

Fiscal Year 2025

Ver. 2026-01-22a



Course number: CSC.T440  
School of Computing,  
Graduate major in Computer Science

# Computer Organization and Architecture

## 6. Thread Level Parallelism: Coherence and Synchronization

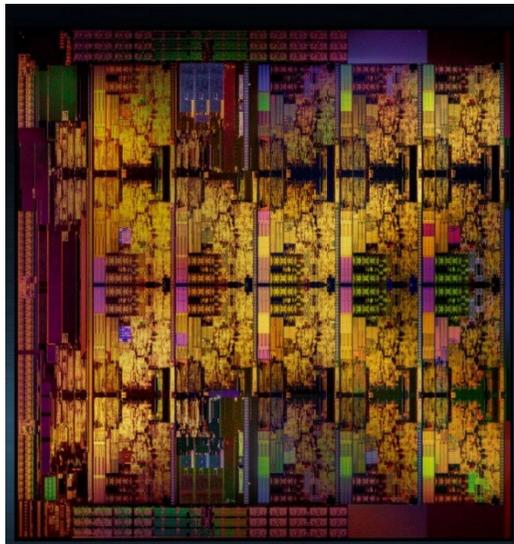


[www.arch.cs.titech.ac.jp/lecture/coa/](http://www.arch.cs.titech.ac.jp/lecture/coa/)  
Room No. M-112(H117), Lecture (Face-to-face)  
Thr 13:30-15:10

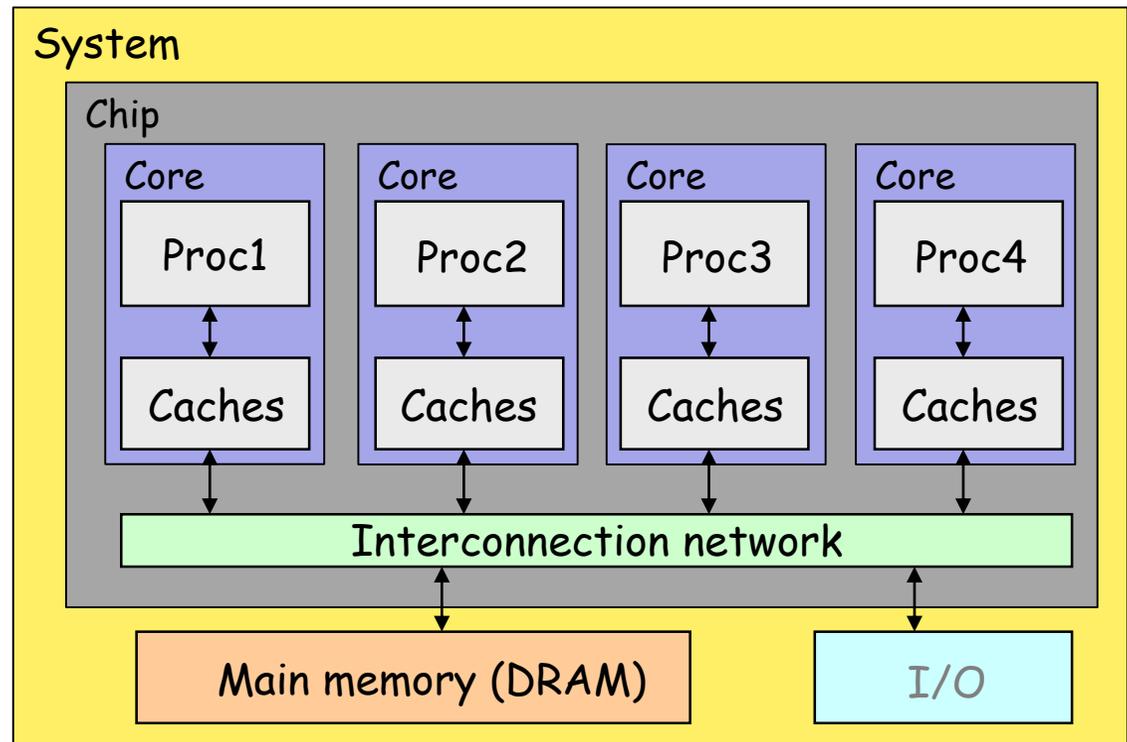
Kenji Kise, Department of Computer Science  
kise[at]comp.isct.ac.jp

# Shared memory many-core architecture

- A single-chip integrates many cores (conventional processors) and an interconnection network.
- The shared memory and **shared address space (SAS)** are used as a means for communication between these cores.

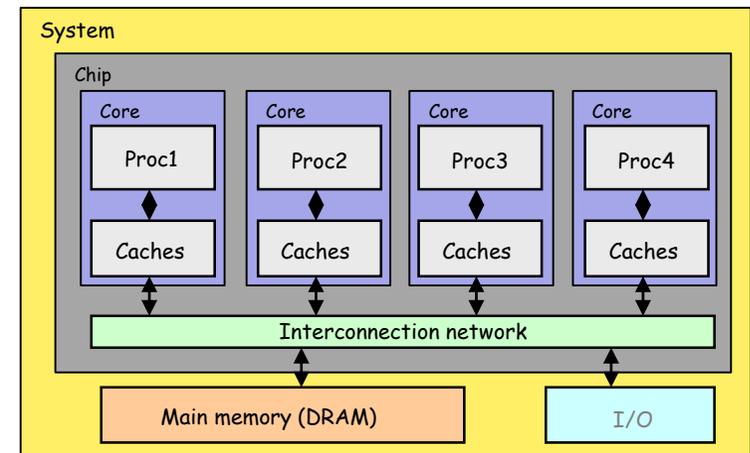


Intel Skylake-X, Core i9-7980XE, 2017



# Key components of many-core processors

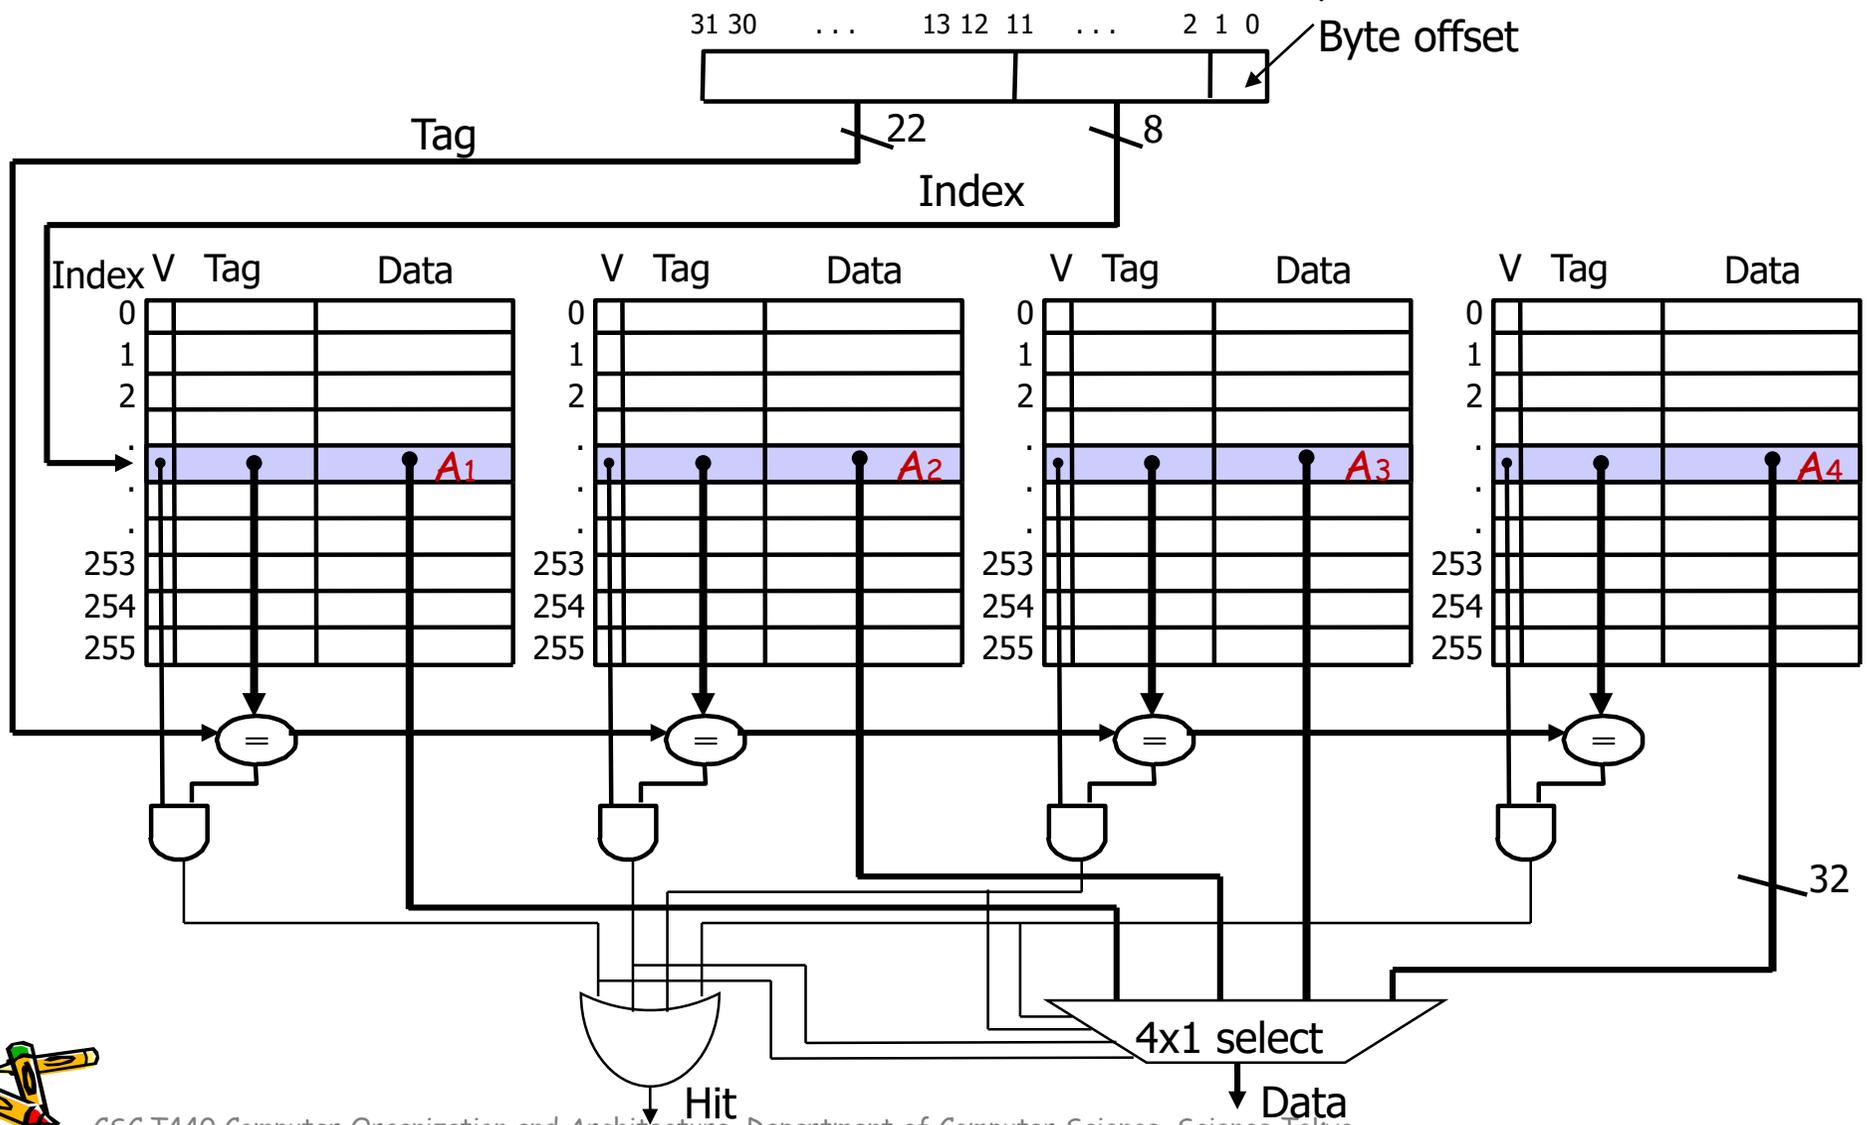
- **Interconnection network**
  - connecting many modules on a chip achieving high throughput and low latency
- **Main memory and caches**
  - Caches are used to reduce latency and to lower network traffic
  - A parallel program has private data and shared data
  - New issues are **cache coherence**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



Shared memory many-core architecture

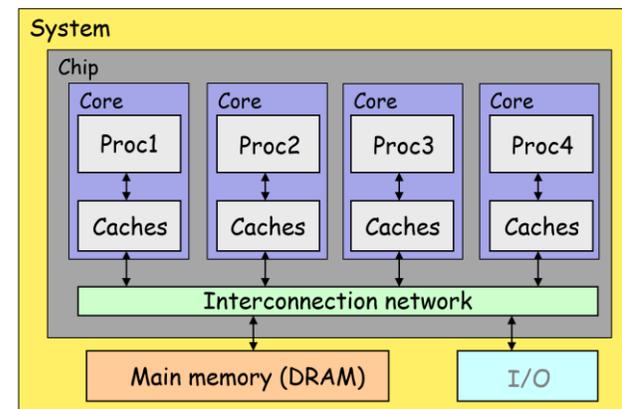
# Four-Way Set Associative Cache

- One word/block,  $2^8 = 256$  sets where each with four ways (each with one block)



# Cache writing policy

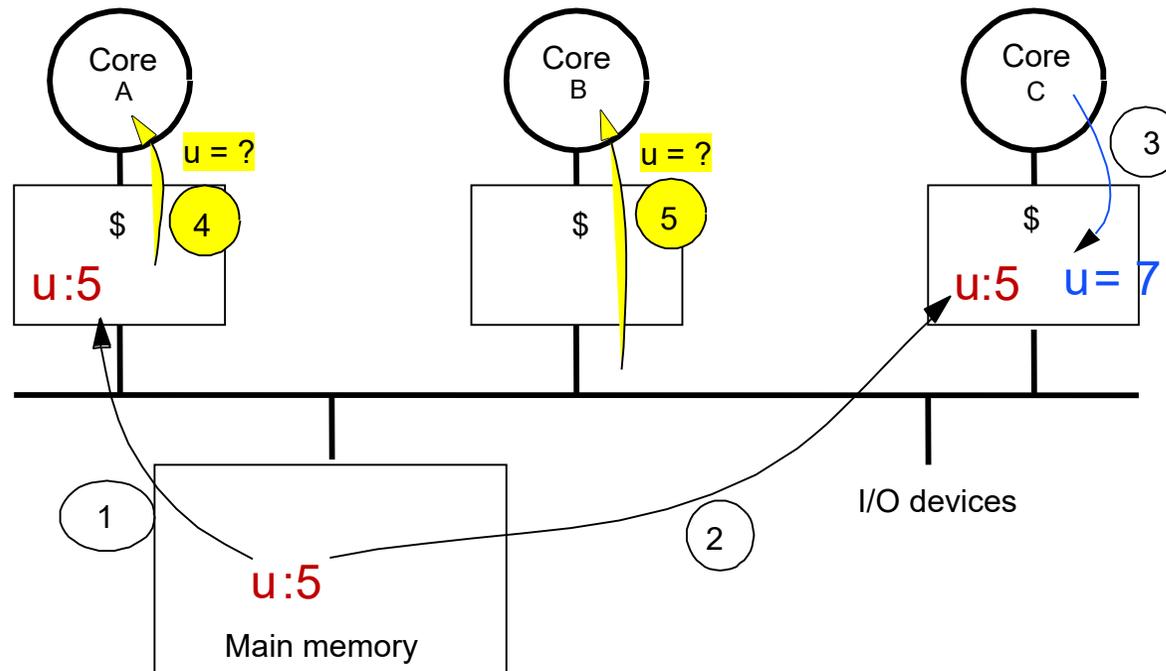
- **Write-through**
  - writing is done synchronously both to the cache and to the main memory. All stores update the main memory and **memory bandwidth becomes a performance bottleneck.**
- **Write-back**
  - initially, writing is done only to the cache. The write to the main memory is postponed until 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 (1/2)

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



# Cache coherence problem (2/2)

- 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 memory contents for <code>u</code>
0				5
1	Core A reads <code>u</code>	5		5
2	Core C reads <code>u</code>	5	5	5
3	Core A stores 7 into <code>u</code>	7	5	5



# 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



# Snooping coherence protocols using bus network

- 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	Contents of core B's cache	Contents of main memory location u
				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 u	Invalidation for u	7		5
Core B reads u	Cache miss for u	7	7	7
				5

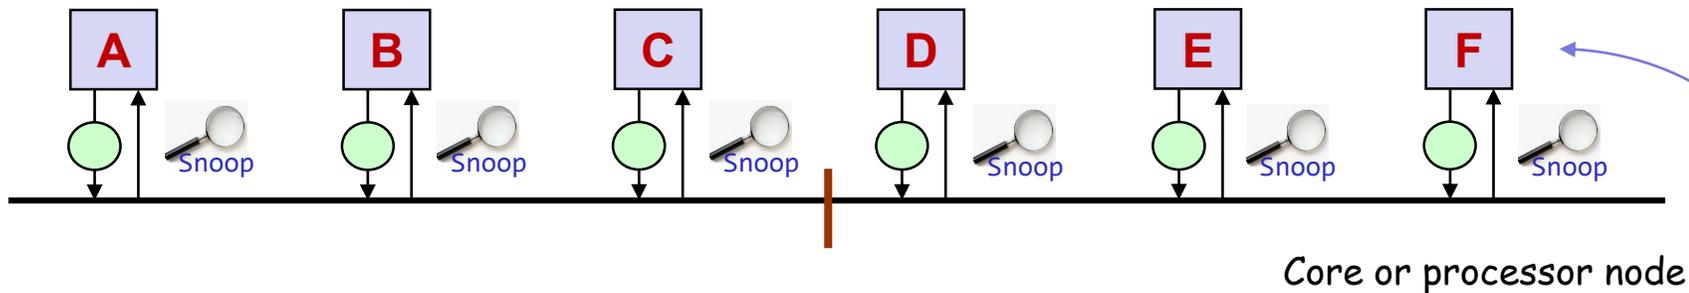
- Write update

- On write, update all copies

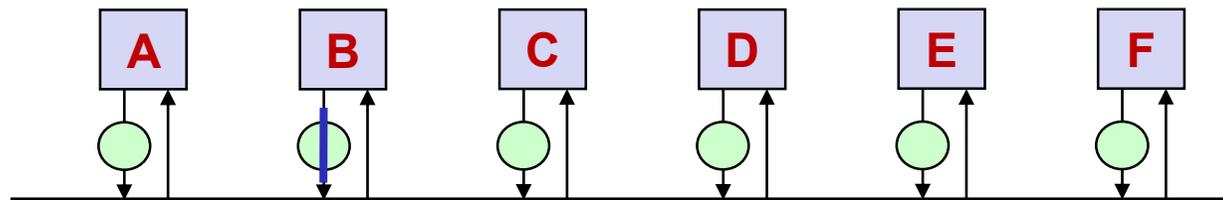


# Bus Network

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

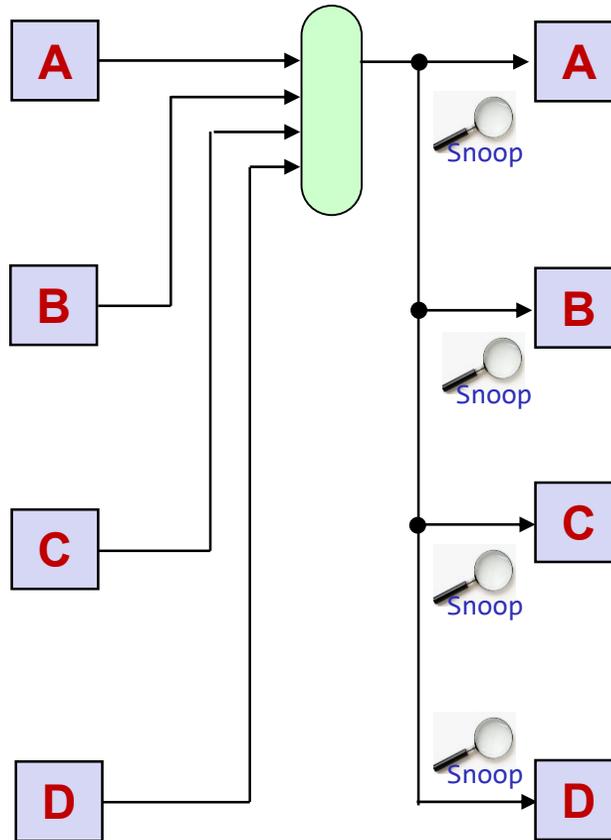


The case where **core B** sends a packet to someone



# Bus Network with multiplexer (mux)

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



The bus network organization of 4 cores using a 4-input mux.



# Snooping coherence protocols using bus network

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

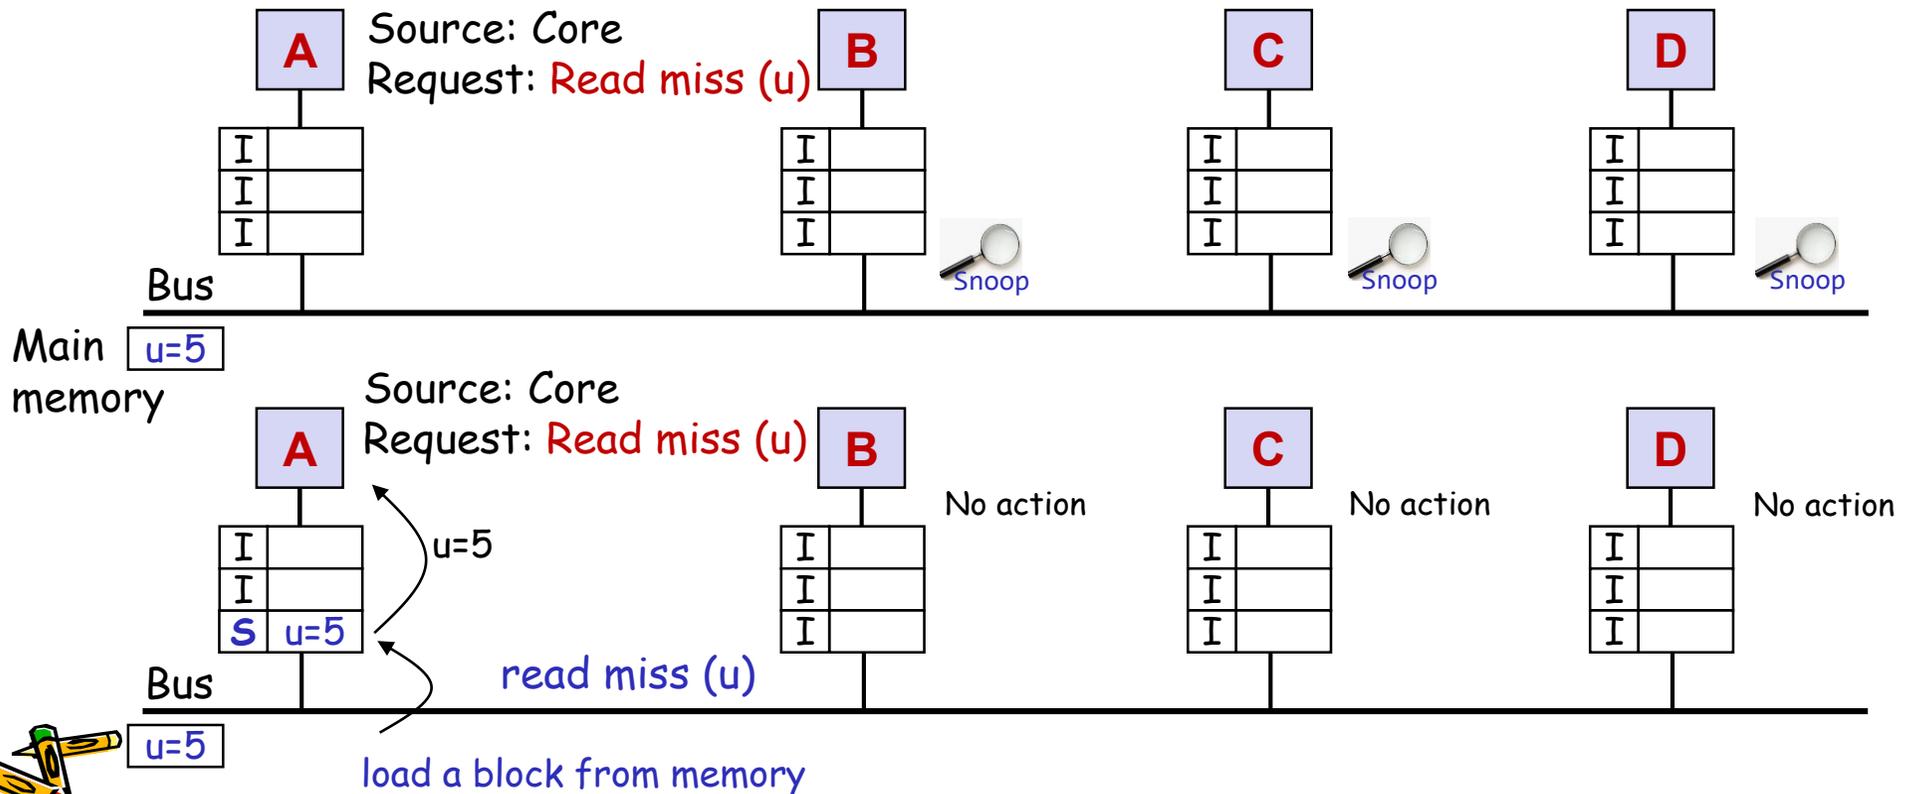
	Request	Source	State of addressed cache block	Type of cache action	Function and explanation
Coh1	Read hit	Processor	Shared or modified	Normal hit	Read data in local cache.
	Read miss	Processor	Invalid	Normal miss	Place read miss on bus.
	Read miss	Processor	Shared	Replacement	Address conflict miss: place 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.
	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.
	Coh2	Read miss	Bus	Shared	No action
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.

# 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

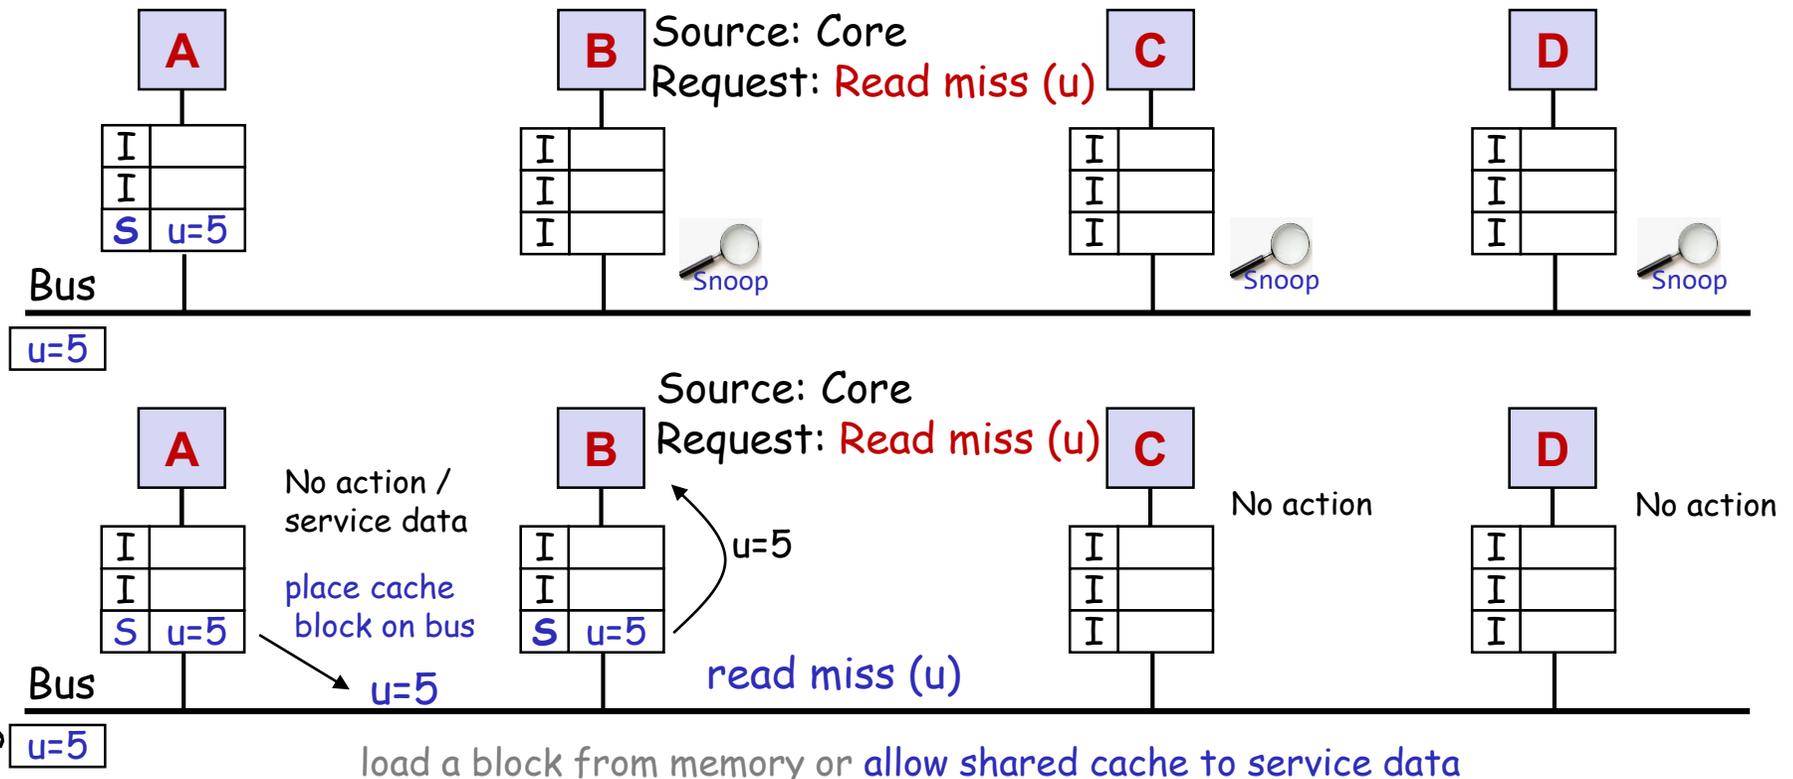


# Cache miss and the addressed block is invalid

- Core B

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

Event: Core B reads u



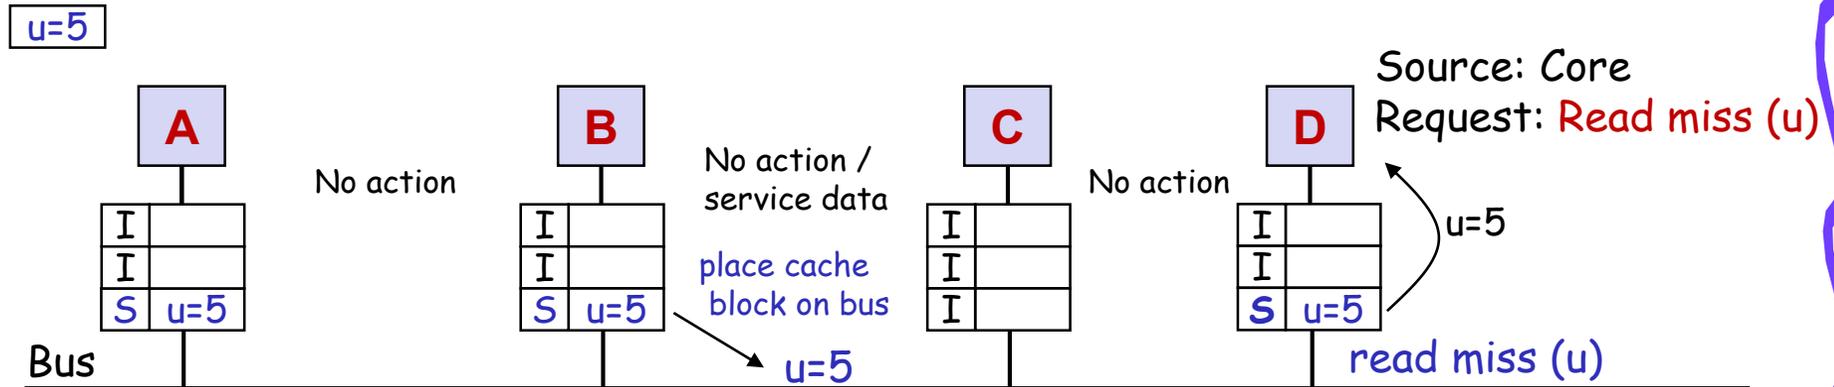
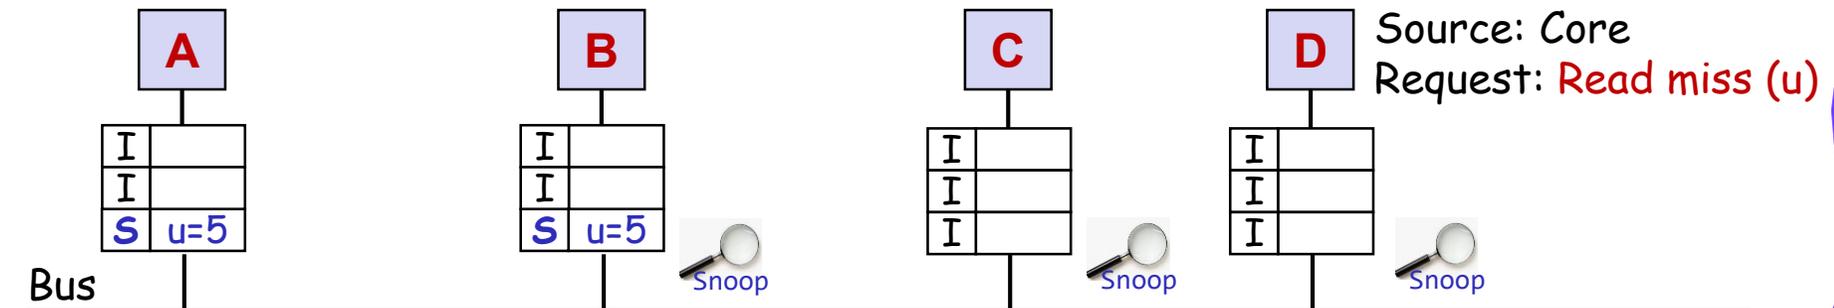
load a block from memory or allow shared cache to service data

# 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



load a block from memory or allow shared cache to service data

# Snooping coherence protocols using bus network

- 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 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.



# Exercise 1

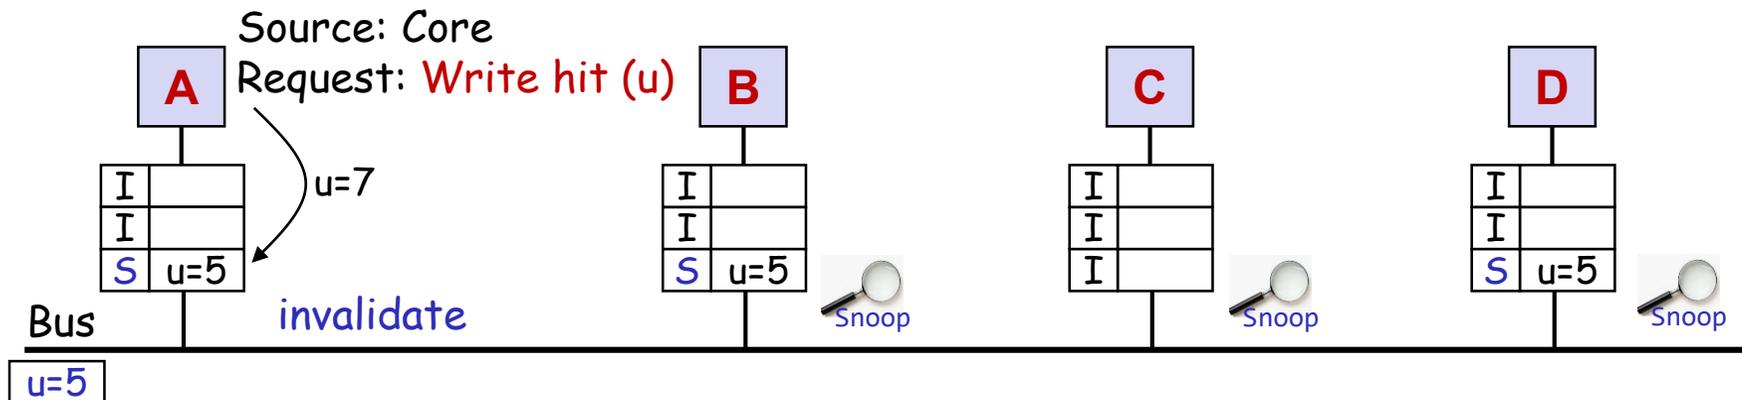
- **Coh1 (Core A)**

- 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

# Coherence 1 (Coh1) and Coherence3 (Coh3)

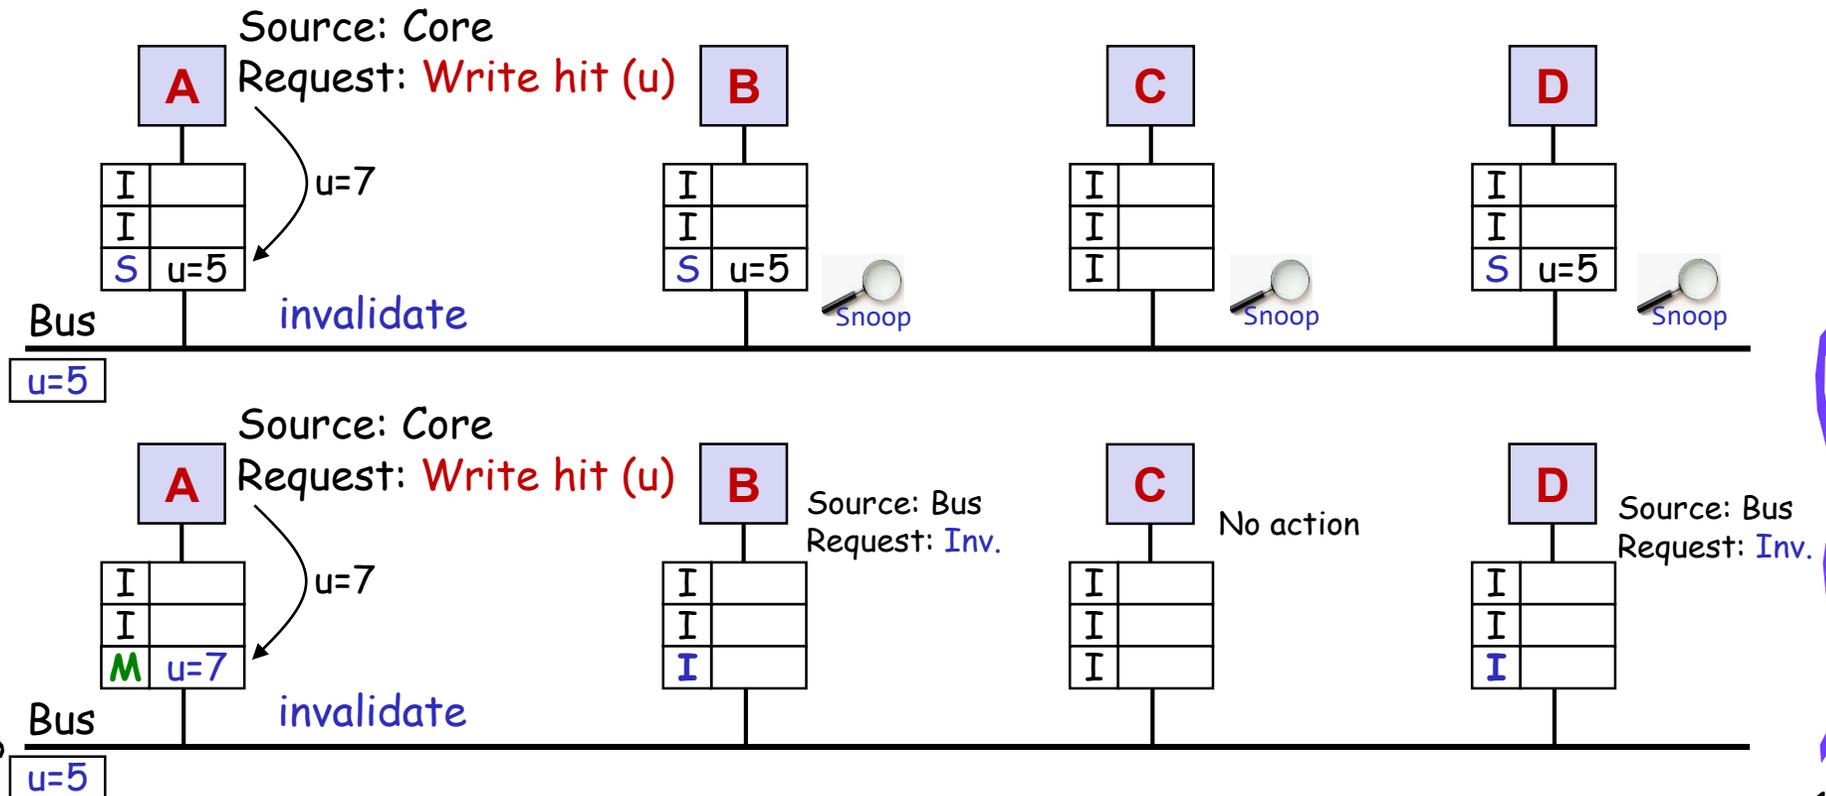
## Coh1 (Core A)

- 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



# Snooping coherence protocols using bus network

- 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.
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.
Read miss	Bus	Modified	Coherence	Attempt to share data: place cache block on bus and change state to shared.
Invalidate	Bus	Shared	Coherence	Attempt to write shared block; invalidate the block.
Write miss	Bus	Shared	Coherence	Attempt to write shared block; invalidate the cache block.
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.

Coh1

Coh2

Coh3

Coh4

Coh5



# Coherence 2 (Coh2)



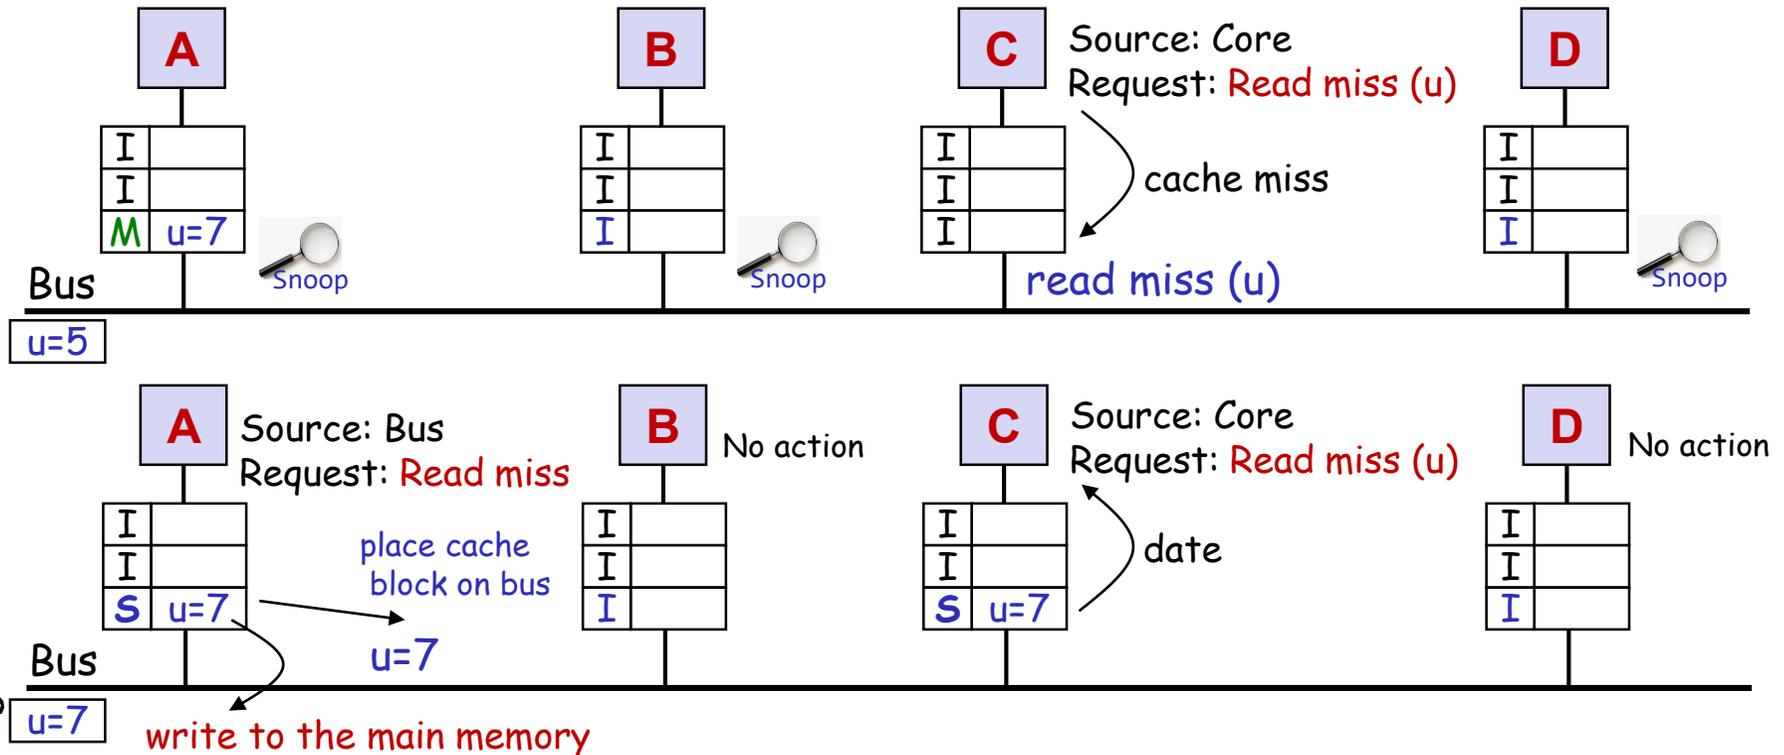
## Core C

- 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 share data; place cache block on bus and change state to shared

Event:  
Core C reads u



# Snooping coherence protocols using bus network

- 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.
	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.



# Coherence 4 (Coh4)

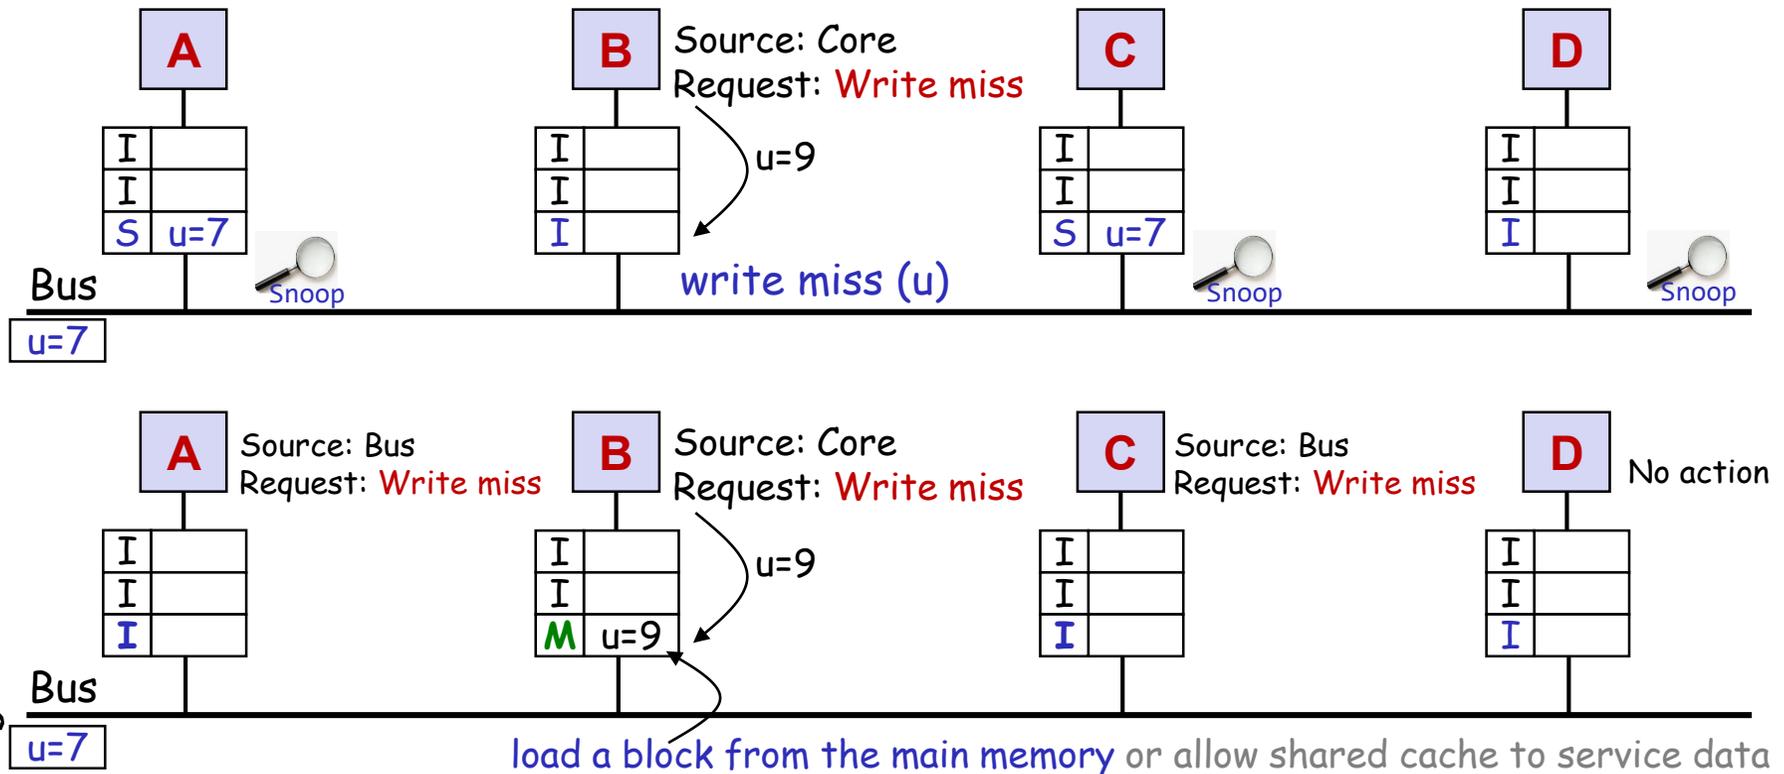
## Core B

- 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



load a block from the main memory or allow shared cache to service data

# Snooping coherence protocols using bus network

- 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.
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.
Read miss	Bus	Modified	Coherence	Attempt to share data: place cache block on bus and change state to shared.
Invalidate	Bus	Shared	Coherence	Attempt to write shared block; invalidate the block.
Write miss	Bus	Shared	Coherence	Attempt to write shared block; invalidate the cache block.
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.

Coh1

Coh2

Coh3

Coh4

Coh5





# Snooping coherence protocols using bus network

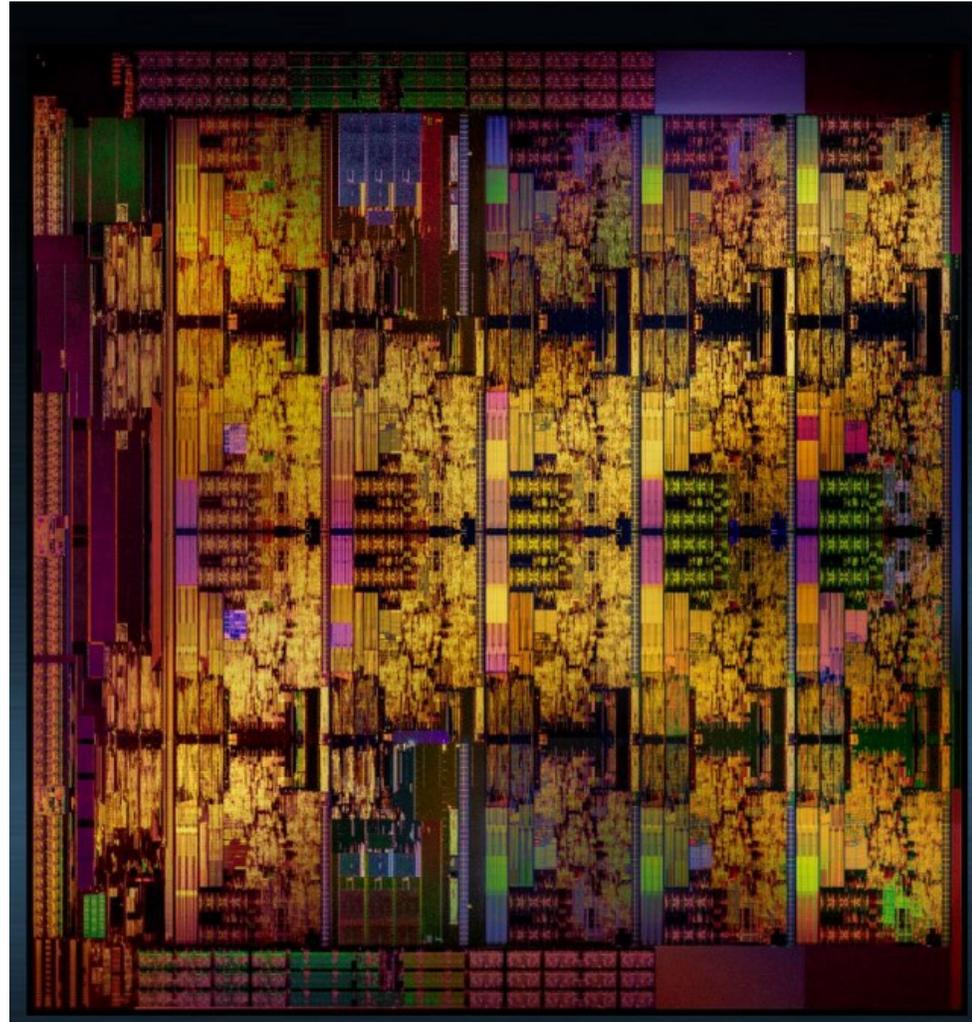


- The basic coherence protocol
  - **MSI** (Modified, Shared, Invalid) protocol
- Extensions
  - **MESI** (Modified, **Exclusive**, Shared, Invalid) protocol
  - **MOESI** (MESI + **Owned**) protocol



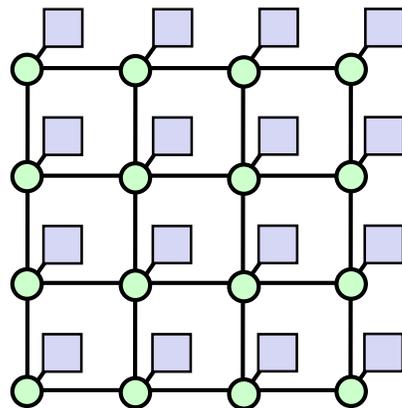
# Intel Skylake-X, Core i9-7980XE (2017)

- 18 core
- 2D mesh topology



# Directory protocols

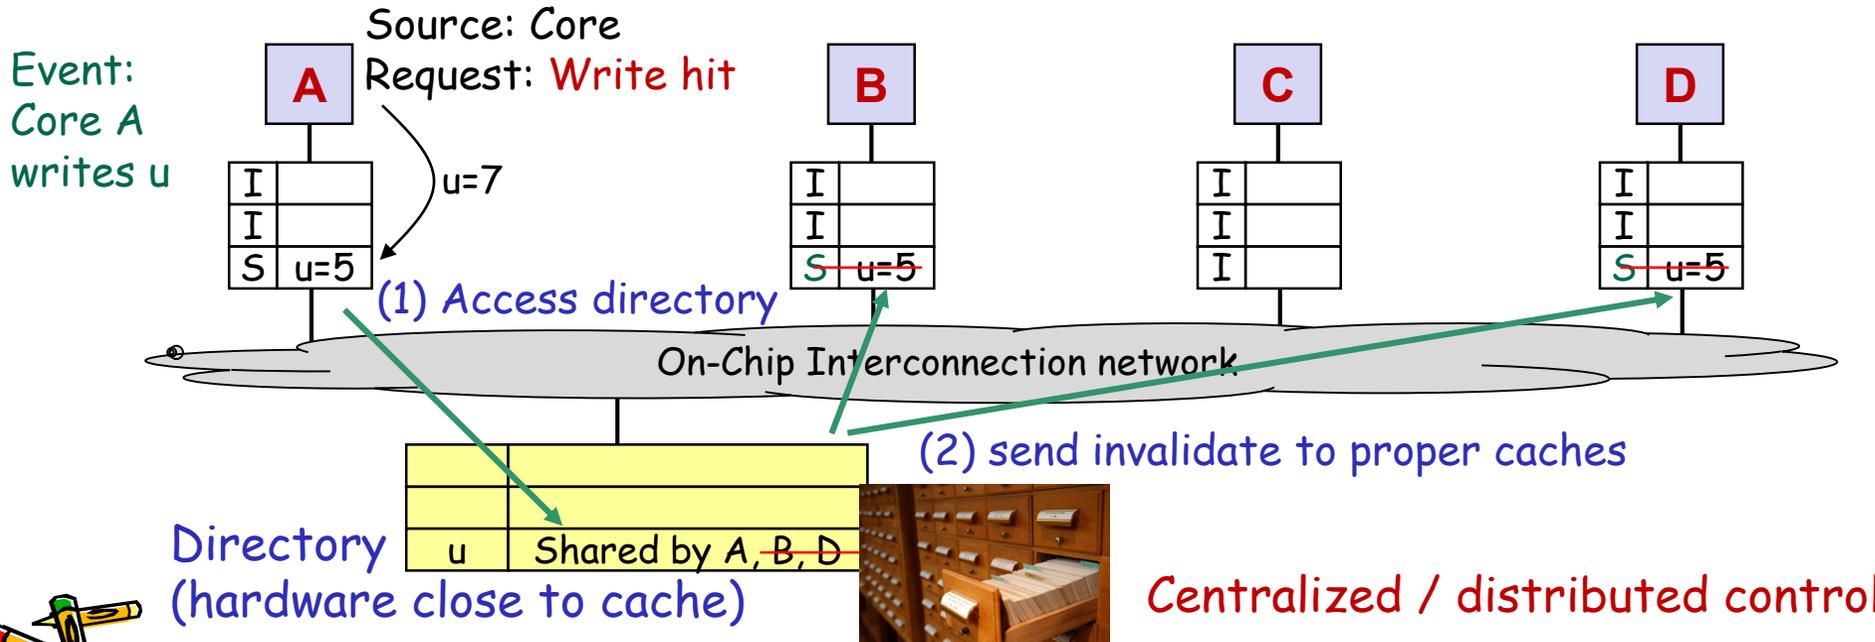
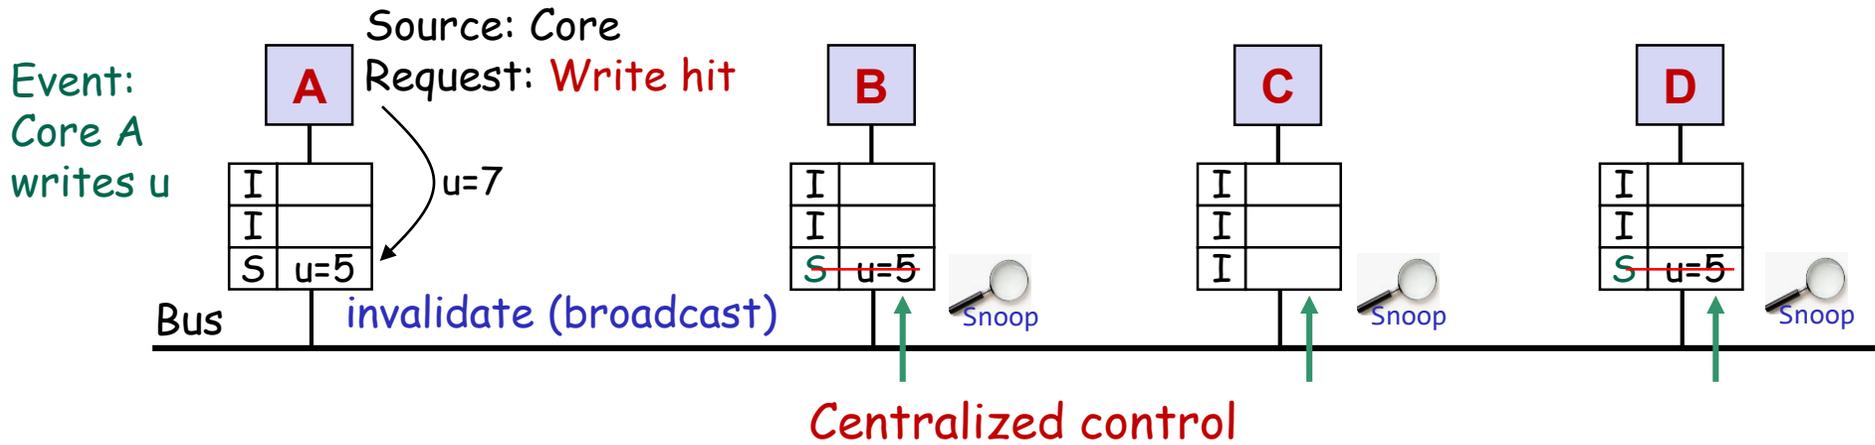
- Snooping coherence protocols are based on the use of bus network.  
What are the protocols for **mesh topology NoC**?
- **Directory protocols**
  - **Sharing status** of each cache line kept in one location
  - A central **directory** keeps track of where the copies of each cache block reside.  
Caches consult this directory to ensure coherence.



Mesh network

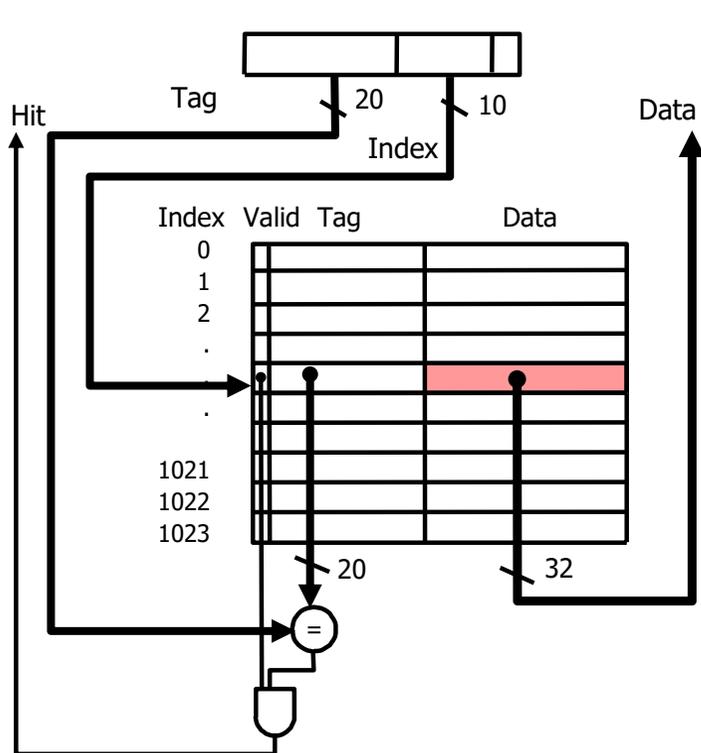


# Snooping coherence protocol and one with directory

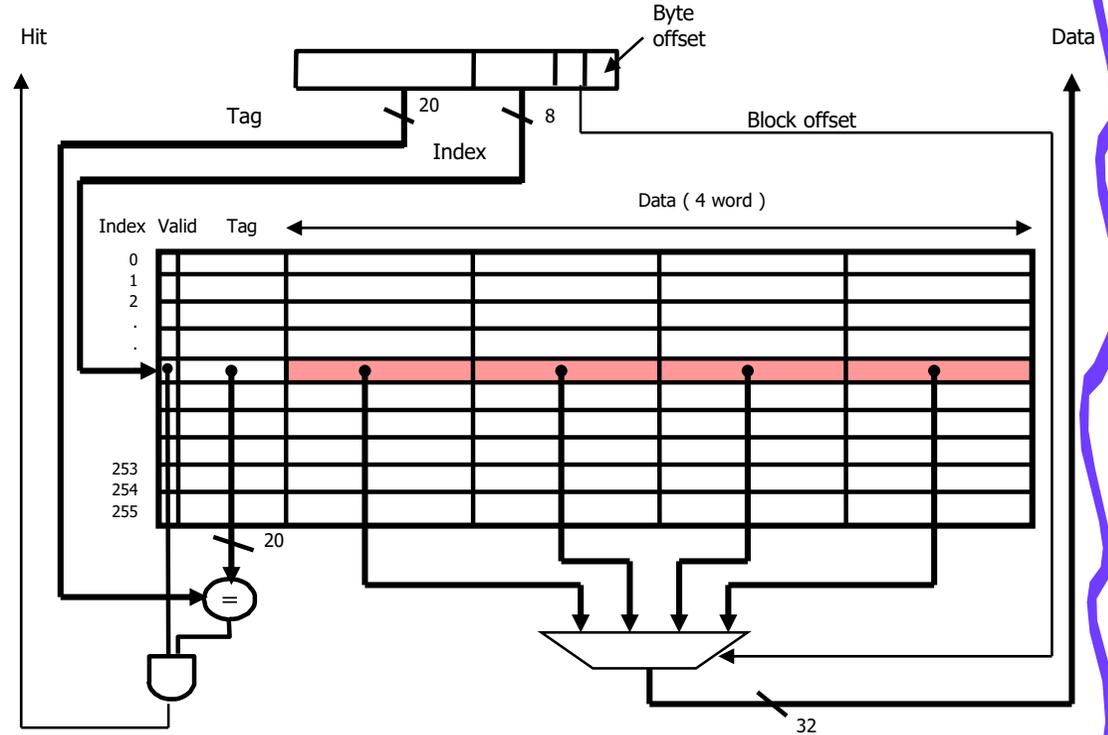


# Two caches of different **block sizes**

- Temporal Locality (Locality in Time):
  - 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



cache line of one word



cache line of **four words (multiword block)**

# Taking Advantage of Spatial Locality

- Let cache block hold more than one word (two words/block)

• 0 1 2 3 4 3 4 15

**0 miss**

00	Mem(1)	Mem(0)

**1 hit**

00	Mem(1)	Mem(0)

**2 miss**

00	Mem(1)	Mem(0)
00	Mem(3)	Mem(2)

**3 hit**

00	Mem(1)	Mem(0)
00	Mem(3)	Mem(2)

**4 miss**

<sup>01</sup> <del>00</del>	<sup>5</sup> <del>Mem(1)</del>	<sup>4</sup> <del>Mem(0)</del>
00	Mem(3)	Mem(2)

**3 hit**

01	Mem(5)	Mem(4)
00	Mem(3)	Mem(2)

**4 hit**

01	Mem(5)	Mem(4)
00	Mem(3)	Mem(2)

**15 miss**

<sup>11</sup> 01	Mem(5)	Mem(4)
<del>00</del>	<del>Mem(3)</del>	<del>Mem(2)</del>

- 8 requests, 4 misses

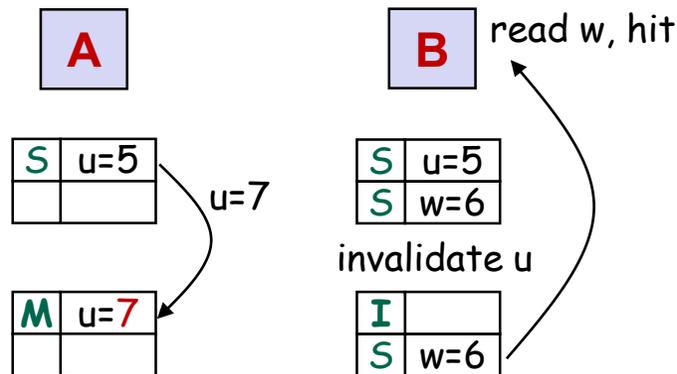
# Coherence influences the cache miss rate

- Coherence misses

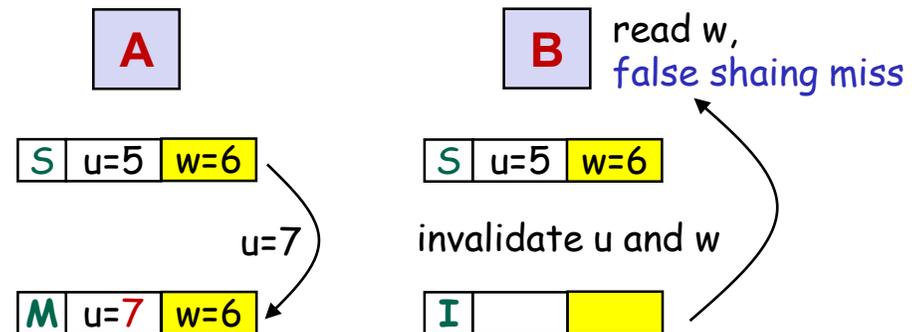
- True sharing misses
- False sharing misses

- When shared data  $u$  and unshared data  $w$  are allocated in the same block, they are both treated as shared data.

Scenario: **A** reads  $u$  -> **B** reads  $u$  -> **B** reads  $w$  -> **A** writes  $u$  -> **B** reads  $w$



cache line of one word



cache line of two words (multiword block)  
 $u$  and  $w$  are in the same cache block

