

---

# Lecture 03: ISA Introduction

**CSE 564 Computer Architecture Summer 2017**

Department of Computer Science and Engineering

Yonghong Yan

[yan@oakland.edu](mailto:yan@oakland.edu)

[www.secs.oakland.edu/~yan](http://www.secs.oakland.edu/~yan)

---

# Contents

---

1. Introduction
2. Classifying Instruction Set Architectures
3. Memory Addressing
4. Type and Size of Operands
5. Operations in the Instruction Set
6. Instructions for Control Flow
7. Encoding an Instruction Set
8. Crosscutting Issues: The Role of Compilers
9. RISC-V ISA

- **Supplement**
  - MIPS ISA
  - RISC vs CISC
  - Compiler compilation stages
  - ISA Historical
    - Appendix L
  - Comparison of ISA
    - Appendix K

# The MIPS Instruction Set

---

- Used as the example as introduction
- Stanford MIPS commercialized by MIPS Technologies ([www.mips.com](http://www.mips.com))
- Large share of embedded core market
  - Applications in consumer electronics, network/storage equipment, cameras, printers, ...
  - Closed one to RISC-V
- Typical of many modern ISAs
  - See MIPS Reference Data tear-out card, and Appendixes B and E in the reference textbook

# Arithmetic Operations

---

- Add and subtract, three operands
  - Two sources and one destinationadd a, b, c # a gets b + c
- All arithmetic operations have this form
- *Design Principle 1*: Simplicity favours regularity
  - Regularity makes implementation simpler
  - Simplicity enables higher performance at lower cost

# Arithmetic Example

---

- C code:

```
f = (g + h) - (i + j);
```

- Compiled MIPS code:

```
add t0, g, h    # temp t0 = g + h
add t1, i, j    # temp t1 = i + j
sub f, t0, t1   # f = t0 - t1
```

# Register Operands

---

- Arithmetic instructions use register operands
- MIPS has a  $32 \times 32$ -bit register file
  - Use for frequently accessed data
  - Numbered 0 to 31
  - 32-bit data called a “word”
- Assembler names
  - $\$t0, \$t1, \dots, \$t9$  for temporary values
  - $\$s0, \$s1, \dots, \$s7$  for saved variables
- *Design Principle 2: Smaller is faster*
  - c.f. main memory: millions of locations

# Register Operand Example

---

- C code:

```
f = (g + h) - (i + j);
```

– f, ..., j in \$s0, ..., \$s4

- Compiled MIPS code:

```
add $t0, $s1, $s2
```

```
add $t1, $s3, $s4
```

```
sub $s0, $t0, $t1
```

# Memory Operands

---

- Main memory used for composite data
  - Arrays, structures, dynamic data
- To apply arithmetic operations
  - Load values from memory into registers
  - Store result from register to memory
- Memory is byte addressed
  - Each address identifies an 8-bit byte
- Words are aligned in memory
  - Address must be a multiple of 4
  
- MIPS is Big Endian
  - Most-significant byte at least address of a word
  - *c.f.* Little Endian: least-significant byte at least address

# Memory Operand Example 1

---

- C code:

```
g = h + A[8];
```

- g in \$s1, h in \$s2, **base address** of A in \$s3

- Compiled MIPS code:

- **Index 8** requires offset of 32

- **4 bytes per word**

```
lw    $t0, 32($s3)    # load word
add   $s1, $s2, $t0
```

offset

base register

# Memory Operand Example 2

---

- C code:

```
A[12] = h + A[8];
```

- $h$  in  $\$s2$ , base address of  $A$  in  $\$s3$

- Compiled MIPS code:

- Index 8 requires offset of 32

```
lw    $t0, 32($s3)    # load word
add   $t0, $s2, $t0
sw    $t0, 48($s3)    # store word
```

# Registers vs. Memory

---

- Registers are faster to access than memory
- Operating on memory data requires loads and stores
  - More instructions to be executed
- Compiler must use registers for variables as much as possible
  - Only **spill** to memory for less frequently used variables
  - Register optimization is important!

# Immediate Operands

---

- Constant data specified in an instruction  
addi \$s3, \$s3, 4
- No subtract immediate instruction
  - Just use a negative constant  
addi \$s2, \$s1, -1
- *Design Principle 3: Make the common case fast*
  - Small constants are common
  - Immediate operand avoids a load instruction

# The Constant Zero

---

- MIPS/RISC-V register 0 (\$zero) is the constant 0
  - Cannot be overwritten
- Useful for common operations
  - E.g., move between registers  
add \$t2, \$s1, \$zero

# Representing Instructions

---

- Instructions are encoded in binary
  - Called machine code
- MIPS instructions
  - Encoded as 32-bit instruction words
  - Small number of formats encoding operation code (opcode), register numbers, ...
  - Regularity!
- Register numbers
  - \$t0 – \$t7 are reg' s 8 – 15
  - \$t8 – \$t9 are reg' s 24 – 25
  - \$s0 – \$s7 are reg' s 16 – 23

# MIPS R-format Instructions

---



- Instruction fields
  - op: operation code (opcode)
  - rs: first source register number
  - rt: second source register number
  - rd: destination register number
  - shamt: shift amount (00000 for now)
  - funct: function code (extends opcode)

# R-format Example

---

|        |        |        |        |        |        |
|--------|--------|--------|--------|--------|--------|
| op     | rs     | rt     | rd     | shamt  | funct  |
| 6 bits | 5 bits | 5 bits | 5 bits | 5 bits | 6 bits |

add \$t0, \$s1, \$s2

|         |       |       |       |       |        |
|---------|-------|-------|-------|-------|--------|
| special | \$s1  | \$s2  | \$t0  | 0     | add    |
| 0       | 17    | 18    | 8     | 0     | 32     |
| 000000  | 10001 | 10010 | 01000 | 00000 | 100000 |

$$00000010001100100100000000100000_2 = 02324020_{16}$$

# Hexadecimal

---

- Base 16
  - Compact representation of bit strings
  - 4 bits per hex digit

|   |      |   |      |   |      |   |      |
|---|------|---|------|---|------|---|------|
| 0 | 0000 | 4 | 0100 | 8 | 1000 | c | 1100 |
| 1 | 0001 | 5 | 0101 | 9 | 1001 | d | 1101 |
| 2 | 0010 | 6 | 0110 | a | 1010 | e | 1110 |
| 3 | 0011 | 7 | 0111 | b | 1011 | f | 1111 |

- Example: eca8 6420
  - 1110 1100 1010 1000 0110 0100 0010 0000

# MIPS I-format Instructions

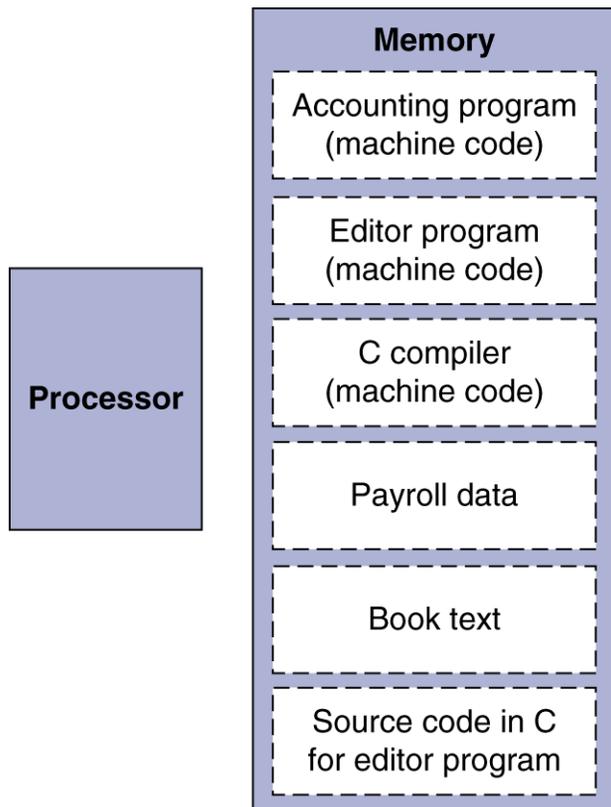
---



- Immediate arithmetic and load/store instructions
  - rt: destination or source register number
  - Constant:  $-2^{15}$  to  $+2^{15} - 1$
  - Address: offset added to base address in rs
- *Design Principle 4*: Good design demands good compromises
  - Different formats complicate decoding, but allow 32-bit instructions uniformly
  - Keep formats as similar as possible

# Stored Program Computers

## The BIG Picture



- Instructions represented in binary, just like data
- Instructions and data stored in memory
- Programs can operate on programs
  - e.g., compilers, linkers, ...
- Binary compatibility allows compiled programs to work on different computers
  - Standardized ISAs

# Logical Operations

- Instructions for bitwise manipulation

| Operation   | C  | Java | MIPS      |
|-------------|----|------|-----------|
| Shift left  | << | <<   | sll       |
| Shift right | >> | >>>  | srl       |
| Bitwise AND | &  | &    | and, andi |
| Bitwise OR  |    |      | or, ori   |
| Bitwise NOT | ~  | ~    | nor       |

- Useful for extracting and inserting groups of bits in a word

# Shift Operations

---



- shamt: how many positions to shift
- Shift left logical
  - Shift left and fill with 0 bits
  - `sll` by  $i$  bits multiplies by  $2^i$
- Shift right logical
  - Shift right and fill with 0 bits
  - `srl` by  $i$  bits divides by  $2^i$  (unsigned only)

# AND Operations

---

- Useful to mask bits in a word
  - Select some bits, clear others to 0

and \$t0, \$t1, \$t2

|      |   |
|------|---|
| \$t2 | 0000 0000 0000 0000 0000 1101 1100 0000 |
| \$t1 | 0000 0000 0000 0000 0011 1100 0000 0000 |
| \$t0 | 0000 0000 0000 0000 0000 1100 0000 0000 |

# OR Operations

---

- Useful to include bits in a word
  - Set some bits to 1, leave others unchanged

or \$t0, \$t1, \$t2

|      |   |
|------|---|
| \$t2 | 0000 0000 0000 0000 0000 1101 1100 0000 |
| \$t1 | 0000 0000 0000 0000 0011 1100 0000 0000 |
| \$t0 | 0000 0000 0000 0000 0011 1101 1100 0000 |

# NOT Operations

---

- Useful to invert bits in a word
  - Change 0 to 1, and 1 to 0
- MIPS has NOR 3-operand instruction
  - $a \text{ NOR } b == \text{NOT} ( a \text{ OR } b )$

```
nor $t0, $t1, $zero
```



Register 0: always read as zero

```
$t1 0000 0000 0000 0000 0011 1100 0000 0000
```

```
$t0 1111 1111 1111 1111 1100 0011 1111 1111
```

# Conditional Operations

---

- Branch to a labeled instruction if a condition is true
  - Otherwise, continue sequentially
- `beq rs, rt, L1`
  - if (`rs == rt`) branch to instruction labeled L1;
- `bne rs, rt, L1`
  - if (`rs != rt`) branch to instruction labeled L1;
- `j L1`
  - unconditional jump to instruction labeled L1

# Compiling If Statements

- C code:

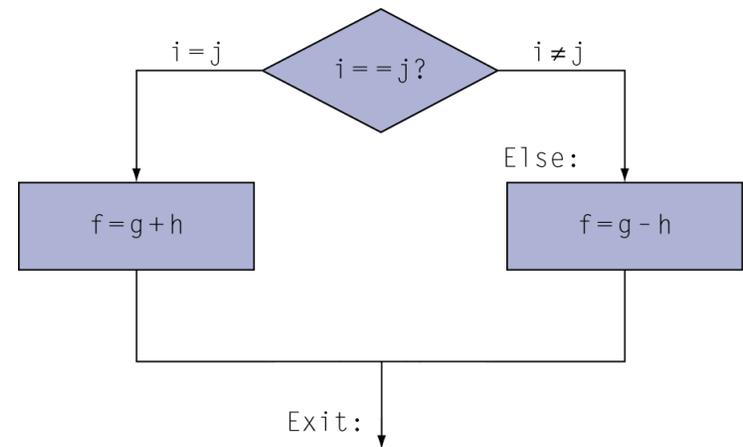
```
if (i==j) f = g+h;  
else f = g-h;
```

– *f, g, ...* in *\$s0, \$s1, ...*

- Compiled MIPS code:

```
        bne $s3, $s4, Else  
        add $s0, $s1, $s2  
        j   Exit  
Else:   sub $s0, $s1, $s2  
Exit:   ...
```

Assembler calculates addresses



# Compiling Loop Statements

---

- C code:

```
while (save[i] == k) i += 1;
```

– *i* in *\$s3*, *k* in *\$s5*, address of *save* in *\$s6*

- Compiled MIPS code:

```
Loop:  sll    $t1, $s3, 2 /* x4 */
       add   $t1, $t1, $s6
       lw    $t0, 0($t1)
       bne   $t0, $s5, Exit
       addi  $s3, $s3, 1
       j     Loop
Exit:  ...
```

# More Conditional Operations

---

- Set result to 1 if a condition is true
  - Otherwise, set to 0
- `slt rd, rs, rt`
  - if ( $rs < rt$ )  $rd = 1$ ; else  $rd = 0$ ;
- `slti rt, rs, constant`
  - if ( $rs < \text{constant}$ )  $rt = 1$ ; else  $rt = 0$ ;
- Use in combination with `beq`, `bne`
  - `slt $t0, $s1, $s2 # if ($s1 < $s2)`
  - `bne $t0, $zero, L # branch to L`

# Branch Addressing

---

- Branch instructions specify
  - Opcode, two registers, target address
- Most branch targets are near branch
  - Forward or backward

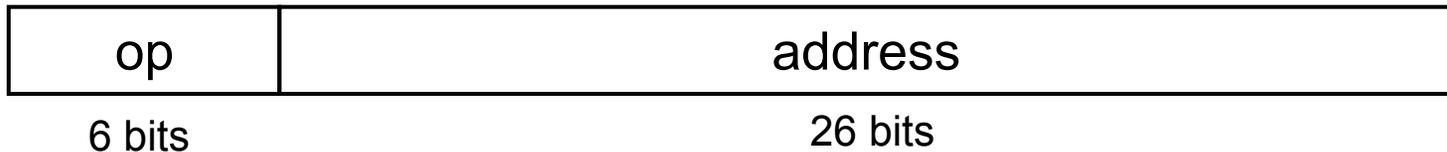


- PC-relative addressing
  - Target address =  $PC + \text{offset} \times 4$
  - PC already incremented by 4 by this time

# Jump Addressing

---

- Jump (j and jal) targets could be anywhere in text segment
  - Encode full address in instruction



- (Pseudo)Direct jump addressing
  - Target address =  $PC_{31..28} : (\text{address} \times 4)$

# Target Addressing Example

- Loop code from earlier example
  - Assume Loop at location 80000

|           |               |                  |       |    |       |    |   |      |
|-----------|---------------|------------------|-------|----|-------|----|---|------|
| Loop: sll | \$t1, \$s3, 2 | 80000            | 0     | 0  | 19    | 9  | 4 | 0    |
|           | add           | \$t1, \$t1, \$s6 | 80004 | 0  | 9     | 22 | 9 | 0 32 |
|           | lw            | \$t0, 0(\$t1)    | 80008 | 35 | 9     | 8  | 0 |      |
|           | bne           | \$t0, \$s5, Exit | 80012 | 5  | 8     | 21 | 2 |      |
|           | addi          | \$s3, \$s3, 1    | 80016 | 8  | 19    | 19 | 1 |      |
|           | j             | Loop             | 80020 | 2  | 20000 |    |   |      |
| Exit: ... |               | 80024            |       |    |       |    |   |      |

# Addressing Mode Summary

---

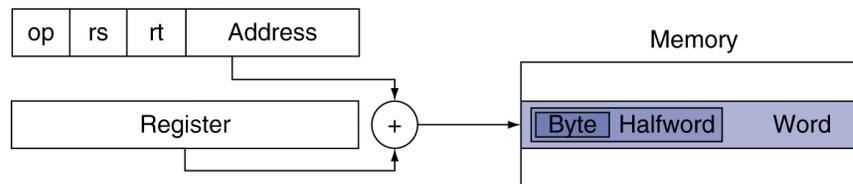
## 1. Immediate addressing



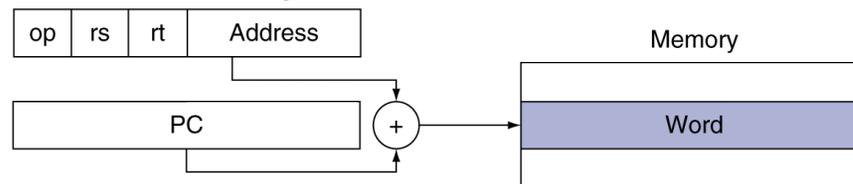
## 2. Register addressing



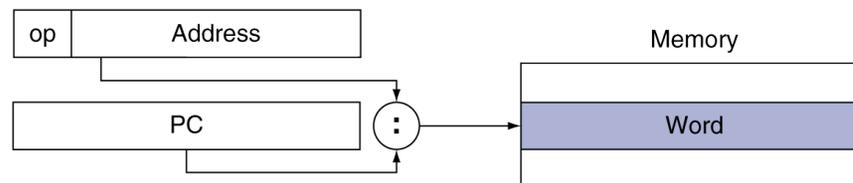
## 3. Base addressing



## 4. PC-relative addressing



## 5. Pseudodirect addressing



# C Sort Example

---

- Illustrates use of assembly instructions for a C bubble sort function
- Swap procedure (leaf)

```
void swap(int v[], int k)
{
    int temp;
    temp = v[k];
    v[k] = v[k+1];
    v[k+1] = temp;
}
```

- v in \$a0, k in \$a1, temp in \$t0

# The Procedure Swap

---

```
swap: sll $t1, $a1, 2    # $t1 = k * 4
      add $t1, $a0, $t1 # $t1 = v+(k*4)
                          # (address of v[k])
      lw  $t0, 0($t1)    # $t0 (temp) = v[k]
      lw  $t2, 4($t1)    # $t2 = v[k+1]
      sw  $t2, 0($t1)    # v[k] = $t2 (v[k+1])
      sw  $t0, 4($t1)    # v[k+1] = $t0 (temp)
      jr  $ra           # return to calling routine
```

```
void swap(int v[], int k)
{
    int temp;
    temp = v[k];
    v[k] = v[k+1];
    v[k+1] = temp;
}
```

# The Sort Procedure in C

---

- Non-leaf (calls swap)

```
void sort (int v[], int n)
{
    int i, j;
    for (i = 0; i < n; i += 1) {
        for (j = i - 1;
            j >= 0 && v[j] > v[j + 1];
            j -= 1) {
            swap(v, j);
        }
    }
}
```

– v in \$a0, k in \$a1, i in \$s0, j in \$s1

# The Procedure Body

```
        move $s2, $a0          # save $a0 into $s2
        move $s3, $a1          # save $a1 into $s3
        move $s0, $zero        # i = 0
for1tst: slt  $t0, $s0, $s3     # $t0 = 0 if $s0 ≥ $s3 (i ≥ n)
        beq  $t0, $zero, exit1  # go to exit1 if $s0 ≥ $s3 (i ≥ n)
        addi $s1, $s0, -1      # j = i - 1
for2tst: slti $t0, $s1, 0       # $t0 = 1 if $s1 < 0 (j < 0)
        bne  $t0, $zero, exit2  # go to exit2 if $s1 < 0 (j < 0)
        sll  $t1, $s1, 2       # $t1 = j * 4
        add  $t2, $s2, $t1     # $t2 = v + (j * 4)
        lw   $t3, 0($t2)       # $t3 = v[j]
        lw   $t4, 4($t2)       # $t4 = v[j + 1]
        slt  $t0, $t4, $t3     # $t0 = 0 if $t4 ≥ $t3
        beq  $t0, $zero, exit2  # go to exit2 if $t4 ≥ $t3
        move $a0, $s2          # 1st param of swap is v (old $a0)
        move $a1, $s1          # 2nd param of swap is j
        jal  swap              # call swap procedure
        addi $s1, $s1, -1      # j -= 1
        j    for2tst           # jump to test of inner loop
exit2:  addi $s0, $s0, 1       # i += 1
        j    for1tst           # jump to test of outer loop
```

Move  
params

Outer loop

Inner loop

Pass  
params  
& call

Inner loop

Outer loop

# The Full Procedure

---

```
sort:    addi $sp,$sp, -20      # make room on stack for 5 registers
        sw $ra, 16($sp)       # save $ra on stack
        sw $s3,12($sp)        # save $s3 on stack
        sw $s2, 8($sp)        # save $s2 on stack
        sw $s1, 4($sp)        # save $s1 on stack
        sw $s0, 0($sp)        # save $s0 on stack
        ...                    # procedure body
        ...
        exit1: lw $s0, 0($sp)  # restore $s0 from stack
        lw $s1, 4($sp)        # restore $s1 from stack
        lw $s2, 8($sp)        # restore $s2 from stack
        lw $s3,12($sp)        # restore $s3 from stack
        lw $ra,16($sp)        # restore $ra from stack
        addi $sp,$sp, 20      # restore stack pointer
        jr $ra                # return to calling routine
```

# Arrays vs. Pointers

---

- Array indexing involves
  - Multiplying index by element size
  - Adding to array base address
- Pointers correspond directly to memory addresses
  - Can avoid indexing complexity

# Example: Clearing and Array

```
clear1(int array[], int size) {
    int i;
    for (i = 0; i < size; i += 1)
        array[i] = 0;
}
```

```
clear2(int *array, int size) {
    int *p;
    for (p = &array[0]; p < &array[size];
        p = p + 1)
        *p = 0;
}
```

```
        move $t0,$zero    # i = 0
loop1:  sll $t1,$t0,2      # $t1 = i * 4
        add $t2,$a0,$t1   # $t2 =
                                # &array[i]
        sw $zero, 0($t2)  # array[i] = 0
        addi $t0,$t0,1    # i = i + 1
        slt $t3,$t0,$a1   # $t3 =
                                # (i < size)
        bne $t3,$zero,loop1 # if (...)
                                # goto loop1
```

```
        move $t0,$a0      # p = & array[0]
        sll $t1,$a1,2      # $t1 = size * 4
        add $t2,$a0,$t1   # $t2 =
                                # &array[size]
loop2:  sw $zero,0($t0)   # Memory[p] = 0
        addi $t0,$t0,4     # p = p + 4
        slt $t3,$t0,$t2   # $t3 =
                                # (p < &array[size])
        bne $t3,$zero,loop2 # if (...)
                                # goto loop2
```

# Comparison of Array vs. Ptr

---

- Multiply “strength reduced” to shift
- Array version requires shift to be inside loop
  - Part of index calculation for incremented  $i$
  - c.f. incrementing pointer
- Compiler can achieve same effect as manual use of pointers
  - Induction variable elimination
  - Better to make program clearer and safer

# Summary

---

| Instruction class | MIPS examples                        | SPEC2006 Int | SPEC2006 FP |
|-------------------|--------------------------------------|--------------|-------------|
| Arithmetic        | add, sub, addi                       | 16%          | 48%         |
| Data transfer     | lw, sw, lb, lbu,<br>lh, lhu, sb, lui | 35%          | 36%         |
| Logical           | and, or, nor, andi,<br>ori, sll, srl | 12%          | 4%          |
| Cond. Branch      | beq, bne, slt,<br>slti, sltiu        | 34%          | 8%          |
| Jump              | j, jr, jal                           | 2%           | 0%          |

---

# Backup

# ARM & MIPS Similarities

- ARM: the most popular embedded core
- Similar basic set of instructions to MIPS

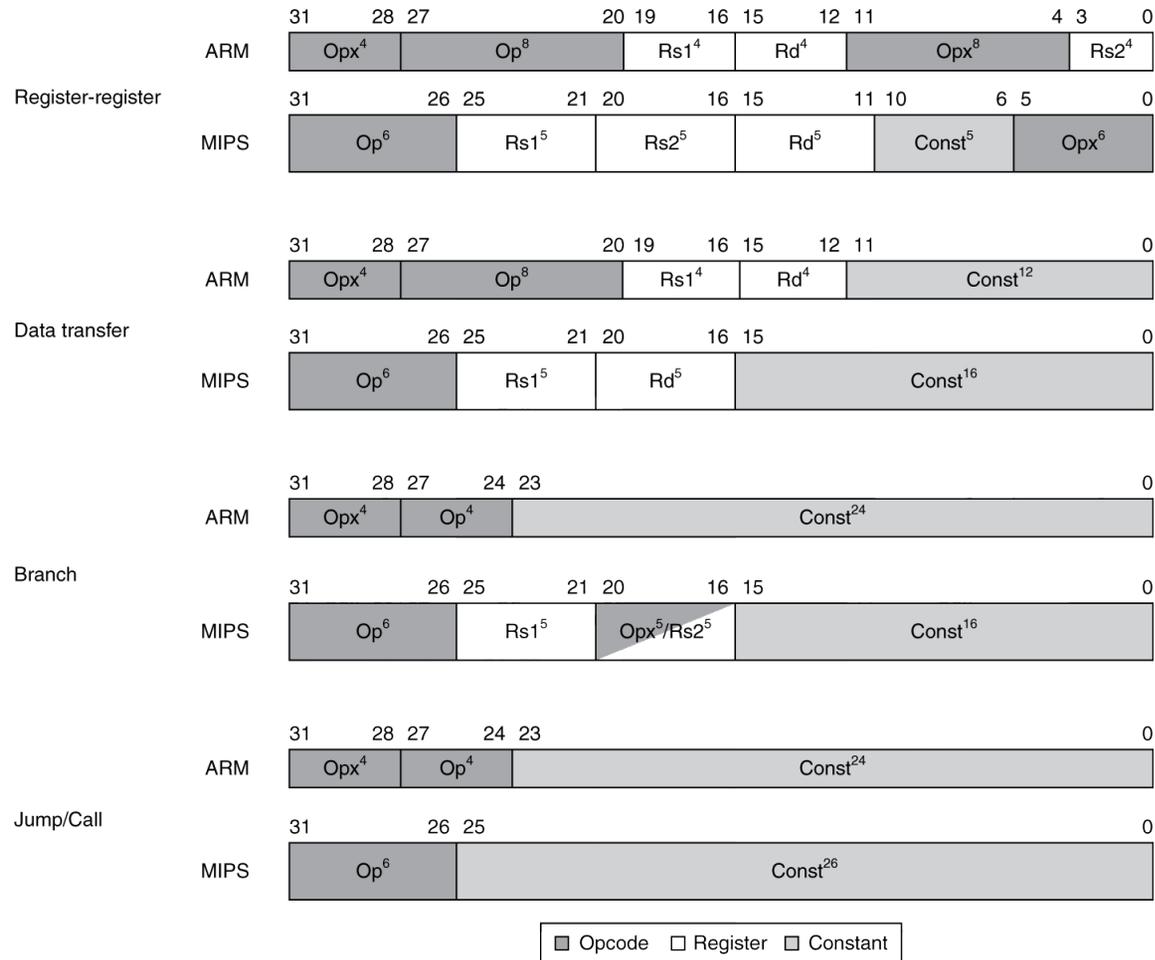
|                       | ARM           | MIPS          |
|-----------------------|---------------|---------------|
| Date announced        | 1985          | 1985          |
| Instruction size      | 32 bits       | 32 bits       |
| Address space         | 32-bit flat   | 32-bit flat   |
| Data alignment        | Aligned       | Aligned       |
| Data addressing modes | 9             | 3             |
| Registers             | 15 × 32-bit   | 31 × 32-bit   |
| Input/output          | Memory mapped | Memory mapped |

# Compare and Branch in ARM

---

- Uses condition codes for result of an arithmetic/logical instruction
  - Negative, zero, carry, overflow
  - Compare instructions to set condition codes without keeping the result
- Each instruction can be conditional
  - Top 4 bits of instruction word: condition value
  - Can avoid branches over single instructions

# Instruction Encoding



# The Intel x86 ISA

---

- Evolution with backward compatibility
  - 8080 (1974): 8-bit microprocessor
    - Accumulator, plus 3 index-register pairs
  - 8086 (1978): 16-bit extension to 8080
    - Complex instruction set (CISC)
  - 8087 (1980): floating-point coprocessor
    - Adds FP instructions and register stack
  - 80286 (1982): 24-bit addresses, MMU
    - Segmented memory mapping and protection
  - 80386 (1985): 32-bit extension (now IA-32)
    - Additional addressing modes and operations
    - Paged memory mapping as well as segments

# The Intel x86 ISA

---

- Further evolution...
  - i486 (1989): pipelined, on-chip caches and FPU
    - Compatible competitors: AMD, Cyrix, ...
  - Pentium (1993): superscalar, 64-bit datapath
    - Later versions added MMX (Multi-Media eXtension) instructions
    - The infamous FDIV bug
  - Pentium Pro (1995), Pentium II (1997)
    - New microarchitecture (see Colwell, *The Pentium Chronicles*)
  - Pentium III (1999)
    - Added SSE (Streaming SIMD Extensions) and associated registers
  - Pentium 4 (2001)
    - New microarchitecture
    - Added SSE2 instructions

# The Intel x86 ISA

---

- And further...
  - AMD64 (2003): extended architecture to 64 bits
  - EM64T – Extended Memory 64 Technology (2004)
    - AMD64 adopted by Intel (with refinements)
    - Added SSE3 instructions
  - Intel Core (2006)
    - Added SSE4 instructions, virtual machine support
  - AMD64 (announced 2007): SSE5 instructions
    - Intel declined to follow, instead...
  - Advanced Vector Extension (announced 2008)
    - Longer SSE registers, more instructions
- If Intel didn't extend with compatibility, its competitors would!
  - Technical elegance ≠ market success

# Basic x86 Registers

---

| Name   | 31    | 0     | Use                                  |
|--------|-------|-------|--------------------------------------|
| EAX    | [Bar] |       | GPR 0                                |
| ECX    | [Bar] |       | GPR 1                                |
| EDX    | [Bar] |       | GPR 2                                |
| EBX    | [Bar] |       | GPR 3                                |
| ESP    | [Bar] |       | GPR 4                                |
| EBP    | [Bar] |       | GPR 5                                |
| ESI    | [Bar] |       | GPR 6                                |
| EDI    | [Bar] |       | GPR 7                                |
|        | CS    | [Bar] | Code segment pointer                 |
|        | SS    | [Bar] | Stack segment pointer (top of stack) |
|        | DS    | [Bar] | Data segment pointer 0               |
|        | ES    | [Bar] | Data segment pointer 1               |
|        | FS    | [Bar] | Data segment pointer 2               |
|        | GS    | [Bar] | Data segment pointer 3               |
| EIP    | [Bar] |       | Instruction pointer (PC)             |
| EFLAGS | [Bar] |       | Condition codes                      |

# Basic x86 Addressing Modes

---

- Two operands per instruction

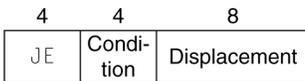
| Source/dest operand | Second source operand |
|---------------------|-----------------------|
| Register            | Register              |
| Register            | Immediate             |
| Register            | Memory                |
| Memory              | Register              |
| Memory              | Immediate             |

- Memory addressing modes

- Address in register
- $\text{Address} = R_{\text{base}} + \text{displacement}$
- $\text{Address} = R_{\text{base}} + 2^{\text{scale}} \times R_{\text{index}}$  (scale = 0, 1, 2, or 3)
- $\text{Address} = R_{\text{base}} + 2^{\text{scale}} \times R_{\text{index}} + \text{displacement}$

# x86 Instruction Encoding

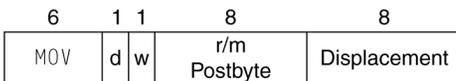
a. JE EIP + displacement



b. CALL



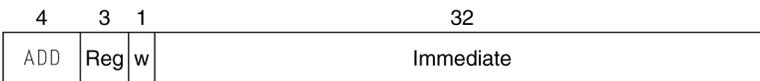
c. MOV EBX, [EDI + 45]



d. PUSH ESI



e. ADD EAX, #6765



f. TEST EDX, #42



- Variable length encoding
  - Postfix bytes specify addressing mode
  - Prefix bytes modify operation
  - Operand length, repetition, locking, ...

# Implementing IA-32

---

- Complex instruction set makes implementation difficult
  - Hardware translates instructions to simpler microoperations
    - Simple instructions: 1–1
    - Complex instructions: 1–many
  - Microengine similar to RISC
  - Market share makes this economically viable
- Comparable performance to RISC
  - Compilers avoid complex instructions