computer Society Press, Los Alamitos, CA, 1996.

3. Campbell, Joe: The RS−232 Solution, 2nd ed., Sybex, 1982.
4. Campbell, Joe: C Programmer's Guide to Serial Communications, 2nd ed., Unknown Publisher, 1993.
5. Levine, Donald: POSIX Programmer's Guide, O'Reilly, 1991.
6. Nelson, Mark: Serial Communications Developer's Guide, 2nd ed., Hungry Minds, 2000.
7. Putnam, Byron W.: RS−232 Simplified, Prentice Hall, 1987.
8. Seyer, Martin D.: RS−232 Made Easy, 2nd ed., Prentice Hall, 1991.
9. Stevens, Richard W.: Advanced Programming in the UNIX Environment, (ISBN 0−201−56317−7; Addison−Wesley)
10. Tischert, Michael & Bruno Jennrich: PC Intern, Abacus 1996. Chapter 7: Serial Ports

Notes re books:

1. "... Complete" has hardware details (including register) but the programming aspect is Window oriented.
2. "Physical Layer ..." covers much more than just EIA−232.

22.2 Serial Software

It's best to use the nearest mirror site, but here's the main sites:
Serial Software for Linux software for the serial ports including getty and port monitors.
Serial Communications for communication programs.

- `irqtune` will give serial port interrupts higher priority to improve performance. Using `hdparm` for hard−disk tuning may help some more.
- `modemstat` and `statserial` show the current state of various modem control lines. See Serial Monitoring/Diagnostics

22.3 Related Linux Documents

- man pages for: `setserial` and `stty`
- Low−Level Terminal Interface part of "GNU C Library Reference manual" (in libc (or glibc) docs package). It covers the detailed meaning of "stty" commands, etc.
- Modem−HOWTO: modems on the serial port
- PPP−HOWTO: help with PPP (using a modem on the serial port)
- Printing−HOWTO: for setting up a serial printer
- Serial−Programming−HOWTO: for some aspects of serial−port programming
- Text−Terminal−HOWTO: how they work and how to install and configure
- UPS−HOWTO: setting up UPS sensors connected to your serial port
- UUCP−HOWTO: for information on setting up UUCP

22.4 Usenet newsgroups:

- comp.os.linux.answers
- comp.os.linux.hardware: Hardware compatibility with the Linux operating system.
- comp.os.linux.networking: Networking and communications under Linux.
- comp.os.linux.setup: Linux installation and system administration.

22.5 Serial Mailing List

The Linux serial mailing list. To join, send email to `majordomo@vger.kernel.org`, 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.

22.6 Internet

- Linux Serial Driver home page Includes info about PCI support.
- Serial Suite by Vern Hoxie is a collection of blurbs about the care and feeding of the Linux serial port plus some simple programs. He also has a Serial-Programming-HOWTO (not yet available from the Linux Documentation Project). Your browser should automatically log you in but if you do it manually login as "anonymous" and use your full e-mail address as the password.
- A white paper discussing serial communications and multiport serial boards was available from Cyclades at `http://www.cyclades.com`.

---

23. Appendix: Obsolete Hardware (prior to 1990) Info

23.1 Replacing obsolete UARTS

Many 486 PCs (old) and all Pentiums (or the like) should have modern 16550As (usually called just 16550's) with FIFOs. 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 built into 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 card on the Internet (few retail stores stock them today) or find a used one.

END OF Serial-HOWTO