

**Computer Architecture**  
(CSC-3501)  
**Lecture 25**  
(24 April 2008)

Seung-Jong Park (Jay)  
<http://www.csc.lsu.edu/~sjpark>

1

CSC3501 - S.J. Park

## Announcement

2

CSC3501 - S.J. Park

## Chapter 9 Objectives

- Learn the properties that often distinguish RISC from CISC architectures.
- Understand how multiprocessor architectures are classified.
- Appreciate the factors that create complexity in multiprocessor systems.
- Become familiar with the ways in which some architectures transcend the traditional von Neumann paradigm.

3

CSC3501 - S.J. Park

## 9.1 Introduction

- We have so far studied only the simplest models of computer systems; classical single-processor von Neumann systems.
- This chapter presents a number of different approaches to computer organization and architecture.
- Some of these approaches are in place in today's commercial systems. Others may form the basis for the computers of tomorrow.

4

CSC3501 - S.J. Park

## 9.2 RISC Machines

- The underlying philosophy of RISC machines is that a system is better able to manage program execution when the program consists of only a few different instructions that are the same length and require the same number of clock cycles to decode and execute.
- RISC systems access memory only with explicit load and store instructions.
- In CISC systems, many different kinds of instructions access memory, making instruction length variable and fetch-decode-execute time unpredictable.

5

CSC3501 - S.J. Park

## 9.2 RISC Machines

- The difference between CISC and RISC becomes evident through the basic computer performance equation:

$$\text{CPU Time} = \frac{\text{seconds}}{\text{program}} = \frac{\text{instructions}}{\text{program}} \times \frac{\text{avg. cycles}}{\text{instruction}} \times \frac{\text{seconds}}{\text{cycle}}$$

- RISC systems shorten execution time by reducing the clock cycles per instruction.
- CISC systems improve performance by reducing the number of instructions per program.

6

CSC3501 - S.J. Park

## 9.2 RISC Machines

- The simple instruction set of RISC machines enables control units to be hardwired for maximum speed.
- The more complex-- and variable-- instruction set of CISC machines requires microcode-based control units that interpret instructions as they are fetched from memory. This translation takes time.
- With fixed-length instructions, RISC lends itself to pipelining and speculative execution.

7

CSC3501 - S.J. Park

## 9.2 RISC Machines

- Consider the the program fragments:
 

<b>CISC</b>	<pre>mov ax, 10 mov bx, 5 mul bx, ax</pre>	<b>RISC</b>	<pre>Begin   add ax, bx   loop Begin</pre>
-------------	--	-------------	--
- The total clock cycles for the CISC version might be:
 
$$(2 \text{ movs} \times 1 \text{ cycle}) + (1 \text{ mul} \times 30 \text{ cycles}) = 32 \text{ cycles}$$
- While the clock cycles for the RISC version is:
 
$$(3 \text{ movs} \times 1 \text{ cycle}) + (5 \text{ adds} \times 1 \text{ cycle}) + (5 \text{ loops} \times 1 \text{ cycle}) = 13 \text{ cycles}$$
- With RISC clock cycle being shorter, RISC gives us much faster execution speeds.

8

CSC3501 - S.J. Park

## 9.2 RISC Machines

- Because of their load-store ISAs, RISC architectures require a large number of CPU registers.
- These register provide fast access to data during sequential program execution.
- They can also be employed to reduce the overhead typically caused by passing parameters to subprograms.
- Instead of pulling parameters off of a stack, the subprogram is directed to use a subset of registers.

9

CSC3501 - S.J. Park

## 9.2 RISC Machines

- It is becoming increasingly difficult to distinguish RISC architectures from CISC architectures.
- Some RISC systems provide more extravagant instruction sets than some CISC systems.
- Some systems combine both approaches.
- The following two slides summarize the characteristics that traditionally typify the differences between these two architectures.

10

CSC3501 - S.J. Park

## 9.2 RISC Machines

<ul style="list-style-type: none"> <li>■ <b>RISC</b> <ul style="list-style-type: none"> <li><input type="checkbox"/> Multiple register sets.</li> <li><input type="checkbox"/> Three operands per instruction.</li> <li><input type="checkbox"/> Parameter passing through register windows.</li> <li><input type="checkbox"/> Single-cycle instructions.</li> <li><input type="checkbox"/> Hardwired control.</li> <li><input type="checkbox"/> Highly pipelined.</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>■ <b>CISC</b> <ul style="list-style-type: none"> <li><input type="checkbox"/> Single register set.</li> <li><input type="checkbox"/> One or two register operands per instruction.</li> <li><input type="checkbox"/> Parameter passing through memory.</li> <li><input type="checkbox"/> Multiple cycle instructions.</li> <li><input type="checkbox"/> Microprogrammed control.</li> <li><input type="checkbox"/> Less pipelined.</li> </ul> </li> </ul>
---	--

Continued....

11

CSC3501 - S.J. Park

## 9.2 RISC Machines

<ul style="list-style-type: none"> <li>■ <b>RISC</b> <ul style="list-style-type: none"> <li><input type="checkbox"/> Simple instructions, few in number.</li> <li><input type="checkbox"/> Fixed length instructions.</li> <li><input type="checkbox"/> Complexity in compiler.</li> <li><input type="checkbox"/> Only <b>LOAD/STORE</b> instructions access memory.</li> <li><input type="checkbox"/> Few addressing modes.</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>■ <b>CISC</b> <ul style="list-style-type: none"> <li><input type="checkbox"/> Many complex instructions.</li> <li><input type="checkbox"/> Variable length instructions.</li> <li><input type="checkbox"/> Complexity in microcode.</li> <li><input type="checkbox"/> Many instructions can access memory.</li> <li><input type="checkbox"/> Many addressing modes.</li> </ul> </li> </ul>
---	---

12

CSC3501 – S.J. Park

### 9.3 Flynn's Taxonomy

- Many attempts have been made to come up with a way to categorize computer architectures.
- *Flynn's Taxonomy* has been the most enduring of these, despite having some limitations.
- Flynn's Taxonomy takes into consideration the number of processors and the number of data paths incorporated into an architecture.
- A machine can have one or many processors that operate on one or many data streams.

13

CSC3501 – S.J. Park

### 9.3 Flynn's Taxonomy

- The four combinations of multiple processors and multiple data paths are described by Flynn as:

Instruction streams	Data streams	Name	Examples
1	1	SISD	Classical Von Neumann machine
1	Multiple	SIMD	Vector supercomputer, array processor
Multiple	1	MISD	Arguably none
Multiple	Multiple	MIMD	Multiprocessor, multicomputer

14

CSC3501 – S.J. Park

### 9.3 Flynn's Taxonomy

- Flynn's Taxonomy falls short in a number of ways:
- First, there appears to be no need for MISD machines.
- Second, parallelism is not homogeneous. This assumption ignores the contribution of specialized processors.
- Third, it provides no straightforward way to distinguish architectures of the MIMD category.
  - One idea is to divide these systems into those that share memory, and those that don't, as well as whether the interconnections are bus-based or switch-based.

15

CSC3501 – S.J. Park

### 9.3 Flynn's Taxonomy

- Symmetric multiprocessors (SMP) and massively parallel processors (MPP) are MIMD architectures that differ in how they use memory.
- SMP systems share the same memory and MPP do not.
- An easy way to distinguish SMP from MPP is:
  - MPP ⇒ many processors + distributed memory + communication via network
  - SMP ⇒ fewer processors + shared memory + communication via memory

16

CSC3501 – S.J. Park

### 9.3 Flynn's Taxonomy

- Other examples of MIMD architectures are found in distributed computing, where processing takes place collaboratively among networked computers.
  - A **network of workstations** (NOW) uses otherwise idle systems to solve a problem.
  - A **collection of workstations** (COW) is a NOW where one workstation coordinates the actions of the others.
  - A **dedicated cluster parallel computer** (DCPC) is a group of workstations brought together to solve a specific problem.
  - A **pile of PCs** (POPC) is a cluster of (usually) heterogeneous systems that form a dedicated parallel system.

17

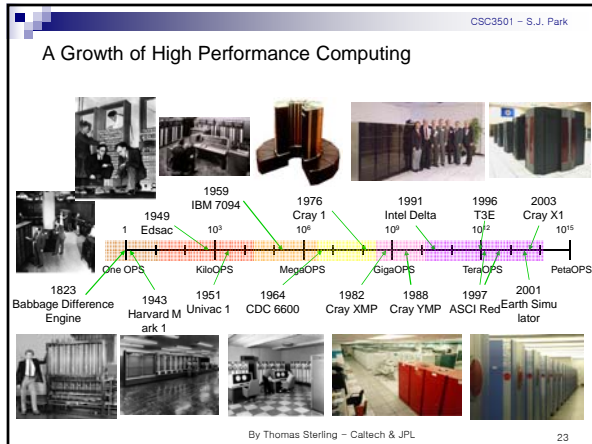
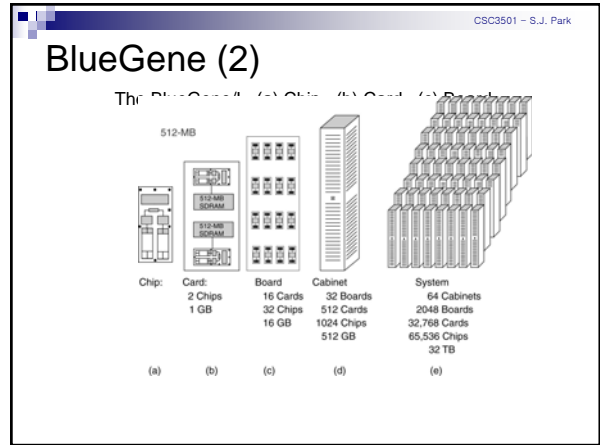
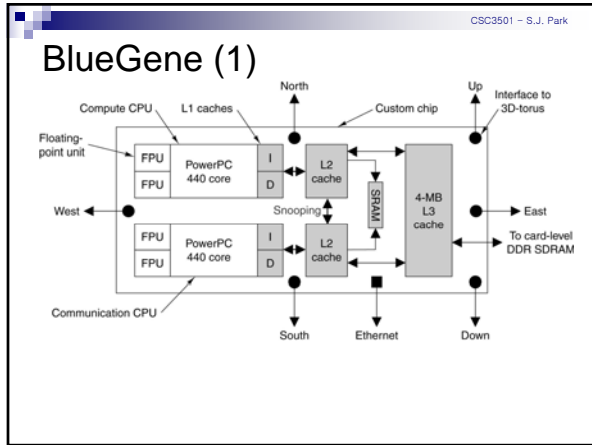
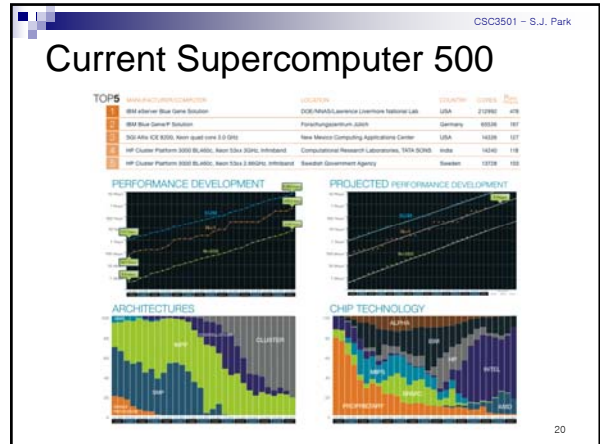
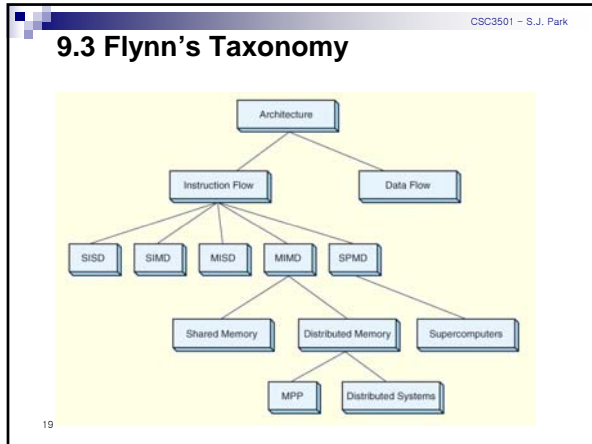
CSC3501 – S.J. Park

### 9.3 Flynn's Taxonomy

- Flynn's Taxonomy has been expanded to include SPMD (single program, multiple data) architectures.
- Each SPMD processor has its own data set and program memory. Different nodes can execute different instructions within the same program using instructions similar to:
  - If myNodeNum = 1 do this, else do that
- Yet another idea missing from Flynn's is whether the architecture is instruction driven or data driven.
 

The next slide provides a revised taxonomy.

18



- CSC3501 - S.J. Park
- ### 9.4 Parallel and Multiprocessor Architectures
- Parallel processing is capable of economically increasing system throughput while providing better fault tolerance.
  - The limiting factor is that no matter how well an algorithm is parallelized, there is always some portion that must be done sequentially.
    - Additional processors sit idle while the sequential work is performed.
  - Thus, it is important to keep in mind that an  $n$ -fold increase in processing power does not necessarily result in an  $n$ -fold increase in throughput.
- 24

CSC3501 - S.J. Park

### 9.4 Superpipelining

- Recall that pipelining divides the fetch-decode-execute cycle into stages that each carry out a small part of the process on a set of instructions.
- Ideally, an instruction exits the pipeline during each tick of the clock.
- Superpipelining** occurs when a pipeline has stages that require less than half a clock cycle to complete.
  - The pipeline is equipped with a separate clock running at a frequency that is at least double that of the main system clock.
- Superpipelining is only one aspect of superscalar design.

25

CSC3501 - S.J. Park

### 9.4 Superscalar (Dynamic multiple-issue processors)

- Superscalar architectures include multiple execution units such as specialized integer and floating-point adders and multipliers.
- A critical component of this architecture is the *instruction fetch unit*, which can simultaneously retrieve several instructions from memory.
- A *decoding unit* determines which of these instructions can be executed in parallel and combines them accordingly.
- This architecture also requires compilers that make optimum use of the hardware.

26

CSC3501 - S.J. Park

### Pentium 4 Block Diagram

AGU = address generation unit  
 BTB = branch target buffer  
 D-TLB = data translation lookaside buffer  
 I-TLB = instruction translation lookaside buffer

27

CSC3501 - S.J. Park

### 9.4 VLIW (Static multiple-issue processors)

- Very long instruction word (VLIW) architectures differ from superscalar architectures because the VLIW compiler, instead of a hardware decoding unit, packs independent instructions into one long instruction that is sent down the pipeline to the execution units.
  - Compiler takes greater responsibility for exploiting parallelism
- One could argue that this is the best approach because the compiler can better identify instruction dependencies.
- However, compilers tend to be conservative and cannot have a view of the run time code.
  - E.g., Intel Itanium and Itanium 2 for the IA-64 ISA - EPIC (Explicit Parallel Instruction Computer)

28

CSC3501 - S.J. Park

### Example of VLIW - IA-64

128 bits (bundle)

Instruction 2 (41 bits)    Instruction 1 (41 bits)    Instruction 0 (41 bits)    Template (5 bits)

Memory (M)    Memory (M)    Integer (I)    (MMI)

29

CSC3501 - S.J. Park

### CISC vs RISC vs SS vs VLIW

	CISC	RISC	Superscalar	VLIW
Instr size	variable size	fixed size	fixed size	fixed size (but large)
Instr format	variable format	fixed format	fixed format	fixed format
Registers	few, some special	many GP	GP and rename (RUU)	many, many GP
Memory reference	embedded in many instr's	load/store	load/store	load/store
Key Issues	decode complexity	data forwarding, hazards	hardware dependency resolution	(compiler) code scheduling
Instruction flow				

30