STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

CSE 520
Computer Architecture II

Cache Coherence

Prof. Michel A. Kinsy

1

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Architecture Taxonomy

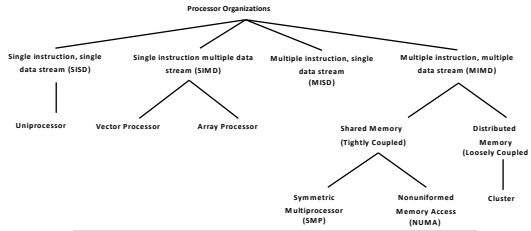

Parallelism Paradigms: Instruction level, Data level and Task level Parallelisms

2

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Architecture Taxonomy

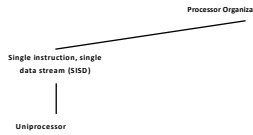

Parallelism Paradigms: Instruction level, Data level and Task level Parallelisms

3

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Parallel Architectures

- Single instruction, single data stream – SISD
  - Single processor
  - Single instruction stream
  - Data stored in single memory

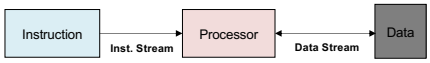

4

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Architecture Taxonomy

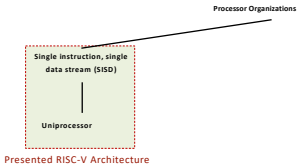

Presented RISC-V Architecture

Parallelism Paradigms: Instruction level, Data level and Task level Parallelisms

5

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Architecture Taxonomy

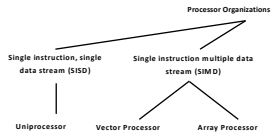

Parallelism Paradigms: Instruction level, Data level and Task level Parallelisms

6

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Parallel Architectures

- Single instruction, multiple data stream – SIMD
- Single machine instruction
  - Each instruction executed on different set of data by different processors
- Number of processing elements
  - Machine controls simultaneous execution
    - Lockstep basis
  - Each processing element has associated data memory
- Application: Vector and array processing

7

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Parallel Architectures

- Single instruction, multiple data stream – SIMD

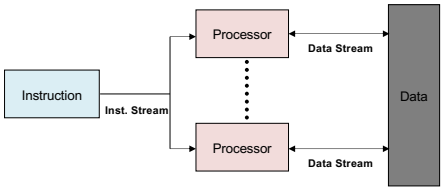

8

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Architecture Taxonomy

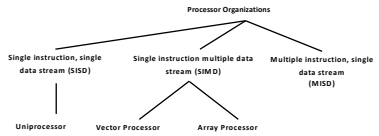

Parallelism Paradigms: Instruction level, Data level and Task level Parallelisms

9

Parallel Architectures

- Multiple instruction, single data stream – MISD
  - Sequence of data
  - Transmitted to set of processors
  - Each processor executes different instruction sequence
- Do not know any implemented case

10

Parallel Architectures

- Multiple instruction, single data stream – MISD

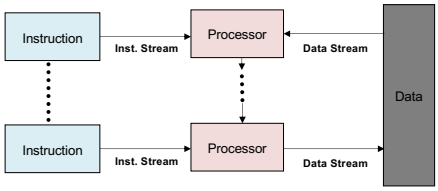

11

Architecture Taxonomy

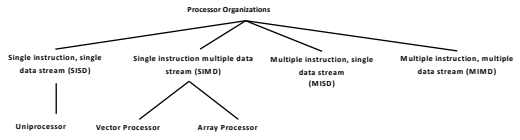

Parallelism Paradigms: Instruction level, Data level and Task level Parallelisms

12

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Parallel Architectures

- Multiple instruction, multiple data stream- MIMD
  - Set of processors
  - Simultaneously executes different instruction sequences
  - Different sets of data
- Examples: SMPs, NUMA systems, and Clusters

13

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Parallel Architectures

- Multiple instruction, multiple data stream- MIMD

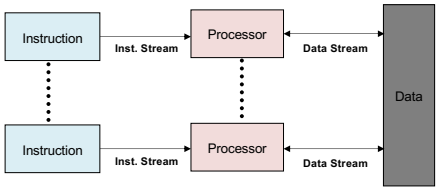

14

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Architecture Taxonomy

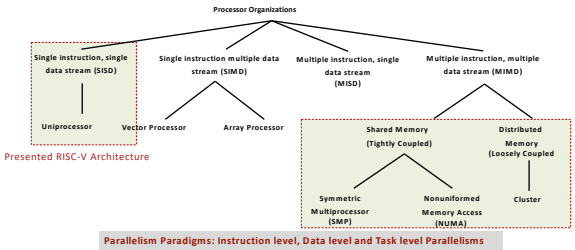

Presented RISC-V Architecture

Parallelism Paradigms: Instruction level, Data level and Task level Parallelisms

15

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Symmetric Multiprocessors (SMP)

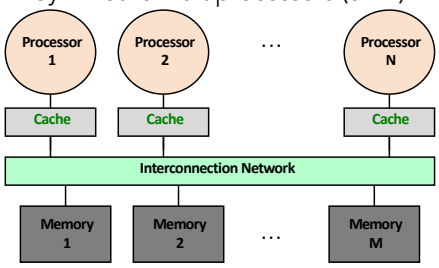

16

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Symmetric Multiprocessors (SMP)

- A collection of processors and a collection of memory connected through some interconnect
- Symmetric because latency for any processor to access any memory is constant – uniform memory access (UMA)

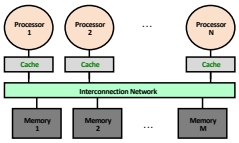

17

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Shared Memory Architectures

- Key differentiating feature: the address space is shared, i.e., any processor can directly address any memory location and access them with load/store instructions

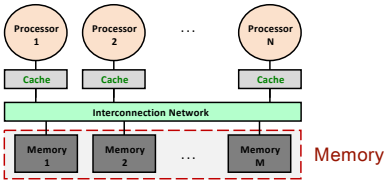

18

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Distributed Memory Multiprocessors

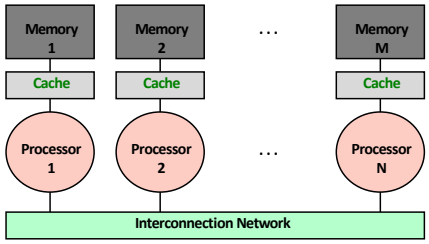

19

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Distributed Memory Multiprocessors

- Each processor has local memory that is accessible through a fast interconnect
- The different nodes are connected as I/O devices with (potentially) slower interconnect
- Local memory access is a lot faster than remote memory
  - Nonuniform memory access (NUMA)

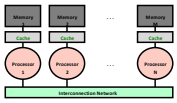

20

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Parallel Architectures

- The parallel computers are classified as
  - Shared memory
  - Distributed memory
- Both shared and distributed memory systems have:
  - Processors: now generally commodity processors
  - Memory: now general commodity DRAM/DDR
  - Network/interconnect: between the processors or memory

21

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Memory Consistency in SMPs

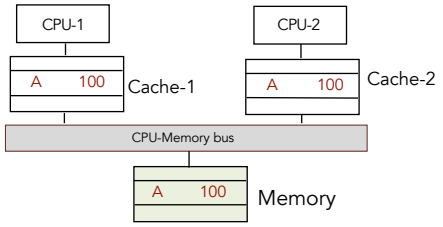

22

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Cache Reads & Writes

- Cache reads
  - With a hit, the data can be retrieved
  - On a miss, we need to take a set of actions
- Cache writes
  - On a hit
    - Write-Through: Every time we write to a variable in cache, we also need to update that variable in memory
    - Write-Back: When we write to a variable in the cache, the modified copy is written back to memory when it gets evicted or replaced in the cache.
      - Write-back can improve performance, but is more complex to implement
        - Dirty Bit: In a write-back cache, cache blocks have a dirty bit
        - Block Replacement: Block is written back to memory if it is dirty and replaced with a new memory block

23

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Cache Reads & Writes

- Cache reads
  - With a hit, the data can be retrieved
  - On a miss, we need to take a set of actions
- Cache writes
  - On a miss
    - Write Allocate - i.e., Fetch on Write
      - Load block into memory as on a read miss
      - Perform Write Hit actions
    - No Write Allocate - i.e., Write Around
      - Modify block in lower memory directly without loading into cache
- Relation of write miss policies to write hit policies
  - Common pairings
    - Write Through with No Write Allocate
    - Write Back with Write Allocate

24

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS — ASU Ira A. Fulton Schools of Engineering, Arizona State University

Memory Consistency in SMPs

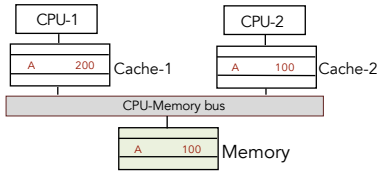

- Suppose CPU-1 updates A to 200
  - Write-back: memory and cache-2 have stale values
  - Write-through: cache-2 has a stale value

25

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS — ASU Ira A. Fulton Schools of Engineering, Arizona State University

Memory Consistency in SMPs

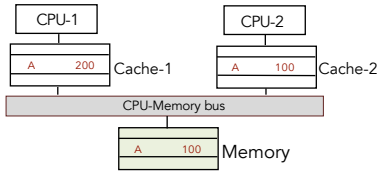

- Suppose CPU-1 updates A to 200
  - Do these stale values matter?
  - What is the view of shared memory for programming?

26

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS — ASU Ira A. Fulton Schools of Engineering, Arizona State University

Write-through Caches & Sequential Consistency

- T1 is executed

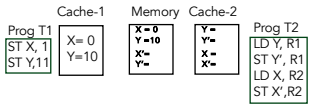

27

Write-through Caches & Sequential Consistency

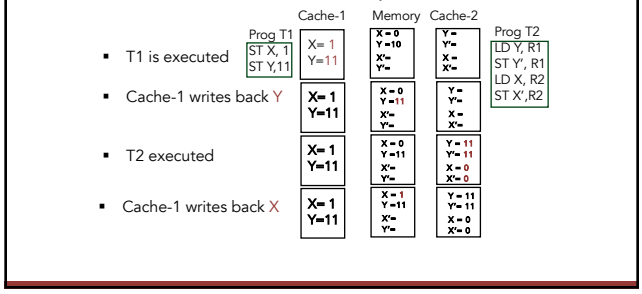

28

Write-through Caches & Sequential Consistency

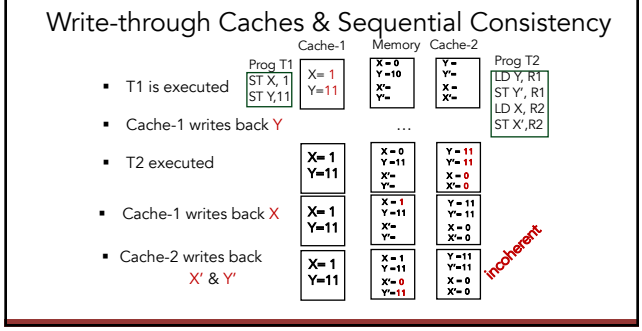

29

Write-through Caches & Sequential Consistency

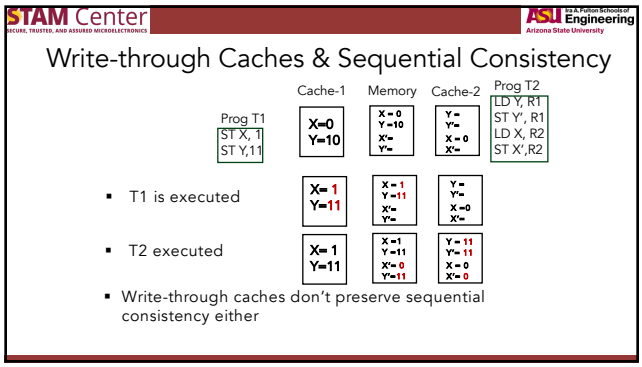

30

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Maintaining Sequential Consistency

- Problem: Multiple copies of a location in various caches can cause SC to break down.
- Hardware support is required such that
  - Only one processor at a time has write permission for a location
  - No processor can load a stale copy of the location after a write
    - Cache coherence protocols

31

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

A System with Multiple Caches

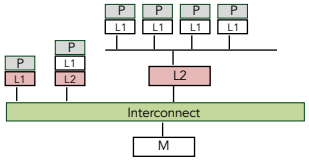

- Modern systems often have hierarchical caches
- Each cache has exactly one parent but can have zero or more children
- Only a parent and its children can communicate directly

32

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

A System with Multiple Caches

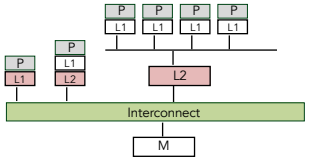

- Inclusion property is maintained between a parent and its children, i.e.,
  - a in L_i implies that a in L_{i+1}

33

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Cache Coherence Protocols for SC

- Write request:
  - The address is invalidated in all other caches before the write is performed, or
  - The address is updated in all other caches after the write is performed
- Read request:
  - If a dirty copy is found in some cache, a write-back is performed before the memory is read
  - We will focus on Invalidation protocols as opposed to Update protocols

34

STAM Center — SECURE, TRUSTED, AND ASSURED MICROELECTRONICS
ASU Ira A. Fulton Schools of Engineering — Arizona State University

Next Learning Module

- Cache Coherence Protocols

35