


计算机组成与设计

硬件/软件接口

(英文版·第3版)



David A. Patterson
加州大学伯克利分校
John L. Hennessy
斯坦福大学 著



机械工业出版社
China Machine Press

经典原版书库

计算机组成与设计

硬件/软件接口

(英文版·第3版)

Computer Organization and Design

The Hardware/Software Interface

(Third Edition)

(美) David A. Patterson
加州大学伯克利分校
John L. Hennessy
斯坦福大学 著



机械工业出版社
China Machine Press

David A. Patterson and John L. Hennessy: Computer Organization and Design: The Hardware/Software Interface, Third Edition (ISBN: 1-55860-604-1 ISBN-13: 978-1-55860-604-3).

Original English language edition copyright © 2005 by Elsevier Inc. All rights reserved.

Authorized English language reprint edition published by the Proprietor.

ISBN: 981-259-715-8 ISBN-13: 978-981-259-715-1

Copyright © 2006 by Elsevier (Singapore) Pte Ltd.

Printed in China by China Machine Press under special arrangement with Elsevier (Singapore) Pte Ltd. This edition is authorized for sale in China only, excluding Hong Kong SAR and Taiwan.

Unauthorized export of this edition is a violation of the Copyright Act. Violation of this Law is subject to Civil and Criminal Penalties.

本书英文影印版由Elsevier (Singapore) Pte Ltd.授权机械工业出版社在中国大陆境内独家发行。本版仅限在中国境内(不包括香港特别行政区及台湾地区)出版及标价销售。未经许可之出口,视为违反著作权法,将受法律之制裁。

版权所有,侵权必究。

本书法律顾问 北京市展达律师事务所

本书版权登记号:图字:01-2006-3122

图书在版编目(CIP)数据

计算机组成与设计:硬件/软件接口(英文版·第3版)/(美)帕特森(Patterson, D. A.)等著.-北京:机械工业出版社,2006.7

(经典原版书库)

书名原文:Computer Organization and Design: The Hardware/Software Interface, Third Edition
ISBN 7-111-19339-3

I. 计… II. 帕… III. ① 计算机体系结构-英文 ② 微型计算机-接口设备-英文
IV. ① TP303 ② TP364

中国版本图书馆CIP数据核字(2006)第062551号

机械工业出版社(北京市西城区百万庄大街22号 邮政编码 100037)

责任编辑:迟振春

北京中兴印刷有限公司印刷·新华书店北京发行所发行

2006年7月第1版第1次印刷

186mm×240mm·41.5印张

定价:85.00元(附光盘)

凡购本书,如有倒页、脱页、缺页,由本社发行部调换
本社购书热线(010)68326294

MIPS Reference Data

①



CORE INSTRUCTION SET

NAME	MNE-MON-FOR-IC MAT	OPERATION (in Verilog)	OPCODE/FUNCT (Hex)
Add	add R	$R[rd] = R[rs] + R[rt]$	(1) 0/20 _{hex}
Add Immediate	addi I	$R[rt] = R[rs] + \text{SignExtImm}$	(1)(2) 8 _{hex}
Add Imm. Unsigned	addiu I	$R[rt] = R[rs] + \text{SignExtImm}$	(2) 9 _{hex}
Add Unsigned	addu R	$R[rd] = R[rs] + R[rt]$	0/21 _{hex}
And	and R	$R[rd] = R[rs] \& R[rt]$	0/24 _{hex}
And Immediate	andi I	$R[rt] = R[rs] \& \text{ZeroExtImm}$	(3) C _{hex}
Branch On Equal	beq I	$\text{if}(R[rs] == R[rt])$ $\text{PC} = \text{PC} + 4 + \text{BranchAddr}$	(4) 4 _{hex}
Branch On Not Equal	bne I	$\text{if}(R[rs] != R[rt])$ $\text{PC} = \text{PC} + 4 + \text{BranchAddr}$	(4) 5 _{hex}
Jump	j J	$\text{PC} = \text{JumpAddr}$	(5) 2 _{hex}
Jump And Link	jal J	$R[31] = \text{PC} + 4; \text{PC} = \text{JumpAddr}$	(5) 3 _{hex}
Jump Register	jr R	$\text{PC} = R[rs]$	0/08 _{hex}
Load Byte Unsigned	lbu I	$R[rt] = \{24'b0, M[R[rs]] + \text{SignExtImm}(7:0)\}$	(2) 0/24 _{hex}
Load Halfword Unsigned	lhu I	$R[rt] = \{16'b0, M[R[rs]] + \text{SignExtImm}(15:0)\}$	(2) 0/25 _{hex}
Load Upper Imm.	lui J	$R[rt] = \{\text{imm}, 16'b0\}$	f _{hex}
Load Word	lw I	$R[rt] = M[R[rs]] + \text{SignExtImm}$	(2) 0/23 _{hex}
Nor	nor R	$R[rd] = \sim(R[rs]) \sim(R[rt])$	0/27 _{hex}
Or	or R	$R[rd] = R[rs] R[rt]$	0/25 _{hex}
Or Immediate	ori I	$R[rt] = R[rs] \text{ZeroExtImm}$	(3) d _{hex}
Set Less Than	slt R	$R[rd] = (R[rs] < R[rt]) ? 1 : 0$	0/2a _{hex}
Set Less Than Imm.	slti I	$R[rt] = (R[rs] < \text{SignExtImm}) ? 1 : 0$	(2) a _{hex}
Set Less Than Imm. Unsigned	sltiu I	$R[rt] = (R[rs] < \text{SignExtImm}) ? 1 : 0$	(2)(6) b _{hex}
Set Less Than Unsigned	sltu R	$R[rd] = (R[rs] < R[rt]) ? 1 : 0$	(6) 0/2b _{hex}
Shift Left Logical	sll R	$R[rd] = R[rs] \ll \text{shamt}$	0/00 _{hex}
Shift Right Logical	srl R	$R[rd] = R[rs] \gg \text{shamt}$	0/02 _{hex}
Store Byte	sb I	$M[R[rs] + \text{SignExtImm}(7:0)] = R[rt](7:0)$	(2) 28 _{hex}
Store Halfword	sh I	$M[R[rs] + \text{SignExtImm}(15:0)] = R[rt](15:0)$	(2) 29 _{hex}
Store Word	sw I	$M[R[rs] + \text{SignExtImm}] = R[rt]$	(2) 2b _{hex}
Subtract	sub R	$R[rd] = R[rs] - R[rt]$	(1) 0/22 _{hex}
Subtract Unsigned	subu R	$R[rd] = R[rs] - R[rt]$	0/23 _{hex}

- May cause overflow exception
- $\text{SignExtImm} = \{16\{\text{immediate}[15]\}, \text{immediate}\}$
- $\text{ZeroExtImm} = \{16\{1'b'0\}, \text{immediate}\}$
- $\text{BranchAddr} = \{14\{\text{immediate}[15]\}, \text{immediate}, 2'b'0\}$
- $\text{JumpAddr} = \{\text{PC}[31:28], \text{address}, 2'b'0\}$
- Operands considered unsigned numbers (vs. 2's comp.)

BASIC INSTRUCTION FORMATS

R	opcode	rs	rt	rd	shamt	funct
31	26-25	21-20	16-15	11-10	6-5	0

I	opcode	rs	rt	immediate
31	26-25	21-20	16-15	

J	opcode	address
31	26-25	

ARITHMETIC CORE INSTRUCTION SET

②

NAME	MNE-MON-FOR-IC MAT	OPERATION	OPCODE/FUNCT (Hex)
Branch On FP True	bclt FI	$\text{if}(\text{FPcond}) \text{PC} = \text{PC} + 4 + \text{BranchAddr}$	(4) 11/8/1-
Branch On FP False	bclf FI	$\text{if}(\sim \text{FPcond}) \text{PC} = \text{PC} + 4 + \text{BranchAddr}$	(4) 11/8/0-
Divide	div R	$\text{Lo} = R[rs]/R[rt]; \text{Hi} = R[rs]\%R[rt]$	0/-/-1a
Divide Unsigned	divu R	$\text{Lo} = R[rs]/R[rt]; \text{Hi} = R[rs]\%R[rt]$	(6) 0/-/-1b
FP Add Single	add.s FR	$F[fd] = F[fs] + F[ft]$	11/10/-/0
FP Add Double	add.d FR	$\{F[fd], F[fd+1]\} = \{F[fs], F[fs+1]\} + \{F[ft], F[ft+1]\}$	11/11/-/0
FP Compare Single	c.x.s* FR	$\text{FPcond} = (F[fs] \text{ op } F[ft]) ? 1 : 0$	11/10/-/y
FP Compare Double	c.x.d* FR	$\text{FPcond} = (\{F[fs], F[fs+1]\} \text{ op } \{F[ft], F[ft+1]\}) ? 1 : 0$	11/11/-/y
* (x is eq, lt, or le) (op is ==, <, or <=) (y is 32, 3c, or 3e)			
FP Divide Single	div.s FR	$F[fd] = F[fs] / F[ft]$	11/10/-/3
FP Divide Double	div.d FR	$\{F[fd], F[fd+1]\} = \{F[fs], F[fs+1]\} / \{F[ft], F[ft+1]\}$	11/11/-/3
FP Multiply Single	mul.s FR	$F[fd] = F[fs] * F[ft]$	11/10/-/2
FP Multiply Double	mul.d FR	$\{F[fd], F[fd+1]\} = \{F[fs], F[fs+1]\} * \{F[ft], F[ft+1]\}$	11/11/-/2
FP Subtract Single	sub.s FR	$F[fd] = F[fs] - F[ft]$	11/10/-/1
FP Subtract Double	sub.d FR	$\{F[fd], F[fd+1]\} = \{F[fs], F[fs+1]\} - \{F[ft], F[ft+1]\}$	11/11/-/1
Load FP Single	lwc1 I	$F[rt] = M[R[rs] + \text{SignExtImm}]$	(2) 31/-/-/1
Load FP Double	ldc1 I	$\{F[rt], F[rt+1]\} = M[R[rs] + \text{SignExtImm}]; \{F[rt+1], F[rt+2]\} = M[R[rs] + \text{SignExtImm} + 4]$	(2) 35/-/-/1
Move From Hi	mfmhi R	$R[rd] = \text{Hi}$	0/-/-/10
Move From Lo	mfmlo R	$R[rd] = \text{Lo}$	0/-/-/12
Move From Control	mfmco R	$R[rd] = \text{CR}[rs]$	16/0/-/0
Multiply	mult R	$\{\text{Hi}, \text{Lo}\} = R[rs] * R[rt]$	0/-/-/18
Multiply Unsigned	multu R	$\{\text{Hi}, \text{Lo}\} = R[rs] * R[rt]$	(6) 0/-/-/19
Store FP Single	swc1 I	$M[R[rs] + \text{SignExtImm}] = F[rt]$	(2) 39/-/-/1
Store FP Double	sdc1 I	$\{M[R[rs] + \text{SignExtImm}] = F[rt]; M[R[rs] + \text{SignExtImm} + 4] = F[rt+1]\}$	(2) 3d/-/-/1

FLOATING POINT INSTRUCTION FORMATS

FR	opcode	fmt	ft	fs	fd	funct
31	26-25	21-20	16-15	11-10	6-5	0

FI	opcode	fmt	ft	immediate
31	26-25	21-20	16-15	

PSEUDO INSTRUCTION SET

NAME	MNEMONIC	OPERATION
Branch Less Than	blt	$\text{if}(R[rs] < R[rt]) \text{PC} = \text{Label}$
Branch Greater Than	bgt	$\text{if}(R[rs] > R[rt]) \text{PC} = \text{Label}$
Branch Less Than or Equal	ble	$\text{if}(R[rs] \leq R[rt]) \text{PC} = \text{Label}$
Branch Greater Than or Equal	bge	$\text{if}(R[rs] \geq R[rt]) \text{PC} = \text{Label}$
Load Immediate	li	$R[rd] = \text{immediate}$
Move	move	$R[rd] = R[rs]$

REGISTER NAME, NUMBER, USE, CALL CONVENTION

NAME	NUMBER	USE	PRESERVED ACROSS A CALL?
\$zero	0	The Constant Value 0	N.A.
\$at	1	Assembler Temporary	No
\$v0-\$v1	2-3	Values for Function Results and Expression Evaluation	No
\$a0-\$a3	4-7	Arguments	No
\$t0-\$t7	8-15	Temporaries	No
\$s0-\$s7	16-23	Saved Temporaries	Yes
\$t8-\$t9	24-25	Temporaries	No
\$k0-\$k1	26-27	Reserved for OS Kernel	No
\$gp	28	Global Pointer	Yes
\$sp	29	Stack Pointer	Yes
\$fp	30	Frame Pointer	Yes
\$ra	31	Return Address	Yes

OPCODES, BASE CONVERSION, ASCII SYMBOLS

MIPS opcode (31:26)	(1) MIPS funct (5:0)	(2) MIPS funct (5:0)	Binary	Decimal	Hexa-decimal	ASCII Character	Decimal	Hexa-decimal	ASCII Character
(1)	sll	add _f	00 0000	0	0	NUL	64	40	@
	sub _f	sub _f	00 0001	1	1	SOH	65	41	A
	mul _f	mul _f	00 0010	2	2	STX	66	42	B
	div _f	div _f	00 0011	3	3	ETX	67	43	C
	sqrt _f	sqrt _f	00 0100	4	4	EOT	68	44	D
	abs _f	abs _f	00 0101	5	5	ENQ	69	45	E
	mov _f	mov _f	00 0110	6	6	ACK	70	46	F
	neg _f	neg _f	00 0111	7	7	BEL	71	47	G
	jr		00 1000	8	8	BS	72	48	H
	jalr		00 1001	9	9	HT	73	49	I
	movz		00 1010	10	a	LF	74	4a	J
	movn		00 1011	11	b	VT	75	4b	K
	syscall	round.w _f	00 1100	12	c	FF	76	4c	L
	break	trunc.w _f	00 1101	13	d	CR	77	4d	M
		ceil.w _f	00 1110	14	e	SO	78	4e	N
	sync	floor.w _f	00 1111	15	f	SI	79	4f	O
(2)	mfi		01 0000	16	10	DLE	80	50	P
	mthi		01 0001	17	11	DC1	81	51	Q
	mflo	movz _f	01 0010	18	12	DC2	82	52	R
	mtlo	movn _f	01 0011	19	13	DC3	83	53	S
			01 0100	20	14	DC4	84	54	T
			01 0101	21	15	NAK	85	55	U
			01 0110	22	16	SYN	86	56	V
			01 0111	23	17	ETB	87	57	W
	mult		01 1000	24	18	CAN	88	58	X
	multu		01 1001	25	19	EM	89	59	Y
	div		01 1010	26	1a	SUB	90	5a	Z
	divu		01 1011	27	1b	ESC	91	5b	{
			01 1100	28	1c	FS	92	5c	
			01 1101	29	1d	GS	93	5d	^
			01 1110	30	1e	RS	94	5e	_
			01 1111	31	1f	US	95	5f	`
	lb	add	10 0000	32	20	Space	96	60	
	lh	addu	10 0001	33	21	1	97	61	a
	lwl	sub	10 0010	34	22	2	98	62	b
	lwr	subu	10 0011	35	23	#	99	63	c
	lbu	and	10 0100	36	24	\$	100	64	d
	lhu	or	10 0101	37	25	%	101	65	e
	lwr	xor	10 0110	38	26	&	102	66	f
		nor	10 0111	39	27	'	103	67	g
	sb		10 1000	40	28	(104	68	h
	sh		10 1001	41	29)	105	69	i
	swl	sll	10 1010	42	2a	*	106	6a	j
	sw	slltu	10 1011	43	2b	+	107	6b	k
			10 1100	44	2c	,	108	6c	l
			10 1101	45	2d	-	109	6d	m
	swr		10 1110	46	2e	.	110	6e	n
	cache		10 1111	47	2f	/	111	6f	o
	ll	tge	11 0000	48	30	0	112	70	p
	lwc1	tgeu	11 0001	49	31	1	113	71	q
	lwc2	tle	11 0010	50	32	2	114	72	r
	pcrf	tltu	11 0011	51	33	3	115	73	s
	teq	c.oilt _f	11 0100	52	34	4	116	74	t
	ldc1	c.oile _f	11 0101	53	35	5	117	75	u
	ldc2	tne	11 0110	54	36	6	118	76	v
		c.oile _f	11 0111	55	37	7	119	77	w
	sc	c.sf _f	11 1000	56	38	8	120	78	x
	sw1	c.ngle _f	11 1001	57	39	9	121	79	y
	sw2	c.seq _f	11 1010	58	3a	*	122	7a	z
		c.ngl _f	11 1011	59	3b	{	123	7b	{
		c.lt _f	11 1100	60	3c	<	124	7c	<
	sdcl	c.ngse _f	11 1101	61	3d	=	125	7d	=
	sdcl	c.le _f	11 1110	62	3e	>	126	7e	>
		c.ngt _f	11 1111	63	3f	?	127	7f	DEL


(1) opcode(31:26) = 0
 (2) opcode(31:26) = 17_{hex} (11_{hex}); if fmt(25:21) = 16_{hex} (10_{hex}) f = s (single);
 if fmt(25:21) = 17_{hex} (11_{hex}) f = d (double)

IEEE 754 FLOATING POINT STANDARD

$$(-1)^S \times (1 + \text{Fraction}) \times 2^{(\text{Exponent} - \text{Bias})}$$

where Single Precision Bias = 127,
 Double Precision Bias = 1023.

IEEE Single Precision and Double Precision Formats:




IEEE 754 Symbols


Exponent	Fraction	Object
0	0	± 0
0	≠0	± Denorm
1 to MAX - 1	anything	± Fl. Pt. Num.
MAX	0	±∞
MAX	≠0	NaN

S.P. MAX = 255, D.P. MAX = 2047


MEMORY ALLOCATION




STACK FRAME



DATA ALIGNMENT



EXCEPTION CONTROL REGISTERS: CAUSE AND STATUS



BD = Branch Delay, UM = User Mode, EL = Exception Level, IE = Interrupt Enable

EXCEPTION CODES

Number	Name	Cause of Exception	Number	Name	Cause of Exception
0	Int	Interrupt (hardware)	9	Bp	Breakpoint Exception
4	AdE	Address Error Exception (load or instruction fetch)	10	RI	Reserved Instruction Exception
5	AdES	Address Error Exception (store)	11	CpU	Coprocessor Unimplemented
6	IBE	Bus Error on Instruction Fetch	12	Ov	Arithmetic Overflow Exception
7	DBE	Bus Error on Load or Store	13	Tr	Trap
8	Sys	System Exception	15	FPE	Floating Point Exception

SIZE PREFIXES (10⁴ for Disk, Communication; 2⁸ for Memory)

SIZE	PRE-FIX	SIZE	PRE-FIX	SIZE	PRE-FIX	SIZE	PRE-FIX
10 ³	Kilo	10 ¹⁵	Peta	10 ⁻³	milli	10 ⁻¹⁵	femto
10 ⁶	Mega	10 ¹⁸	Exa	10 ⁻⁶	micro	10 ⁻¹⁸	atto
10 ⁹	Giga	10 ²¹	Zetta	10 ⁻⁹	nano	10 ⁻²¹	zepto
10 ¹²	Tera	10 ²⁴	Yotta	10 ⁻¹²	pico	10 ⁻²⁴	yocto

The symbol for each prefix is just its first letter, except μ is used for micro.

1. Pull along perforation to separate card. 2. Fold bottom side (columns 3 and 4) together.

出版者的话

文艺复兴以降，源远流长的科学精神和逐步形成的学术规范，使西方国家在自然科学的各个领域取得了垄断性的优势；也正是这样的传统，使美国在信息技术发展的六十多年间名家辈出、独领风骚。在商业化的进程中，美国的产业界与教育界越来越紧密地结合，计算机学科中的许多泰山北斗同时身处科研和教学的最前线，由此而产生的经典科学著作，不仅擘划了研究的范畴，还揭橥了学术的源变，既遵循学术规范，又自有学者个性，其价值并不会因年月的流逝而减退。

近年，在全球信息化大潮的推动下，我国的计算机产业发展迅猛，对专业人才的需求日益迫切。这对计算机教育界和出版界都既是机遇，也是挑战；而专业教材的建设在教育战略上显得举足轻重。在我国信息技术发展时间较短、从业人员较少的现状下，美国等发达国家在其计算机科学发展的几十年间积淀的经典教材仍有许多值得借鉴之处。因此，引进一批国外优秀计算机教材将对我国计算机教育事业的发展起积极的推动作用，也是与世界接轨、建设真正的世界一流大学的必由之路。

机械工业出版社华章图文信息有限公司较早意识到“出版要为教育服务”。自1998年开始，华章公司就将工作重点放在了遴选、移译国外优秀教材上。经过几年的不懈努力，我们与Prentice Hall, Addison-Wesley, McGraw-Hill, Morgan Kaufmann等世界著名出版公司建立了良好的合作关系，从它们现有的数百种教材中甄选出Tanenbaum, Stroustrup, Kernighan, Jim Gray等大师名家的一批经典作品，以“计算机科学丛书”为总称出版，供读者学习、研究及收藏。大理石纹理的封面，也正体现了这套丛书的品位和格调。

“计算机科学丛书”的出版工作得到了国内外学者的鼎力襄助，国内的专家不仅提供了中肯的选题指导，还不辞劳苦地担任了翻译和审校的工作；而原书的作者也相当关注其作品在中国的传播，有的还专程为其书的中译本作序。迄今，“计算机科学丛书”已经出版了近百个品种，这些书籍在读者中树立了良好的口碑，并被许多高校采用为正式教材和参考书籍，为进一步推广与发展打下了坚实的基础。

随着学科建设的初步完善和教材改革的逐渐深化，教育界对国外计算机教材的需求和应用都步入一个新的阶段。为此，华章公司将加大引进教材的力度，在“华章教育”的总规划之下出版三个系列的计算机教材：除“计算机科学丛书”之外，对影印版的教材，则单独开辟出“经典原版书库”；同时，引进全美通行的教学辅导书“Schaum's Outlines”系列组成“全美经典学习指导系列”。为了保证这三套丛书的权威性，同时也为了更好地为学校和老师服务，华章公司聘请了中国科学院、北京大学、清华大学、国防科技大学、复旦大学、上海交通大学、南京大学、浙江大学、中国科技大学、哈尔滨工业大学、西安交通大学、中国人民大学、北京航空航天大学、北京邮电大学、中山大学、解放军理工大学、郑州大学、湖北工学院、中国国

家信息安全测评认证中心等国内重点大学和科研机构在计算机的各个领域的著名学者组成“专家指导委员会”，为我们提供选题意见和出版监督。

这三套丛书是响应教育部提出的使用外版教材的号召，为国内高校的计算机及相关专业的教学度身订造的。其中许多教材均已为M. I. T., Stanford, U.C. Berkeley, C. M. U. 等世界名牌大学所采用。不仅涵盖了程序设计、数据结构、操作系统、计算机体系结构、数据库、编译原理、软件工程、图形学、通信与网络、离散数学等国内大学计算机专业普遍开设的核心课程，而且各具特色——有的出自语言设计者之手、有的历经三十年而不衰、有的已被全世界的几百所高校采用。在这些圆熟通博的名师大作的指引之下，读者必将在计算机科学的宫殿中由登堂而入室。

权威的作者、经典的教材、一流的译者、严格的审校、精细的编辑，这些因素使我们的图书有了质量的保证，但我们的目标是尽善尽美，而反馈的意见正是我们达到这一终极目标的重要帮助。教材的出版只是我们的后续服务的起点。华章公司欢迎老师和读者对我们的工作提出建议或给予指正，我们的联系方式如下：

电子邮件：hzjsj@hzbook.com

联系电话：(010) 68995264

联系地址：北京市西城区百万庄南街1号

邮政编码：100037

专家指导委员会

(按姓氏笔画顺序)

尤晋元	王 珊	冯博琴	史忠植	史美林
石教英	吕 建	孙玉芳	吴世忠	吴时霖
张立昂	李伟琴	李师贤	李建中	杨冬青
邵维忠	陆丽娜	陆鑫达	陈向群	周伯生
周克定	周傲英	孟小峰	岳丽华	范 明
郑国梁	施伯乐	钟玉琢	唐世渭	袁崇义
高传善	梅 宏	程 旭	程时端	谢希仁
裘宗燕	戴 葵			

Preface

*The most beautiful thing we can experience is the mysterious.
It is the source of all true art and science.*

Albert Einstein, *What I Believe*, 1930

About This Book

We believe that learning in computer science and engineering should reflect the current state of the field, as well as introduce the principles that are shaping computing. We also feel that readers in every specialty of computing need to appreciate the organizational paradigms that determine the capabilities, performance, and, ultimately, the success of computer systems.

Modern computer technology requires professionals of every computing specialty to understand both hardware and software. The interaction between hardware and software at a variety of levels also offers a framework for understanding the fundamentals of computing. Whether your primary interest is hardware or software, computer science or electrical engineering, the central ideas in computer organization and design are the same. Thus, our emphasis in this book is to show the relationship between hardware and software and to focus on the concepts that are the basis for current computers.

The audience for this book includes those with little experience in assembly language or logic design who need to understand basic computer organization as well as readers with backgrounds in assembly language and/or logic design who want to learn how to design a computer or understand how a system works and why it performs as it does.

About the Other Book

Some readers may be familiar with *Computer Architecture: A Quantitative Approach*, popularly known as Hennessy and Patterson. (This book in turn is called Patterson and Hennessy.) Our motivation in writing that book was to describe the principles of computer architecture using solid engineering funda-

mentals and quantitative cost/performance trade-offs. We used an approach that combined examples and measurements, based on commercial systems, to create realistic design experiences. Our goal was to demonstrate that computer architecture could be learned using quantitative methodologies instead of a descriptive approach. It is intended for the serious computing professional who wants a detailed understanding of computers.




























































A majority of the readers for this book do not plan to become computer architects. The performance of future software systems will be dramatically affected, however, by how well software designers understand the basic hardware techniques at work in a system. Thus, compiler writers, operating system designers, database programmers, and most other software engineers need a firm grounding in the principles presented in this book. Similarly, hardware designers must understand clearly the effects of their work on software applications.






Thus, we knew that this book had to be much more than a subset of the material in *Computer Architecture*, and the material was extensively revised to match the different audience. We were so happy with the result that the subsequent editions of *Computer Architecture* were revised to remove most of the introductory material; hence, there is much less overlap today than with the first editions of both books.

Changes for the Third Edition

We had six major goals for the third edition of *Computer Organization and Design*: make the book work equally well for readers with a software focus or with a hardware focus; improve pedagogy in general; enhance understanding of program performance; update the technical content to reflect changes in the industry since the publication of the second edition in 1998; tie the ideas from the book more closely to the real world *outside* the computing industry; and reduce the size of this book.

First, the table on the next page shows the hardware and software paths through the material. Chapters 1, 4, and 7 are found on both paths, no matter what the experience or the focus. Chapters 2 and 3 are likely to be review material for the hardware-oriented, but are essential reading for the software-oriented, especially for those readers interested in learning more about compilers and object-oriented programming languages. The first sections of Chapters 5 and 6 give overviews for those with a software focus. Those with a hardware focus, however, will find that these chapters present core material; they may also, depending on background, want to read Appendix B on logic design first and the sections on microprogramming and how to use hardware description languages to specify control. Chapter 8 on input/output is key to readers with a software focus and should be read if time permits by others. The last chapter on multiprocessors and clusters is again a question of time for the reader. Even the history sections show this balanced focus; they include short histories of programming languages, compilers, numerical software, operating systems, networking protocols, and databases.

Chapter or Appendix	Sections	Software Focus	Hardware Focus
1. Computer Abstractions and Technology	1.1 to 1.6		
	■ 1.7 (History)		
2. Instructions: Language of the Computer	2.1 to 2.11		
	■ 2.12 (Compilers)		
	2.13 (C sort)		
	■ 2.14 (Java)		
	2.15 to 2.18		
3. Arithmetic for Computers	3.1 to 3.9		
	■ 3.10 (History)		
D. RISC instruction set architectures	■ D.1 to D.19		
4. Assessing and Understanding Performance	4.1 to 4.6		
	■ 4.7 (History)		
B. The Basics of Logic Design	■ B.1 to B.13		
5. The Processor: Datapath and Control	5.1 (Overview)		
	5.2 to 5.7		
	■ 5.8 (Microcode)		
	■ 5.9 (Verilog)		
	5.10 to 5.12		
C. Mapping Control to Hardware	■ 5.13 (History)		
	■ C.1 to C.6		
6. Enhancing Performance with Pipelining	6.1 (Overview)		
	6.2 to 6.6		
	■ 6.7 (verilog)		
	6.8 to 6.9		
	6.10 to 6.12		
7. Large and Fast: Exploiting Memory Hierarchy	■ 6.13 (History)		
	7.1 to 7.8		
8. Storage, Networks, and Other Peripherals	■ 7.9 (History)		
	8.1 to 8.2		
	■ 8.3 (Networks)		
	8.4 to 8.10		
9. Multiprocessors and Clusters	■ 8.13 (History)		
	■ 9.1 to 9.10		
A. Assemblers, Linkers, and the SPIM Simulator	■ 9.11 (History)		
	■ A.1 to A.12		
Computers in the Real World	Between Chapters		

Read carefully Read if have time Reference Review or read Read for culture 

The next goal was to improve the exposition of the ideas in the book, based on difficulties mentioned by readers of the second edition. We added five new book elements to help. To make the book work better as a reference, we placed definitions of new terms in the margins at their first occurrence. We hope this will help readers find the sections when they want to refer back to material they have already read. Another change was the insertion of the “Check Yourself” sections, which we added to help readers to check their comprehension of the material on the first time through it. A third change is that added extra exercises in the “For More Practice” section. Fourth, we added the answers to the “Check Yourself” sections and to the For More Practice exercises to help readers see for themselves if they understand the material by comparing their answers to the book. The final new book element was inspired by the “Green Card” of the IBM System/360. We believe that you will find that the MIPS Reference Data Card will be a handy reference when writing MIPS assembly language programs. Our idea is that you will remove the card from the front of the book, fold it in half, and keep it in your pocket, just as IBM S/360 programmers did in the 1960s.

Third, computers are so complex today that understanding the performance of a program involves understanding a good deal about the underlying principles and the organization of a given computer. Our goal is that readers of this book should be able to understand the performance of their programs and how to improve it. To aid in that goal, we added a new book element called “Understanding Program Performance” in several chapters. These sections often give concrete examples of how ideas in the chapter affect performance of real programs.

Fourth, in the interval since the second edition of this book, Moore’s law has marched onward so that we now have processors with 200 million transistors, DRAM chips with a billion transistors, and clock rates of multiple gigahertz. The “Real Stuff” examples have been updated to describe such chips. This edition also includes AMD64/IA-32e, the 64-bit address version of the long-lived 80x86 architecture, which appears to be the nemesis of the more recent IA-64. It also reflects the transition from parallel buses to serial networks and switches. Later chapters describe Google, which was born after the second edition, in terms of its cluster technology and in novel uses of search.

Fifth, although many computer science and engineering students enjoy information technology for technology’s sake, some have more altruistic interests. This latter group tends to have more women and underrepresented minorities. Consequently, we have added a new book element, “Computers in the Real World,” two-page layouts found between each chapter. Our perspective is that information technology is more valuable for humanity than most other topics you could study—whether it is preserving our art heritage, helping the Third World, saving our environment, or even changing political systems—and so we demonstrate our view with concrete examples of nontraditional applications. We think readers of these segments will have a greater appreciation of the computing culture beyond

the inherently interesting technology, much like those who read the history sections at the end of each chapter

Finally, books are like people: they usually get larger as they get older. By using technology, we have managed to do all the above and yet shrink the page count by hundreds of pages. As the table illustrates, the core portion of the book for hardware and software readers is on paper, but sections that some readers would value more than others are found on the companion CD. This technology also allows your authors to provide longer histories and more extensive exercises without concerns about lengthening the book. Once we added the CD to the book, we could then include a great deal of free software and tutorials that many instructors have told us they would like to use in their courses. This hybrid paper-CD publication weighs about 30% less than it did six years ago—an impressive goal for books as well as for people.

Instructor Support

We have collected a great deal of material to help instructors teach courses using this book. Solutions to exercises, figures from the book, lecture notes, lecture slides, and other materials are available to adopters from the publisher. Check the publisher's Web site for more information:

www.mkp.com/companions/1558606041

Concluding Remarks

If you read the following acknowledgments section, you will see that we went to great lengths to correct mistakes. Since a book goes through many printings, we have the opportunity to make even more corrections. If you uncover any remaining, resilient bugs, please contact the publisher by electronic mail at cod3bugs@mkp.com or by low-tech mail using the address found on the copyright page. The first person to report a technical error will be awarded a \$1.00 bounty upon its implementation in future printings of the book!

This book is truly collaborative, despite one of us running a major university. Together we brainstormed about the ideas and method of presentation, then individually wrote about one-half of the chapters and acted as reviewer for every draft of the other half. The page count suggests we again wrote almost exactly the same number of pages. Thus, we equally share the blame for what you are about to read.

Acknowledgments for the Third Edition

We'd like to again express our appreciation to **Jim Larus** for his willingness in contributing his expertise on assembly language programming, as well as for welcoming readers of this book to use the simulator he developed and maintains. Our

exercise editor **Dan Sorin** took on the Herculean task of adding new exercises and answers. **Peter Ashenden** worked similarly hard to collect and organize the companion CD.

We are grateful to the many instructors who answered the publisher's surveys, reviewed our proposals, and attended focus groups to analyze and respond to our plans for this edition. They include the following individuals: Michael Anderson (University of Hartford), David Bader (University of New Mexico), Rusty Baldwin (Air Force Institute of Technology), John Barr (Ithaca College), Jack Briner (Charleston Southern University), Mats Brorsson (KTH, Sweden), Colin Brown (Franklin University), Lori Carter (Point Loma Nazarene University), John Casey (Northeastern University), Gene Chase (Messiah College), George Cheney (University of Massachusetts, Lowell), Daniel Citron (Jerusalem College of Technology, Israel), Albert Cohen (INRIA, France), Lloyd Dickman (PathScale), Jose Duato (Universidad Politécnica de Valencia, Spain), Ben Dugan (University of Washington), Derek Eager (University of Saskatchewan, Canada), Magnus Ekman (Chalmers University of Technology, Sweden), Ata Elahi (Southern Connecticut State University), Soundararajan Ezekiel (Indiana University of Pennsylvania), Ernest Ferguson (Northwest Missouri State University), Michael Fry (Lebanon Valley College, Pennsylvania), R. Gaede (University of Arkansas at Little Rock), Jean-Luc Gaudiot (University of California, Irvine), Thomas Gendreau (University of Wisconsin, La Crosse), George Georgiou (California State University, San Bernardino), Paul Gillard (Memorial University of Newfoundland, Canada), Joe Grimes (California Polytechnic State University, SLO), Max Hailperin (Gustavus Adolphus College), Jayantha Herath (St. Cloud State University, Minnesota), Mark Hill (University of Wisconsin, Madison), Michael Hsaio (Virginia Tech), Richard Hughey (University of California, Santa Cruz), Tony Jebara (Columbia University), Elizabeth Johnson (Xavier University), Peter Kogge (University of Notre Dame), Morris Lancaster (BAH), Doug Lawrence (University of Montana), David Lilja (University of Minnesota), Nam Ling (Santa Clara University, California), Paul Lum (Agilent Technologies), Stephen Mann (University of Waterloo, Canada), Diana Marculescu (Carnegie Mellon University), Margaret McMahon (U.S. Naval Academy Computer Science), Uwe Meyer-Baese (Florida State University), Chris Milner (University of Virginia), Tom Pittman (Southwest Baptist University), Jalel Rejeb (San Jose State University, California), Bill Siever (University of Missouri, Rolla), Kevin Skadron (University of Virginia), Pam Smallwood (Regis University, Colorado), K. Stuart Smith (Rocky Mountain College), William J. Taffe (Plymouth State University), Michael E. Thomodakis (Texas A&M University), Rупpa K. Thulasiram (University of Manitoba, Canada), Ye Tung (University of South Alabama), Steve VanderLeest (Calvin College), Neal R. Wagner (University of Texas at San Antonio), and Kent Wilken (University of California, Davis).

We are grateful too to those who carefully read our draft manuscripts; some read successive drafts to help ensure new errors didn't creep in as we revised. They include Krste Asanovic (Massachusetts Institute of Technology), Jean-Loup Baer (University of Washington), David Brooks (Harvard University), Doug Clark (Princeton University), Dan Connors (University of Colorado at Boulder), Matt Farrens (University of California, Davis), Manoj Franklin (University of Maryland College Park), John Greiner (Rice University), David Harris (Harvey Mudd College), Paul Hilfinger (University of California, Berkeley), Norm Jouppi (Hewlett-Packard), David Kaeli (Northeastern University), David Oppenheimer (University of California, Berkeley), Timothy Pinkston (University of Southern California), Mark Smotherman (Clemson University), and David Wood (University of Wisconsin, Madison).

To help us meet our goal of creating 70% new exercises and solutions for this edition, we recruited several graduate students recommended to us by their professors. We are grateful for their creativity and persistence: Michael Black (University of Maryland), Lei Chen (University of Rochester), Nirav Dave (Massachusetts Institute of Technology), Wael El Essawy (University of Rochester), Nikil Mehta (Brown University), Nicholas Nelson (University of Rochester), Aaron Smith (University of Texas, Austin), and Charlie Wang (Duke University).

We would like to especially thank **Mark Smotherman** for making a careful final pass to find technical and writing glitches that significantly improved the quality of this edition.

We wish to thank the extended Morgan Kaufmann family for agreeing to publish this book again under the able leadership of **Denise Penrose**. She developed the vision of the hybrid paper-CD book and recruited the many people above who played important roles in developing the book.

Simon Crump managed the book production process, and **Summer Block** coordinated the surveying of users and their responses. We thank also the many freelance vendors who contributed to this volume, especially **Nancy Logan** and **Dartmouth Publishing, Inc.**, our compositors.

The contributions of the nearly 100 people we mentioned here have made this third edition our best book yet. Enjoy!

David A. Patterson

John L. Hennessy

Contents

Preface vii

CHAPTERS

1 Computer Abstractions and Technology 2

- 1.1 Introduction 3
- 1.2 Below Your Program 11
- 1.3 Under the Covers 15
- 1.4 Real Stuff: Manufacturing Pentium 4 Chips 28
- 1.5 Fallacies and Pitfalls 33
- 1.6 Concluding Remarks 35
- 1.7 Historical Perspective and Further Reading 36
- 1.8 Exercises 36

COMPUTERS IN THE REAL WORLD Information Technology for the 4 Billion without IT 44

2 Instructions: Language of the Computer 46

- 2.1 Introduction 48
- 2.2 Operations of the Computer Hardware 49
- 2.3 Operands of the Computer Hardware 52
- 2.4 Representing Instructions in the Computer 60
- 2.5 Logical Operations 68
- 2.6 Instructions for Making Decisions 72
- 2.7 Supporting Procedures in Computer Hardware 79
- 2.8 Communicating with People 90
- 2.9 MIPS Addressing for 32-Bit immediates and addresses 95
- 2.10 Translating and Starting a Program 106
- 2.11 How Compilers Optimize 116
- 2.12 How Compilers Work: An Introduction 121

- 2.13 A C Sort Example to Put It All Together 121
- 2.14 Implementing an Object-Oriented Language 130
- 2.15 Arrays versus Pointers 130
- 2.16 Real Stuff: IA-32 Instructions 134
- 2.17 Fallacies and Pitfalls 143
- 2.18 Concluding Remarks 145
- 2.19 Historical Perspective and Further Reading 147
- 2.20 Exercises 147

COMPUTERS IN THE REAL WORLD
Helping Save Our Environment with Data 156

3 Arithmetic for Computers 158

- 3.1 Introduction 160
- 3.2 Signed and Unsigned Numbers 160
- 3.3 Addition and Subtraction 170
- 3.4 Multiplication 176
- 3.5 Division 183
- 3.6 Floating Point 189
- 3.7 Real Stuff: Floating Point in the IA-32 217
- 3.8 Fallacies and Pitfalls 220
- 3.9 Concluding Remarks 225
- 3.10 Historical Perspective and Further Reading 229
- 3.11 Exercises 229

COMPUTERS IN THE REAL WORLD
Reconstructing the Ancient World 236

4 Assessing and Understanding Performance 238

- 4.1 Introduction 240
- 4.2 CPU Performance and Its Factors 246
- 4.3 Evaluating Performance 254
- 4.4 Real Stuff: Two SPEC Benchmarks and the Performance of Recent Intel Processors 259
- 4.5 Fallacies and Pitfalls 266
- 4.6 Concluding Remarks 270
- 4.7 Historical Perspective and Further Reading 272
- 4.8 Exercises 272

COMPUTERS IN THE REAL WORLD
Moving People Faster and More Safely 280