EECS 361 Computer Architecture Lecture 3 – Instruction Set Architecture

Prof. Alok N. Choudhary

choudhar@ece.northwestern.edu

Today's Lecture

Quick Review of Last Week

Classification of Instruction Set Architectures

Instruction Set Architecture Design Decisions

• Operands

Annoucements

- Operations
- Memory Addressing
- Instruction Formats

Instruction Sequencing

Language and Compiler Driven Decisions

Summary of Lecture 2

Two Notions of "Performance"

Plane	DC to Paris	Speed	Passengers	Throughput (pmph)
Boeing 747	6.5 hours	610 mph	470	286,700
Concorde	3 hours	1350 mph	132	178,200

Which has higher performance?

Execution time (response time, latency, ...)

• Time to do a task

Throughput (bandwidth, ...)

• Tasks per unit of time

Response time and throughput often are in opposition EECS 361

Definitions

Performance is typically in units-per-second

• bigger is better

If we are primarily concerned with response time

 performance = 1 execution_time

" X is n times faster than Y" means

 $\frac{ExecutionTime_{y}}{ExecutionTime_{x}} = \frac{Performance_{x}}{Performance_{y}} = n$

Organizational Trade-offs



Instruction Set Architecture with the Implementation of that architecture, and the program measured

Principal Design Metrics: CPI and Cycle Time

$$Performance = \frac{1}{ExecutionTime}$$

$$Performance = \frac{1}{CPI \times CycleTime}$$

$$Performance = \frac{1}{\frac{Cycles}{Instruction} \times \frac{Seconds}{Cycle}} = \frac{Instructions}{Seconds}$$

Amdahl's "Law": Make the Common Case Fast



Suppose that enhancement E accelerates a fraction F of the task

by a factor S and the remainder of the task is unaffected then,

ExTime(with E) = ((1-F) + F/S) X ExTime(without E)

Performance improvement is limited by how much the improved feature is used \rightarrow Invest resources where time is spent.

$$Speedup_{overall} = \frac{ExecTime_{old}}{ExecTime_{new}} = \frac{1}{\frac{Fraction_{enhanced}}{Speedup_{enhanced}} + (1 - Fraction_{enhanced})}$$

EECS 361

Classification of Instruction Set Architectures

Instruction Set Design



ISAs evolve: MIPS-I, MIPS-II, MIPS-II, MIPS-IV, MIPS, MDMX, MIPS-32, MIPS-64

Typical Processor Execution Cycle



Obtain instruction from program storage

Determine required actions and instruction size

Locate and obtain operand data

Compute result value or status

Deposit results in register or storage for later use

Determine successor instruction

Instruction and Data Memory: Unified or Separate



Princeton (Von Neumann) Architecture

- --- Data and Instructions mixed in same unified memory
- --- Program as data
- --- Storage utilization
- --- Single memory interface



Harvard Architecture

- --- Data & Instructions in separate memories
- --- Has advantages in certain high performance implementations
- --- Can optimize each memory

Basic Addressing Classes



Declining cost of registers

© 2003 Elsevier Science (USA). All rights reserved.

Stack Architectures

- Stack: First-In Last-Out data structure (FILO)
- Instruction operands
 - None for ALU operations
 - One for push/pop
- Advantages:
 - Short instructions
 - Compiler is easy to write
- Disadvantages
 - Code is inefficient
 - Fix: random access to stacked values
 - Stack size & access latency
 - Fix : register file or cache for top entries
- Examples
 - 60s: Burroughs B5500/6500, HP 3000/70
 - Today: Java VM



Accumulator Architectures

- Single register (accumulator)
- Instructions
 - ALU (Acc \leftarrow Acc + *M)
 - Load to accumulator (Acc $\leftarrow *M$)
 - Store from accumulator (*M \leftarrow Acc)
- Instruction operands
 - One explicit (memory address)
 - One implicit (accumulator)
- Attributes:
 - Short instructions
 - Minimal internal state; simple design
 - Many loads and stores
- Examples:
 - Early machines: IBM 7090, DEC PDP-8
 - Today: DSP architectures



Register-Set Architectures

- General Purpose Registers (GPRs)
- Registers:
 - Explicitly managed memory for holding recently used values
- The dominant architecture: CDC 6600, IBM 360/370, PDP-11, 68000, RISC
- Advantages:
 - Allows fast access to temporary values
 - Permits clever compiler optimization
 - Reduced traffic to memory
- Disadvantages:
 - Longer instructions (than accumulator designs)



Register-to-Register: Load-Store Architectures



Register-to-Memory Architectures

- One memory address in ALU ops
- Typically 2-operand ALU ops
- Advantages
 - Small instruction count
 - Dense encoding
- Disadvantages
 - Result destroys an operand
 - Instruction length varies
 - Clocks per instruction varies
 - Harder to pipeline
- Examples
 - IBM 360/370, VAX



Memory-to-Memory Architectures

- All ALU operands from memory addresses
- Advantages
 - No register wastage
 - Lowest instruction count
- Disadvantages
 - Large variation in instruction length
 - Large variation in clocks per instructions
 - Huge memory traffic
- Examples
 - VAX

	D=	B+ (C*D)	
mul	D	<-	C*D	
add	D	<-	D+B	

Instruction Set Architecture Design Decisions

Basic Issues in Instruction Set Design

What data types are supported. What size.

What operations (and how many) should be provided

- LD/ST/INC/BRN sufficient to encode any computation, or just Sub and Branch!
- But not useful because programs too long!

How (and how many) operands are specified

Most operations are dyadic (eg, A <- B + C)

• Some are monadic (eg, A <- ~B)

Location of operands and result

- where other than memory?
- how many explicit operands?
- how are memory operands located?
- which can or cannot be in memory?
- How are they addressed

How to encode these into consistent instruction formats

- Instructions should be multiples of basic data/address widths
- Encoding Driven by static measurement and dynamic tracing of selected benchmarks and workloads.

Typical instruction set:

32 bit word
basic operand addresses are 32 bits long
basic operands, like integers, are 32 bits long
in general case, instruction could reference 3 operands (A := B + C)

Typical challenge:

•encode operations in a small number of bits

3-21

Operands

Comparing Number of Instructions

Code sequence for (C = A + B) for four classes of instruction sets:

Stack	Accumulator	Register (register-memory)	Register (load-store)
Push A	Load A	Load R1,A	Load R1,A
Push B	Add B	Add R1,B	Load R2,B
Add	Store C	Store C, R1	Add R3,R1,R2
Рор С			Store C,R3

$$ExecutionTime = \frac{1}{Performance} = Instructions \times \frac{Cycles}{Instruction} \times \frac{Seconds}{Cycle}$$

Examples of Register Usage

Number of memory addresses per typical ALU instruction

	Maxi	mum number of operands per typical ALU instruction
		Examples
0	3	SPARC, MIPS, Precision Architecture, Power PC
1	2	Intel 80x86, Motorola 68000
2	2	VAX (also has 3-operand formats)
3	3	VAX (also has 2-operand formats)

General Purpose Registers Dominate

1975-2002 all machines use general purpose registers

Advantages of registers

- Registers are faster than memory
- Registers compiler technology has evolved to efficiently generate code for register files
 - E.g., (A*B) (C*D) (E*F) can do multiplies in any order vs. stack
- Registers can hold variables
 - Memory traffic is reduced, so program is sped up (since registers are faster than memory)
- Code density improves (since register named with fewer bits than memory location)
- Registers imply operand locality

Operand Size Usage



Frequency of reference by size

 Support for these data sizes and types: 8-bit, 16-bit, 32-bit integers and 32-bit and 64-bit IEEE 754 floating point numbers

Announcements

Next lecture

• MIPS Instruction Set

Operations

Typical Operations (little change since 1960)



EECS 361

Top 10 80x86 Instructions

° Rank	instruction I	nteger Average Percent total executed
1	load	22%
2	conditional branc	h 20%
3	compare	16%
4	store	12%
5	add	8%
6	and	6%
7	sub	5%
8	move register-reg	jister 4%
9	call	1%
10	return	1%
	Total	96%

° Simple instructions dominate instruction frequency

Memory Addressing

Since 1980, almost every machine uses addresses to level of 8-bits (byte)

Two questions for design of ISA:

- Since could read a 32-but word as four loads of bytes from sequential byte address of as one load word from a single byte address, how do byte addresses map onto words?
- Can a word be placed on any byte boundary?

Mapping Word Data into a Byte Addressable Memory:



Intel 80x86, DEC Vax, DEC Alpha (Windows NT)

Mapping Word Data into a Byte Addressable Memory: Alignment



Alignment: require that objects fall on address that is multiple of their size.

Addressing Modes

• Addressing modes specify a constant, a register, or a location in memory

_	Register	add r1,	r2	r1 <- r1+r2
_	Immediate	add r1,	#5	r1 <- r1+5
_	Direct	add r1,	(0x200)	r1 <- r1+M[0x200]
_	Register indirect	add r1,	(r2)	r1 <- r1+M[r2]
_	Displacement	add r1,	100(r2)	r1 <- r1+M[r2+100]
_	Indexed	add r1,	(r2+r3)	r1 <- r1+M[r2+r3]
_	Scaled	add r1,	(r2+r3*4)	r1 <- r1+M[r2+r3*4]
_	Memory indirect	add r1,	@(r2)	r1 <- r1+M[M[r2]]
_	Auto-increment	add r1,	(r2)+	r1 <- r1+M[r2], r2++
_	Auto-decrement	add r1,	-(r2)	r2, r1 <- r1+M[r2]

• Complicated modes reduce instruction count at the cost of complex implementations

Common Memory Addressing Modes



Measured on the VAX-11

Register operations account for 51% of all references

~75% - displacement and immediate

EECS 361 ~85% - displacement, immediate and register indirect

Displacement Address Size



Average of 5 SPECint92 and 5 SPECfp92 programs

~1% of addresses > 16-bits

12 ~ 16 bits of displacement cover most usage (+ and -)

EECS 361

Frequency of Immediates (Instruction Literals)



~25% of all loads and ALU operations use immediates

15~20% of all instructions use immediates

Size of Immediates



50% to 60% fit within 8 bits

75% to 80% fit within 16 bits

EECS 361

Addressing Summary

Data Addressing modes that are important:

• Displacement, Immediate, Register Indirect

Displacement size should be 12 to 16 bits

Immediate size should be 8 to 16 bits

Instruction Formats

Instruction Format

Specify

- Operation / Data Type
- Operands

Stack and Accumulator architectures have implied operand addressing

If have many memory operands per instruction and/or many addressing modes:

• Need one address specifier per operand

If have load-store machine with 1 address per instruction and one or two addressing modes:

• Can encode addressing mode in the opcode



If code size is most important, use variable length instructions

If performance is most important, use fixed length instructions

Recent embedded machines (ARM, MIPS) added optional mode to execute subset of 16bit wide instructions (Thumb, MIPS16); per procedure decide performance or density

Some architectures actually exploring on-the-fly decompression for more density. EECS 361

Support these simple instructions, since they will dominate the number of instructions executed:

```
load,
store,
add,
subtract,
move register-register,
and,
shift,
compare equal, compare not equal,
branch,
jump,
call,
return;
```

Example: MIPS Instruction Formats and Addressing Modes

All instructions 32 bits wide



Instruction Set Design Metrics

Static Metrics

• How many bytes does the program occupy in memory?

Dynamic Metrics

- How many instructions are executed?
- How many bytes does the processor fetch to execute the program?
- How many clocks are required per instruction?



Instruction Sequencing

Instruction Sequencing

The next instruction to be executed is typically implied

- Instructions execute sequentially
- Instruction sequencing increments a Program Counter



Sequencing flow is disrupted conditionally and unconditionally

• The ability of computers to test results and conditionally instructions is one of the reasons computers have become so useful



Dynamic Frequency



Condition Testing

° Condition Codes

Processor status bits are set as a side-effect of arithmetic instructions (possibly on Moves) or explicitly by compare or test instructions.

ex: add r1, r2, r3 bz label

° Condition Register

- Ex: cmp r1, r2, r3 bgt r1, label
- ° Compare and Branch
 - Ex: bgt r1, r2, label

Condition Codes

Setting CC as side effect can reduce the # of instructions

But also has disadvantages:

- --- not all instructions set the condition codes which do and which do not often confusing! e.g., shift instruction sets the carry bit
- --- dependency between the instruction that sets the CC and the one that tests it



Branches

--- Conditional control transfers

V overflow C carry

Sixteen combinations of the basic four conditions:

Always	Unconditional
Never	NOP
Not Equal	~Z
Equal	Z
Greater	~[Z + (N ⊕ V)]
Less or Equal	Z + (N ⊕ V)
Greater or Equal	~(N + V)
Less	NOV
Greater Unsigned	~(C + Z)
Less or Equal Unsigned	C + Z
Carry Clear	~C
Carry Set	С
Positive	~N
Negative	Ν
Overflow Clear	~V
Overflow Set	V

Conditional Branch Distance



PC-relative (+-)

25% of integer branches are 2 to 4 instructions

At least 8 bits suggested (± 128 instructions)

Language and Compiler Driven Facilities

Calls: Why Are Stacks So Great?

Stacking of Subroutine Calls & Returns and Environments:



Some machines provide a memory stack as part of the architecture (e.g., VAX)

Sometimes stacks are implemented via software convention (e.g., MIPS)

Memory Stacks

Useful for stacked environments/subroutine call & return even if operand stack not part of architecture

Stacks that Grow Up vs. Stacks that Grow Down:



How is empty stack represented?

Little --> Big/Last Full

- POP: Read from Mem(SP) Decrement SP
- PUSH: Increment SP Write to Mem(SP)

inf. Big	0 Little	
grows up	grows down	Memory Addresses
0 Little	inf. Big	

Little --> Big/Next Empty

- POP: Decrement SP Read from Mem(SP)
- PUSH: Write to Mem(SP) Increment SP

Call-Return Linkage: Stack Frames



Many variations on stacks possible (up/down, last pushed /next)

Compilers normally keep scalar variables in registers, not memory!

Compilers and Instruction Set Architectures

Ease of compilation

- Orthogonality: no special registers, few special cases, all operand modes available with any data type or instruction type
- Completeness: support for a wide range of operations and target applications
- Regularity: no overloading for the meanings of instruction fields
- Streamlined: resource needs easily determined

Register Assignment is critical too

• Easier if lots of registers

Provide at least 16 general purpose registers plus separate floating-point registers

Be sure all addressing modes apply to all data transfer instructions

Aim for a minimalist instruction set

Summary

Quick Review of Last Week

Classification of Instruction Set Architectures

Instruction Set Architecture Design Decisions

- Operands
- Operations
- Memory Addressing
- Instruction Formats

Instruction Sequencing

Language and Compiler Driven Decisions