Aula 01: Introdução

- **Arquitetura de Computadores**

Mario F. Montenegro Campos
ICEx 4013 - mario@dcc.ufmg.br

O que é arquitetura de computadores?

"os atributos de um sistema de computação na visão do programador, i.e., a estrutura conceitual e o comportamento funcional, … em oposição à implementação física."

Amdahl, Blaaw, and Brooks, 1964

O que é este curso?

Entender as técnicas de projeto, arquitetura das máquinas correntes, fatores da tecnologia, métodos de avaliação que vão determinar a estrutura dos computadores do próximo século

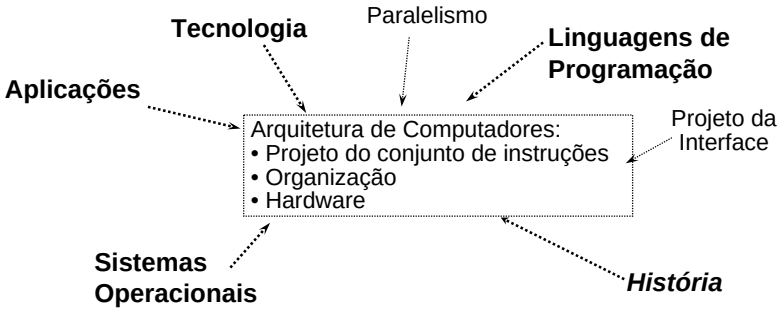

Pré-requisitos

- Conceitos em organização de computadores
  - Organização de Computadores I
- Projeto de sistemas lógicos
  - Introdução aos Sistemas Lógicos
- Linguagens: C e assembler
- Conceitos básicos de sistemas operacionais
- Conceitos básicos de compiladores

Material do curso

- Livro texto e papers
- Listas de exercícios
- Trabalhos de implementação de uma arquitetura RISC
- Projeto de final de curso, envolvendo um tópico avançado

Cursos relacionados

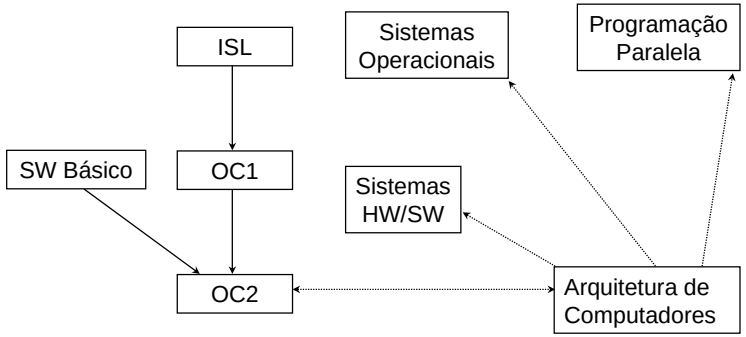

Programação

Hennessy and Patterson, *Computer Architecture: A Quantitative Approach*, 4th Ed., Morgan Kaufman, 2006.

1 Fundamentals of Computer Design

2 Instruction Level Parallelism and Its Exploitation

3 Advanced Techniques for Exploiting Instruction-Level Parallelism and Their Limits

4 Multiprocessors and Thread-Level Parallelism

5 Memory Hierarchy Design

6 Storage Systems

Appendix A: Pipelining: Basic and Intermediate Concepts

Appendix B: Instruction Set Principles and Examples

Appendix C: Introduction to Memory Hierarchy

Livros de referência complementar

- Patterson and Hennessy, *Computer Organization and Design: the Hardware/Software Interface*, Segunda edição, Morgan Kaufmann, 1997.
- Sailer, P.M. and Kaeli, D.R., *The DLX Instruction Set Architecture Handbook*, Morgan Kaufmann, 1996.
- R. Lipsett, C. Schaefer, C. Ussery, *VHDL: Hardware Description Language and Design*, Kluwer Academic Publishers, 1989.
- Peter J Ashenden, *The Designer's Guide to VHDL*, Morgan Kaufmann, 1995.
- manuais (SCI, SCSI, etc…) e papers

Tópicos a serem estudados

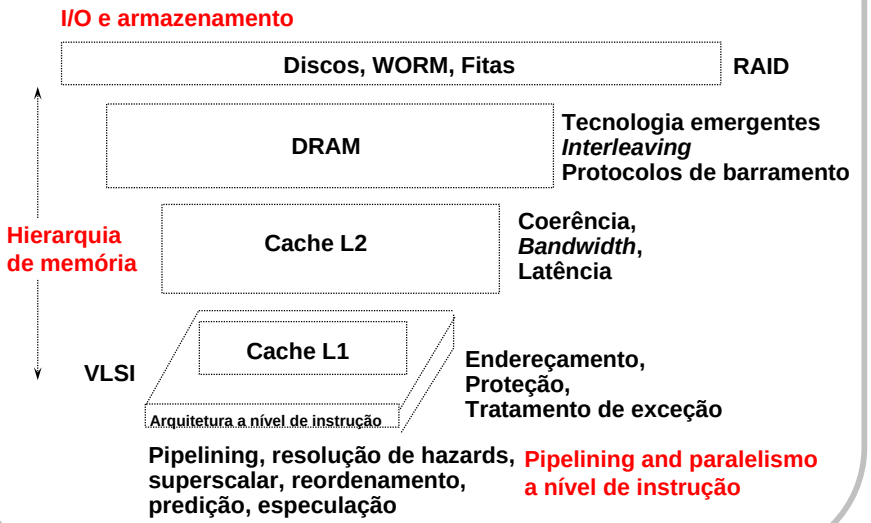

Tópicos a serem estudados

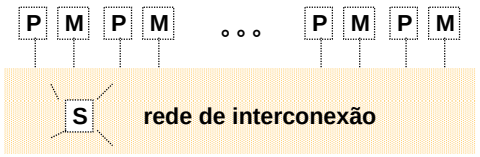

switch entre processador e memória

Multiprocessadores, redes e interconexão

Memória compartilhada, Troca de mensagens,

Interfaces de redes

Topologias, roteamento, *bandwith*, latência, *reliability*

Metodologia de projeto de arquiteturas de computadores

Metodologia de projeto de arquiteturas de computadores (Idéia básica)

- Regra número 1: Nada vem de graça
- Regra número 2: Uma nova arquitetura ou uma nova idéia só é implementada após um estudo da sua viabilidade
- Regra número 3: Nem sempre o melhor e o mais bem desenvolvido vai ganhar o mercado
  - Microsoft Windows vs. Mac OS + OS2
  - 68000 vs. x86
  - Qualquer implementação nova tenta diminuir sua dependência com a sorte o máximo o possível (isso é um esforço multi-disciplinar)

Evolução Tecnológica

	Ano	Componente	Armazen.	Linguagens	O/S
	54	Tubes	core (8 ms)		
	58	Transistor (10µs)			Fortran
	60			Algol, Cobol	Batch
	64	Hybrid (1µs)	thin films	Lisp, APL, Basic	
	66	IC (100ns)	(200ns)	PL1, Simula,C	
	67				Multiprog.
Gerações	71	LSI (10ns)	1k DRAM	O.O.	V.M.
Evoluções	73	(8-bit µP)			
	75	(16-bit µP)	4k DRAM		
	78	VLSI (10ns)	16k DRAM		Redes
	80		64k DRAM		
	84	(32-bit µP)	256k DRAM	ADA	
Paralelismo	87	ULSI	1M DRAM		
	89	GAs	4M DRAM	C++	
	92	(64-bit µP)	16M DRAM	Fortran90	

Evolução Tecnológica Perspectiva da Intel

Ano	Processador	# Xtors
1971	4004	2300
1972	8008	3500
1974	8080	6000
1978	8086/8088	29000
1982	80286	134K
1985	80386	275K
1989	80486	1.2M
1993	Pentium	3.1M
1995	Pentium Pro	5.5M
1997	Pentium II	7.5M
1999	Pentium III	9.5M

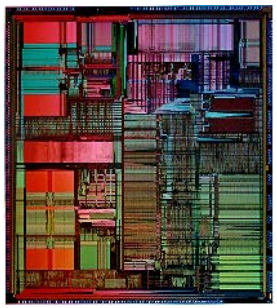

Evolução Tecnológica *Moore's Law*

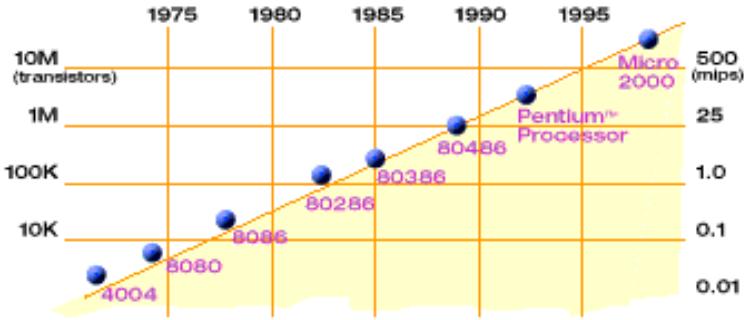

Projeto de Novas Arquiteturas

- Área de aplicação
  - *Special Purpose* (e.g., DSP) / *General Purpose*
  - Científico (intenso em FP) / Comercial (Mainframe)
- Nível de compatibilidade de Software
  - Compatibilidade de código objeto/binário (custo HW vs. SW, x86)
  - Linguagem de máquina (modificações no código objeto/binário são possíveis no projeto da arquitetura)
  - Linguagens de programação (por que não?)

Projeto de Novas Arquiteturas

- Requisitos do sistema operacional para aplicações do tipo *general purpose*
  - Tamanho do espaço de endereçamento (*Address Space*)
  - Gerenciamento de memória e proteção
  - Trocas de contexto
  - Interrupções e *Traps*
- Padrões: inovação vs. competição
  - IEEE 754 (Ponto flutuante)
  - Barramentos de I/O
  - Redes
  - Sistemas operacionais / Linguagens de programação ...

What is "Computer Architecture"

Computer Architecture =

Instruction Set Architecture +

Machine Organization

Instruction Set Architecture (subset of Computer Arch.)

... the attributes of a [computing] system as seen by the programmer, *i.e.* the conceptual structure and functional behavior, as distinct from the organization of the data flows and controls the logic design, and the physical implementation.

– Amdahl, Blaaw, and Brooks, 1964

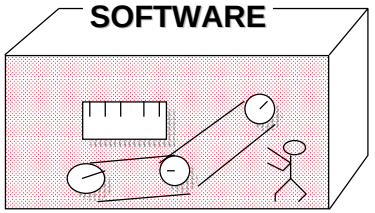

-- Organization of Programmable Storage

-- Data Types & Data Structures: Encodings & Representations

-- Instruction Set

-- Instruction Formats

-- Modes of Addressing and Accessing Data Items and Instructions

-- Exceptional Conditions

The Instruction Set: a Critical Interface

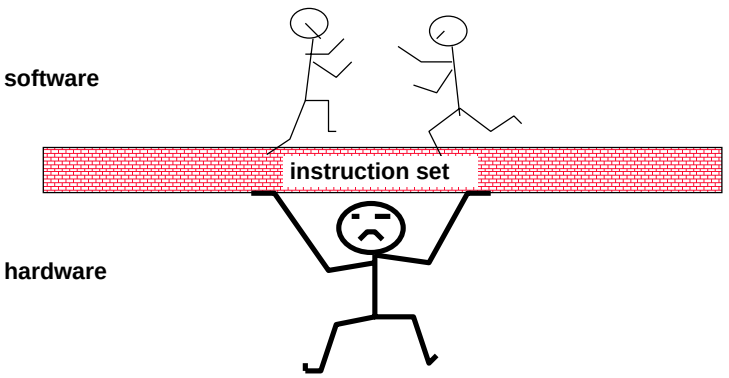

Example ISAs (Instruction Set Architectures)

° Digital Alpha	(v1, v3)	1992-97
° HP PA-RISC	(v1.1, v2.0)	1986-96
° Sun Sparc	(v8, v9)	1987-95
° SGI MIPS	(MIPS I, II, III, IV, V)	1986-96
° Intel	(8086,80286,80386, 80486,Pentium, MMX, ...)	1978-96

MIPS R3000 Instruction Set Architecture (Summary)

° **Instruction Categories**
- Load/Store
- Computational
- Jump and Branch
- Floating Point
  - coprocessor
- Memory Management
- Special

Registers

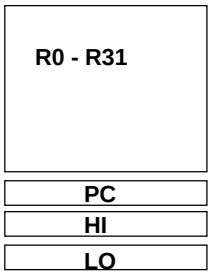

3 Instruction Formats: all 32 bits wide

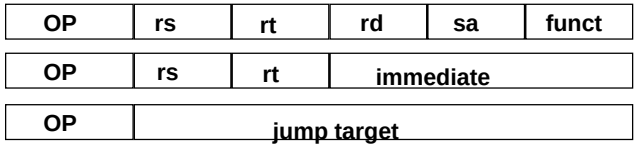

Q: How many already familiar with MIPS ISA?

Organization

° **Capabilities & Performance Characteristics of Principal Functional Units**
- (e.g., Registers, ALU, Shifters, Logic Units, ...)

° **Ways in which these components are interconnected**

° **Information flows between components**

° **Logic and means by which such information flow is controlled.**

° **Choreography of FUs to realize the ISA**

° **Register Transfer Level (RTL) Description**

Logic Designer's View

ISA Level

FUs & Interconnect

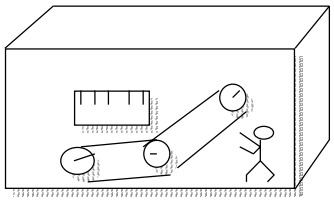

Example Organization

° TI SuperSPARC^{tm} TMS390Z50 in Sun SPARCstation20

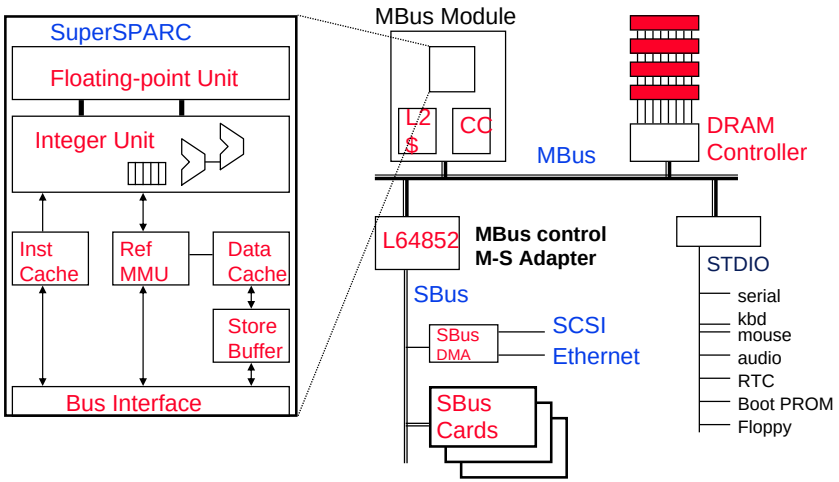

What is "Computer Architecture"?

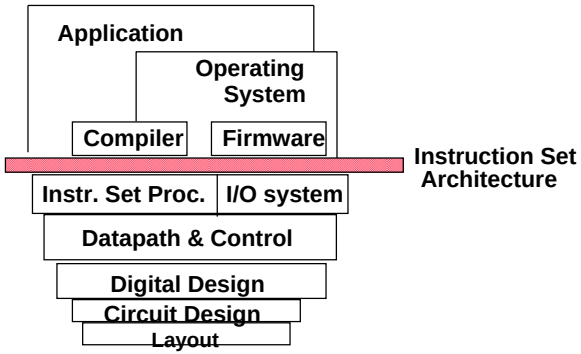

° Coordination of many *levels of abstraction*

° Under a rapidly changing set of forces

° Design, Measurement, *and* Evaluation

Technology

DRAM chip capacity

DRAM	
Year	Size
1980	64 Kb
1983	256 Kb
1986	1 Mb
1989	4 Mb
1992	16 Mb
1996	64 Mb
1999	256 Mb
2002	1 Gb

Microprocessor Logic Density

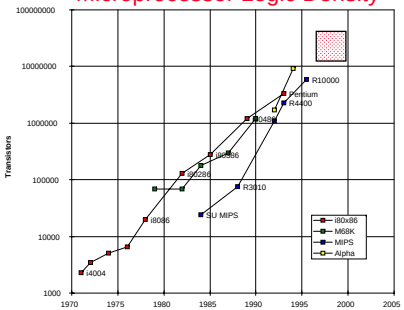

° In ~1985 the single-chip processor (32-bit) and the single-board computer emerged

- => workstations, personal computers, multiprocessors have been riding this wave since

° In the 2002+ timeframe, these may well look like mainframes compared single-chip computer (maybe 2 chips)

Technology => dramatic change

° Processor

- logic capacity: about 30% per year
- clock rate: about 20% per year

° Memory

- DRAM capacity: about 60% per year (4x every 3 years)
- Memory speed: about 10% per year
- Cost per bit: improves about 25% per year

° Disk

- capacity: about 60% per year

Performance Trends

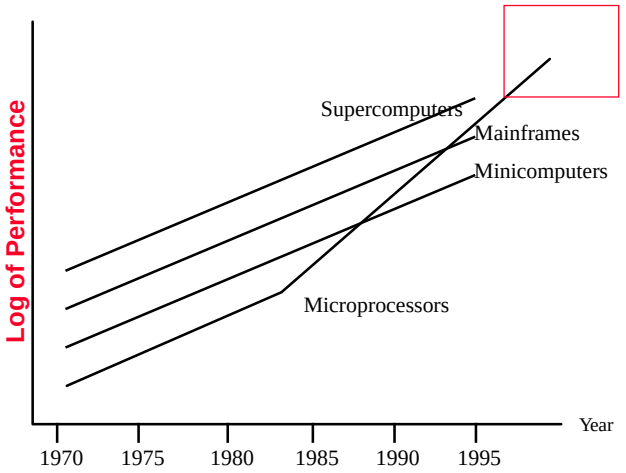

Processor Performance (SPEC)

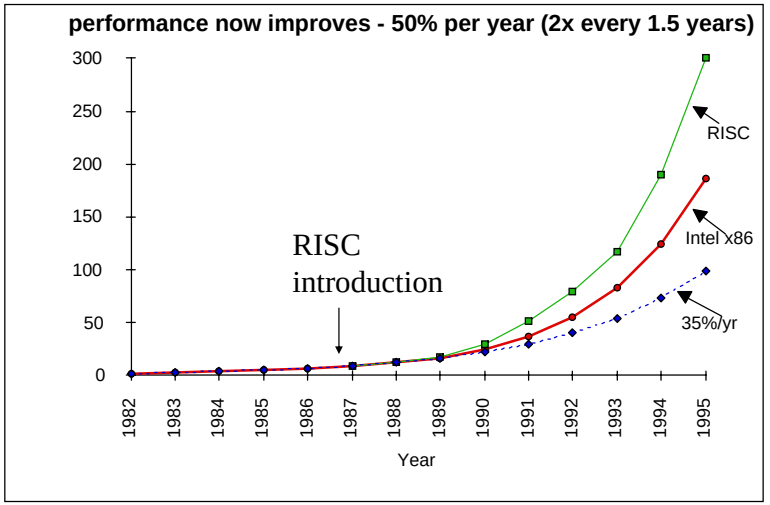

Did RISC win the technology battle and lose the market war?

Applications and Languages

- CAD, CAM, CAE, . . .
- Lotus, DOS, . . .
- Multimedia, . . .
- The Web, . . .
- JAVA, . . .
- ???

Measurement and Evaluation

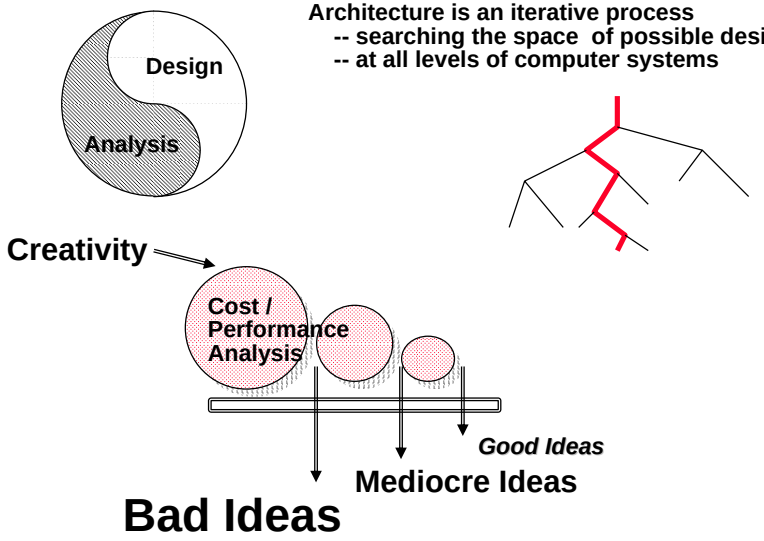

Computer Architecture and Engineering

Instruction Set Design	Computer Organization
Interfaces	Hardware Components
Compiler/System View	Logic Designer's View
-"Building Architect"	-"Construction Engineer"

Conceptual tool box?

- Evaluation Techniques
- Levels of translation (e.g., Compilation)
- Levels of Interpretation (e.g., Microprogramming)
- Hierarchy (e.g, registers, cache, mem,disk,tape)
- Pipelining and Parallelism
- Static / Dynamic Scheduling
- Indirection and Address Translation
- Synchronous and Asynchronous Control Transfer
- Timing, Clocking, and Latching
- CAD Programs, Hardware Description Languages, Simulation
- Physical Building Blocks (e.g., CLA)
- Understanding Technology Trends

Levels of Representation (61C Review)

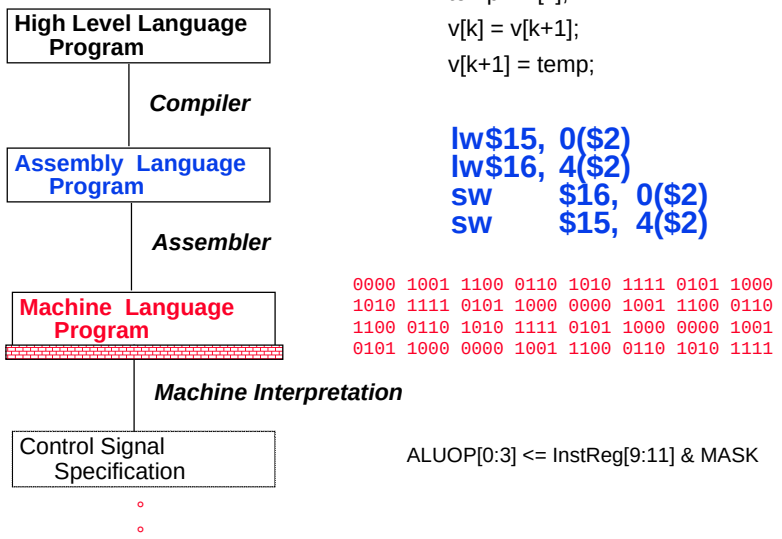

What are "Machine Structures"?

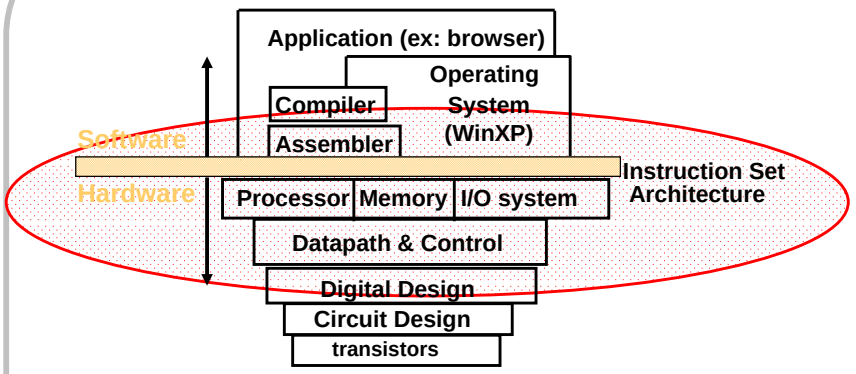

* Coordination of many

levels (layers) of abstraction

Levels of Representation

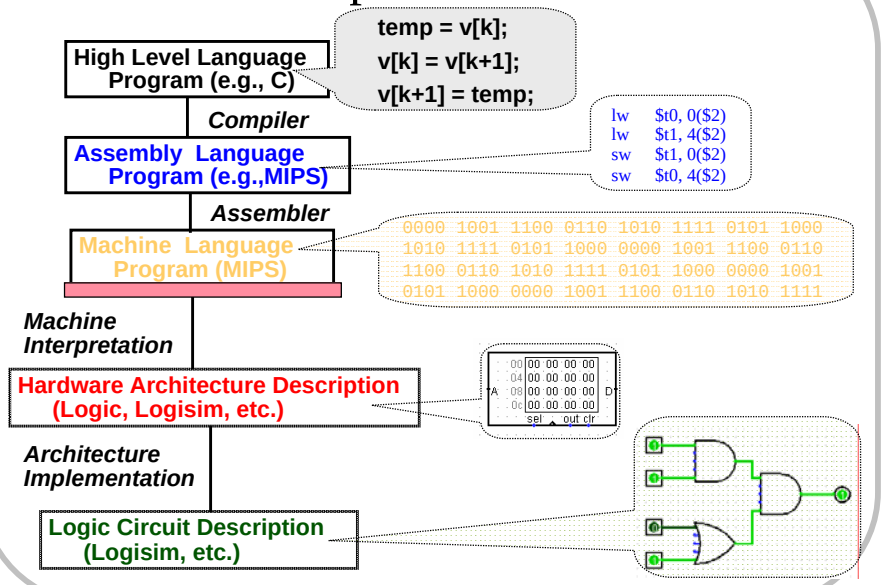

Anatomy: 5 components of any computer

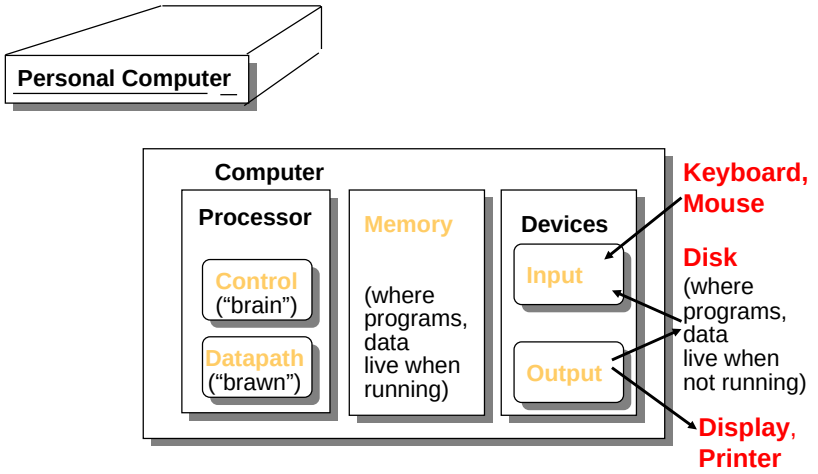

Overview of Physical Implementations

The hardware out of which we make systems.

- Integrated Circuits (ICs)
  - Combinational logic circuits, memory elements, analog interfaces.
- Printed Circuits (PC) boards
  - substrate for ICs and interconnection, distribution of CLK, Vdd, and GND signals, heat dissipation.
- Power Supplies
  - Converts line AC voltage to regulated DC low voltage levels.
- Chassis (rack, card case, ...)
  - holds boards, power supply, provides physical interface to user or other systems.
- Connectors and Cables.

Integrated Circuits (2003 state-of-the-art)

Bare Die

- Primarily Crystalline Silicon
- 1mm - 25mm on a side
- 2003 - feature size ~ 0.13µm = 0.13 x 10^{-6} m
- 100 - 400M transistors
- (25 - 100M "logic gates")
- 3 - 10 conductive layers
- "CMOS" (complementary metal oxide semiconductor) - most common.

Chip in Package

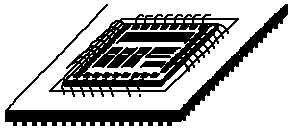

- Package provides:
  - spreading of chip-level signal paths to board-level
  - heat dissipation.
- Ceramic or plastic with gold wires.

Printed Circuit Boards

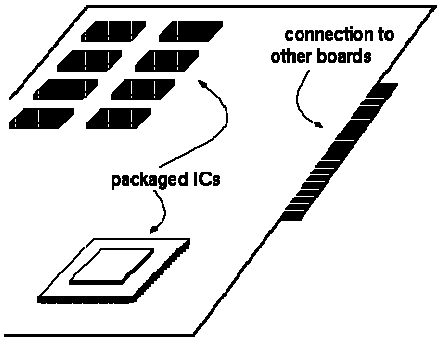

- fiberglass or ceramic
- 1-20 conductive layers
- 1-20in on a side
- IC packages are soldered down.

Technology Trends: Memory Capacity (Single-Chip DRAM)

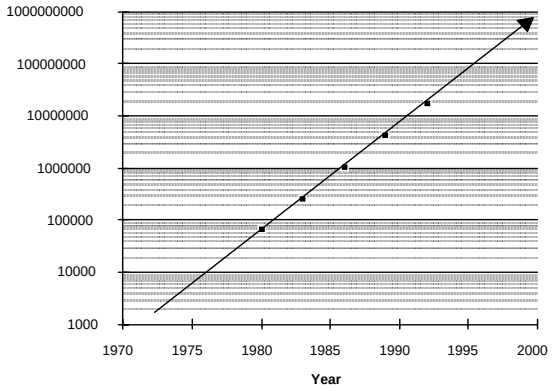

year	size (Mbit)
1980	0.0625
1983	0.25
1986	1
1989	4
1992	16
1996	64
1998	128
2000	256
2002	512

- **Now 1.4X/yr, or 2X every 2 years.**
- **8000X since 1980!**

Technology Trends: Microprocessor Complexity

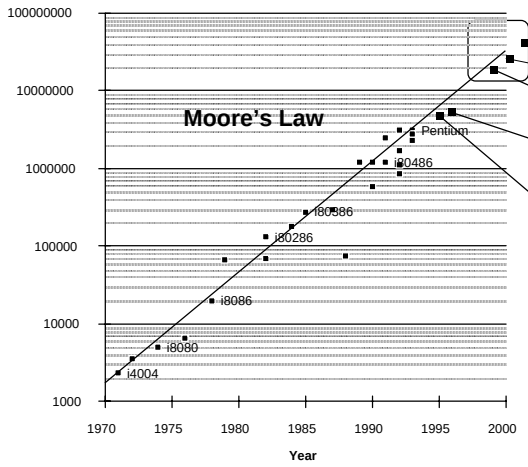

Itanium 2: 41 Million
Athlon (K7): 22 Million
Alpha 21264: 15 million
Pentium Pro: 5.5 million
PowerPC 620: 6.9 million
Alpha 21164: 9.3 million
Sparc Ultra: 5.2 million

2X transistors/Chip Every 1.5 years

Called "Moore's Law"

Technology Trends: Processor Performance

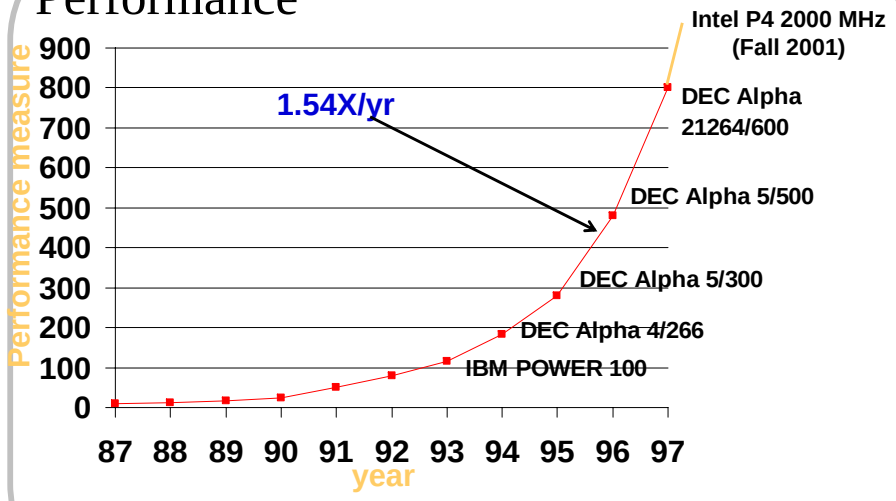

We'll talk about processor performance later on…

Computer Technology - Dramatic Change

- Memory
  - DRAM capacity: 2x / 2 years (since ‘96);
    64x size improvement in last decade.
- Processor
  - Speed 2x / 1.5 years (since ‘85);
    100X performance in last decade.
- Disk
  - Capacity: 2x / 1 year (since ‘97)
    250X size in last decade.

Computer Technology - Dramatic Change

- State-of-the-art PC when you graduate:
  (at least…)
  - Processor clock speed: 5000 MegaHertz (5.0 GigaHertz)
  - Memory capacity: 4000 MegaBytes (4.0 GigaBytes)
  - Disk capacity: 2000 GigaBytes (2.0 TeraBytes)
  - New units! Mega => Giga, Giga => Tera

  (Tera => Peta, Peta => Exa, Exa => Zetta
  Zetta => Yotta = 10^{24})

Levels of Organization

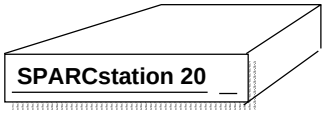

Workstation Design Target:
25% of cost on Processor
25% of cost on Memory
(minimum memory size)
Rest on I/O devices,
power supplies, box

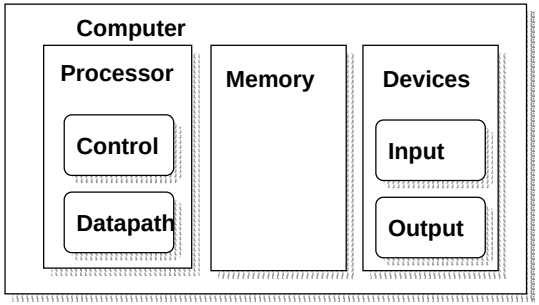

Execution Cycle

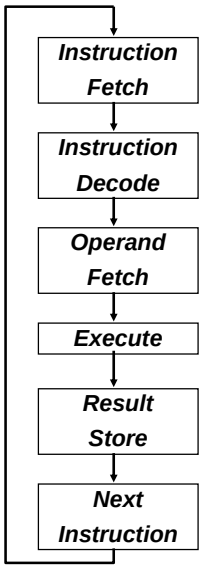

Obtain instruction from program storage

Determine required actions and instruction size

Locate and obtain operand data

Compute result value or status

Deposit results in storage for later use

Determine successor instruction

The SPARCstation 20

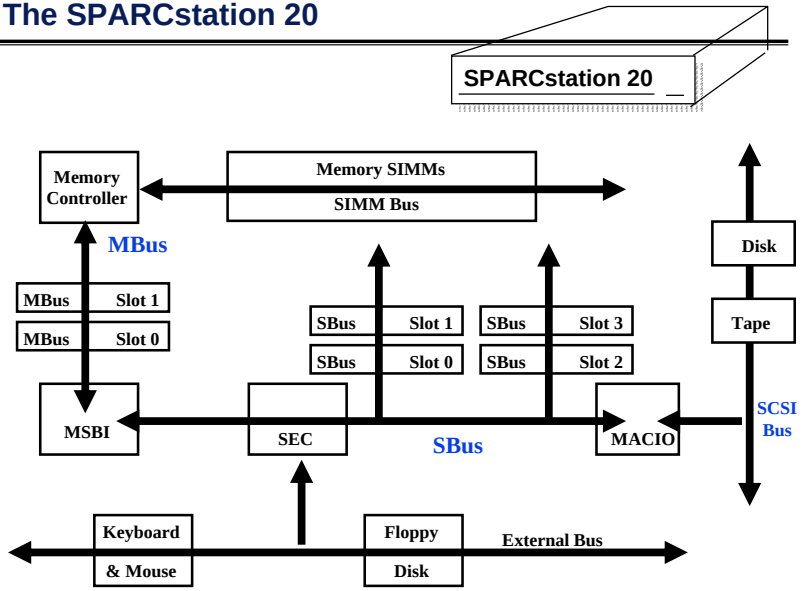

The Underlying Interconnect

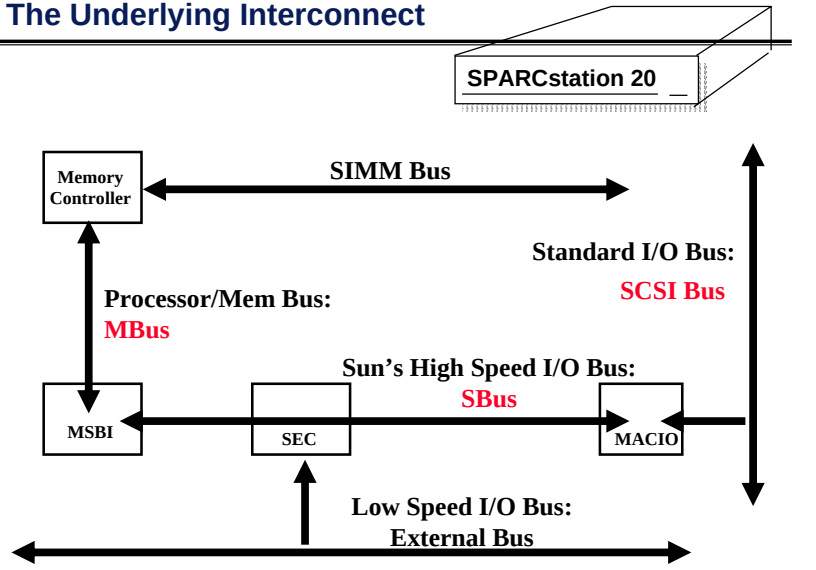

Processor and Caches

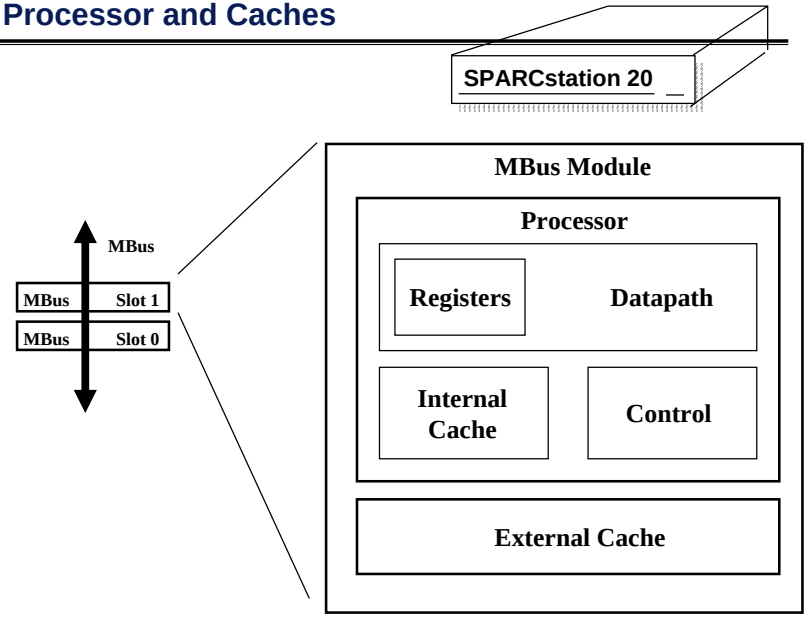

Memory

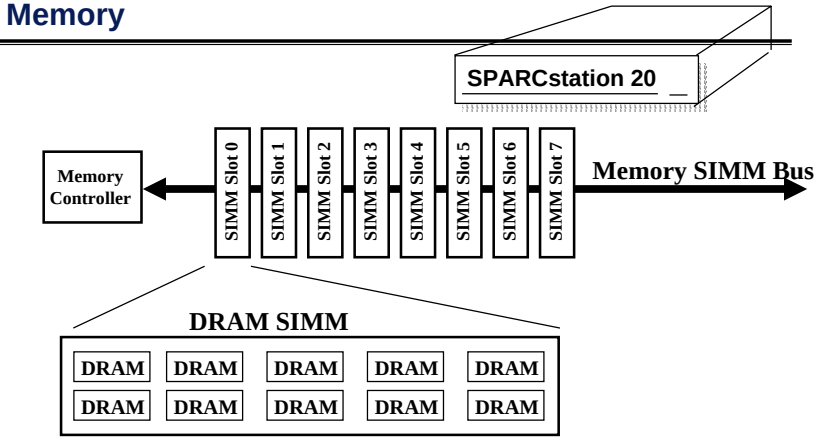

Input and Output (I/O) Devices

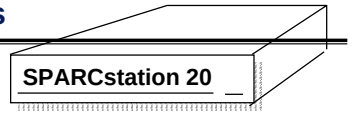

- SCSI Bus: Standard I/O Devices
- SBus: High Speed I/O Devices
- External Bus: Low Speed I/O Device

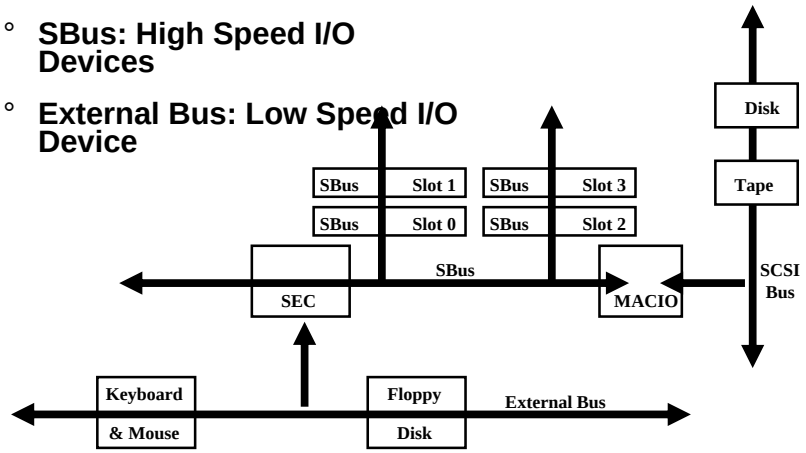

Standard I/O Devices

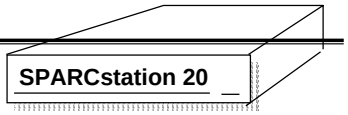

- SCSI = Small Computer Systems Interface
- A standard interface (IBM, Apple, HP, Sun ... etc.)
- Computers and I/O devices communicate with each other
- The hard disk is one I/O device resides on the SCSI Bus

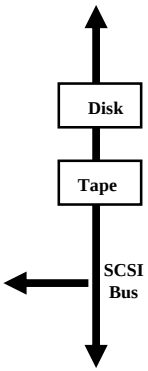

High Speed I/O Devices

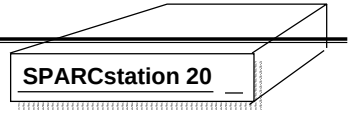

- SBus is SUN's own high speed I/O bus
- SS20 has four SBus slots where we can plug in I/O devices
- Example: graphics accelerator, video adaptor, ... etc.
- High speed and low speed are relative terms

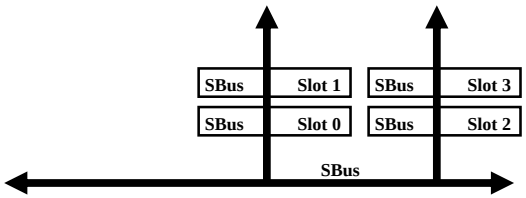

Slow Speed I/O Devices

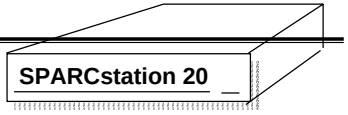

- The are only four SBus slots in SS20--"seats" are expensive
- The speed of some I/O devices is limited by human reaction time--very very slow by computer standard
- Examples: Keyboard and mouse
- No reason to use up one of the expensive SBus slot