Fiscal Year 2024

Ver. 2025-01-30a

Course number: CSC.T433 School of Computing, Graduate major in Computer Science

Advanced Computer Architecture

12. Thread Level Parallelism: Cache Coherence Protocol

www.arch.cs.titech.ac.jp/lecture/ACA/ Room No. W8E-308, Lecture (Face-to-face) Mon 13:30-15:10, Thr 13:30-15:10

Kenji Kise, Department of Computer Science kise _at_ c.titech.ac.jp

Key components of many-core processors

- Interconnection network
 - connecting many modules on a chip achieving high throughput and low latency
- Main memory and caches
 - Caches are used to reduce latency and to lower network traffic
 - A parallel program has private data and shared data
 - New issues are cache coherence and memory consistency
- Core
 - High-performance superscalar processor providing a hardware mechanism to support thread synchronization (lock, unlock, barrier)



Shared memory many-core architecture

Four-Way Set Associative Cache

One word/block, 2⁸ = 256 sets where each with four ways (each with one block) • 13 12 11 ... 31 30 ^{2 1 0} / Byte offset . . . _22 8 Tag Index Index V Tag V Tag Tag V Tag Data V Data Data Data 2 253 253 253 253 254 254 254 254 255 255 255 255 _ 32 4x1 select CSC.T433 Advanced Computer Architecture, Department of Computer Science, Science Tokyo Data

Cache writing policy

- Write-through
 - writing is done synchronously both to the cache and to the main memory. All stores update the main memory and memory bandwidth becomes a performance bottleneck.
- Write-back
 - initially, writing is done only to the cache. The write to the main memory is postponed until the modified content is about to be replaced by another cache block.
 - reduces the required network and memory bandwidth.
 - preferable for manycore.
- we assume the use of write-back



Shared memory many-core architecture

Cache coherence problem

- In this example behavior of five events,
- Cores see different values for shared data **u** after event 3
- With write-back caches, value written back to memory depends on which cache line flushes or writes back
 - main memory may see stale (out-of-date) value for a long time
- Unacceptable for programming, and its frequent!



Cache coherence problem

- Cores may see different values through their caches
 - assuming a write-back cache
 - after the value 7 of u has been written by core A, core A's cache contains the new value, but core C's cache and the main memory do not

| Time | Event | Cache contents for core A | Cache contents for core C | Main memo contents | ory for U |
|------|-----------------------------|------------------------------|------------------------------|-----------------------|--------------|
| 0 | | | | 5 | - |
| 1 | Core A reads <mark>u</mark> | 5 | | 5 | - |
| 2 | Core C reads <mark>u</mark> | 5 | 5 | 5 | - |
| 3 | Core A stores 7 | 7 | 5 | 5 | |
| | into <mark>u</mark> | | | | _ |



Cache coherence and enforcing coherence

- Cache coherence
 - All reads by any core must return the most recently written value
 - Writes to the same location by any two cores are seen in the same order by all cores
- Cache coherence protocols
 - (1) Snooping (write invalidate / write update)
 - Each cache tracks sharing status of each cache line
 - (2) Directory based
 - Sharing status of each cache line kept in one location

- Write invalidate
 - On write, invalidate all other copies by an invalidate broadcast
 - Use bus itself to serialize
 - Write cannot complete until bus access is obtained

| Processor activity | Bus activity | Contents of core A's cache | 1 | Contents of core B's cache | Contents of main memory location <mark>u</mark> |
|--|-------------------------------|-------------------------------|---|-------------------------------|--|
| | | | | | 5 |
| Core A reads u | Cache miss for u | 5 | | | 5 |
| Core B reads u | Cache miss for u | 5 | | 5 | 5 |
| Core A writes a 7 to <mark>u</mark> | Invalidation for | u 7 | | | 5 |
| Core B reads u | Cache miss for <mark>u</mark> | 7 | | 7 | 7 |

5

- Write update
 - On write, update all copies

Bus Network

- N cores (), N switch (), 1 link (the bus)
- Only 1 simultaneous transfer at a time
 - NB (best case) = link (bus) bandwidth x 1
 - BB (worst case) = link (bus) bandwidth x 1
- All processors can snoop the bus





The case where core B sends a packet to someone



Bus Network with multiplexer (mux)

- One N-input multiplexer for N cores
- Arbitration, node ID, centralized control







• A write invalidate, cache coherence protocol for a private write-back cache showing the states and state transitions for each block in the cache



MSI (Modified, Shared, Invalid) protocol

Cache miss and the addressed block is invalid

- Core A
 - Source: Core
 - State: Invalid
 - Request: Read miss (u)
 - Function: Place read miss on bus

Event: Core A reads u



Cache miss and the addressed block is invalid

- Core B
 - Source: Core
 - State: Invalid
 - Request: Read miss (u)
 - Function: Place read miss on bus

Event: Core B reads u



Cache miss and the addressed block is invalid

- Core D
 - Source: Core
 - State: Invalid
 - Request: Read miss (u)
 - Function: Place read miss on bus

Event: Core D reads u



The coherence mechanism of a private cache (using word processor for core).

| | Request | Source | State of addressed cache block | Type of cache action | Function and explanation |
|------|------------|-----------|--------------------------------------|-------------------------|---|
| | Read hit | Processor | Shared or modified | Normal hit | Read data in local cache. |
| | Read miss | Processor | Invalid | Normal miss | Place read miss on bus. |
| | Read miss | Processor | Shared | Replacement | Address conflict miss: place read miss on bus. |
| | Read miss | Processor | Modified | Replacement | Address conflict miss: write-back block, then place read miss on bus. |
| | Write hit | Processor | Modified | Normal hit | Write data in local cache. |
| Coh1 | | | | | |
| | Write miss | Processor | Invalid | Normal miss | Place write miss on bus. |
| | Write miss | Processor | Shared | Replacement | Address conflict miss: place write miss on bus. |
| | Write miss | Processor | Modified | Replacement | Address conflict miss: write-back block, then place write miss on bus. |
| | Read miss | Bus | Shared | No action | Allow shared cache or memory to service read miss. |
| Coh2 | Read miss | Bus | Modified | Coherence | Attempt to share data: place cache block on bus and change state to shared. |
| Coh3 | | | | | |
| Coh4 | Write miss | Bus | Shared | Coherence | Attempt to write shared block; invalidate the cache block. |
| Coh5 | Write miss | Bus | Modified | Coherence | Attempt to write block that is exclusive elsewhere; write-back the cache block and make its state invalid in the local cache. |

CSC.T433 Advanced Computer Architecture, Department of Computer Science, Science Tokyo

Exercise 1

- Coh1 (Core A)
 - Source: Core
 - State: Shared
 - Request: Write hit (u)
 - Function: Place invalidate on bus

- Coh3 (Core B, D)
 - Source: Bus
 - State: Shared
 - Request: Invalidate
 - Function: attempt to write shared block; invalidate the block

Event:

Core A writes u



Draw the behavior of this request

u=5

Coherence 1 (Coh1) and Coherence3 (Coh3)

- Coh1 (Core A) ٠
 - Source: Core
 - State: Shared
 - Request: Write hit (u) ٠
 - Function: Place invalidate on bus •

- Coh3 (Core B, D) •
 - Source: Bus
 - State: Shared
 - Request: Invalidate
 - Function: attempt to write shared block; • invalidate the block

Event:

Core A writes u



• The coherence mechanism of a private cache (using word processor for core).

| | Request Read hit | Source Processor | State of addressed cache block Shared or modified | Type of cache action Normal hit | Function and explanation Read data in local cache. |
|------|---------------------|---------------------|---|---------------------------------------|--|
| | | | niounicu | | |
| | Write hit | Processor | Modified | Normal hit | Write data in local cache. |
| Coh1 | Write hit | Processor | Shared | Coherence | Place invalidate on bus. These operations are often called upgrade or <i>ownership</i> misses, since they do not fetch the data but only change the state. |
| | Write miss | Processor | Invalid | Normal miss | Place write miss on bus. |
| | Write miss | Processor | Shared | Replacement | Address conflict miss: place write miss on bus. |
| | Write miss | Processor | Modified | Replacement | Address conflict miss: write-back block, then place write miss on bus. |
| | Read miss | Bus | Shared | No action | Allow shared cache or memory to service read miss. |
| Coh2 | | | | | |
| Coh3 | Invalidate | Bus | Shared | Coherence | Attempt to write shared block; invalidate the block. |
| Coh4 | Write miss | Bus | Shared | Coherence | Attempt to write shared block; invalidate the cache block. |
| Coh5 | Write miss | Bus | Modified | Coherence | Attempt to write block that is exclusive elsewhere; write-back the cache block and make its state invalid in the local cache. |

Coherence 2 (Coh2)

Core C

- Source: Core
- State: Invalid
- Request: Read miss (u)
- Function: Place read miss on bus

- Coh2 (Core A)
 - Source: Bus
 - State: Modified
 - Request: Read miss (u)
 - Function: attempt to shared data; place cache block on bus and change state to shared

Event:

Core C reads u



• The coherence mechanism of a private cache

| niss on bus. |
|---|
| lock, then place read miss on |
| |
| operations are often called they do not fetch the data but |
| |
| service read miss. |
| block on bus and change state |
| lidate the block. |
| |
| ive elsewhere; write-back the lid in the local cache. |
| |

Coherence 4 (Coh4)

Core B

- Source: Core
- State: Invalid
- Request: Write miss (u)
- Function: Place write miss on bus

- Coh4 (Core A, C)
 - Source: Bus
 - State: Shared
 - Request: Write miss (u)
 - Function: attempt to write shared block; invalidate the cache block

Event:

Core B writes u

21



• The coherence mechanism of a private cache

| | Request | Source | State of addressed cache block | Type of cache action | Function and explanation |
|--------|------------|-----------|--------------------------------------|-------------------------|--|
| | Read hit | Processor | Shared or modified | Normal hit | Read data in local cache. |
| | Read miss | Processor | Invalid | Normal miss | Place read miss on bus. |
| | Read miss | Processor | Shared | Replacement | Address conflict miss: place read miss on bus. |
| | Read miss | Processor | Modified | Replacement | Address conflict miss: write-back block, then place read miss on bus. |
| | Write hit | Processor | Modified | Normal hit | Write data in local cache. |
| Coh1 | Write hit | Processor | Shared | Coherence | Place invalidate on bus. These operations are often called upgrade or <i>ownership</i> misses, since they do not fetch the data but only change the state. |
| | | | | | |
| | Read miss | Bus | Shared | No action | Allow shared cache or memory to service read miss. |
| Coh2 | Read miss | Bus | Modified | Coherence | Attempt to share data: place cache block on bus and change state to shared. |
| Coh3 | Invalidate | Bus | Shared | Coherence | Attempt to write shared block; invalidate the block. |
| Coh4 | Write miss | Bus | Shared | Coherence | Attempt to write shared block; invalidate the cache block. |
| o Coh5 | | | | | |

• A write invalidate, cache coherence protocol for a private write-back cache showing the states and state transitions for each block in the cache



MSI (Modified, Shared, Invalid) protocol

- The basic coherence protocol
 - MSI (Modified, Shared, Invalid) protocol
- Extensions
 - MESI (Modified, Exclusive, Shared, Invalid) protocol
 - MOESI (MESI + Owned) protocol

Intel Single-Chip Cloud Computer (2009)

• To research multi-core processors and parallel processing.





A many-core architecture with 2D Mesh NoC



Intel Single-Chip Cloud Computer (48 Core)

Directory protocols

- Snooping coherence protocols are based on the use of bus network.
 What are the protocols for mesh topology NoC?
- Directory protocols
 - A logically-central directory keeps track of where the copies of each cache block reside.

Caches consult this directory to ensure coherence.





Snooping coherence protocol and one with directory



Two caches of different block sizes

- Temporal Locality (Locality in Time): ٠
 - Keep most recently accessed data items closer to the processor
- Spatial Locality (Locality in Space) •
 - Move blocks consisting of contiguous words to the upper levels



Coherence influences the cache miss rate

- Coherence misses
 - True sharing misses
 - False sharing misses
 - When shared data u and unshared data w are allocated in the same block, they are both treated as shared data.

Senario: A reads u -> B reads u -> B reads w -> A writes u -> B reads w



Key components of many-core processors

- Interconnection network
 - connecting many modules on a chip achieving high throughput and low latency
- Main memory and caches
 - Caches are used to reduce latency and to lower network traffic
 - A parallel program has private data and shared data
 - New issues are cache coherence and memory consistency
- Core
 - High-performance superscalar processor providing a hardware mechanism to support thread synchronization (lock, unlock, barrier)



Shared memory many-core architecture

• A write invalidate, cache coherence protocol for a private write-back cache showing the states and state transitions for each block in the cache



MSI (Modified, Shared, Invalid) protocol

• The coherence mechanism of a private cache (using word processor for core).

| | Request | Source | State of addressed cache block | Type of cache action | Function and explanation | |
|------|------------|-----------|--------------------------------------|-------------------------|--|--|
| | Read hit | Processor | Shared or modified | Normal hit | Read data in local cache. | |
| | Read miss | Processor | Invalid | Normal miss | Place read miss on bus. | |
| | Read miss | Processor | Shared | Replacement | Address conflict miss: place read miss on bus. | |
| | Read miss | Processor | Modified | Replacement | Address conflict miss: write-back block, then place read miss on bus. | |
| Coh1 | Write hit | Processor | Modified | Normal hit | Write data in local cache. | |
| | Write hit | Processor | Shared | Coherence | Place invalidate on bus. These operations are often called upgrade or <i>ownership</i> misses, since they do not fetch the data but only change the state. | |
| | Write miss | Processor | Invalid | Normal miss | Place write miss on bus. | |
| | Write miss | Processor | Shared | Replacement | Address conflict miss: place write miss on bus. | |
| | Write miss | Processor | Modified | Replacement | Address conflict miss: write-back block, then place write miss on bus. | |
| | Read miss | Bus | Shared | No action | Allow shared cache or memory to service read miss. | |
| Coh2 | Read miss | Bus | Modified | Coherence | Attempt to share data: place cache block on bus and change state to shared. | |
| Coh3 | Invalidate | Bus | Shared | Coherence | Attempt to write shared block; invalidate the block. | |
| Coh4 | Write miss | Bus | Shared | Coherence | Attempt to write shared block; invalidate the cache block. | |
| Coh5 | Write miss | Bus | Modified | Coherence | Attempt to write block that is exclusive elsewhere; write-back the cache block and make its state invalid in the local cache. | |

Orchestration

