

Contents

Foreword	xxi
Preface	xxiii
Acknowledgments	xxix
About the Author	xxx
Chapter 1 Introduction	1
Evolution	2
The GNU Copyleft	3
Kernel.org	4
Mailing Lists and Forums	4
Linux Distributions	5
Looking at the Sources	6
Building the Kernel	10
Loadable Modules	12
Before Starting	14
Chapter 2 A Peek Inside the Kernel	17
Booting Up	18
Kernel Mode and User Mode	30
Process Context and Interrupt Context	30
Kernel Timers	31
HZ and Jiffies	31
Long Delays	33

Short Delays	36
Pentium Time Stamp Counter	36
Real Time Clock	37
Concurrency in the Kernel	39
Spinlocks and Mutexes	39
Atomic Operators	45
Reader-Writer Locks	46
Debugging	48
Process Filesystem	49
Allocating Memory	49
Looking at the Sources	52
Chapter 3 Kernel Facilities	55
Kernel Threads	56
Creating a Kernel Thread	56
Process States and Wait Queues	61
User Mode Helpers	63
Helper Interfaces	65
Linked Lists	65
Hash Lists	72
Work Queues	72
Notifier Chains	74
Completion Interface	78
Kthread Helpers	81
Error-Handling Aids	83
Looking at the Sources	85
Chapter 4 Laying the Groundwork	89
Introducing Devices and Drivers	90
Interrupt Handling	92
Interrupt Context	92
Assigning IRQs	94

Device Example: Roller Wheel	94
Softirqs and Tasklets	99
The Linux Device Model	103
Udev	103
Sysfs, Kobjects, and Device Classes	106
Hotplug and Coldplug	110
Microcode Download	111
Module Autoload	112
Memory Barriers	114
Power Management	114
Looking at the Sources	115
Chapter 5 Character Drivers	119
Char Driver Basics	120
Device Example: System CMOS	121
Driver Initialization	122
Open and Release	127
Exchanging Data	129
Seek	136
Control	137
Sensing Data Availability	139
Poll	139
Fasync	142
Talking to the Parallel Port	145
Device Example: Parallel Port LED Board	146
RTC Subsystem	156
Pseudo Char Drivers	157
Misc Drivers	160
Device Example: Watchdog Timer	160
Character Caveats	166
Looking at the Sources	167

Chapter 6 Serial Drivers	171
Layered Architecture	173
UART Drivers	176
Device Example: Cell Phone	178
RS-485	191
TTY Drivers	192
Line Disciplines	194
Device Example: Touch Controller	195
Looking at the Sources	205
Chapter 7 Input Drivers	207
Input Event Drivers	210
The Evdev Interface	210
Input Device Drivers	216
Serio	217
Keyboards	217
Mice	220
Touch Controllers	227
Accelerometers	228
Output Events	228
Debugging	230
Looking at the Sources	231
Chapter 8 The Inter-Integrated Circuit Protocol	233
What's I ² C/SMBus?	234
I ² C Core	235
Bus Transactions	237
Device Example: EEPROM	238
Initializing	238
Probing the Device	241
Checking Adapter Capabilities	244

Accessing the Device	244
More Methods	246
Device Example: Real Time Clock	247
I2C-dev	251
Hardware Monitoring Using LM-Sensors	251
The Serial Peripheral Interface Bus	251
The 1-Wire Bus	254
Debugging	254
Looking at the Sources	255
Chapter 9 PCMCIA and Compact Flash	257
What's PCMCIA/CF?	258
Linux-PCMCIA Subsystem	260
Host Controller Drivers	262
PCMCIA Core	263
Driver Services	263
Client Drivers	264
Data Structures	264
Device Example: PCMCIA Card	267
Tying the Pieces Together	271
PCMCIA Storage	272
Serial PCMCIA	272
Debugging	273
Looking at the Sources	275
Chapter 10 Peripheral Component Interconnect	277
The PCI Family	278
Addressing and Identification	281
Accessing PCI Regions	285
Configuration Space	285
I/O and Memory	286
Direct Memory Access	288

Device Example: Ethernet-Modem Card	292
Initializing and Probing	293
Data Transfer	301
Debugging	308
Looking at the Sources	308
Chapter 11 Universal Serial Bus	311
USB Architecture	312
Bus Speeds	314
Host Controllers	315
Transfer Types	315
Addressing	316
Linux-USB Subsystem	317
Driver Data Structures	317
The usb_device Structure	318
USB Request Blocks	319
Pipes	321
Descriptor Structures	322
Enumeration	324
Device Example: Telemetry Card	324
Initializing and Probing	325
Accessing Registers	332
Data Transfer	335
Class Drivers	338
Mass Storage	339
USB-Serial	345
Human Interface Devices	348
Bluetooth	348
Gadget Drivers	348
Debugging	349
Looking at the Sources	351

Chapter 12 Video Drivers	355
Display Architecture	356
Linux-Video Subsystem	359
Display Parameters	361
The Frame Buffer API	362
Frame Buffer Drivers	365
Device Example: Navigation System	365
Console Drivers	380
Device Example: Cell Phone Revisited	382
Boot Logo	387
Debugging	387
Looking at the Sources	388
Chapter 13 Audio Drivers	391
Audio Architecture	392
Linux-Sound Subsystem	394
Device Example: MP3 Player	396
Driver Methods and Structures	399
ALSA Programming	409
Debugging	412
Looking at the Sources	412
Chapter 14 Block Drivers	415
Storage Technologies	416
Linux Block I/O Layer	421
I/O Schedulers	422
Block Driver Data Structures and Methods	423
Device Example: Simple Storage Controller	426
Initialization	427
Block Device Operations	430
Disk Access	432

Advanced Topics	434
Debugging	436
Looking at the Sources	437
Chapter 15 Network Interface Cards	439
Driver Data Structures	440
Socket Buffers	441
The Net Device Interface	443
Activation	444
Data Transfer	444
Watchdog	445
Statistics	445
Configuration	446
Bus Specific	448
Talking with Protocol Layers	448
Receive Path	448
Transmit Path	449
Flow Control	449
Buffer Management and Concurrency Control	450
Device Example: Ethernet NIC	451
ISA Network Drivers	457
Asynchronous Transfer Mode	458
Network Throughput	459
Driver Performance	459
Protocol Performance	461
Looking at the Sources	461
Chapter 16 Linux Without Wires	465
Bluetooth	467
BlueZ	469
Device Example: CF Card	471

Device Example: USB Adapter	471
RFCOMM	473
Networking	475
Human Interface Devices	477
Audio	477
Debugging	478
Looking at the Sources	478
Infrared	478
Linux-IrDA	480
Device Example: Super I/O Chip	482
Device Example: IR Dongle	483
IrComm	486
Networking	486
IrDA Sockets	487
Linux Infrared Remote Control	488
Looking at the Sources	489
WiFi	489
Configuration	490
Device Drivers	494
Looking at the Sources	496
Cellular Networking	496
GPRS	496
CDMA	498
Current Trends	500
Chapter 17 Memory Technology Devices	503
What's Flash Memory?	504
Linux-MTD Subsystem	505
Map Drivers	506
Device Example: Handheld	506
NOR Chip Drivers	511

NAND Chip Drivers	513
User Modules	516
Block Device Emulation	516
Char Device Emulation	517
JFFS2	517
YAFFS2	518
MTD-Utills	518
Configuring MTD	519
eXecute In Place	520
The Firmware Hub	520
Debugging	524
Looking at the Sources	524
Chapter 18 Embedding Linux	527
Challenges	528
Component Selection	530
Tool Chains	531
Embedded Bootloaders	531
Memory Layout	535
Kernel Porting	537
Embedded Drivers	538
Flash Memory	538
UART	539
Buttons and Wheels	539
PCMCIA/CF	540
SD/MMC	540
USB	540
RTC	541
Audio	541
Touch Screen	541

Video	541
CPLD/FPGA	542
Connectivity	542
Domain-Specific Electronics	542
More Drivers	543
The Root Filesystem	544
NFS-Mounted Root	544
Compact Middleware	546
Test Infrastructure	548
Debugging	548
Board Rework	549
Debuggers	550
Chapter 19 Drivers in User Space	551
Process Scheduling and Response Times	553
The Original Scheduler	553
The O(1) Scheduler	553
The CFS Scheduler	555
Response Times	555
Accessing I/O Regions	558
Accessing Memory Regions	562
User Mode SCSI	565
User Mode USB	567
User Mode I ² C	571
UIO	573
Looking at the Sources	574
Chapter 20 More Devices and Drivers	577
ECC Reporting	578
Device Example: ECC-Aware Memory Controller	579
Frequency Scaling	583

Embedded Controllers	584
ACPI	585
ISA and MCA	587
FireWire	588
Intelligent Input/Output	589
Amateur Radio	590
Voice over IP	590
High-Speed Interconnects	591
InfiniBand	592
RapidIO	592
Fibre Channel	592
iSCSI	593
Chapter 21 Debugging Device Drivers	595
Kernel Debuggers	596
Entering a Debugger	597
Kernel Debugger (kdb)	598
Kernel GNU Debugger (kgdb)	600
GNU Debugger (gdb)	604
JTAG Debuggers	605
Downloads	609
Kernel Probes	609
Kprobes	609
Jprobes	614
Return Probes	617
Limitations	619
Looking at the Sources	620
Kexec and Kdump	620
Kexec	620
Kexec with Kdump	621

Kdump	622
Looking at the Sources	629
Profiling	629
Kernel Profiling with OProfile	629
Application Profiling with Gprof	633
Tracing	634
Linux Trace Toolkit	634
Linux Test Project	638
User Mode Linux	638
Diagnostic Tools	638
Kernel Hacking Config Options	639
Test Equipment	640
Chapter 22 Maintenance and Delivery	641
Coding Style	642
Change Markers	642
Version Control	643
Consistent Checksums	643
Build Scripts	645
Portable Code	647
Chapter 23 Shutting Down	649
Checklist	650
What Next?	651
Appendix A Linux Assembly	653
Debugging	659
Appendix B Linux and the BIOS	661
Real Mode Calls	662
Protected Mode Calls	665
BIOS and Legacy Drivers	666

Appendix C Seq Files	669
The Seq File Advantage	670
Updating the NVRAM Driver	677
Looking at the Sources	679
Index	681