

TECH TODAY




x64 ASSEMBLY LANGUAGE

STEP-BY-STEP

Programming with Linux[®]

4TH EDITION



JEFF DUNTEMANN

WILEY

x64 Assembly Language Step-by-Step

4TH Edition



x64 Assembly Language Step-by-Step

Programming with Linux®

4TH Edition

Jeff Duntemann

WILEY

Copyright © 2024 by Jeff Duntemann. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
Published simultaneously in Canada and the United Kingdom.

ISBNs: 9781394155248 (Hardback), 9781394155545 (epdf), 9781394155255 (epub)

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at www.wiley.com/go/permission.

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries and may not be used without written permission. Linux is a registered trademark of Linus Torvalds. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

If you believe you've found a mistake in this book, please bring it to our attention by emailing our reader support team at wileysupport@wiley.com with the subject line "Possible Book Errata Submission."

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Control Number: 2023944290

Cover image: © CSA-Printstocks/Getty Images
Cover design: Wiley

To the eternal memory of my father,

Frank W. Duntemann, Engineer

1922–1978

Who said, "When you build 'em right, they fly."

You did. And I do.



About the Author

Jeff Duntemann has had his technical nonfiction and science-fiction work professionally published since 1974. He worked as a programmer for Xerox Corporation and as a technical editor for Ziff-Davis Publishing and Borland International. He launched and edited two print magazines for programmers and has 20 technical books to his credit, including the bestselling *Assembly Language Step-by-Step*. He wrote the “Structured Programming” column in *Dr. Dobbs’s Journal* for four years and has published dozens of technical articles in many magazines. He has a longstanding interest in “strong” artificial intelligence, and most of his fiction explores the possibilities and consequences of strong AI. His other interests include electronics and amateur radio (callsign K7JPD), telescopes, and kites. Jeff lives in Phoenix, Arizona, with his wife Carol.

About the Technical Editor

David Stafford is an enthusiast of low-level programming in assembly language, from 8-bit processors to modern 64-bit multicore architectures. He lives in the Seattle area and works in the field of artificial intelligence for robotics.



Acknowledgments

Thanks are due to a number of people who helped me out as this edition took shape, in various ways. First, thanks to Jim Minatel and Pete Gaughan of Wiley, who got the project underway and made sure it went to completion. Also thanks to David Stafford, who acted as technical editor and provided a constant stream of invaluable advice.

The event shocked me to the core, but antony-jr of GitHub managed to create a working Linux appimage of the quirky and ancient but very accessible Insight debugger, which was pulled from Linux repositories soon after the third edition of this book hit print in 2009. A very big thanks to him for what was likely a *very* peculiar project. You can find his Insight appimage here: <https://appimage.github.io/Insight>.

Abundant thanks also to Dmitriy Manushin, who created SASM, a free assembly language IDE targeted at beginners: <https://dman95.github.io/SASM/english.html>.

When I ran into a weirdness in `glibc`, my wizard crew on Contrapositive Diary helped me figure it out:

- Jim Strickland
- Bill Buhler
- Jason Bucata
- Jonathan O'Neal
- Bruce and Keith (last names not given, and that's OK—the advice was golden)

Finally, and as always, a toast to Carol for her support and sacramental friendship that have enlivened me now for 54 years and enabled me to take on difficult projects like this and see them through to the end, no matter how nuts they made me along the way!



Contents at a Glance

Introduction	xxix
Chapter 1 It's All in the Plan	1
Chapter 2 Alien Bases	11
Chapter 3 Lifting the Hood	41
Chapter 4 Location, Location, Location	73
Chapter 5 The Right to Assemble	103
Chapter 6 A Place to Stand, with Access to Tools	143
Chapter 7 Following Your Instructions	175
Chapter 8 Our Object All Sublime	213
Chapter 9 Bits, Flags, Branches, and Tables	251
Chapter 10 Dividing and Conquering	299
Chapter 11 Strings and Things	377
Chapter 12 Heading Out to C	423
Conclusion: Not the End, But Only the Beginning	489
Appendix A The Return of the Insight Debugger	493
Appendix B Partial x64 Instruction Reference	501
Appendix C Character Set Charts	575
Index	579



Contents

Introduction		xxix
Chapter 1	It's All in the Plan	1
	Another Pleasant Valley Saturday	1
	Steps and Tests	3
	More Than Two Ways?	3
	Computers Think Like Us	4
	Had This Been the Real Thing . . .	5
	Assembly Language Programming As a Square Dance	5
	Assembly Language Programming As a Board Game	6
	Code and Data	8
	Addresses	8
	Metaphor Check!	9
Chapter 2	Alien Bases	11
	The Return of the New Math Monster	11
	Counting in Martian	12
	Dissecting a Martian Number	14
	The Essence of a Number Base	16
	Octal: How the Grinch Stole Eight and Nine	16
	Who Stole Eight and Nine?	17
	Hexadecimal: Solving the Digit Shortage	20
	From Hex to Decimal and from Decimal to Hex	24
	From Hex to Decimal	24
	From Decimal to Hex	25
	Practice. Practice! PRACTICE!	27
	Arithmetic in Hex	28

	Columns and Carries	30
	Subtraction and Borrows	31
	Borrows Across Multiple Columns	33
	What's the Point?	33
	Binary	34
	Values in Binary	36
	Why Binary?	38
	Hexadecimal as Shorthand for Binary	38
	Prepare to Compute	40
Chapter 3	Lifting the Hood	41
	RAXie, We Hardly Knew Ye	41
	Gus to the Rescue	42
	Switches, Transistors, and Memory	43
	One If by Land...	43
	Transistor Switches	44
	The Incredible Shrinking Bit	46
	Random Access	47
	Memory Access Time	49
	Bytes, Words, Double Words, and Quad Words	50
	Pretty Chips All in a Row	51
	The Shop Supervisor and the Assembly Line	54
	Talking to Memory	55
	Riding the Data Bus	56
	The Shop Supervisor's Pockets	57
	The Assembly Line	58
	The Box That Follows a Plan	58
	Fetch and Execute	60
	The Supervisor's Innards	61
	Changing Course	62
	What vs. How: Architecture and Microarchitecture	63
	Evolving Architectures	64
	The Secret Machinery in the Basement	65
	Enter the Plant Manager	67
	Operating Systems: The Corner Office	67
	BIOS: Software, Just Not as Soft	68
	Multitasking Magic	68
	Promotion to Kernel	70
	The Core Explosion	70
	The Plan	72
Chapter 4	Location, Location, Location	73
	The Joy of Memory Models	73
	16 Bits'll Buy You 64 KB	75

The Nature of a Megabyte	78
Backward Compatibility and Virtual 86 Mode	79
16-Bit Blinders	79
The Nature of Segments	80
A Horizon, Not a Place	84
Making 20-Bit Addresses Out of 16-Bit Registers	84
Segment Registers	87
Segment Registers and x64	88
General-Purpose Registers	88
Register Halves	91
The Instruction Pointer	92
The Flags Register	94
Math Coprocessors and Their Registers	94
The Four Major Assembly Programming Models	95
Real-Mode Flat Model	95
Real-Mode Segmented Model	97
32-Bit Protected Mode Flat Model	99
64-Bit Long Mode	101
Chapter 5 The Right to Assemble	103
The Nine and Sixty Ways to Code	103
Files and What’s Inside Them	104
Binary vs. Text Files	105
Looking at Binary File Internals with the GHex Hex Editor	106
Interpreting Raw Data	110
“Endianness”	111
Text In, Code Out	115
Assembly Language	116
Comments	118
Beware “Write-Only” Source Code!	119
Object Code, Linkers, and Libraries	120
Relocatability	123
The Assembly Language Development Process	123
The Discipline of Working Directories	125
Editing the Source Code File	126
Assembling the Source Code File	126
Assembler Errors	127
Back to the Editor	128
Assembler Warnings	129
Linking the Object Code File	130
Linker Errors	131
Testing the EXE File	131
Errors vs. Bugs	132

Are We There Yet?	133	
Debuggers and Debugging	133	
Taking a Trip Down Assembly Lane	134	
Installing the Software	134	
Step 1: Edit the Program in an Editor	137	
Step 2: Assemble the Program with NASM	138	
Step 3: Link the Program with ld	140	
Step 4: Test the Executable File	141	
Step 5: Watch It Run in the Debugger	141	
Chapter 6	A Place to Stand, with Access to Tools	143
Integrated Development Environments	143	
Introducing SASM	146	
Configuring SASM	146	
SASM's Fonts	147	
Using a Compiler to Link	148	
A Quick Tour of SASM	149	
SASM's Editor	152	
What SASM Demands of Your Code	152	
Linux and Terminals	153	
The Linux Console	153	
Character Encoding in Konsole	154	
The Three Standard Unix Files	156	
I/O Redirection	158	
Simple Text Filters	159	
Using Standard Input and Standard Output from Inside SASM	161	
Terminal Control with Escape Sequences	161	
So Why Not GUI Apps?	163	
Using Linux Make	164	
Dependencies	165	
When a File Is Up-to-Date	167	
Chains of Dependencies	167	
Invoking Make	169	
Creating a Custom Key Binding for Make	170	
Using Touch to Force a Build	172	
Debugging with SASM	172	
Pick up Your Tools. . .	174	
Chapter 7	Following Your Instructions	175
Build Yourself a Sandbox	176	
A Minimal NASM Program for SASM	176	
Instructions and Their Operands	178	
Source and Destination Operands	178	

Immediate Data	179
Register Data	181
Memory Data and Effective Addresses	184
Confusing Data and Its Address	185
The Size of Memory Data	185
The Bad Old Days	186
Rally Round the Flags, Boys!	186
Flag Etiquette	190
Watching Flags from SASM	190
Adding and Subtracting One with INC and DEC	191
How Flags Change Program Execution	192
How to Inspect Variables in SASM	194
Signed and Unsigned Values	195
Two's Complement and NEG	196
Sign Extension and MOVSX	198
Implicit Operands and MUL	200
MUL and the Carry Flag	202
Unsigned Division with DIV	203
MUL and DIV Are Slowpokes	204
Reading and Using an Assembly Language Reference	205
Memory Joggers for Complex Memories	205
An Assembly Language Reference for Beginners	206
Flags	207
NEG Negate (Two's Complement; i.e., Multiply by -1)	208
Flags Affected	208
Legal Forms	208
Examples	208
Notes	208
Legal Forms	209
Operand Symbols	209
Examples	210
Notes	210
What's Not Here. . .	210
Chapter 8 Our Object All Sublime	213
The Bones of an Assembly Language Program	213
The Initial Comment Block	215
The .data Section	216
The .bss Section	216
The .text Section	217
Labels	217
Variables for Initialized Data	218
String Variables	219

Deriving String Length with EQU and \$	221
Last In, First Out via the Stack	223
Five Hundred Plates an Hour	223
Stacking Things Upside Down	225
Push-y Instructions	226
POP Goes the Opcode	227
PUSHA and POPA Are Gone	228
Pushing and Popping in Detail	229
Storage for the Short Term	231
Using Linux Kernel Services Through Syscall	231
X64 Kernel Services via the SYSCALL Instruction	232
ABI vs. API?	232
The ABI's Register Parameter Scheme	233
Exiting a Program via SYSCALL	234
Which Registers Are Trashed by SysCall?	235
Designing a Nontrivial Program	235
Defining the Problem	235
Starting with Pseudocode	236
Successive Refinement	237
Those Inevitable "Whoops!" Moments	241
Scanning a Buffer	242
"Off by One" Errors	244
From Pseudocode to Assembly Code	246
The SASM Output Window Gotcha	248
Going Further	248
Chapter 9 Bits, Flags, Branches, and Tables	251
Bits Is Bits (and Bytes Is Bits)	251
Bit Numbering	252
"It's the Logical Thing to Do, Jim. . ."	252
The AND Instruction	253
Masking Out Bits	254
The OR Instruction	255
The XOR Instruction	256
The NOT Instruction	257
Segment Registers Don't Respond to Logic!	258
Shifting Bits	258
Shift by What?	258
How Bit Shifting Works	259
Bumping Bits into the Carry Flag	260
The Rotate Instructions	260
Rotating Bits Through the Carry Flag	261
Setting a Known Value into the Carry Flag	262

Bit-Bashing in Action	262
Splitting a Byte into Two Nybbles	264
Shifting the High Nybble into the Low Nybble	265
Using a Lookup Table	266
Multiplying by Shifting and Adding	267
Flags, Tests, and Branches	270
Unconditional Jumps	271
Conditional Jumps	271
Jumping on the Absence of a Condition	272
Flags	273
Comparisons with CMP	274
A Jungle of Jump Instructions	275
“Greater Than” Versus “Above”	275
Looking for 1-Bits with TEST	277
Looking for 0-Bits with BT	279
x64 Long Mode Memory Addressing in Detail	279
Effective Address Calculations	281
Displacements	282
The x64 Displacement Size Problem	283
Base Addressing	283
Base + Displacement Addressing	283
Base + Index Addressing	284
Index X Scale + Displacement Addressing	285
Other Addressing Schemes	287
LEA: The Top-Secret Math Machine	289
Character Table Translation	290
Translation Tables	291
Translating with MOV or with XLAT	293
Tables Instead of Calculations	298
Chapter 10 Dividing and Conquering	299
Boxes within Boxes	300
Procedures as Boxes for Code	301
Calling and Returning	309
Calls Within Calls	311
The Dangers of Accidental Recursion	313
A Flag Etiquette Bug to Beware Of	314
Procedures and the Data They Need	315
Saving the Caller’s Registers	316
Preserving Registers Across Linux System Calls	317
PUSHAD and POPAD Are Gone	319
Local Data	321
Placing Constant Data in Procedure Definitions	322

More Table Tricks	323
Local Labels and the Lengths of Jumps	325
“Forcing” Local Label Access	328
Short, Near, and Far Jumps	329
Building External Procedure Libraries	330
When Tools Reach Their Limits	330
Using Include Files in SASM	331
Where SASM’s Include Files Must Be Stored	337
The Best Way to Create an Include File Library	338
Separate Assembly and Modules	339
Global and External Declarations	339
The Mechanics of Globals and Externals	342
Linking Libraries into Your Programs	351
The Dangers of Too Many Procedures and Too Many Libraries	352
The Art of Crafting Procedures	352
Maintainability and Reuse	353
Deciding What Should Be a Procedure	354
Use Comment Headers!	355
Simple Cursor Control in the Linux Console	356
Console Control Cautions	363
Creating and Using Macros	364
The Mechanics of Macro Definition	366
Defining Macros with Parameters	371
The Mechanics of Invoking Macros	372
Local labels Within Macros	373
Macro Libraries as Include Files	374
Macros vs. Procedures: Pros and Cons	375
Chapter 11 Strings and Things	377
The Notion of an Assembly Language String	378
Turning Your “String Sense” Inside-Out	378
Source Strings and Destination Strings	379
A Text Display Virtual Screen	379
REP STOSB, the Software Machine Gun	387
Machine-Gunning the Virtual Display	388
Executing the STOSB Instruction	389
STOSB and the Direction Flag DF	390
Defining Lines in the Display Buffer	391
Sending the Buffer to the Linux Console	391
The Semiautomatic Weapon: STOSB Without REP	392
Who Decrements RCX?	392
The LOOP Instructions	393
Displaying a Ruler on the Screen	394
MUL Is Not IMUL	395

Ruler's Lessons	396
The Four Sizes of STOS	396
Goodbye, BCD Math	397
MOVSB: Fast Block Copies	397
DF and Overlapping Block Moves	398
Single-Stepping REP String Instructions	401
Storing Data to Discontinuous Strings	402
Displaying an ASCII Table	402
Nested Instruction Loops	404
Jumping When RCX Goes to 0	405
Closing the Inner Loop	406
Closing the Outer Loop	407
Showchar Recap	408
Command-Line Arguments, String Searches, and the Linux Stack	408
Displaying Command-Line Arguments from SASM	408
String Searches with SCASB	411
REPNE vs. REPE	413
You Can't Pass Command-Line Arguments to Programs Within SASM	413
The Stack, Its Structure, and How to Use It	414
Accessing the Stack Directly	417
Program Prologs and Epilogs	419
Addressing Data on the Stack	420
Don't Pop!	422
Chapter 12 Heading Out to C	423
What's GNU?	424
The Swiss Army Compiler	425
Building Code the GNU Way	425
SASM Uses GCC	427
How to Use gcc in Assembly Work	427
Why Not gas?	428
Linking to the Standard C Library	429
C Calling Conventions	431
Callers, Callees, and Clobbers	431
Setting Up a Stack Frame	433
Destroying a Stack Frame in the Epilog	434
Stack Alignment	435
Characters Out Via puts()	437
Formatted Text Output with printf()	438
Passing Parameters to printf()	440
Printf() Needs a Preceding 0 in RAX	442
You Shall Have -No-Pie	442
Data In with fgets() and scanf()	442

Using scanf() for Entry of Numeric Values	445
Be a Linux Time Lord	448
The C Library's Time Machine	449
Fetching time_t Values from the System Clock	451
Converting a time_t Value to a Formatted String	451
Generating Separate Local Time Values	452
Making a Copy of glibc's tm Struct with MOVSD	453
Understanding AT&T Instruction Mnemonics	456
AT&T Mnemonic Conventions	457
AT&T Memory Reference Syntax	459
Generating Random Numbers	460
Seeding the Generator with srand()	461
Generating Pseudorandom Numbers	461
Some Bits Are More Random Than Others	467
Calls to Addresses in Registers	469
Using puts() to Send a Naked Linefeed to the Console	470
How to Pass a libc Function More Than Six Parameters	470
How C Sees Command-Line Arguments	472
Simple File I/O	474
Converting Strings into Numbers with sscanf()	475
Creating and Opening Files	477
Reading Text from Files with fgets()	478
Writing Text to Files with fprintf()	481
Notes on Gathering Your Procedures into Libraries	482
Conclusion: Not the End, But Only the Beginning	489
Appendix A The Return of the Insight Debugger	493
Insight's Shortcomings	494
Opening a Program Under Insight	495
Setting Command-Line Arguments with Insight	496
Running and Stepping a Program	496
The Memory Window	497
Showing the Stack in Insight's Memory View	498
Examining the Stack with Insight's Memory View	498
Learn gdb!	500
Appendix B Partial x64 Instruction Reference	501
What's Been Removed from x64	502
Flag Results	502
Size Specifiers	503
Instruction Index	505
ADC: Arithmetic Addition with Carry	507
Flags Affected	507

Legal Forms	507
Examples	507
Notes	507
ADD: Arithmetic Addition	509
Flags Affected	509
Legal Forms	509
Examples	509
Notes	509
AND: Logical AND	511
Flags Affected	511
Legal Forms	511
Examples	511
Notes	511
BT: Bit Test	513
Flags Affected	513
Legal Forms	513
Examples	513
Notes	513
CALL: Call Procedure	515
Flags Affected	515
Legal Forms	515
Examples	515
Notes	515
CLC: Clear Carry Flag (CF)	517
Flags Affected	517
Legal Forms	517
Examples	517
Notes	517
CLD: Clear Direction Flag (DF)	518
Flags Affected	518
Legal Forms	518
Examples	518
Notes	518
CMP: Arithmetic Comparison	519
Flags Affected	519
Legal Forms	519
Examples	519
Notes	519
DEC: Decrement Operand	521
Flags Affected	521
Legal Forms	521
Examples	521
Notes	521

DIV: Unsigned Integer Division	522
Flags Affected	522
Legal Forms	522
Examples	522
Notes	522
INC: Increment Operand	524
Flags Affected	524
Legal Forms	524
Examples	524
Notes	524
J?: Jump If Condition Is Met	525
Flags Affected	525
Examples	525
Notes	525
JECXZ: Jump if ECX=0	527
Flags Affected	527
Legal Forms	527
Examples	527
Notes	527
JRCXZ: Jump If RCX=0	528
Flags Affected	528
Legal Forms	528
Examples	528
Notes	528
JMP: Unconditional Jump	529
Flags Affected	529
Legal Forms	529
Examples	529
Notes	529
LEA: Load Effective Address	531
Flags Affected	531
Legal Forms	531
Examples	531
Notes	531
LOOP: Loop Until CX/ECX/RCX=0	532
Flags Affected	532
Legal Forms	532
Examples	532
Notes	532
LOOPNZ/LOOPNE: Loop Until CX/ECX/RCX=0 and ZF=0	534
Flags Affected	534
Legal Forms	534

Examples	534
Notes	534
LOOPZ/LOOPE: Loop Until CX/ECX/RCX=0 and ZF=1	535
Flags Affected	535
Legal Forms	535
Examples	535
Notes	535
MOV: Copy Right Operand into Left Operand	536
Flags Affected	536
Legal Forms	536
Examples	536
Notes	536
MOVS: Move String	538
Flags Affected	538
Legal Forms	538
Examples	538
Notes	538
MOVSX: Copy with Sign Extension	540
Flags Affected	540
Legal Forms	540
Examples	540
Notes	540
MUL: Unsigned Integer Multiplication	542
Flags Affected	542
Legal Forms	542
Examples	542
Notes	542
NEG: Negate (Two's Complement; i.e., Multiply by -1)	544
Flags Affected	544
Legal Forms	544
Examples	544
Notes	544
NOP: No Operation	546
Flags Affected	546
Legal Forms	546
Examples	546
Notes	546
NOT: Logical NOT (One's Complement)	547
Flags Affected	547
Legal Forms	547
Examples	547
Notes	547

OR: Logical OR	548
Flags Affected	548
Legal Forms	548
Examples	548
Notes	548
POP: Copy Top of Stack into Operand	550
Flags Affected	550
Legal Forms	550
Examples	550
Notes	550
POPF/D/Q: Copy Top of Stack into Flags Register	552
Flags Affected	552
Legal Forms	552
Examples	552
Notes	552
PUSH: Push Operand onto Top of Stack	553
Flags Affected	553
Legal Forms	553
Examples	553
Notes	553
PUSHF/D/Q: Push Flags Onto the Stack	555
Flags Affected	555
Legal Forms	555
Examples	555
Notes	555
RET: Return from Procedure	556
Flags Affected	556
Legal Forms	556
Examples	556
Notes	556
ROL/ROR: Rotate Left/Rotate Right	558
Flags Affected	558
Legal Forms	558
Examples	558
Notes	558
SBB: Arithmetic Subtraction with Borrow	560
Flags Affected	560
Legal Forms	560
Examples	560
Notes	560
SHL/SHR: Shift Left/Shift Right	562
Flags Affected	562
Legal Forms	562
Examples	562

Notes	562
STC: Set Carry Flag (CF)	564
Flags Affected	564
Legal Forms	564
Examples	564
Notes	564
STD: Set Direction Flag (DF)	565
Flags Affected	565
Legal Forms	565
Examples	565
Notes	565
STOS/B/W/D/Q: Store String	566
Flags Affected	566
Legal Forms	566
Examples	566
Notes	566
SUB: Arithmetic Subtraction	568
Flags Affected	568
Legal Forms	568
Examples	568
Notes	569
SYSCALL: Fast System Call into Linux	570
Flags Affected	570
Legal Forms	570
Examples	570
Notes	570
XCHG: Exchange Operands	571
Flags Affected	571
Legal Forms	571
Examples	571
Notes	571
XLAT: Translate Byte Via Table	572
Flags Affected	572
Legal Forms	572
Examples	572
Notes	572
XOR: Exclusive OR	573
Flags Affected	573
Legal Forms	573
Examples	573
Notes	573
Appendix C Character Set Charts	575
Index	579



Introduction

“Why Would You Want to Do That?”

It was 1985, and I was in a chartered bus in New York City, heading for a press reception with a bunch of other restless media egomaniacs. I was only beginning my tech journalist career (as technical editor for *PC Tech Journal*), and my first book was still months in the future. I happened to be sitting next to an established programming writer/guru, with whom I was impressed and to whom I was babbling about one thing or another. I would like to eliminate this statement; it adds little to the book, and as annoying as he is, even though we don't name him, I now understand why he's so annoying: He lives and works in a completely different culture than I do.

During our chat, I happened to let slip that I was a Turbo Pascal fanatic, and what I really wanted to do was learn how to write Turbo Pascal programs that made use of the brand new Microsoft Windows user interface. He wrinkled his nose and grimaced wryly, before speaking the Infamous Question:

“Why would you want to do that?”

I had never heard the question before (though I would hear it many times thereafter), and it took me aback. Why? Because, well, because. . . I wanted to know how it *worked*.

“Heh. That's what C is for.”

Further discussion got me nowhere in a Pascal direction. But some probing led me to understand that you *couldn't* write Windows apps in Turbo Pascal. It was impossible. Or. . .the programming writer/guru didn't know how. Maybe both. I never learned the truth as it stood in 1985. (Delphi answered the question once and for all in 1995.) But I did learn the meaning of the Infamous Question.

Note well: When somebody asks you, *“Why would you want to do that?”* what it really means is this: *“You've asked me how to do something that is either*

impossible using tools that I favor or completely outside my experience, but I don't want to lose face by admitting it. So. . .how 'bout those Blackhawks?"

I heard it again and again over the years:

Q: How can I set up a C string so that I can read its length without scanning it?

A: Why would you want to do *that*?

Q: How can I write an assembly language subroutine callable from Turbo Pascal?

A: Why would you want to do *that*?

Q: How can I write Windows apps in assembly language?

A: Why would you want to do *that*?

You get the idea. The answer to the Infamous Question is always the same, and if the weasels ever ask it of you, snap back as quickly as possible: *because I want to know how it works.*

That is a completely sufficient answer. It's the answer I've used every single time, except for one occasion a considerable number of years ago, when I put forth that I wanted to write a book that taught people how to program in assembly language as their *first* experience in programming.

Q: Good grief, why would you want to do *that*?

A: Because it's the best way there is to build the skills required to understand how *all the rest* of the programming universe works.

Being a programmer is one thing above all else: It is understanding how things work. Learning to be a programmer, furthermore, is almost entirely a process of learning how things work. This can be done at various levels, depending on the tools you're using. If you're programming in Visual Basic, you have to understand how certain things work, but those things are by and large confined to Visual Basic itself. A great deal of machinery is hidden by the layer that Visual Basic places between the programmer and the computer. (The same is true of Delphi, Lazarus, Java, Python, and many other very high-level programming environments.) If you're using a C compiler, you're a lot closer to the machine, so you see a lot more of that machinery—and must, therefore, understand how it works to be able to use it. However, quite a bit remains hidden, even from the hardened C programmer.

If, on the other hand, you're working in assembly language, you're as close to the machine as you can get. Assembly language hides *nothing*, and withholds no power. The flipside, of course, is that no magical layer between you and the machine will absolve any ignorance and "take care of" things for you. If you don't understand how something works, you're dead in the water—unless you know enough to be able to figure it out on your own.

That's a key point: My goal in creating this book is not entirely to teach you assembly language *per se*. If this book has a prime directive at all, it is to impart a certain disciplined curiosity about the underlying machine, along with some basic context from which you can begin to explore the machine at its very lowest levels—that, and the confidence to give it your best shot. This is difficult stuff, but it's nothing you can't master given some concentration, patience, and the time it requires—which, I caution, may be considerable.

In truth, what I'm really teaching you here is how to learn.

What You'll Need

To program as I intend to teach, you're going to need a 64-bit Intel computer running a 64-bit distribution of Linux. The one I used in preparing this book is Linux Mint Cinnamon V20.3 Una. "Una" here is a code name for this version of Linux Mint. It's nothing more than a short way of saying "Linux Mint 20.3." I recommend Mint; it's thrown me fewer curves than any other distro I've ever used—and I've used Linux here and there ever since it first appeared. I don't think which graphical shell you use matters a great deal. I like Cinnamon, but you can use whatever you like or are familiar with.

You need to be reasonably proficient with Linux at the user level. I can't teach you how to install, configure, and run Linux in this book. If you're not already familiar with Linux, get a tutorial text and work through it. There are many such online.

You'll need a piece of free software called SASM, which is a simple interactive development environment (IDE) for programming in assembly. Basically, it consists of an editor, a build system, and a front end to the standard Linux debugger `gdb`. You'll also need a free assembler called NASM.

You don't have to know how to download, install, and configure these tools in advance because, at the appropriate times, I'll cover all necessary tool installation and configuration.

Do note that other Unix implementations not based on the Linux kernel may not function precisely the same way under the hood. BSD Unix uses different conventions for making system calls, for example, and other Unix versions like Solaris are outside my experience.

Remember that *this book is about the x64 architecture*. To the extent that x64 contains x86, I will also be teaching elements of the x86 architecture. The gulf between 32-bit x86 and 64-bit x64 is a *lot* narrower than the gulf between 16-bit x86 and 32-bit x86. If you already have a firm grounding in 32-bit x86, you'll breeze through most of this book at a gallop. If you can do that, cool—just please remember that the book is for those who are just starting out in programming on Intel CPUs.

Also remember that this book is limited in size by its publisher: Paper, ink, and cover stock aren't free. That means I have to narrow the scope of what I teach and explain within those limits. I wish I had the space to cover the AVX math subsystem. I don't. But I'll bet that once you go through this book, you can figure much of it out by yourself.

The Master Plan

This book starts at the beginning, and I mean the *beginning*. Maybe you're already there, or well past it. I respect that. I still think that it wouldn't hurt to start at the first chapter and read through all the chapters in order. Review is useful, and hey—you may realize that you didn't know *quite* as much as you thought you did. (Happens to me all the time!)

But if time is at a premium, here's the cheat sheet:

- If you already understand the fundamental ideas of computer programming, skip Chapter 1.
- If you already understand the ideas behind number bases other than decimal (especially hexadecimal and binary), skip Chapter 2.
- If you already have a grip on the nature of computer internals (memory, CPU architectures, and so on) skip Chapter 3.
- If you already understand x64 memory addressing, skip Chapter 4.
- No. Stop. Scratch that. Even if you already understand x64 memory addressing, *read Chapter 4*.

The last bullet is there, and emphatic, for a reason: *Assembly language programming is about memory addressing*. If you don't understand memory addressing, nothing else you learn in assembly will help you one. . . bit. So, don't skip Chapter 4 no matter what else you know or think you know. Start from there, and see it through to the end. Memory addressing comes up regularly throughout the rest of the book. It's really the heart of the topic.

Load every example program, assemble each one, and run them all. Strive to understand every single line in every program. Take nothing on faith. Furthermore, don't stop there. Change the example programs as things begin to make sense to you. Try different approaches. Try things that I don't mention. Be audacious. Nay, go nuts—bits don't have feelings, and the worst thing that can happen is that Linux throws a segmentation fault, which may hurt your program but does not hurt Linux. The only catch is that when you do try something, strive to understand why it *doesn't* work as clearly as you understand all the other things that do. Single-step your way through a program in the SASM debugger, even when the program works. *Take notes*.

That is, ultimately, what I'm after: to show you the way to understand what every however distant corner of your machine is doing and how all its many pieces work together. This doesn't mean I'll explain every corner of it myself—no one will live long enough to do that because computing isn't simple anymore—but if you develop the discipline of patient research and experimentation, you can probably work it out for yourself. Ultimately, that's the only way to learn it: by yourself. The guidance you find—in friends, on the Net, in books like this—is only guidance and grease on the axles. You have to decide who's to be the master, you or the machine, and make it so. Assembly programmers are the only programmers who can truly claim to be masters, which is a truth worth meditating on.

A Note on Capitalization Conventions

Assembly language is peculiar among programming languages in that there is no universal standard for case-sensitivity. In the C language, all identifiers are case-sensitive, and I have seen assemblers that do not recognize differences in case at all. NASM, the assembler I'm presenting in this book, is case-sensitive only for programmer-defined identifiers. The instruction mnemonics and the names of registers, however, are *not* case sensitive.

There are customs in the literature on assembly language, and one of those customs is to treat CPU instruction mnemonics as uppercase in the chapter text and in lowercase in source code files and code snippets interspersed within the text. I'll be following that custom here. Within discussion text, I'll speak of `MOV` and `CALL` and `CMP`. In example code, it will be `mov` and `call` and `cmp`. Code snippets and listings will be in a monospace Courier-style font. When mentioned in the text, registers will be in uppercase but not in the Courier font and lowercase in snippets and listings.

There are two reasons for this:

- In text discussions, the mnemonics need to stand out. It's too easy to lose track of them amid a torrent of ordinary mixed-case words.
- To read and learn from existing documents and source code outside of this one book, you need to be able to easily read assembly language whether it's in uppercase, lowercase, or mixed case. Getting comfortable with different ways of expressing the same things is important.

Remember Why You're Here

Anyway. Wherever you choose to start the book, it's time to get underway. Just remember that whatever gets in your face, be it the weasels, the machine, or

your own inexperience, the thing to keep in the forefront of your mind is this:
You're in it to figure out how it works.

Let's go.

Jeff Duntemann
Scottsdale, Arizona
May 24, 2023

x64 Assembly Language Step-by-Step

4TH Edition

It's All in the Plan

Understanding What Computers Really Do

Another Pleasant Valley Saturday

“Quick, Mike, get your sister and brother up; it’s past 7. Nicky’s got Little League at 9, and Dione’s got ballet at 10. Give Max his heartworm pill! (We’re out of them, Mom, remember?) Your father picked a great weekend to go fishing Here, let me give you 10 bucks and go get more pills at the vet’s My God, that’s right, Hank needed gas money and left me broke. There’s a teller machine over by Kmart, and if I go there, I can take that stupid toilet seat back and get the right one. “I guess I’d better make a list”

It’s another Pleasant Valley Saturday, and 30-odd million suburban homemakers sit down with a pencil and pad at the kitchen table to try to make sense of a morning that would kill and pickle any lesser being. In her mind she thinks of the dependencies and traces the route:

“Drop Nicky at Rand Park, go back to Dempster, and it’s about 10 minutes to Golf Mill Mall. Do I have gas? I’d better check first—if not, stop at Del’s Shell or I won’t make it to Milwaukee Avenue. Milk the teller machine at Golf Mill; then cross the parking lot to Kmart to return the toilet seat that Hank

bought last weekend without checking what shape it was. Gotta remember to throw the toilet seat in back of the van—write that at the top of the list.

“By then it’ll be half past, maybe later. Ballet is all the way down Greenwood in Park Ridge. No left turn from Milwaukee—but there’s the sneak path around behind the mall. I have to remember not to turn right onto Milwaukee like I always do—jot that down. While I’m in Park Ridge, I can check and see if Hank’s new glasses are in—should call, but they won’t even be open until 9:30. Oh, and groceries—can do that while Dione dances. On the way back I can cut over to Oakton and get the dog’s pills.”

In about 90 seconds flat the list is complete:

- Throw toilet seat in van.
- Check gas—if empty, stop at Del’s Shell.
- Drop Nicky at Rand Park.
- Stop at Golf Mill teller machine.
- Return toilet seat at Kmart.
- Drop Dione at ballet (remember the sneak path to Greenwood).
- See if Hank’s glasses are at Pearle Vision—if they are, make double sure they remembered the extra scratch coating.
- Get groceries at Jewel.
- Pick up Dione.
- Stop at vet for heartworm pills.
- Drop off groceries at home.
- If it’s time, pick up Nicky. If not, collapse for a few minutes; then pick up Nicky.
- Collapse!

What we often call a “laundry list” (whether it involves laundry or not) is the perfect metaphor for a computer program. Without realizing it, our intrepid homemaker has written herself a computer program and then set out (with herself acting as the computer) to execute it and be done before noon.

Computer programming is nothing more than this: you the programmer write a list of steps and tests. The computer then performs each step and test in sequence. When the list of steps has been executed, the computer stops.

A computer program is a list of steps and tests, nothing more.

Steps and Tests

Think for a moment about what I call a test in the preceding laundry list. A *test* is the sort of either/or decision we make dozens or hundreds of times on even the most placid of days, sometimes nearly without thinking about it.

Our homemaker performed a test when she jumped into the van to get started on her adventure. She looked at the gas gauge. The gas gauge would tell her one of two things: (1) she has enough gas, or (2) she doesn't. If she has enough gas, she takes a right and heads for Rand Park. If she doesn't have enough gas, she takes a left down to the corner and fills the tank at Del's Shell. (Del takes credit cards.) Then, with a full tank, she continues the program by making a U-turn and heading for Rand Park.

In the abstract, a test consists of these two parts:

- First, you take a look at something that can go one of two ways.
- Then you do one of two things, depending on what you saw when you took a look.

Toward the end of the program, our homemaker got home, took the groceries out of the van, and looked at the clock. If it isn't time to get Nicky back from Little League, she has a moment to collapse on the couch in a nearly empty house. If it *is* time to get Nicky, there's no rest for the ragged: she sprints for the van and heads back to Rand Park.

(Any guesses as to whether she really gets to rest when the program finishes running?)

More Than Two Ways?

You might object, saying that many or most tests involve more than two alternatives. Sorry, you're wrong—in every case. Read this twice: except for totally impulsive or psychotic behavior, every human decision comes down to the choice between two alternatives.

What you have to do is look a little more closely at what goes through your mind when you make decisions. The next time you buzz down to Chow Now for fast Chinese, observe yourself while you're poring over the menu. The choice might seem, at first, to be of one item out of 26 Cantonese main courses. Not so—the choice, in fact, is between choosing one item and *not* choosing that one item. Your eyes rest on chicken with cashews. Naw, too bland. *That was a test.* You slide down to the next item. Chicken with black mushrooms. Hmmm, no, had that last week. *That was another test.* Next item: kung pao chicken. Yeah, that's it! *That was a third test.*

The choice was not among chicken with cashews, chicken with black mushrooms, and chicken with kung pao. Each dish had its moment, poised before the critical eye of your mind, and you turned thumbs up or thumbs down on it, individually. Eventually, one dish won, but it won in that same game of “to eat or not to eat.”

Let me give you another example. Many of life's most complicated decisions come about because 99.99867 percent of us are not nudists. You've been there: you're standing in the clothes closet in your underwear, flipping through your rack of pants. The tests come thick and fast. This one? No. This one? No. This one? No. This one? Yeah. You pick a pair of blue pants, say. (It's a Monday, after all, and blue would seem an appropriate color.) Then you stumble over to your sock drawer and take a look. Whoops, no blue socks. *That was a test.* So you stumble back to the clothes closet, hang your blue pants back on the pants rack, and start over. This one? No. This one? No. This one? Yeah. This time it's brown pants, and you toss them over your arm and head back to the sock drawer to take another look. Nertz, out of brown socks, too. So it's back to the clothes closet . . .

What you might consider a single decision, or perhaps two decisions inextricably tangled (like picking pants and socks of the same color, given stock on hand), is actually a series of small decisions, always binary in nature: pick 'em or don't pick 'em. Find 'em or don't find 'em. The Monday morning episode in the clothes closet is a good analogy of a programming structure called a *loop*: you keep doing a series of things until you get it right, and then you stop (assuming you're not the kind of geek who wears blue socks with brown pants). But whether you get everything right always comes down to a sequence of simple either/or decisions.

Computers Think Like Us

I can almost hear what you're thinking: “Sure, it's a computer book, and he's trying to get me to think like a computer.” Not at all. Computers think like *us*. We designed them; how else could they think? No, what I'm trying to do is get you to take a long, hard look at how *you* think. We run on automatic for so much of our lives that we literally do most of our thinking without really thinking about it.

The best model for the logic of a computer program is the same logic we use to plan and manage our daily affairs. No matter what we do, it comes down to a matter of confronting two alternatives and picking one. What we might think of as a single large and complicated decision is nothing more than a messy tangle of many smaller decisions. The skill of looking at a complex decision and seeing all the little decisions in its tummy will serve you well in learning how to

program. Observe yourself the next time you have to decide something. Count up the little decisions that make up the big one. You'll be surprised.

And, surprise! You'll be a programmer.

Had This Been the Real Thing . . .

Do not be alarmed. What you have just experienced was a metaphor. It was not the real thing. (The real thing comes later.)

I use metaphors a lot in this book. A metaphor is a loose comparison drawn between something familiar (such as a Saturday morning laundry list) and something unfamiliar (such as a computer program). The idea is to anchor the unfamiliar in terms of the familiar so that when I begin tossing facts at you, you'll have someplace comfortable to lay them down.

The most important thing for you to do right now is keep an open mind. If you know a little bit about computers or programming, don't pick nits. Yes, there are important differences between a homemaker following a scribbled laundry list and a computer executing a program. I'll mention those differences all in good time.

For now, it's still Chapter 1. Take these initial metaphors on their own terms. Later, they'll help a lot.

Assembly Language Programming As a Square Dance

Carol and I have a certain fondness for "called" dances, the most prevalent type being square dances. There are others, like New England contra dances, which are a lot like square dances but with better music. In a called dance, the caller person at the front of the hall calls out movements, and the dancers perform those movements. The music provides a beat, like the ticking of a clock. The sequence of movements taken together is the dance, and the dance usually has a name.

The first time Carol and I attended a contra dance, I was poleaxed: *this was like assembly language programming!* The caller called out "allemande left," and we performed the movement known as "allemande left." The caller called out "forward and back," and we executed the "forward and back" movement. The caller called out "box the gnat," and, well, we boxed the gnat. (I am not making this up!) There are a reasonable number of movements, and to be good at that sort of dancing, you have to memorize them all by name. Otherwise, if the caller calls a movement that you don't know, the dance might stumble or grind to a halt. (Bluescreen!)

At its deepest level, a computer understands a collection of individual operations called *instructions*. These perform arithmetic, execute logic like AND and OR, move data around, and do many other things. Each instruction is performed inside the CPU chip. Just as a set of dance movements are the individual atoms of motion making up a square dance, instructions are the atoms of a computer program. The program is like the dance as a whole: a sequence of instructions executed in order. The couples taking part in the dance execute the dance/program as the caller moves down the list of movements, calling out each one in turn. The couples, then, are the computer on which the dance runs.

That's about as far as the square dance metaphor goes. Once you get the knack of assembly language, hey, go take square dance or contra dance lessons somewhere and see if you don't come to the same conclusion that I did.

Assembly Language Programming As a Board Game

Board games were a really big deal when I was a kid, when board games were actually printed on a species of board. (OK, cardboard.) Monopoly was one that almost everybody had. There was a sort of pathway around the edge of the board divided into squares. You had a game piece that advanced from square to square according to dice throws, and when your piece landed on a square, you could do one of several things: buy property that hadn't been bought yet, pay rent on property owned by other players, pull a card from the Chance stack, or—eek!—go to jail. You had a pile of Monopoly money to spend, and when another player had to pay rent, you got more.

The specifics of the Monopoly game aren't important here. What matters is that you progress through a series of steps, and at each step, something happens. Your pile of money grows or shrinks. Assembly language is a little like that: a program is like the game board. Each step in the program does something. There are places where you can store numbers. The numbers change as you move through the program.

Now that you're thinking in terms of board games, take a look at Figure 1.1. What I've drawn is actually a fair approximation of assembly language as it was used on some of our simpler computers 50 or 60 years ago. The column marked "Program Instructions" is the main path around the edge of the board, of which only a portion can be shown here. This is the assembly language computer program, the actual series of steps and tests that, when executed, causes the computer to do something useful. Setting up this series of program instructions is what programming in assembly language actually *is*.

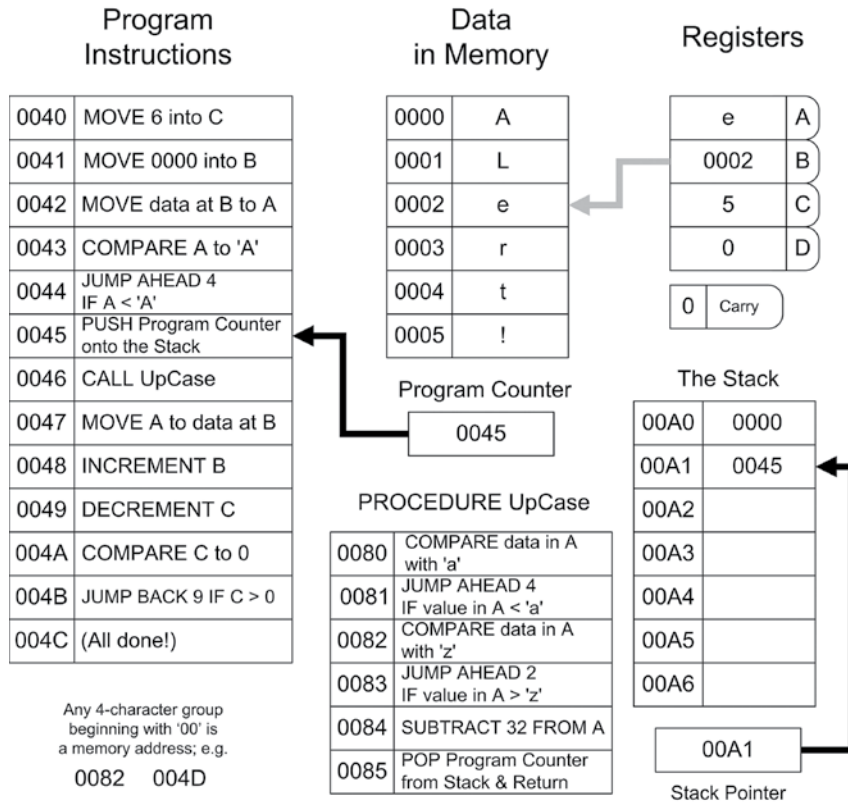


Figure 1.1: The Game of Assembly Language

Everything else is odds and ends in the middle of the board that serve the game in progress. Most of these are storage locations that contain your data. You're probably noticing (perhaps with sagging spirits) that there are a *lot* of numbers involved. (They're weird numbers, too. What, for example, does "004B" mean? I deal with that issue in Chapter 2, "Alien Bases.") I'm sorry, but that's simply the way the game is played. Assembly language, at its deepest level, is nothing *but* numbers, and if you hate numbers the way most people hate anchovies, you're going to have a rough time of it. (I like anchovies, which is part of my legend. Learn to like numbers. They're not as salty.) Higher-level programming languages such as Pascal or Python disguise the numbers by treating them symbolically. But assembly language, well, it's just you and the numbers.

I should caution you that the Game of Assembly Language in Figure 1.1 represents no real computer processor, like the Intel Core i5. Also, I've made the names of instructions more clearly understandable than the names of the instructions in Intel assembly language actually are. In the real world, instruction names

are typically short things like **LAHF**, **STC**, **INC**, **SHRX**, and other crypticisms that cannot be understood without considerable explanation. We're easing into this stuff sidewise, and in this chapter I have to sugarcoat certain things a little to draw the metaphors clearly.

Code and Data

Like most board games, the assembly language board game consists of two broad categories of elements: game steps and places to store things. The "game steps" are the steps and tests I've been speaking of all along. The places to store things are just that: cubbyholes into which you can place numbers, with the confidence that those numbers will remain where you put them until you take them out or change them somehow.

In programming terms, the game steps are called *code*, and the numbers in their cubbyholes (as distinct from the cubbyholes themselves) are called *data*. The cubbyholes themselves are usually called *storage*. (The difference between the places you store information and the information you store in them is crucial. Don't confuse them.) Consider an instruction in the Game of Assembly Language that says **ADD 32 to A**. An **ADD** instruction in the code alters a data value stored in a cubbyhole named Register A.

Code and data are two very different kinds of critters, but they interact in ways that make the game interesting. The code includes steps that place data into storage (**MOVE** instructions) and steps that alter data that is already in storage (**INCREMENT** and **DECREMENT** instructions, and **ADD** instructions, among others). Most of the time you'll think of code as being the master of data, in that the code writes data values into storage. Data does influence code as well, however. Among the tests that the code makes are tests that examine data in storage, the **COMPARE** instructions. If a given data value exists in storage, the code may do one thing; if that value does not exist in storage, the code will do something else, as in the **JUMP BACK** and **JUMP AHEAD** instructions.

The short block of instructions marked **PROCEDURE** is a detour off the main stream of instructions. At any point in the program you can duck out into the procedure, perform its steps and tests, and then return to the very place from which you left. This allows a sequence of steps and tests that is generally useful and used frequently to exist in only one place rather than exist as separate copies everywhere it's needed.

Addresses

Another critical concept lies in the funny numbers at the left side of the program step locations and data locations. Each number is unique, in that a location tagged with that number appears only *once* inside the computer. This location is called

an *address*. Data is stored and retrieved by specifying the data's address in the machine. Procedures are called by specifying the address at which they begin.

The little box (which is also a storage location) marked "Program Counter" keeps the address of the next instruction to be performed. The number inside the program counter is increased by one (*incremented*) each time an instruction is performed *unless the instruction tells the program counter to do something else*. For example, notice the **JUMP BACK 9** instruction at address 004B. When this instruction is performed, the program counter will "back up" by nine locations. This is analogous to the "go back three spaces" concept in most board games.

Metaphor Check!

That's about as much explanation of the Game of Assembly Language as I'm going to offer for now. This is still Chapter 1, and we're still in metaphor territory. People who have had some exposure to computers will recognize and understand some of what Figure 1.1 is doing. People with no exposure to computer innards at all shouldn't feel left behind for being utterly lost. I created the Game of Assembly Language solely to put across the following points:

- *The individual steps are very simple.* One single instruction rarely does more than move a single value from one storage cubbyhole to another, perform very elementary arithmetic like addition or subtraction, or compare the value contained in one storage cubbyhole to a value contained in another. This is good news, because it allows you to concentrate on the simple task accomplished by a single instruction without being overwhelmed by complexity. The bad news, however, is the next point.
- *It takes a lot of steps to do anything useful.* You can often write a useful program in such languages as Pascal or BASIC in five or six lines. You can actually create useful programs in visual programming systems like Visual Basic, Delphi, or Lazarus *without writing any code at all*. (The code is still there . . . but the code is "canned" and all you're really doing is choosing which chunks of canned code in a collection of many such chunks will run.) A useful assembly language program cannot be implemented in fewer than about 50 lines, and anything challenging takes hundreds or thousands—or tens of thousands—of lines. The skill of assembly language programming lies in structuring these hundreds or thousands of instructions so that the program both operates correctly and can still be read and understood by other programmers—and yourself—six months later.
- *The key to assembly language is understanding memory addresses.* In such languages as Pascal and BASIC, the compiler takes care of where something is located—you simply have to give that something a symbolic name and call it by that name whenever you want to look at it or change it.