

# Other Instruction Set Architectures

ARM

x86

MIPS instructions	Name	Format	Pseudo MIPS	Name	Format
add	add	R	move	move	R
subtract	sub	R	multiply	mult	R
add immediate	addi	I	multiply immediate	mult i	I
load word	lw	I	load immediate	li	I
store word	sw	I	branch less than	blt	I
load half	lh	I	branch less than or equal	ble	I
load half unsigned	lhu	I	branch greater than	bgt	I
store half	sh	I	branch greater than or equal	bge	I
load byte	lb	I			
load byte unsigned	lbu	I			
store byte	sb	I			
load linked	ll	I			
store conditional	sc	I			
load upper immediate	lui	I			
and	and	R			
or	or	R			
nor	nor	R			
and immediate	andi	I			
or immediate	ori	I			
shift left logical	sll	R			
shift right logical	srl	R			
branch on equal	beq	I			
branch on not equal	bne	I			
set less than	slt	R			
set less than immediate	slti	I			
set less than immediate unsigned	sltiu	I			
jump	j	J			
jump register	jr	R			
jump and link	jal	J			

**FIGURE 2.44** The MIPS instruction set covered so far, with the real MIPS instructions on the left and the pseudoinstructions on the right. Appendix B (Section B.10) describes the full MIPS architecture. Figure 2.1 shows more details of the MIPS architecture revealed in this chapter. The information given here is also found in Columns 1 and 2 of the MIPS Reference Data Card at the front of the book. Copyright © 2009 Elsevier, Inc. All rights reserved.

Instruction class	MIPS examples	HLL correspondence	Frequency	
			Integer	Ft. pt.
Arithmetic	add, sub, addi	Operations in assignment statements	16%	48%
Data transfer	lw, sw, lb, lbu, lh, lhu, sb, lui	References to data structures, such as arrays	35%	36%
Logical	and, or, nor, andi, ori, sll, srl	Operations in assignment statements	12%	4%
Conditional branch	beq, bne, slt, slti, sltiu	<i>If</i> statements and loops	34%	8%
Jump	j, jr, jal	Procedure calls, returns, and <i>case/switch</i> statements	2%	0%

**FIGURE 2.45 MIPS instruction classes, examples, correspondence to high-level program language constructs, and percent age of MIPS instructions executed by category for the average SPEC2006 benchmarks.** Figure 3.26 in Chapter 3 shows average percent age of the individual MIPS instructions executed. Copyright © 2009 Elsevier, Inc. All rights reserved.

	ARM	MIPS
Date announced	1985	1985
Instruction size (bits)	32	32
Address space (size, model)	32 bits, flat	32 bits, flat
Data alignment	Aligned	Aligned
Data addressing modes	9	3
Integer registers (number, model, size)	15 GPR × 32 bits	31 GPR × 32 bits
I/O	Memory mapped	Memory mapped

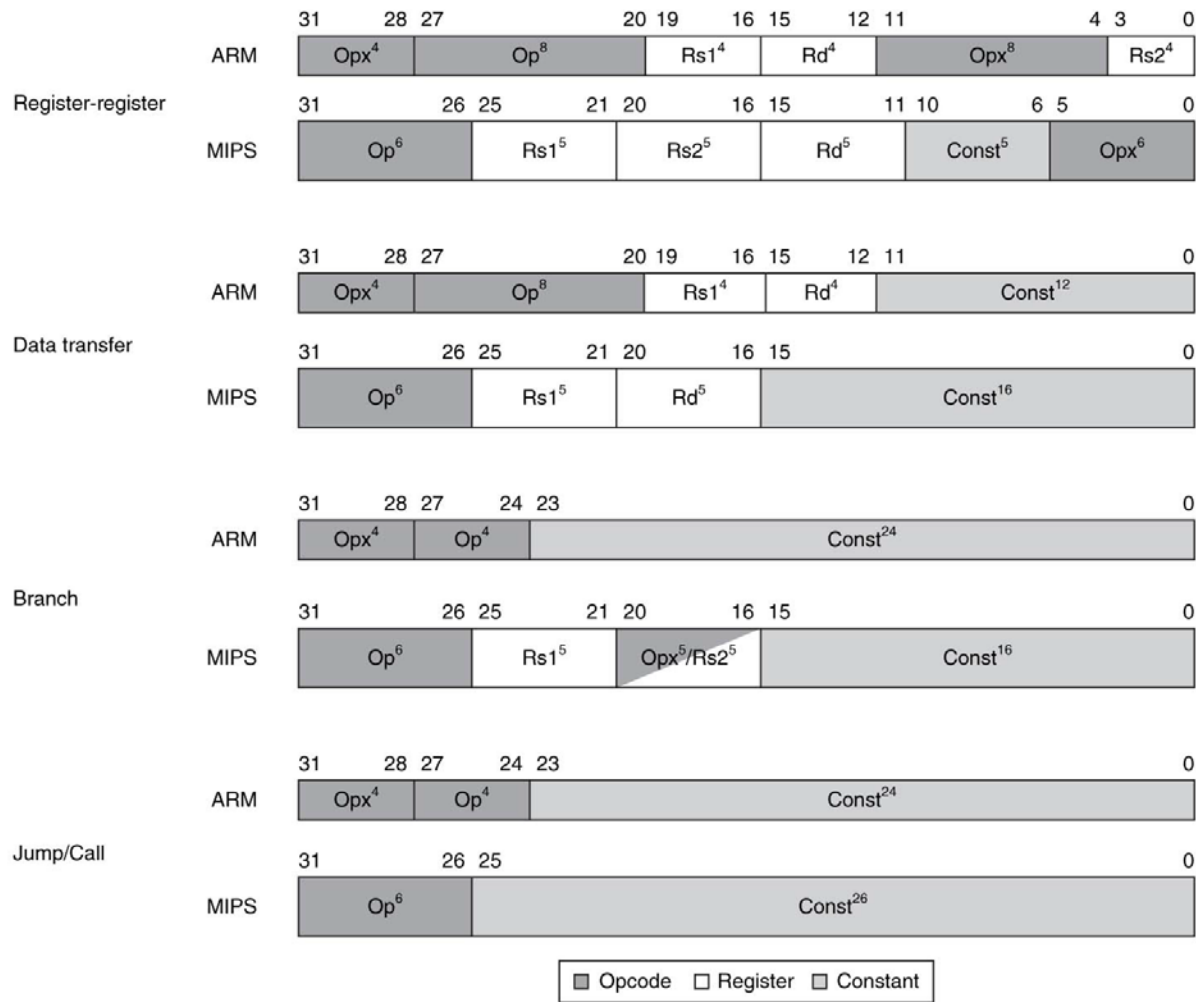
**FIGURE 2.31 Similarities in ARM and MIPS instruction sets.** Copyright © 2009 Elsevier, Inc. All rights reserved.

	Instruction name	ARM	MIPS
Register-register	Add	add	addu, addiu
	Add (trap if overflow)	adds; swivs	add
	Subtract	sub	subu
	Subtract (trap if overflow)	subs; swivs	sub
	Multiply	mul	mult, multu
	Divide	—	div, divu
	And	and	and
	Or	orr	or
	Xor	eor	xor
	Load high part register	—	lui
	Shift left logical	lsl <sup>1</sup>	sllv, sll
	Shift right logical	lsr <sup>1</sup>	srlv, srl
	Shift right arithmetic	asr <sup>1</sup>	srav, sra
	Compare	cmp, cmn, tst, teq	slt/i, slt/iu
Data transfer	Load byte signed	ldrsb	lb
	Load byte unsigned	ldrb	lbu
	Load halfword signed	ldrsh	lh
	Load halfword unsigned	ldrh	lhu
	Load word	ldr	lw
	Store byte	strb	sb
	Store halfword	strh	sh
	Store word	str	sw
	Read, write special registers	mrs, msr	move
	Atomic Exchange	swp, swpb	ll;sc

**FIGURE 2.32 ARM register-register and data transfer instructions equivalent to MIPS core.** Dashes mean the operation is not available in that architecture or not synthesized in a few instructions. If there are several choices of instructions equivalent to the MIPS core, they are separated by commas. ARM includes shifts as part of every data operation instruction, so the shifts with superscript 1 are just a variation of a move instruction, such as lsr<sup>1</sup>. Note that ARM has no divide instruction. Copyright © 2009 Elsevier, Inc. All rights reserved.

Addressing mode	ARM v.4	MIPS
Register operand	X	X
Immediate operand	X	X
Register + offset (displacement or based)	X	X
Register + register (indexed)	X	—
Register + scaled register (scaled)	X	—
Register + offset and update register	X	—
Register + register and update register	X	—
Autoincrement, autodecrement	X	—
PC-relative data	X	—

**FIGURE 2.33 Summary of data addressing modes.** ARM has separate register indirect and register + offset addressing modes, rather than just putting 0 in the offset of the latter mode. To get greater addressing range, ARM shifts the offset left 1 or 2 bits if the data size is halfword or word. Copyright © 2009 Elsevier, Inc. All rights reserved.



**FIGURE 2.34 Instruction formats, ARM, and MIPS.** The differences result from whether the architecture has 16 or 32 registers. Copyright © 2009 Elsevier, Inc. All rights reserved.

Name	Definition	ARM v.4	MIPS
Load immediate	$Rd = Imm$	mov	addi, \$0,
Not	$Rd = \sim(Rs1)$	mvn	nor, \$0,
Move	$Rd = Rs1$	mov	or, \$0,
Rotate right	$Rd = Rs\ i \gg i$ $Rd_{0 \dots i-1} = Rs_{31-i \dots 31}$	ror	
And not	$Rd = Rs1 \& \sim(Rs2)$	bic	
Reverse subtract	$Rd = Rs2 - Rs1$	rsb, rsc	
Support for multiword integer add	CarryOut, $Rd = Rd + Rs1 + OldCarryOut$	adcs	—
Support for multiword integer sub	CarryOut, $Rd = Rd - Rs1 + OldCarryOut$	sbc	—

**FIGURE 2.35 ARM arithmetic/logical instructions not found in MIPS.** Copyright © 2009 Elsevier, Inc. All rights reserved.



# Alternative Architectures


- Design alternative:
  - provide more powerful operations
  - goal is to reduce number of instructions executed
  - danger is a slower cycle time and/or a higher CPI

*–“The path toward operation complexity is thus fraught with peril. To avoid these problems, designers have moved toward simpler instructions”*
  
- Let’s look (briefly) at IA-32

# IA-32 (a.k.a. x86)

- 1978: The Intel 8086 is announced (16 bit architecture)
- 1980: The 8087 floating point coprocessor is added
- 1982: The 80286 increases address space to 24 bits, +instructions
- 1985: The 80386 extends to 32 bits, new addressing modes
- 1989-1995: The 80486, Pentium, Pentium Pro add a few instructions (mostly designed for higher performance)
- 1997: 57 new “MMX” instructions are added, Pentium II
- 1999: The Pentium III added another 70 instructions (SSE – streaming SIMD extensions)
- 2001: Another 144 instructions (SSE2)
- 2003: AMD extends the architecture to increase address space to 64 bits, widens all registers to 64 bits and makes other changes (AMD64)
- 2004: Intel capitulates and embraces AMD64 (calls it EM64T) and adds more media extensions
- 2006: Intel adds 54 instructions (SSE4) and supports virtual machines
- 2007: AMD adds 170 instructions (SSE5)
- 2008: Intel adds Advanced Vector Extension to expand SSE register width from 128 to 256 bits and adds 128 new instructions.

*“This history illustrates the impact of the “golden handcuffs” of compatibility: “adding new features as someone might add clothing to a packed bag” and “an architecture that is difficult to explain and impossible to love”*

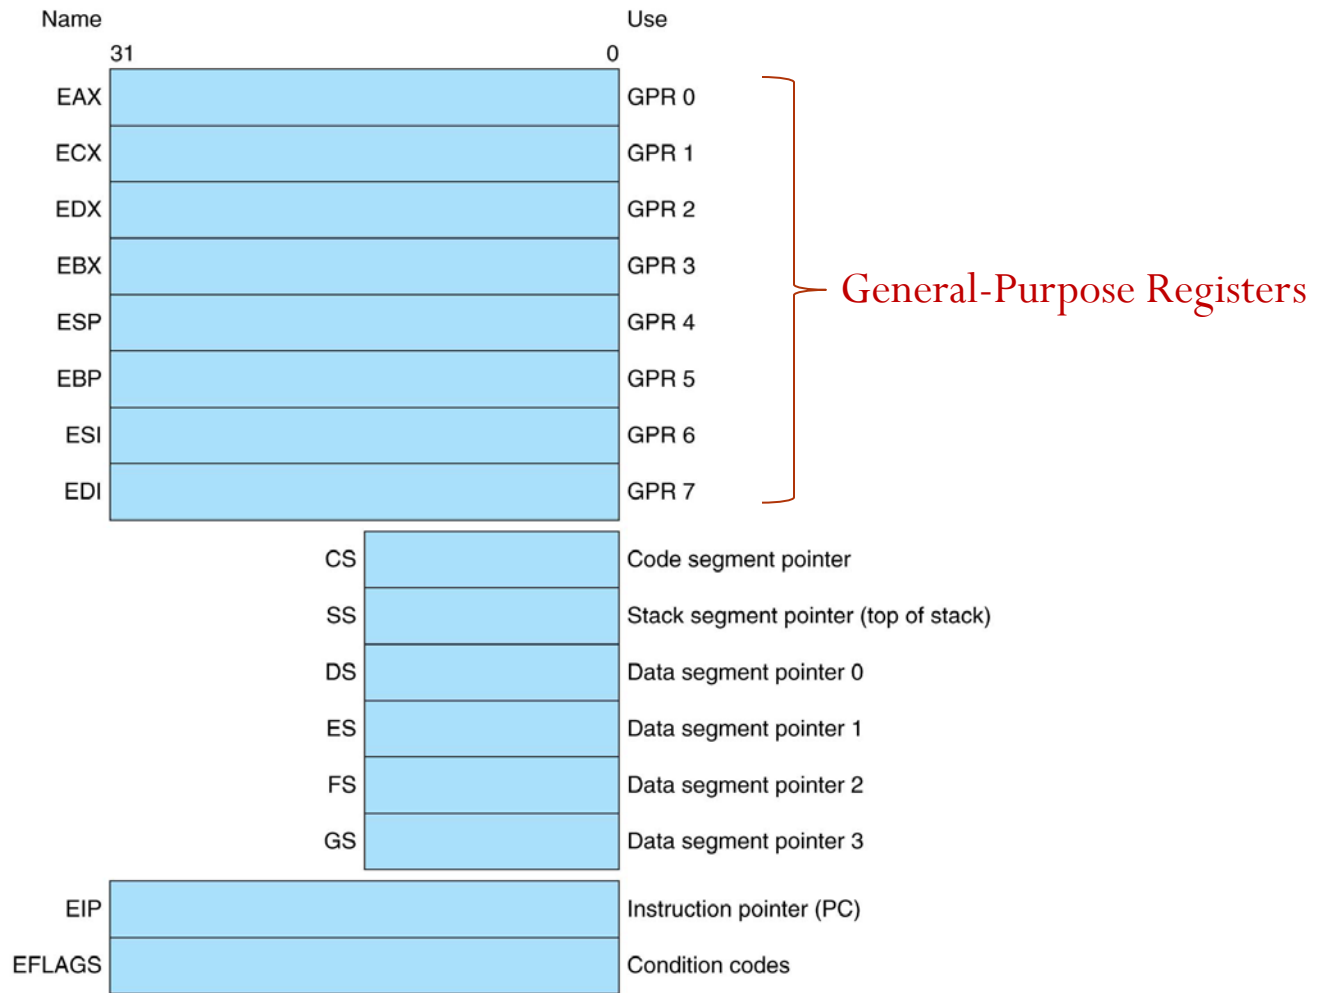


**FIGURE 2.43 Growth of x86 instruction set over time.** While there is clear technical value to some of these extensions, this rapid change also increases the difficulty for other companies to try to build compatible processors. Copyright © 2009 Elsevier, Inc. All rights reserved.

# IA-32 Overview

- Complexity:
  - Instructions from 1 to 17 bytes long
  - one operand must act as both a source and destination
  - one operand can come from memory
  - complex addressing modes
    - e.g., “base or scaled index with 8 or 32 bit displacement”
- Saving grace:
  - the most frequently used instructions are not too difficult to build
  - compilers avoid the portions of the architecture that are slow

*“what the x86 lacks in style is made up in quantity,  
making it beautiful from the right perspective”*



**FIGURE 2.36 The 80386 register set.** Starting with the 80386, the top eight registers were extended to 32 bits and could also be used as general-purpose registers. Copyright © 2009 Elsevier, Inc. All rights reserved.

Instruction	Meaning
<b>Control</b>	<b>Conditional and unconditional branches</b>
jnz, jz	Jump if condition to EIP + 8-bit offset; JNE (for JNZ), JE (for JZ) are alternative names
jmp	Unconditional jump—8-bit or 16-bit offset
call	Subroutine call—16-bit offset; return address pushed onto stack
ret	Pops return address from stack and jumps to it
loop	Loop branch—decrement ECX; jump to EIP + 8-bit displacement if ECX ≠ 0
<b>Data transfer</b>	<b>Move data between registers or between register and memory</b>
move	Move between two registers or between register and memory
push, pop	Push source operand on stack; pop operand from stack top to a register
les	Load ES and one of the GPRs from memory
<b>Arithmetic, logical</b>	<b>Arithmetic and logical operations using the data registers and memory</b>
add, sub	Add source to destination; subtract source from destination; register-memory format
cmp	Compare source and destination; register-memory format
shl, shr, rcr	Shift left; shift logical right; rotate right with carry condition code as fill
cbw	Convert byte in eight rightmost bits of EAX to 16-bit word in right of EAX
test	Logical AND of source and destination sets condition codes
inc, dec	Increment destination, decrement destination
or, xor	Logical OR; exclusive OR; register-memory format
<b>String</b>	<b>Move between string operands; length given by a repeat prefix</b>
movs	Copies from string source to destination by incrementing ESI and EDI; may be repeated
lods	Loads a byte, word, or doubleword of a string into the EAX register

**FIGURE 2.40 Some typical operations on the x86.** Many operations use register-memory for mat, where either the source or the destination may be memory and the other may be a register or immediate operand. Copyright © 2009 Elsevier, Inc. All rights reserved.

Source/destination operand type	Second source operand
Register	Register
Register	Immediate
Register	Memory
Memory	Register
Memory	Immediate

**FIGURE 2.37 Instruction types for the arithmetic, logical, and data transfer instructions.** The x86 allows the combinations shown. The only restriction is the absence of a memory- memory mode. Immediates may be 8, 16, or 32 bits in length; a register is any one of the 14 major registers in Figure 2.36 (not EIP or EFLAGS). Copyright © 2009 Elsevier, Inc. All rights reserved.

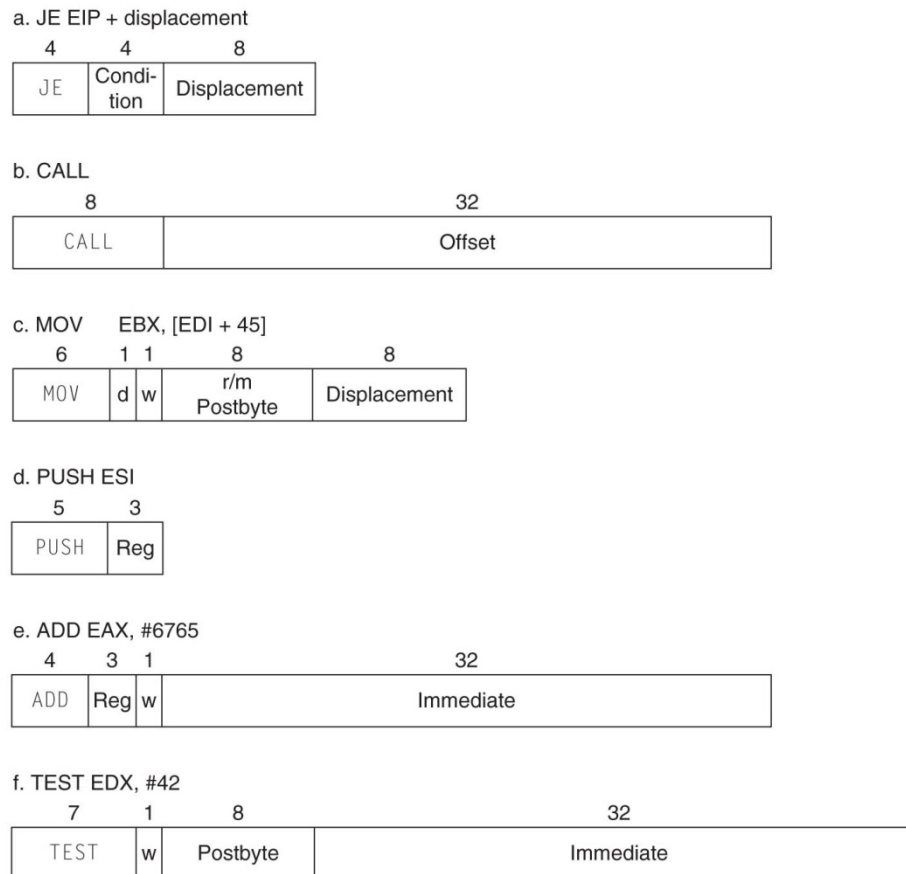
Mode	Description	Register restrictions	MIPS equivalent
Register indirect	Address is in a register.	Not ESP or EBP	lw \$s0,0(\$s1)
Based mode with 8- or 32-bit displacement	Address is contents of base register plus displacement.	Not ESP	lw \$s0,100(\$s1) # <= 16-bit # displacement
Base plus scaled index	The address is Base + (2 <sup>Scale</sup> x Index) where Scale has the value 0, 1, 2, or 3.	Base: any GPR Index: not ESP	mul \$t0,\$s2,4 add \$t0,\$t0,\$s1 lw \$s0,0(\$t0)
Base plus scaled index with 8- or 32-bit displacement	The address is Base + (2 <sup>Scale</sup> x Index) + displacement where Scale has the value 0, 1, 2, or 3.	Base: any GPR Index: not ESP	mul \$t0,\$s2,4 add \$t0,\$t0,\$s1 lw \$s0,100(\$t0) # 16-bit # displacement

**FIGURE 2.38 x86 32-bit addressing modes with register restrictions and the equivalent MIPS code.** The Base plus Scaled Index addressing mode, not found in ARM or MIPS, is included to avoid the multiplies by 4 (scale factor of 2) to turn an index in a register into a byte address (see Figures 2.25 and 2.27). A scale factor of 1 is used for 16-bit data, and a scale factor of 3 for 64-bit data. A scale factor of 0 means the address is not scaled. If the displacement is longer than 16 bits in the second or fourth modes, then the MIPS equivalent mode would need two more instructions: a lui to load the upper 16 bits of the displacement and an add to sum the upper address with the base register \$s1. (Intel gives two different names to what is called Based addressing mode—Based and Indexed—but they are essentially identical and we combine them here.) Copyright © 2009 Elsevier, Inc. All rights reserved.



Instruction	Function
je name	if equal(condition code) {EIP=name}; EIP-128 <= name < EIP+128
jmp name	EIP=name
call name	SP=SP-4; M[SP]=EIP+5; EIP=name;
movw EBX,[EDI+45]	EBX=M[EDI+45]
push ESI	SP=SP-4; M[SP]=ESI
pop EDI	EDI=M[SP]; SP=SP+4
add EAX,#6765	EAX= EAX+6765
test EDX,#42	Set condition code (flags) with EDX and 42
movsl	M[EDI]=M[ESI]; EDI=EDI+4; ESI=ESI+4

**FIGURE 2.39 Some typical x86 instructions and their functions.** A list of frequent operations appears in Figure 2.40. The CALL saves the EIP of the next instruction on the stack. (EIP is the Intel PC.) Copyright © 2009 Elsevier, Inc. All rights reserved.



**FIGURE 2.41 Typical x86 instruction formats.** Figure 2.42 shows the encoding of the postbyte. Many instructions contain the 1-bit field *w*, which says whether the operation is a byte or a double word. The *d* field in MOV is used in instructions that may move to or from memory and shows the direction of the move. The ADD instruction requires 32 bits for the immediate field, because in 32-bit mode, the immediates are either 8 bits or 32 bits. The immediate field in the TEST is 32 bits long because there is no 8-bit immediate for test in 32-bit mode. Overall, instructions may vary from 1 to 17 bytes in length. The long length comes from extra 1-byte prefixes, having both a 4-byte immediate and a 4-byte displacement address, using an opcode of 2 bytes, and using the scaled index mode specifier, which adds another byte. Copyright © 2009 Elsevier, Inc. All rights reserved.

reg	w = 0		w = 1		r/m	mod = 0		mod = 1		mod = 2		mod = 3
	16b	32b	16b	32b		16b	32b	16b	32b			
0	AL	AX	EAX	0	addr=BX+SI	=EAX	same	same	same	same	same	same
1	CL	CX	ECX	1	addr=BX+DI	=ECX	addr as	addr as	addr as	addr as	addr as	as
2	DL	DX	EDX	2	addr=BP+SI	=EDX	mod=0	mod=0	mod=0	mod=0	mod=0	reg
3	BL	BX	EBX	3	addr=BP+SI	=EBX	+ disp8	+ disp8	+ disp16	+ disp32	+ disp32	field
4	AH	SP	ESP	4	addr=SI	=(sib)	SI+disp8	(sib)+disp8	SI+disp8	(sib)+disp32	(sib)+disp32	“
5	CH	BP	EBP	5	addr=DI	=disp32	DI+disp8	EBP+disp8	DI+disp16	EBP+disp32	EBP+disp32	“
6	DH	SI	ESI	6	addr=disp16	=ESI	BP+disp8	ESI+disp8	BP+disp16	ESI+disp32	ESI+disp32	“
7	BH	DI	EDI	7	addr=BX	=EDI	BX+disp8	EDI+disp8	BX+disp16	EDI+disp32	EDI+disp32	“

**FIGURE 2.42 The encoding of the first address specifier of the x86: mod, reg, r/m.** The first four columns show the encoding of the 3-bit reg field, which depends on the w bit from the opcode and whether the machine is in 16-bit mode (8086) or 32-bit mode (80386). The remaining columns explain the mod and r/m fields. The meaning of the 3-bit r/m field depends on the value in the 2-bit mod field and the address size. Basically, the registers used in the address calculation are listed in the sixth and seventh columns, under mod = 0, with mod = 1 adding an 8-bit displacement and mod = 2 adding a 16-bit or 32-bit displacement, depending on the address mode. The exceptions are 1) r/m = 6 when mod = 1 or mod = 2 in 16-bit mode selects BP plus the displacement; 2) r/m = 5 when mod = 1 or mod = 2 in 32-bit mode selects EBP plus displacement; and 3) r/m = 4 in 32-bit mode when mod does not equal 3, where (sib) means use the scaled index mode shown in Figure 2.38. When mod = 3, the r/m field indicates a register, using the same encoding as the reg field combined with the w bit. Copyright © 2009 Elsevier, Inc. All rights reserved.

# Some IA-32 Instructions

- PUSH            5-bit opcode, 3-bit register operand


5-b | 3-b

- JE                4-bit opcode, 4-bit condition, 8-bit jump offset

4-b | 4-b | 8-b

# Some IA-32 Instructions

- **MOV** 6-bit opcode, 8-bit register/mode\*, 8-bit offset



bit indicates byte or double word operation

bit indicates move to or from memory

- **XOR** 8-bit opcode, 8-bit reg/mode\*, 8-bit base, 8-b index



\*8-bit register/mode: See Figure 2.42, page 174.

# Some IA-32 Instructions

- **ADD** 4-bit opcode, 3-bit register, 32-bit immediate

4-b | 3-b |w| 32-b

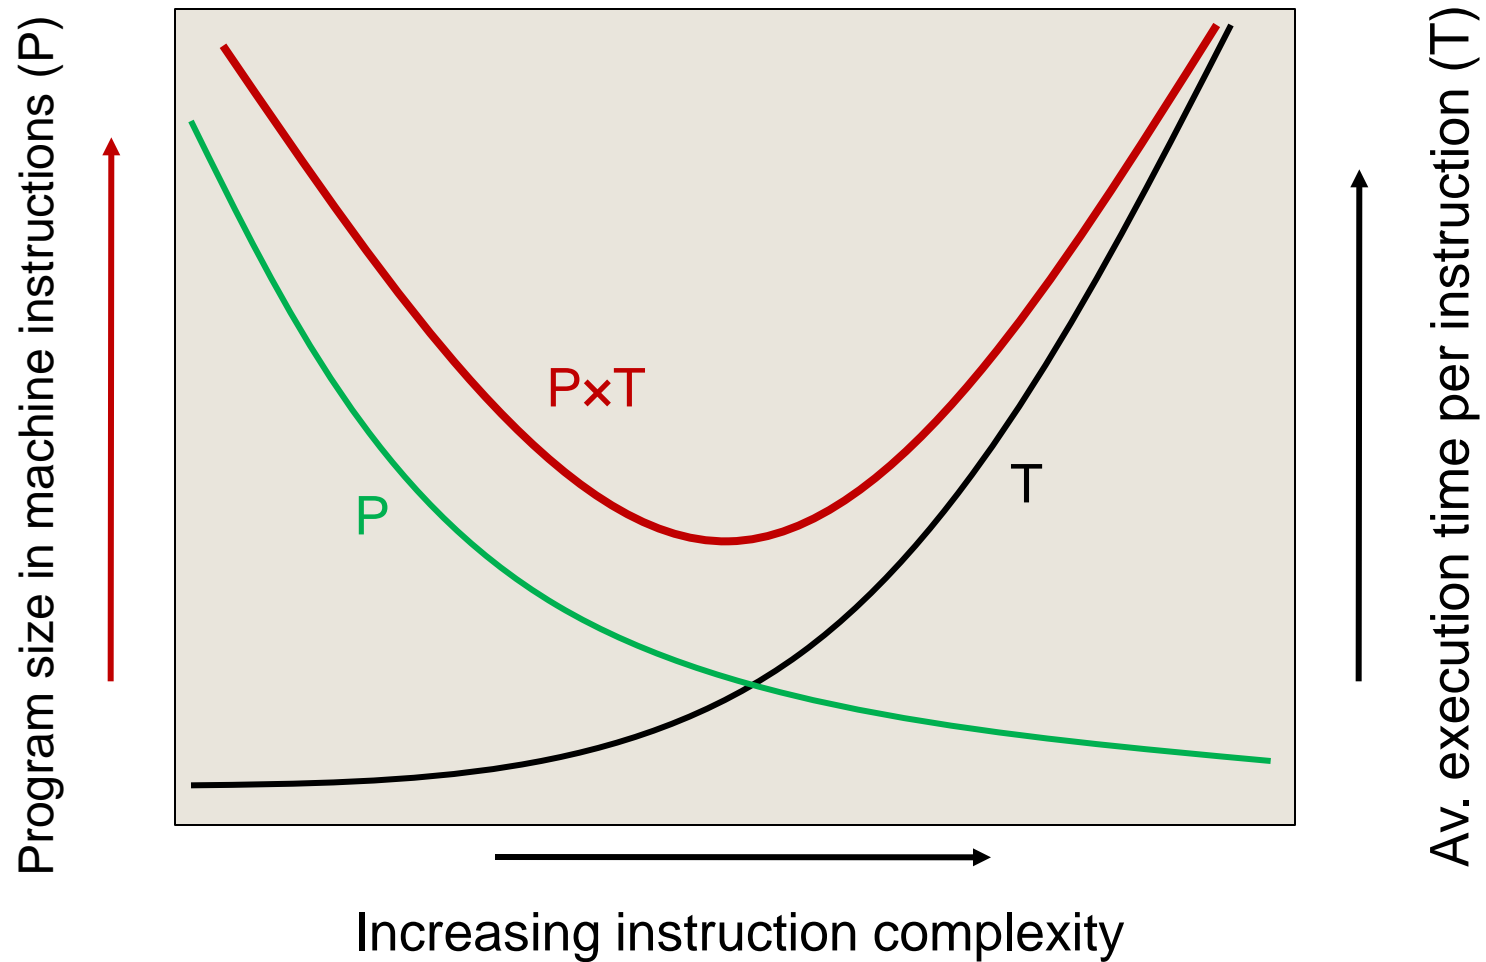
- **TEST** 7-bit opcode, 8-bit reg/mode, 32-bit immediate

7-b |w| 8-b | 32-b

# Additional References

- IA-32, IA-64 (CISC)
  - A. S. Tanenbaum, *Structured Computer Organization, Fifth Edition*, Upper Saddle River, New Jersey: Pearson Prentice-Hall, 2006, Chapter 5.
- ARM (RISC)
  - D. Seal, *ARM Architecture Reference Manual, Second Edition*, Addison-Wesley Professional, 2000.
- SPARC (Scalable Processor Architecture)
- PowerPC
  - V. C. Hamacher, Z. G. Vranesic and S. G. Zaky, *Computer Organization, Fourth Edition*, New York: McGraw-Hill, 1996.

# Instruction Complexity





# Summary

- Instruction complexity is only one variable
  - lower instruction count vs. higher CPI / lower clock rate – *we will see performance measures later*
- Design Principles:
  - simplicity favors regularity
  - smaller is faster
  - good design demands compromise
  - make the common case fast
- Instruction set architecture
  - a very important abstraction indeed!