

# ECE/CS 757: Advanced Computer Architecture II

Instructor: Mikko H Lipasti

Spring 2009  
University of Wisconsin-Madison

Lecture notes based on slides created by John Shen, Mark Hill, David Wood, Guri Sohi, and Jim Smith, Natalie Enright Jerger, and probably others

## Midterm 1 Review

- Introductory material
- Multiprocessor software
- Cores and multithreading
- Multicore
- Coherence
- Consistency

2

## Introduction

- Thread-level parallelism
- Multiprocessor Systems
- Cache Coherence Basics
  - Snoopy
  - Scalable
- Flynn Taxonomy (SISD/SIMD/MIMD)
- UMA vs. NUMA

3

## Introduction - Readings

- Smith ch 1
- Amdahl
- Multicore Amdahl
- Olukotun CMP

4

## MP Software

- Important Multiprocessor Interfaces
  - API
  - ABI
  - ISA
- Programming Models
- Major Abstractions
  - Processes & threads
  - Communication
  - Synchronization
- Shared Memory
  - API description
  - Implementation at ABI, ISA levels
  - ISA support
- Message Passing ( cursory coverage only )
  - API description
  - Implementation at ABI, ISA levels
  - ISA support

01/07

ECE/CS 757; copyright J.E. Smith, 2007

## MP Software - Readings

- Smith Ch. 2
- Lamport
- Sutter/Larus ACM Queue
- Barroso, Commercial Workloads
- MCS Lock/Barrier

6

## Cores (Review of 752)

- ✓ Iron law
- ✓ Beyond pipelining
- ✓ Superscalar challenges
  - ✓ Instruction flow
  - ✓ Register data flow
  - ✓ Memory Dataflow
- ✓ Modern memory interface
- What was not covered
  - Memory hierarchy (caches, DRAM)
  - Virtual memory
  - Power
  - Many implementation/design details

## Multithreading

- Historical Multi-threaded Cores
  - 6600, HEP
- In-Order Multi-threaded Cores
- Out-of-Order Multi-threaded Cores
  - Resource Sharing
  - Thread Scheduling
- Case Studies
  - IBM RS64 IV
  - SUN Niagara
  - IBM Power5

(c) 2007 J. E. Smith

8

## Cores/Multithreading - Readings

- Smith ch. 3
- Marr, Pentium 4 SMT
- Borkenhagen, PowerPC coarse-grained MT

9

## Multicore Summary

- Objective:
  - resource sharing
- Design Issues
  - Where to connect
  - Shared vs private caches
  - Coherence
    - L1, L2 protocols, policies
  - Interconnect
    - Bus, crossbar, ring, mesh

10

## Multicore - Readings

- CMP design space exploration (thermal vs. power)
- Heterogenous CMP
- Piranha
- Multicore CPUs for the masses
- Victim replication

11

## Cache Coherence

- Coherence States
- Snoopy bus-based Invalidate Protocols
- Invalidate protocol optimizations
- Update Protocols (Dragon/Firefly)
- Directory protocols
- Implementation issues

12

## Coherence Readings

- Smith Ch. 4
- Firefly
- Archibald
- Sweazey/Smith
- Natarajan, memory controller study

13

## Consistency Models

- **UNDERSTANDING CONSISTENCY MODELS**
  - Atomicity
  - Program Ordering
  - Visibility/Causality
- **POPULAR CONSISTENCY MODELS**
  - Sequential Consistency
  - IBM/370
  - Processor Consistency
  - SPARC TSO/PSO/RMO
  - Weak Ordering
  - PowerPC Weak Consistency
- **VISIBILITY**
- **MEMORY REFERENCE REORDERING**
- **IMPLEMENTATION ISSUES (not on midterm 1)**

© 2005 Mikko Lijesti

14

## Consistency - Readings

- Adve/Gharachorloo tutorial
- Hill: position paper on simple models

15

## Midterm 1 Review

- Introductory material
- Multiprocessor software
- Cores and multithreading
- Multicore
- Coherence
- Consistency

16