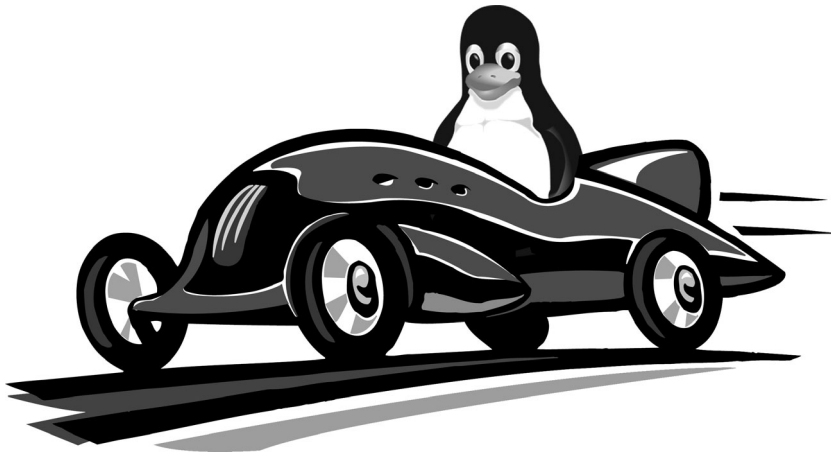


See MIPS[®] Run

Second Edition



See **MIPS**[®] Run

Second Edition

Dominic Sweetman



AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO
Morgan Kaufmann Publishers is an imprint of Elsevier



Publisher: Denise E. M. Penrose
Publishing Services Manager: George Morrison
Senior Project Manager: Brandy Lilly
Editorial Assistant: Kimberlee Honjo
Cover Design: Alisa Andreola and Hannus Design
Composition: diacriTech
Technical Illustration: diacriTech
Copyeditor: Denise Moore
Proofreader: Katherine Antonsen
Indexer: Steve Rath
Interior Printer: The Maple-Vail Book Manufacturing Group, Inc.
Cover Printer: Phoenix Color

Morgan Kaufmann Publishers is an imprint of Elsevier.
500 Sansome Street, Suite 400, San Francisco, CA 94111

This book is printed on acid-free paper.

© 2007 by Elsevier Inc. All rights reserved.

MIPS, MIPS I, MIPS II, MIPS III, MIPS IV, MIPS V, MIPS16, MIPS16e, MIPS-3D, MIPS32, MIPS64, 4K, 4KE, 4KEc, 4KSc, 4Ksd, M4K, 5K, 20Kc, 24K, 24KE, 24Kf, 25Kf, 34K, R3000, R4000, R5000, R10000, CorExtend, MDMX, PDtrace and SmartMIPS are trademarks or registered trademarks of MIPS Technologies, Inc. in the United States and other countries, and used herein under license from MIPS Technologies, Inc. MIPS, MIPS16, MIPS32, MIPS64, MIPS-3D and SmartMIPS, among others, are registered in the U.S. Patent and Trademark Office.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Designations used by companies to distinguish their products are often claimed as trademarks or registered trademarks. In all instances in which Morgan Kaufmann Publishers is aware of a claim, the product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

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, scanning, or otherwise—without prior written permission of the publisher.

Permissions may be sought directly from Elsevier's Science & Technology Rights Department in Oxford, UK: phone: (+44) 1865 843830, fax: (+44) 1865 853333, E-mail: permissions@elsevier.com. You may also complete your request on-line via the Elsevier homepage (<http://elsevier.com>), by selecting "Support & Contact" then "Copyright and Permission" and then "Obtaining Permissions."

Library of Congress Cataloging-in-Publication Data

Application Submitted

ISBN 13: 978-0-12-088421-6

ISBN 10: 0-12-088421-6

For information on all Morgan Kaufmann publications,
visit our Web site at www.mkp.com or www.books.elsevier.com

Printed in the United States of America

06 07 08 09 10 5 4 3 2 1

Working together to grow
libraries in developing countries

www.elsevier.com | www.bookaid.org | www.sabre.org

ELSEVIER

BOOK AID
International

Sabre Foundation

Foreword

The MIPS architecture was born in the early 1980s from the work done by John Hennessy and his students at Stanford University. They were exploring the architectural concept of RISC (Reduced Instruction Set Computing), which theorized that relatively simple instructions, combined with excellent compilers and hardware that used pipelining to execute the instructions, could produce a faster processor with less die area. The concept was so successful that MIPS Computer Systems was formed in 1984 to commercialize the MIPS architecture.

Over the course of the next 14 years, the MIPS architecture evolved in a number of ways and its implementations were used very successfully in workstation and server systems. Over that time, the architecture and its implementations were enhanced to support 64-bit addressing and operations, support for complex memory-protected operating systems such as UNIX, and very high performance floating point. Also in that period, MIPS Computer Systems was acquired by Silicon Graphics and MIPS processors became the standard for Silicon Graphics computer systems. With 64-bit processors, high-performance floating point, and the Silicon Graphics heritage, MIPS processors became the solution of choice in high-volume gaming consoles.

In 1998, MIPS Technologies emerged from Silicon Graphics as a stand-alone company focused entirely on intellectual property for embedded markets. As a result, the pace of architecture development has increased to address the unique needs of these markets: high-performance computation, code compression, geometry processing for graphics, security, signal processing, and multi-threading. Each architecture development has been matched by processor core implementations of the architecture, making MIPS-based processors the standard for high-performance, low-power applications.

The MIPS legacy in complex systems such as workstations and servers directly benefits today's embedded systems, which have, themselves, become very complex. A typical embedded system is composed of multiple processing elements, high-performance memory, and one or more operating systems.

When compared with other embedded architectures, which are just now learning what is required to build a complex system, the MIPS architecture provides a proven base on which to implement such systems.

In many ways, the first edition of *See MIPS Run* was a ground-breaking book on the MIPS architecture and its implementations. While other books covered similar material, *See MIPS Run* focused on what the programmer needed to understand of the architecture and the software environment in order to effectively program a MIPS chip.

Increasing complexity of embedded systems has been matched by enhancements to the MIPS architecture to address the needs of such systems. The second edition of this book is required reading for any current developer of MIPS-based embedded systems. It adds significant new material, including the architectural standardization of the MIPS32 and MIPS64 architectures, brand new application-specific extensions such as multithreading, and a very nice treatment of the implementation of the popular Linux operating system on the MIPS architecture. Short of the MIPS architecture specifications, the second edition of *See MIPS Run* is the most current description of the state of the art of the architecture and is, bar none, the most readable.

I hope that you will find this as worthwhile and as entertaining to read as I did.

Michael Uhler,
Chief Technology Officer, MIPS Technologies, Inc.
Mountain View, CA
May 2006

Contents

<i>Foreword</i>	v
<i>Preface</i>	xv
Style and Limits	xviii
Conventions	xviii
Acknowledgments	xix
<i>Chapter 1</i> RISCs and MIPS Architectures	1
<hr/>	
1.1 Pipelines	2
1.1.1 <i>What Makes a Pipeline Inefficient?</i>	3
1.1.2 <i>The Pipeline and Caching</i>	4
1.2 The MIPS Five-Stage Pipeline	5
1.3 RISC and CISC	7
1.4 Great MIPS Chips of the Past and Present	8
1.4.1 <i>R2000 to R3000 Processors</i>	8
1.4.2 <i>The R6000 Processor: A Diversion</i>	9
1.4.3 <i>The First CPU Cores</i>	11
1.4.4 <i>The R4000 Processor: A Revolution</i>	12
1.4.5 <i>The Rise and Fall of the ACE Consortium</i>	12
1.4.6 <i>SGI Acquires MIPS</i>	13
1.4.7 <i>QED: Fast MIPS Processors for Embedded Systems</i>	13
1.4.8 <i>The R10000 Processor and its Successors</i>	14
1.4.9 <i>MIPS Processors in Consumer Electronics</i>	15
1.4.10 <i>MIPS in Network Routers and Laser Printers</i>	15
1.4.11 <i>MIPS Processors in Modern Times</i>	17
1.4.12 <i>The Rebirth of MIPS Technologies</i>	20
1.4.13 <i>The Present Day</i>	21
1.5 MIPS Compared with CISC Architectures	23
1.5.1 <i>Constraints on MIPS Instructions</i>	23
1.5.2 <i>Addressing and Memory Accesses</i>	24
1.5.3 <i>Features You Won't Find</i>	25
1.5.4 <i>Programmer-Visible Pipeline Effects</i>	27

<i>Chapter 2</i>	MIPS Architecture	29
2.1	A Flavor of MIPS Assembly Language	33
2.2	Registers	34
2.2.1	<i>Conventional Names and Uses of General-Purpose Registers</i>	35
2.3	Integer Multiply Unit and Registers	38
2.4	Loading and Storing: Addressing Modes	39
2.5	Data Types in Memory and Registers	39
2.5.1	<i>Integer Data Types</i>	39
2.5.2	<i>Unaligned Loads and Stores</i>	40
2.5.3	<i>Floating-Point Data in Memory</i>	41
2.6	Synthesized Instructions in Assembly Language	42
2.7	MIPS I to MIPS64 ISAs: 64-Bit (and Other) Extensions	43
2.7.1	<i>To 64 Bits</i>	45
2.7.2	<i>Who Needs 64 Bits?</i>	45
2.7.3	<i>Regarding 64 Bits and No Mode Switch: Data in Registers</i>	46
2.8	Basic Address Space	47
2.8.1	<i>Addressing in Simple Systems</i>	49
2.8.2	<i>Kernel versus User Privilege Level</i>	49
2.8.3	<i>The Full Picture: The 64-Bit View of the Memory Map</i>	50
2.9	Pipeline Visibility	50
<i>Chapter 3</i>	Coprocessor 0: MIPS Processor Control	53
3.1	CPU Control Instructions	55
3.2	Which Registers Are Relevant When?	58
3.3	CPU Control Registers and Their Encoding	59
3.3.1	<i>Status Register (SR)</i>	60
3.3.2	<i>Cause Register</i>	64
3.3.3	<i>Exception Restart Address (EPC) Register</i>	65
3.3.4	<i>Bad Virtual Address (BadVAddr) Register</i>	67
3.3.5	<i>Count/Compare Registers: The On-CPU Timer</i>	68
3.3.6	<i>Processor ID (PRId) Register</i>	68
3.3.7	<i>Config Registers: CPU Resource Information and Configuration</i>	69
3.3.8	<i>EBase and IntCtl: Interrupt and Exception Setup</i>	73
3.3.9	<i>SRSCtl and SRSSMap: Shadow Register Setup</i>	74
3.3.10	<i>Load-Linked Address (LLAddr) Register</i>	75
3.4	CPU Hazards—A Trap for the Unwary	75
3.4.1	<i>Hazard Barrier Instructions</i>	76
3.4.2	<i>Instruction Hazards and User Hazards</i>	77
3.4.3	<i>Hazards between CPU Instructions</i>	77
<i>Chapter 4</i>	How Caches Work on MIPS Processors	79
4.1	Caches and Cache Management	79
4.2	How Caches Work	80
4.3	Write-Through Caches in Early MIPS CPUs	83

4.4	Write-Back Caches in MIPS CPUs	84
4.5	Other Choices in Cache Design	84
4.6	Managing Caches	86
4.7	L2 and L3 Caches	88
4.8	Cache Configurations for MIPS CPUs	88
4.9	Programming MIPS32/64 Caches	90
4.9.1	<i>The Cache Instruction</i>	91
4.9.2	<i>Cache Initialization and Tag/Data Registers</i>	92
4.9.3	<i>CacheErr, ERR, and ErrorEPC Registers: Memory/Cache Error Handling</i>	94
4.9.4	<i>Cache Sizing and Figuring Out Configuration</i>	95
4.9.5	<i>Initialization Routines</i>	96
4.9.6	<i>Invalidating or Writing Back a Region of Memory in the Cache</i>	97
4.10	Cache Efficiency	98
4.11	Reorganizing Software to Influence Cache Efficiency	100
4.12	Cache Aliases	102
Chapter 5	Exceptions, Interrupts, and Initialization	105
5.1	Precise Exceptions	107
5.1.1	<i>Nonprecise Exceptions—The Multiplier in Historic MIPS CPUs</i>	108
5.2	When Exceptions Happen	109
5.3	Exception Vectors: Where Exception Handling Starts	109
5.4	Exception Handling: Basics	113
5.5	Returning from an Exception	114
5.6	Nesting Exceptions	114
5.7	An Exception Routine	115
5.8	Interrupts	115
5.8.1	<i>Interrupt Resources in MIPS CPUs</i>	116
5.8.2	<i>Implementing Interrupt Priority in Software</i>	118
5.8.3	<i>Atomicity and Atomic Changes to SR</i>	120
5.8.4	<i>Critical Regions with Interrupts Enabled: Semaphores the MIPS Way</i>	121
5.8.5	<i>Vectored and EIC Interrupts in MIPS32/64 CPUs</i>	123
5.8.6	<i>Shadow Registers</i>	124
5.9	Starting Up	124
5.9.1	<i>Probing and Recognizing Your CPU</i>	126
5.9.2	<i>Bootstrap Sequences</i>	127
5.9.3	<i>Starting Up an Application</i>	128
5.10	Emulating Instructions	128
Chapter 6	Low-level Memory Management and the TLB	131
6.1	The TLB/MMU Hardware and What It Does	131
6.2	TLB/MMU Registers Described	132
6.2.1	<i>TLB Key Fields—EntryHi and PageMask</i>	134
6.2.2	<i>TLB Output Fields—EntryLo0-1</i>	136

6.2.3	<i>Selecting a TLB Entry—Index, Random, and Wired Registers</i>	137
6.2.4	<i>Page-Table Access Helpers—Context and XContext</i>	138
6.3	TLB/MMU Control Instructions	140
6.4	Programming the TLB	141
6.4.1	<i>How Refill Happens</i>	142
6.4.2	<i>Using ASIDs</i>	143
6.4.3	<i>The Random Register and Wired Entries</i>	143
6.5	Hardware-Friendly Page Tables and Refill Mechanism	143
6.5.1	<i>TLB Miss Handling</i>	145
6.5.2	<i>XTLB Miss Handler</i>	146
6.6	Everyday Use of the MIPS TLB	147
6.7	Memory Management in a Simpler OS	149
 <i>Chapter 7</i> Floating-Point Support		 151
7.1	A Basic Description of Floating Point	151
7.2	The IEEE 754 Standard and Its Background	152
7.3	How IEEE Floating-Point Numbers Are Stored	154
7.3.1	<i>IEEE Mantissa and Normalization</i>	155
7.3.2	<i>Reserved Exponent Values for Use with Strange Values</i>	155
7.3.3	<i>MIPS FP Data Formats</i>	156
7.4	MIPS Implementation of IEEE 754	158
7.4.1	<i>Need for FP Trap Handler and Emulator in All MIPS CPUs</i>	159
7.5	Floating-Point Registers	159
7.5.1	<i>Conventional Names and Uses of Floating-Point Registers</i>	160
7.6	Floating-Point Exceptions/Interrupts	161
7.7	Floating-Point Control: The Control/Status Register	161
7.8	Floating-Point Implementation Register	165
7.9	Guide to FP Instructions	166
7.9.1	<i>Load/Store</i>	167
7.9.2	<i>Move between Registers</i>	168
7.9.3	<i>Three-Operand Arithmetic Operations</i>	169
7.9.4	<i>Multiply-Add Operations</i>	170
7.9.5	<i>Unary (Sign-Changing) Operations</i>	170
7.9.6	<i>Conversion Operations</i>	170
7.9.7	<i>Conditional Branch and Test Instructions</i>	171
7.10	Paired-Single Floating-Point Instructions and the MIPS-3D ASE	173
7.10.1	<i>Exceptions on Paired-Single Instructions</i>	174
7.10.2	<i>Paired-Single Three-Operand Arithmetic, Multiply-Add, Sign-Changing, and Nonconditional Move Operations</i>	174
7.10.3	<i>Paired-Single Conversion Operations</i>	175
7.10.4	<i>Paired-Single Test and Conditional Move Instructions</i>	176
7.10.5	<i>MIPS-3D Instructions</i>	176
7.11	Instruction Timing Requirements	179
7.12	Instruction Timing for Speed	179
7.13	Initialization and Enabling on Demand	180
7.14	Floating-Point Emulation	181

<i>Chapter 8</i>	Complete Guide to the MIPS Instruction Set	183
8.1	A Simple Example	183
8.2	Assembly Instructions and What They Mean	185
8.2.1	<i>U and Non-U Mnemonics</i>	186
8.2.2	<i>Divide Mnemonics</i>	187
8.2.3	<i>Inventory of Instructions</i>	188
8.3	Floating-Point Instructions	210
8.4	Differences in MIPS32/64 Release 1	216
8.4.1	<i>Regular Instructions Added in Release 2</i>	216
8.4.2	<i>Privileged Instructions Added in Release 2</i>	218
8.5	Peculiar Instructions and Their Purposes	218
8.5.1	<i>Load Left/Load Right: Unaligned Load and Store</i>	218
8.5.2	<i>Load-Linked/Store-Conditional</i>	223
8.5.3	<i>Conditional Move Instructions</i>	224
8.5.4	<i>Branch-Likely</i>	225
8.5.5	<i>Integer Multiply-Accumulate and Multiply-Add Instructions</i>	226
8.5.6	<i>Floating-Point Multiply-Add Instructions</i>	227
8.5.7	<i>Multiple FP Condition Bits</i>	228
8.5.8	<i>Prefetch</i>	228
8.5.9	<i>Sync: A Memory Barrier for Loads and Stores</i>	229
8.5.10	<i>Hazard Barrier Instructions</i>	231
8.5.11	<i>Synci: Cache Management for Instruction Writers</i>	232
8.5.12	<i>Read Hardware Register</i>	232
8.6	Instruction Encodings	233
8.6.1	<i>Fields in the Instruction Encoding Table</i>	233
8.6.2	<i>Notes on the Instruction Encoding Table</i>	251
8.6.3	<i>Encodings and Simple Implementation</i>	251
8.7	Instructions by Functional Group	252
8.7.1	<i>No-op</i>	252
8.7.2	<i>Register/Register Moves</i>	252
8.7.3	<i>Load Constant</i>	253
8.7.4	<i>Arithmetical/Logical</i>	253
8.7.5	<i>Integer Multiply, Divide, and Remainder</i>	255
8.7.6	<i>Integer Multiply-Accumulate</i>	256
8.7.7	<i>Loads and Stores</i>	257
8.7.8	<i>Jumps, Subroutine Calls, and Branches</i>	259
8.7.9	<i>Breakpoint and Trap</i>	260
8.7.10	<i>CPO Functions</i>	260
8.7.11	<i>Floating Point</i>	261
8.7.12	<i>Limited User-Mode Access to “Under the Hood” Features</i>	261
<i>Chapter 9</i>	Reading MIPS Assembly Language	263
9.1	A Simple Example	264
9.2	Syntax Overview	268
9.2.1	<i>Layout, Delimiters, and Identifiers</i>	268
9.3	General Rules for Instructions	269

9.3.1	<i>Computational Instructions: Three-, Two-, and One-Register</i>	269
9.3.2	<i>Immediates: Computational Instructions with Constants</i>	270
9.3.3	<i>Regarding 64-Bit and 32-Bit Instructions</i>	271
9.4	Addressing Modes	271
9.4.1	<i>Gp-Relative Addressing</i>	273
9.5	Object File and Memory Layout	274
9.5.1	<i>Practical Program Layout, Including Stack and Heap</i>	277
 <i>Chapter 10</i> Porting Software to the MIPS Architecture		279
<hr/>		
10.1	Low-Level Software for MIPS Applications: A Checklist of Frequently Encountered Problems	280
10.2	Endianness: Words, Bytes, and Bit Order	281
10.2.1	<i>Bits, Bytes, Words, and Integers</i>	281
10.2.2	<i>Software and Endianness</i>	284
10.2.3	<i>Hardware and Endianness</i>	287
10.2.4	<i>Bi-endian Software for a MIPS CPU</i>	293
10.2.5	<i>Portability and Endianness-Independent Code</i>	295
10.2.6	<i>Endianness and Foreign Data</i>	295
10.3	Trouble with Visible Caches	296
10.3.1	<i>Cache Management and DMA Data</i>	298
10.3.2	<i>Cache Management and Writing Instructions: Self-Modifying Code</i>	299
10.3.3	<i>Cache Management and Uncached or Write-Through Data</i>	300
10.3.4	<i>Cache Aliases and Page Coloring</i>	301
10.4	Memory Access Ordering and Reordering	301
10.4.1	<i>Ordering and Write Buffers</i>	304
10.4.2	<i>Implementing <code>wbflush</code></i>	304
10.5	Writing it in C	305
10.5.1	<i>Wrapping Assembly Code with the GNU C Compiler</i>	305
10.5.2	<i>Memory-Mapped I/O Registers and “Volatile”</i>	307
10.5.3	<i>Miscellaneous Issues When Writing C for MIPS Applications</i>	308
 <i>Chapter 11</i> MIPS Software Standards (ABIs)		311
<hr/>		
11.1	Data Representations and Alignment	312
11.1.1	<i>Sizes of Basic Types</i>	312
11.1.2	<i>Sizes of “long” and Pointer Types</i>	313
11.1.3	<i>Alignment Requirements</i>	313
11.1.4	<i>Memory Layout of Basic Types and How It Changes with Endianness</i>	313
11.1.5	<i>Memory Layout of Structure and Array Types and Alignment</i>	315
11.1.6	<i>Bitfields in Structures</i>	315
11.1.7	<i>Unaligned Data from C</i>	318
11.2	Argument Passing and Stack Conventions for MIPS ABIs	319
11.2.1	<i>The Stack, Subroutine Linkage, and Parameter Passing</i>	320
11.2.2	<i>Stack Argument Structure in <code>o32</code></i>	320
11.2.3	<i>Using Registers to Pass Arguments</i>	321

11.2.4	<i>Examples from the C Library</i>	322
11.2.5	<i>An Exotic Example: Passing Structures</i>	323
11.2.6	<i>Passing a Variable Number of Arguments</i>	324
11.2.7	<i>Returning a Value from a Function</i>	325
11.2.8	<i>Evolving Register-Use Standards: SGIs n32 and n64</i>	326
11.2.9	<i>Stack Layouts, Stack Frames, and Helping Debuggers</i>	329
11.2.10	<i>Variable Number of Arguments and stdargs</i>	337
Chapter 12	Debugging MIPS Designs—Debug and Profiling Features	339
12.1	The “EJTAG” On-chip Debug Unit	341
12.1.1	<i>EJTAG History</i>	343
12.1.2	<i>How the Probe Controls the CPU</i>	343
12.1.3	<i>Debug Communications through JTAG</i>	344
12.1.4	<i>Debug Mode</i>	344
12.1.5	<i>Single-Stepping</i>	346
12.1.6	<i>The dseg Memory Decode Region</i>	346
12.1.7	<i>EJTAG CP0 Registers, Particularly Debug</i>	348
12.1.8	<i>The DCR (Debug Control) Memory-Mapped Register</i>	351
12.1.9	<i>EJTAG Breakpoint Hardware</i>	352
12.1.10	<i>Understanding Breakpoint Conditions</i>	355
12.1.11	<i>Imprecise Debug Breaks</i>	356
12.1.12	<i>PC Sampling with EJTAG</i>	356
12.1.13	<i>Using EJTAG without a Probe</i>	356
12.2	Pre-EJTAG Debug Support—Break Instruction and CP0 Watchpoints	358
12.3	PDtrace	359
12.4	Performance Counters	360
Chapter 13	GNU/Linux from Eight Miles High	363
13.1	Components	364
13.2	Layering in the Kernel	368
13.2.1	<i>MIPS CPU in Exception Mode</i>	368
13.2.2	<i>MIPS CPU with Some or All Interrupts off</i>	369
13.2.3	<i>Interrupt Context</i>	370
13.2.4	<i>Executing the Kernel in Thread Context</i>	370
Chapter 14	How Hardware and Software Work Together	371
14.1	The Life and Times of an Interrupt	371
14.1.1	<i>High-Performance Interrupt Handling and Linux</i>	374
14.2	Threads, Critical Regions, and Atomicity	375
14.2.1	<i>MIPS Architecture and Atomic Operations</i>	376
14.2.2	<i>Linux Spinlocks</i>	377
14.3	What Happens on a System Call	378
14.4	How Addresses Get Translated in Linux/MIPS Systems	380

14.4.1	<i>What's Memory Translation For?</i>	382
14.4.2	<i>Basic Process Layout and Protection</i>	384
14.4.3	<i>Mapping Process Addresses to Real Memory</i>	385
14.4.4	<i>Paged Mapping Preferred</i>	386
14.4.5	<i>What We Really Want</i>	387
14.4.6	<i>Origins of the MIPS Design</i>	389
14.4.7	<i>Keeping Track of Modified Pages (Simulating "Dirty" Bits)</i>	392
14.4.8	<i>How the Kernel Services a TLB Refill Exception</i>	393
14.4.9	<i>Care and Maintenance of the TLB</i>	397
14.4.10	<i>Memory Translation and 64-Bit Pointers</i>	397
Chapter 15	MIPS Specific Issues in the Linux Kernel	399
15	Explicit Cache Management	399
15.1.1	<i>DMA Device Accesses</i>	399
15.1.2	<i>Writing Instructions for Later Execution</i>	401
15.1.3	<i>Cache/Memory Mapping Problems</i>	401
15.1.4	<i>Cache Aliases</i>	402
15.2	CP0 Pipeline Hazards	403
15.3	Multiprocessor Systems and Coherent Caches	403
15.4	Demon Tweaks for a Critical Routine	406
Chapter 16	Linux Application Code, PIC, and Libraries	409
16.1	How Link Units Get into a Program	411
16.2	Global Offset Table (GOT) Organization	412
Appendix A	MIPS Multithreading	415
A.1	What Is Multithreading?	415
A.2	Why Is MT Useful?	417
A.3	How to Do Multithreading for MIPS	417
A.4	MT in Action	421
Appendix B	Other Optional Extensions to the MIPS Instruction Set	425
B.1	MIPS16 and MIPS16e ASEs	425
B.1.1	<i>Special Encodings and Instructions in the MIPS16 ASE</i>	426
B.1.2	<i>The MIPS16 ASE Evaluated</i>	427
B.2	The MIPS DSP ASE	428
B.3	The MDMX ASE	429
	<i>MIPS Glossary</i>	431
	<i>References</i>	477
	<i>Books and Articles</i>	477
	<i>Online Resources</i>	478
	<i>Index</i>	481

Preface

This book is about MIPS, the cult hit from the mid-1980s' crop of RISC CPU designs. These days MIPS is not the highest-volume 32-bit architecture, but it is in a comfortable second place. Where it wins, hands down, is its range of applications. A piece of equipment built around a MIPS CPU might have cost you \$35 for a wireless router or hundreds of thousands of dollars for an SGI supercomputer (though with SGI's insolvency, those have now reached the end of the line). Between those extremes are Sony and Nintendo games machines, many Cisco routers, TV set-top boxes, laser printers, and so on.

The first edition of this book has sold close to 10,000 English copies over the years and has been translated into Chinese. I'm pleased and surprised; I didn't know there were so many MIPS programmers out there.

This second edition is *See MIPS Run... Linux*. The first edition struggled to motivate some features of the MIPS architecture, because they don't make sense unless you can see how they help out inside an OS kernel. But now a lot of you have some sense of how Linux works, and I can quote its source code; more importantly, I can refer to it knowing that those of you who get interested can read the source code and find out how it's *really* done.

So this is a book about the MIPS architecture, but the last three chapters stroll through the Linux kernel and application-programming system to cast light on what those weird features do. I hope Linux experts will forgive my relative ignorance of Linux details, but the chance to go for a description of a real OS running on a real architecture was too good to pass up.

MIPS is a RISC: a useful acronym, well applied to the common features of a number of computer architectures invented in the 1980s, to realize efficient pipelined implementation. The acronym CISC is vaguer. I'll use it in a narrow sense, for the kind of features found in x86 and other pre-1982 architectures, designed with microcoded implementations in mind.

Some of you may be up in arms: He's confusing implementation with architecture! But while computer architecture is supposed to be a contract with the

programmer about what programs will run correctly, it's also an engineering design in its own right. A computer architecture is designed to make for good CPUs. As chip design becomes more sophisticated, the trade-offs change.

This book is for programmers, and that's the test we've used to decide what gets included—if a programmer might see it, or is likely to be interested, it's here. That means we don't get to discuss, for example, the strange system interfaces with which MIPS has tortured two generations of hardware design engineers. And your operating system may hide many of the details we talk about here; there is many an excellent programmer who thinks that C is quite low level enough, portability a blessing, and detailed knowledge of the architecture irrelevant. But sometimes you do need to get down to the nuts and bolts—and human beings are born curious as to how bits of the world work.

A result of this orientation is that we'll tend to be rather informal when describing things that may not be familiar to a software engineer—particularly the inner workings of the CPU—but we'll get much more terse and technical when we're dealing with the stuff programmers have met before, such as registers, instructions, and how data is stored in memory.

We'll assume some familiarity and comfort with the C language. Much of the reference material in the book uses C fragments as a way of compressing operation descriptions, particularly in the chapters on the details of the instruction set and assembly language.

Some parts of the book are targeted at readers who've seen some assembly language: the ingenuity and peculiarity of the MIPS architecture shows up best from that viewpoint. But if assembly is a closed book to you, that's probably not a disaster.

This book aims to tell you everything you need to know about programming generic MIPS CPUs. More precisely, it describes the architecture as it's defined by MIPS Technologies' MIPS32 and MIPS64—specifically, the second release of those specifications from 2003. We'll shorten that to “MIPS32/64.” But this is not just a reference manual: To keep an architecture in your head means coming to understand it in the round. I also hope the book will interest students of programming (at college or enrolled in the school of life) who want to understand a modern CPU architecture all the way through.

If you plan to read this book straight through from front to back, you will expect to find a progression from overview to detail, and you won't be disappointed. But you'll also find some progression through history; the first time we talk about a concept we'll usually focus on its first version. Hennessy and Patterson call this “learning through evolution,” and what's good enough for them is certainly good enough for me.

We start in Chapter 1 with some history and background, and set MIPS in context by discussing the technological concerns and ideas that were uppermost in the minds of its inventors. Then in Chapter 2 we discuss the characteristics of the MIPS machine language that follow from their approach.

To help you see the big picture, we leave out the details of processor control until Chapter 3, which introduces the ugly but eminently practical system that allows MIPS CPUs to deal with their caches, exceptions and startup, and memory management. Those last three topics, respectively, become the subjects of Chapters 4 through 6.

The MIPS architecture has been careful to separate out the part of the instruction set that deals with floating-point numbers. That separation allows MIPS CPUs to be built with various levels of floating-point support, from none at all through partial implementations to somewhere near the state of the art. So we have also separated out the floating-point functions, and we keep them back until Chapter 7.

Up to this point, the chapters follow a reasonable sequence for getting to know MIPS. The following chapters change gear and are more like reference manuals or example-based tutorials.

In Chapter 8, we go through the whole machine instruction set; the intention is to be precise but much more terse than the standard MIPS reference works—we cover in 10 pages what takes a hundred in other sources.¹ Chapter 9 is a brief introduction to reading and writing assembly, and falls far short of an assembly programming manual.

Chapter 10 is a checklist with helpful hints for those of you who have to port software between another CPU and a MIPS CPU. The longest section tackles the troublesome problem of endianness in CPUs, software, and systems.

Chapter 11 is a bare-bones summary of the software conventions (register use, argument passing, etc.) necessary to produce interworking software with different toolkits. Chapter 12 introduces the debug and profiling features standardized for MIPS CPUs.

Then we're on to seeing how MIPS runs GNU/Linux. We describe relationship between the Linux kernel and a computer architecture in Chapter 13; then Chapters 14 and 15 dig down into some of the detail as to how the MIPS architecture does what the Linux kernel needs. Chapter 16 gives you a quick look at the dynamic linking magic that makes GNU/Linux applications work.

Appendix A covers the MIPS MT (multithreading) extension, probably the most important addition to the architecture in many years. And Appendix B describes the more important add-ons: MIPS16, the new MIPS DSP extensions, and MDMX.

You will also find at the end of this book a glossary of terms—a good place to look for specialized or unfamiliar usage and acronyms—and a list of books, papers, and online references for further reading.

1. I have taken considerable care in the generation of these tables, and they are mostly right. But if your system depends on it, be sure to cross-check this information. An excellent source of fairly reliable information can be found in the behavior and source code of the GNU tool collection—but I referred to that too, so it's not completely independent.

Style and Limits

Every book reflects its author, so we'd better make a virtue of it.

Since some of you will be students, I wondered whether I should distinguish general use from MIPS use. I decided not to; I aim to be specific except where it costs the reader nothing to be general. I also try to be concrete rather than abstract. I don't worry overmuch about whatever meaning terms like "TLB" have in the wider industry, but I explain them in a MIPS context. Human beings are great generalizers, and this is unlikely to damage your learning much.

It's 20 years since I started working with MIPS CPUs in the fall of 1986. Some of the material in this book goes back as far as 1988, when I started giving training courses on MIPS architecture. In 1993, I gathered them together to make a software manual focused on IDT's R3051 family CPUs. It took quite a lot of extra material to create the first edition, published in 1999.

A lot has happened since 1999. MIPS is now at the very end of its life in servers with SGI but has carved out a significant niche in embedded systems. Linux has emerged as the most-used OS for embedded MIPS, but there's still a lot of diversity in the embedded market. The MIPS specifications have been reorganized around MIPS32 and MIPS64 (which this edition regards as the baseline). This second edition has been in the works for about three years.

The MIPS story continues; if it did not, we'd only be writing this book for historians, and Morgan Kaufmann wouldn't be very interested in publishing it. MIPS developments that weren't announced by the end of 2005 are much too late for this edition.

Conventions

A quick note on the typographical conventions used in this book:

- Type in this font (Minion) is running text.
- Type in this font (Futura) is a sidebar.
- **Type in this font (Courier bold) is used for assembly code and MIPS register names.**
- Type in this font (Courier) is used for C code and hexadecimals.
- *Type in this font (Minion italic, small) is used for hardware signal names.*

Acknowledgments

The themes in this book have followed me through my computing career. Mike Cole got me excited about computing, and I've been trying to emulate his skill in picking out good ideas ever since. In the brief but exciting life of Whitechapel Workstations (1983–1988), many colleagues taught me something about computer architecture and about how to design hardware—Bob Newman and Rick Filipkiewicz probably the most. I also have to thank Whitechapel's salesperson, Dave Gravell, for originally turning me on to MIPS. My fellow engineers during the lifetime of Algorithmics Ltd. (Chris Dearman, Rick Filipkiewicz, Gerald Onions, Nigel Stephens, and Chris Shaw) have to be doubly thanked, both for all I've learned through innumerable discussions, arguments, and designs and for putting up with the book's competition for my time.

Many thanks are due to the reviewers who've read chapters over a long period of time: Phil Bourekas of Integrated Device Technology, Inc.; Thomas Daniel of the LSI Logic Corporation; Mike Murphy of Silicon Graphics, Inc.; and David Nagle of Carnegie Mellon University.

On the second edition: I've known Paul Cobb for a long time, as we both worked around MIPS companies. Paul contributed material updating the historical survey of MIPS CPUs and the programming chapter and cleaning up the references. In all cases, though, I've had a final edit—so any errors are mine.

During the preparation of this edition I've been employed by MIPS Technologies Inc. It's dangerous to pick out some colleagues and not others, but I'll do it anyway.

Ralf Baechle runs the `www.linux-mips.org` site, which coordinates MIPS work on the Linux kernel. He's been very helpful in dispelling some of the illusions I'd formed about Linux: I started off thinking it was like other operating systems . . . (Robert Love's *Linux Kernel Development* book helped too; I warmly recommend it to anyone who wants a more educated guidebook to the kernel). Thanks to MIPS Technologies and my various managers for being flexible about my time, and to many colleagues at MIPS Technologies (too many to name) who have read and commented on drafts.

Todd Bezenek has been my most persistent colleague/reviewer of this edition. Reviewers outside MIPS Technologies did it for their love and respect for the field: notable contributors were Steven Hill (Reality Diluted, Inc.), Jun Sun (DoCoMo USA Labs), Eric DeVolder, and Sophie Wilson.

Denise Penrose is easily the Best Editor Ever. Not many people in Finsbury Park (my home in North London) can say they're just flying to San Francisco for brunch with their publisher.

Last but not least, thanks to Carol O'Brien, who was rash enough to marry me in the middle of this rewrite.