

# CS6354: Snooping Cache Coherency

7 October 2016

1

## To read more...

This day's papers:


Goodman, "Using cache memory to reduce processor-memory traffic"  
Archibald and Baer, "Cache Coherence Models: Evaluation Using a Multiprocessor Simulation Model"

Supplementary readings:

Hennessy and Patterson, section 5.3

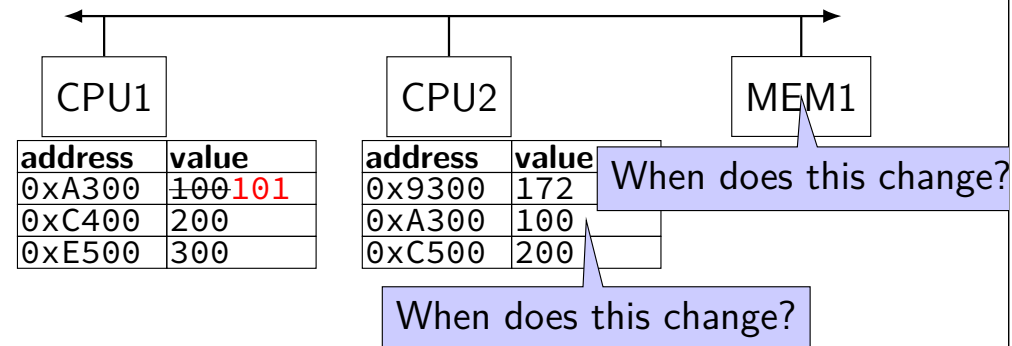
1

## caching shared memories



2

## caching shared memories



CPU1 writes 101 to 0xA300?

2

## cache coherency states

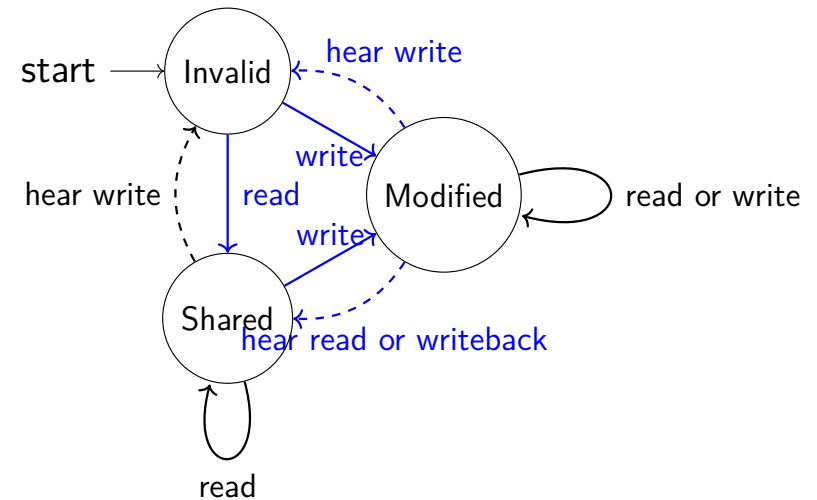
extra information for **each cache block**  
overlaps with valid, dirty bits

stored in **each cache**

different caches may have different states for same block

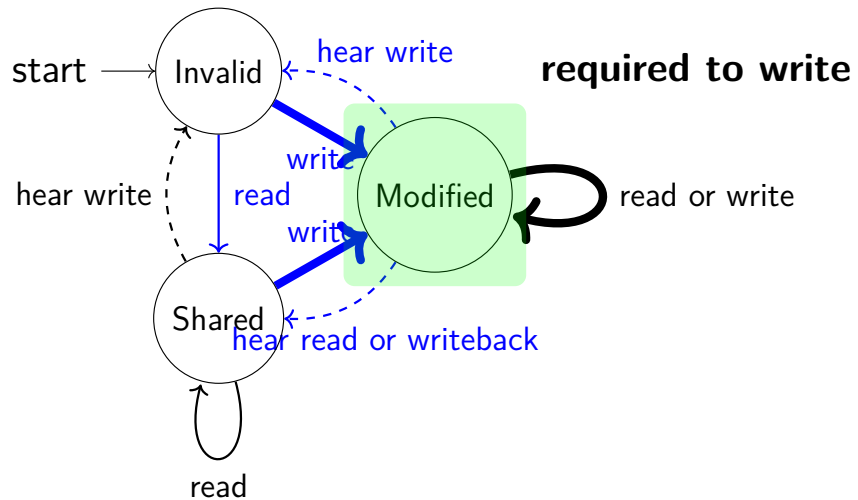
3

## scheme 1: MSI



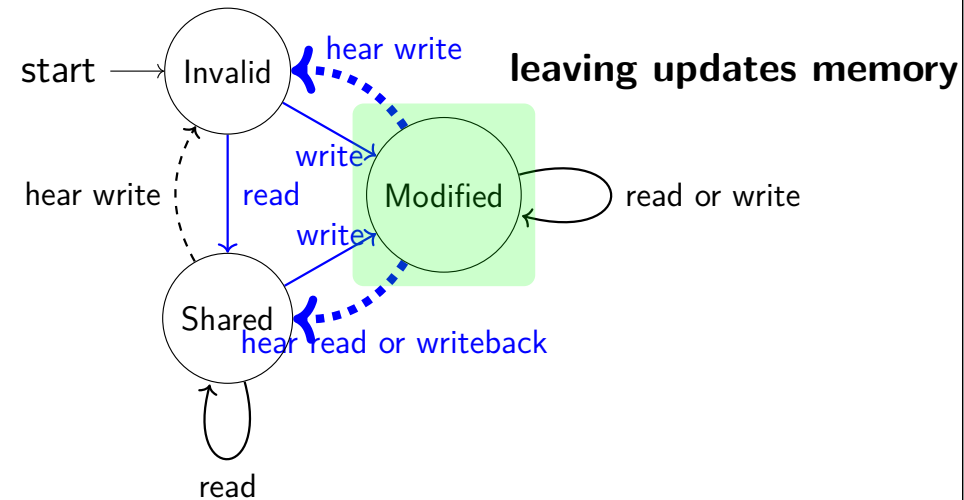
4

## scheme 1: MSI



4

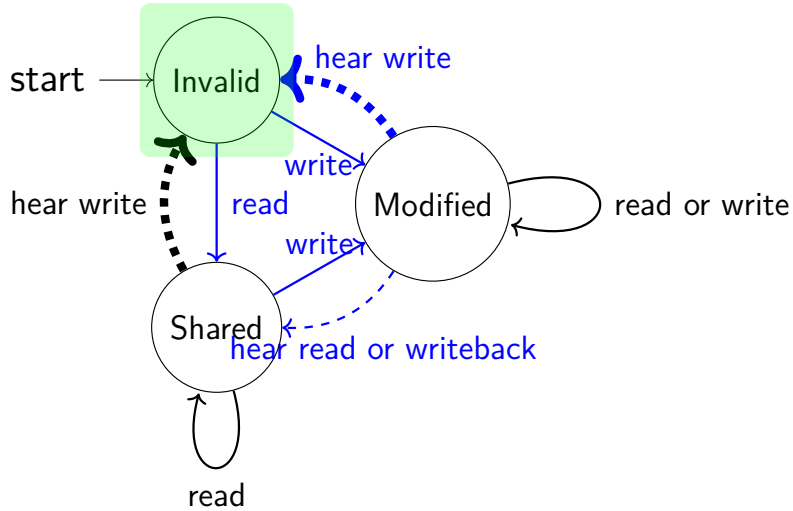
## scheme 1: MSI



4

## scheme 1: MSI

triggered by others writing



4


## scheme 1: MSI

| State    | hear read | hear write | read     | write    |
|----------|-----------|------------|----------|----------|
| Invalid  | —         | —          | Shared   | Modified |
| Shared   | —         | to Invalid | Modified |          |
| Modified | Shared    | Invalid    | —        | —        |

blue: transition sends bus signal

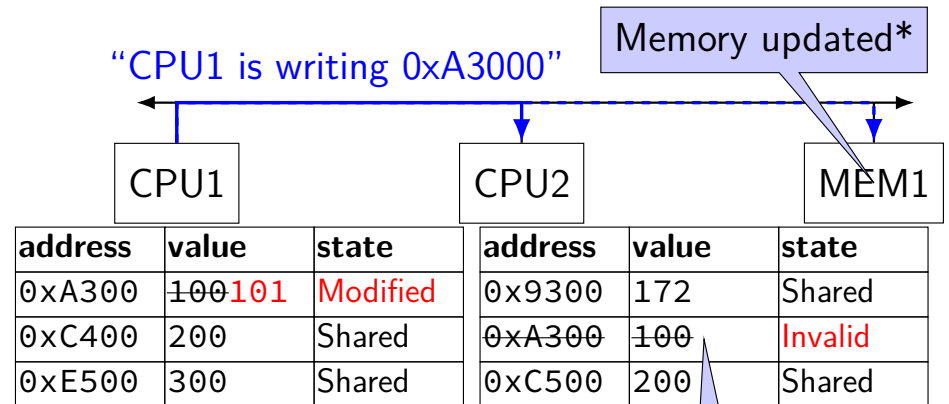
5

## MSI example



6

## MSI example

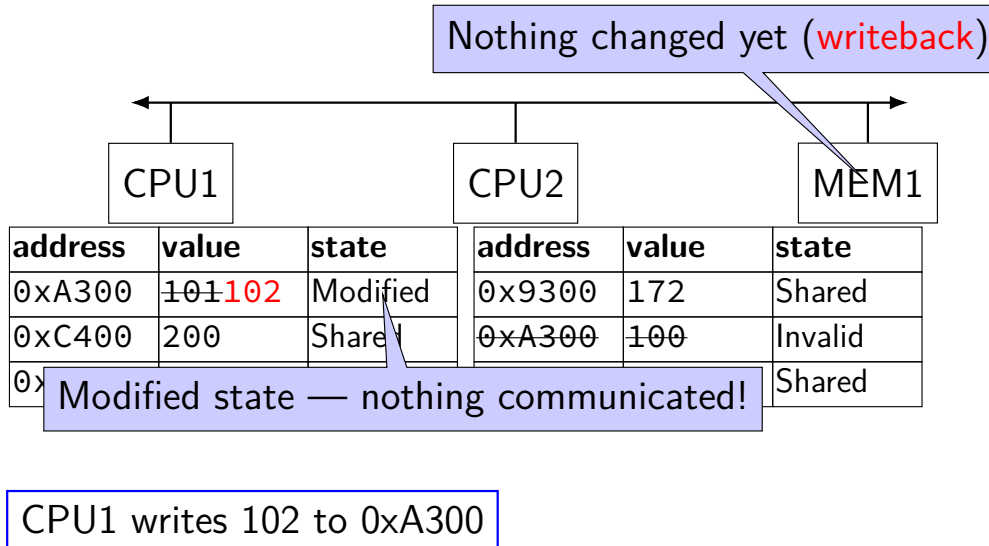


CPU1 writes 101 to 0xA300

Cache sees write:  
invalidate 0xA300

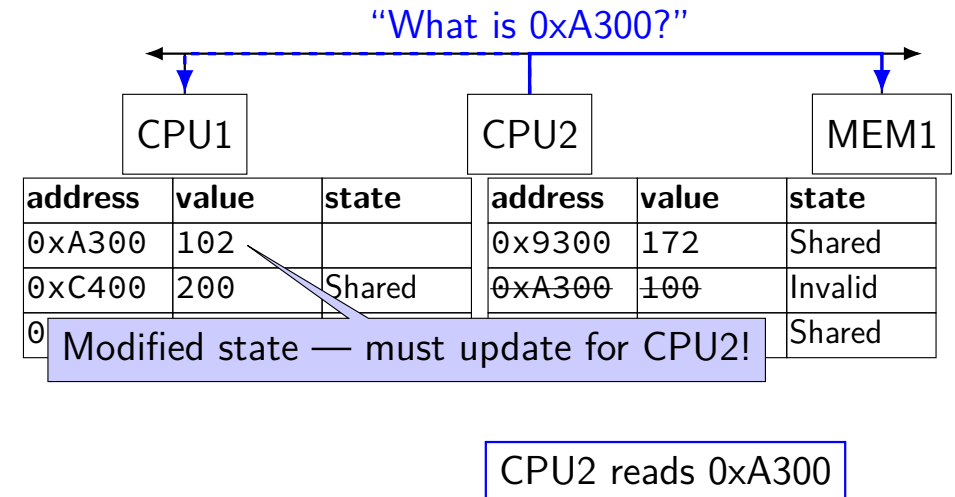
6

## MSI example



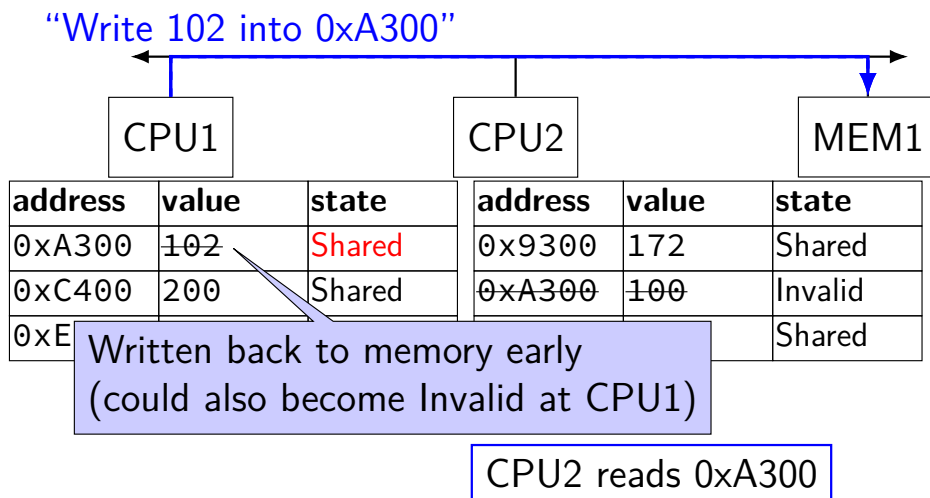
6

## MSI example



6

## MSI example



6

## update memory

to write value (enter modified state), only need to **invalidate** others

more efficient: shorter bus message

7

## on cache replacement/writeback

still happens — e.g. want to store something else  
changes state to **invalid**  
requires writeback if modified (= dirty bit)

8

## scheme 1: MSI

**Modified** value is **different than memory** *and*  
I am the only one who has it

**Shared** value is the **same as memory**

**Invalid** I don't have the value; I will need  
to ask for it

9

## MSI complaints

**modifying** (read then write then write) a value often  
three messages:

initial read from memory

invalidate other caches (and maybe write to  
memory) on initial write

final writeback

10

## scheme 2: MESI

**Modified** value is **different than memory** *and*  
I am the only one who has it

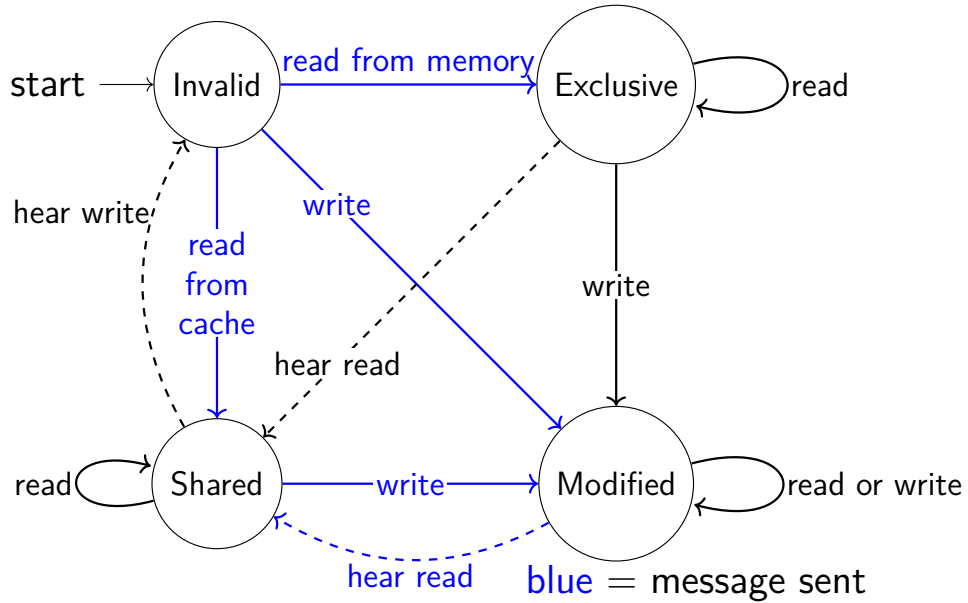
**Exclusive** value is **same as memory** *and* I am  
the only one who has it

**Shared** value is the **same as memory**

**Invalid** I don't have the value; I will need  
to ask for it

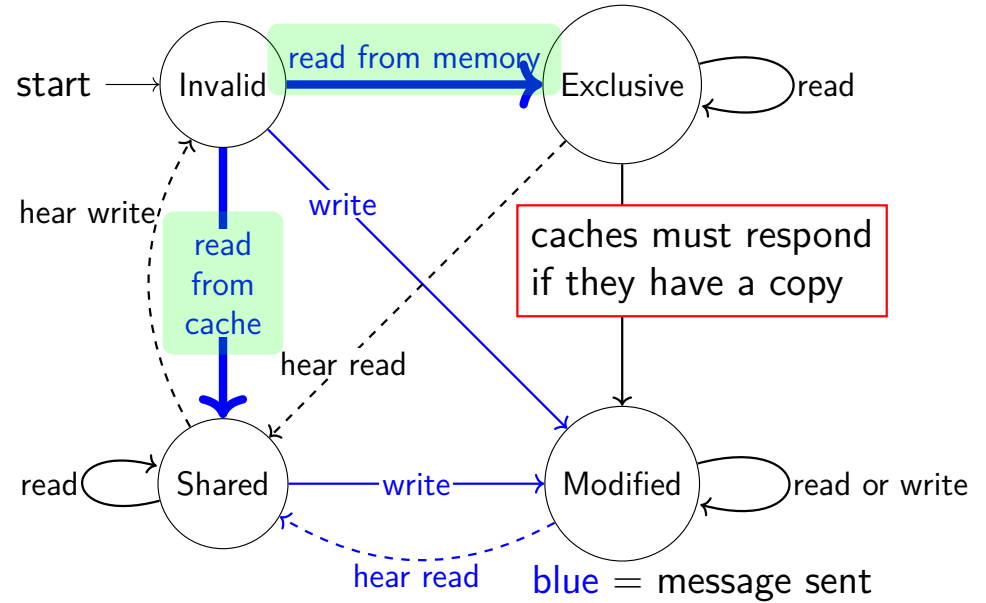
11

## scheme 2: MESI



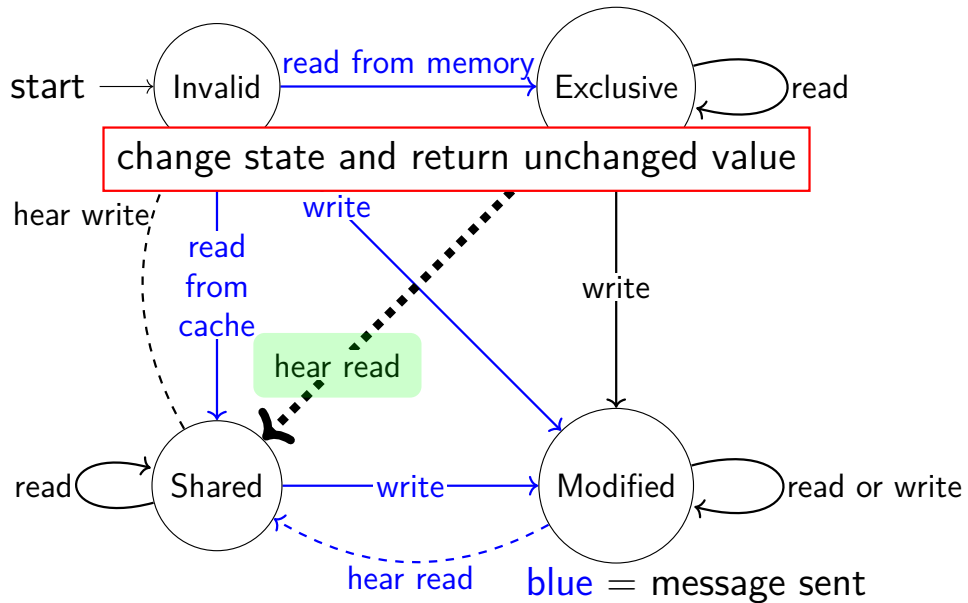
12

## scheme 2: MESI



12

## scheme 2: MESI



12

## read for ownership

reading to modify a value soon?

read into Exclusive state even if reading from cache

invalidate and read

second way to enter exclusive state

13

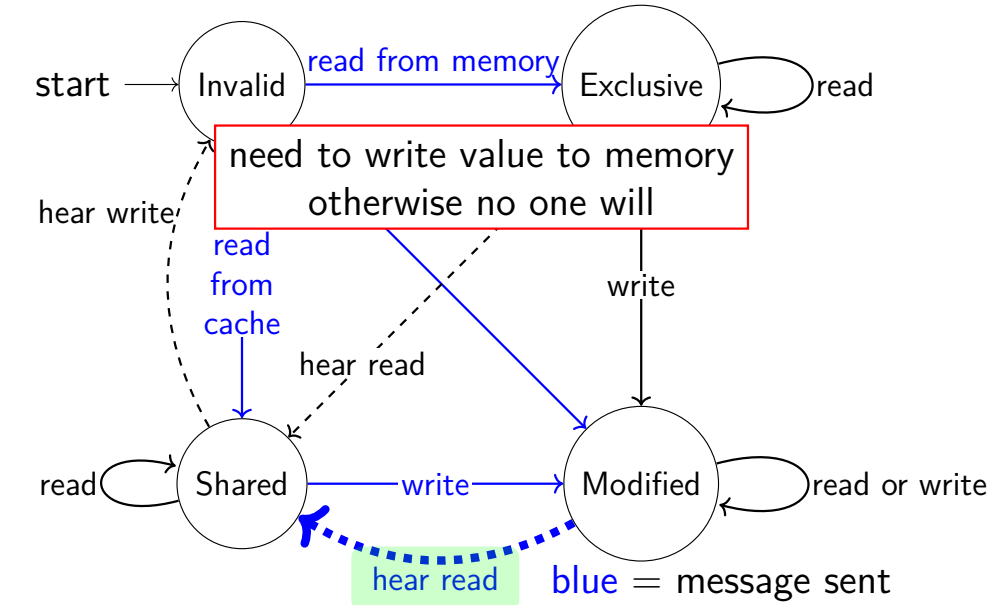
## MESI complaints

have to update memory to share a modified value ...  
even though caches **read from other caches**

read from which cache?

14

## scheme 2: MESI



15

## scheme 3: MOESI

**Modified** value is different than memory *and*  
I am the only one who has it

**Owned** value is different than memory *and*  
I must update memory

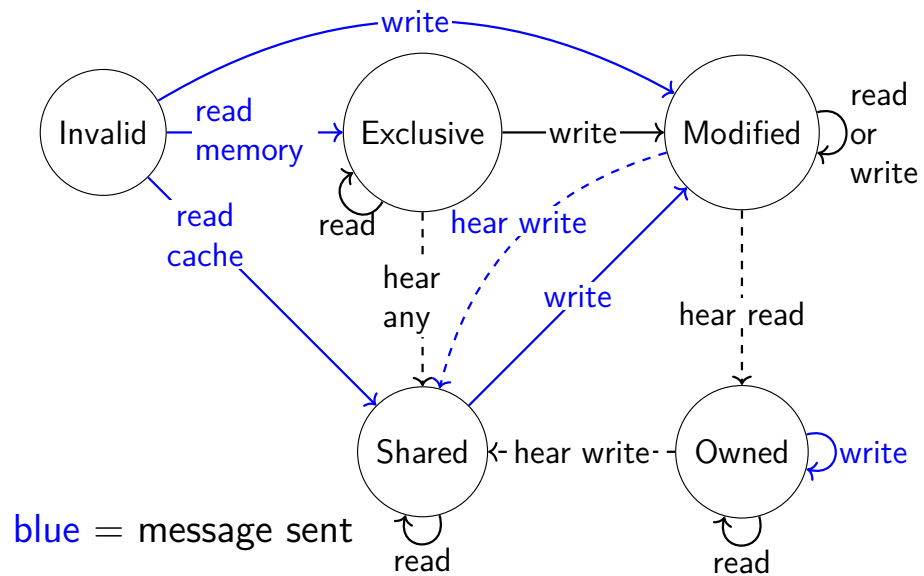
**Exclusive** value is same as memory *and* I am the only one who has it

**Shared** value is same as memory **or cache**  
**in Owned state**

**Invalid** I don't have the value

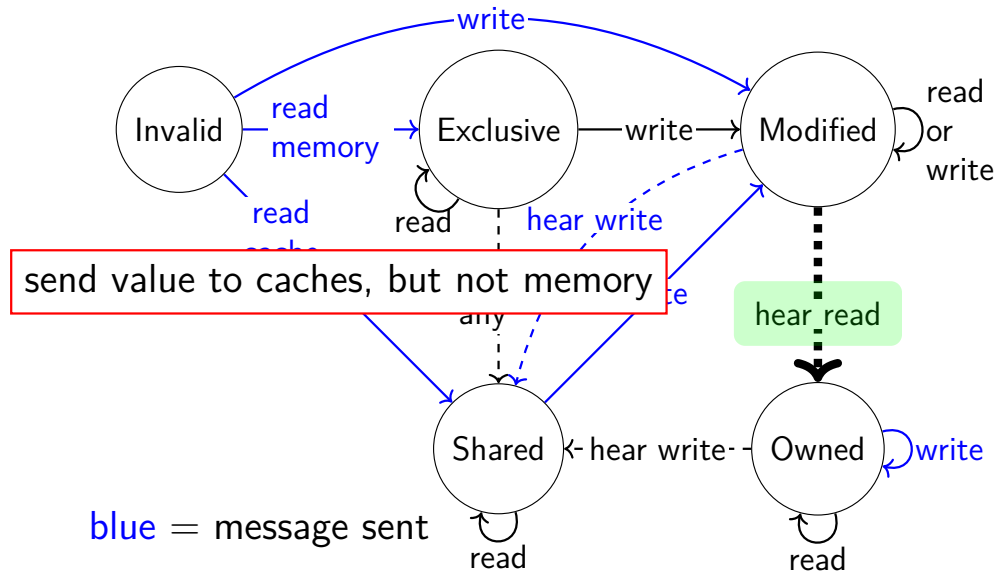
16

## scheme 3: MOESI



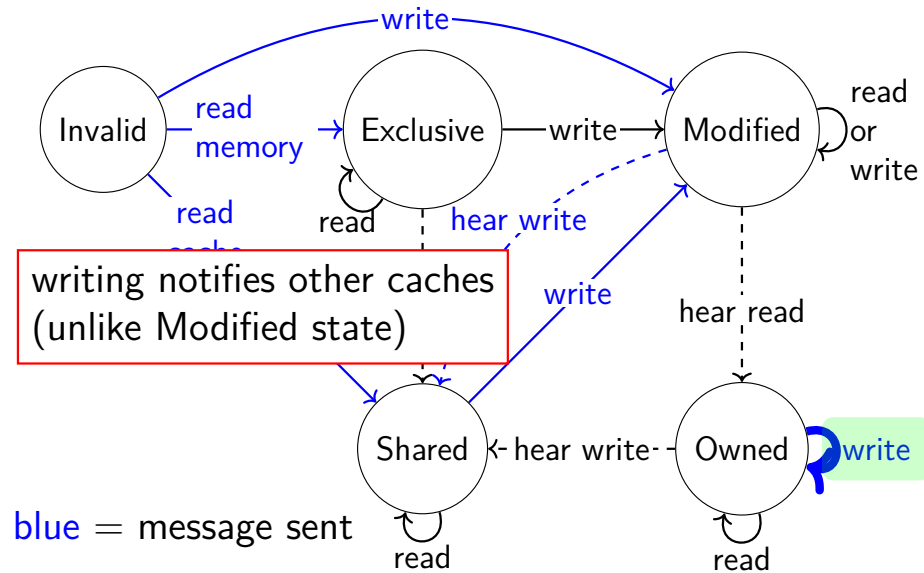
17

## scheme 3: MOESI



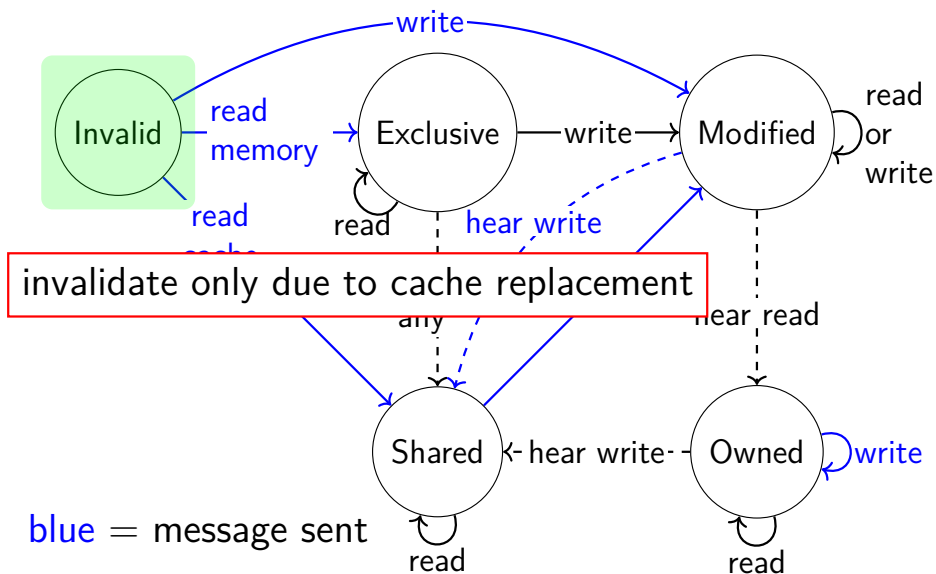
17

## scheme 3: MOESI



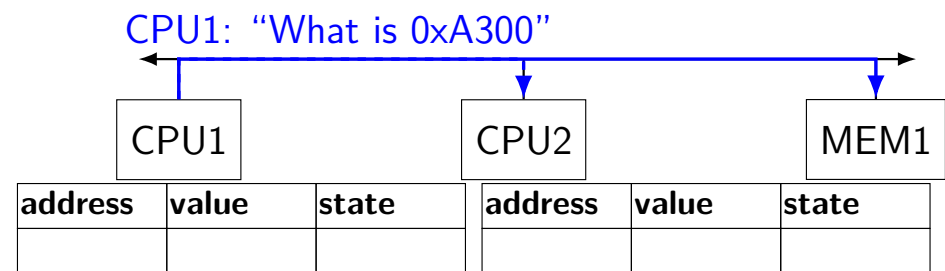
17

## scheme 3: MOESI



17

## MOESI example

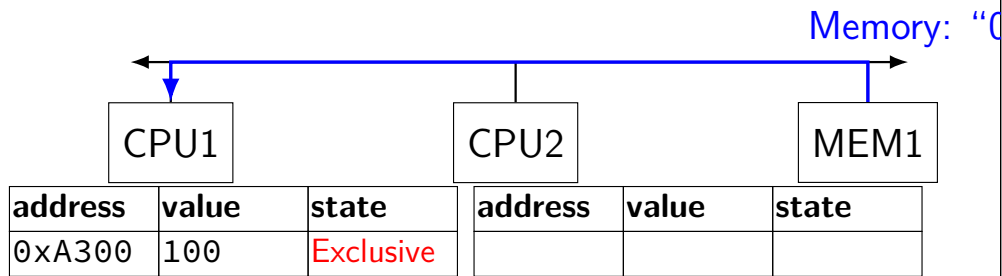


```
CPU1: read 0xA300
CPU1: write 0xA300
CPU1: read 0xA300
CPU2: read 0xA300
CPU2: write 0xA300
```

18




## MOESI example



CPU1: read 0xA300  
 CPU1: write 0xA300  
 CPU1: read 0xA300  
 CPU2: read 0xA300  
 CPU2: write 0xA300

18


## MOESI example



CPU1: read 0xA300  
 CPU1: write 0xA300  
 CPU1: read 0xA300  
 CPU2: read 0xA300  
 CPU2: write 0xA300

18

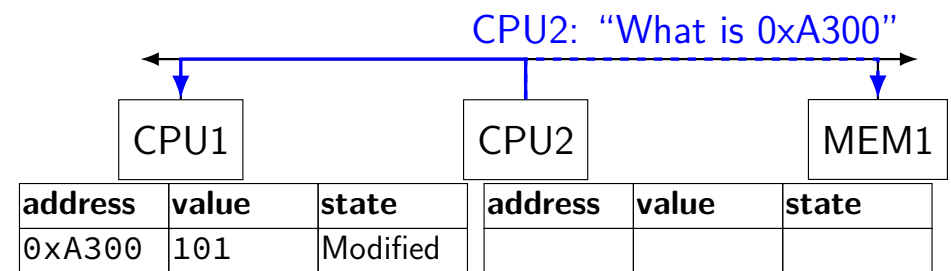
## MOESI example



CPU1: read 0xA300  
 CPU1: write 0xA300  
 CPU1: read 0xA300  
 CPU2: read 0xA300  
 CPU2: write 0xA300

18

## MOESI example




CPU1: read 0xA300  
 CPU1: write 0xA300  
 CPU1: read 0xA300  
 CPU2: read 0xA300  
 CPU2: write 0xA300

18

## MOESI example

CPU1: "0xA300 = 101"




CPU1: read 0xA300  
 CPU1: write 0xA300  
 CPU1: read 0xA300  
 CPU2: read 0xA300  
 CPU2: write 0xA300

18

## MOESI example

CPU2: "I'm changing 0xA300"



CPU1: read 0xA300  
 CPU1: write 0xA300  
 CPU1: read 0xA300  
 CPU2: read 0xA300  
 CPU2: write 0xA300

18

## MSI versus MESI versus MOESI

CPU1: read 0xA300  
 CPU1: write 0xA300 MSI: invalidate  
 CPU1: read 0xA300  
 CPU2: read 0xA300 MSI/MESI: memory write  
 CPU2: write 0xA300 MSI: invalidate

19

## Other cache coherency options

can **invalidate** instead of updating other caches on write

invalidation message faster to send than new value

tradeoff: faster **if** other cache won't use value

20

## Dropping states from MOESI

**Modified** value is different than memory *and* I am the only one who has it

**Owned** value is different than memory *and* I must update memory

**Exclusive** value is same as memory *and* I am the only one who has it

**Shared** value is same as memory or cache in Owned state

**Invalid** I don't have the value

21

## Dropping states from MOESI

**Modified** value is different than memory *and* I am the only one who has it

**Owned** value is different than memory *and* I must update memory

**Exclusive** ~~value is same as memory *and* I am the only one who has it~~

**Shared** value is same as memory or cache in Owned state

**Invalid** I don't have the value

21

## Mapping to the paper

MSI + reread to get in Modified: Synapse

MESI + full-write-to-invalidate: write-once

MOSI + forward-on-write: Berkeley

MESI + forward-on-write: Illinois

MESI + invalidate-on-write: Firefly

MOESI + forward-on-write: Dragon

22

## “System Power”


sum of **processor utilizations**

how much time are CPUs spending waiting for bus

what about overlapping cache accesses and computation??

23

## overhead if almost no shared data



24

## overheads without sharing data

sending invalidation signals no other cache needs

reloading value from memory no cache needs  
(Synapse)

25

## simulation caveats

workloads?

variation in hardware?

26

## false sharing

cache blocks are shared **even if you are accessing different parts**

huge performance problem with writes

27

## Present-day snooping cache coherency

AMD processors use MOESI

Intel uses something called **MESIF**

plus some techniques we'll talk about next time


28

## MESIF states

|                   |  |
|-------------------|--|
| <b>Modified</b>   | value is different than memory <i>and</i> I am the only one who has it     |
| <b>Exclusive</b>  | value is same as memory <i>and</i> I am the only one who has it            |
| <b>Shared</b>     | value is same as memory  |
| <b>Invalid</b>    | I don't have the value   |
| <b>Forwarding</b> | value is same as memory <i>and</i> I should <b>provide it if requested</b> |

29

## Forwarding state: lower traffic



30

## Non-bus topologies


**necessary** to connect large numbers of caches

higher bandwidth — if you don't broadcast everything

next time: avoiding broadcast

31

## timing trickiness



32

## compare-and-swap

```
compare-and-swap(address, expect-old-value, new-v  
    atomically {  
        if (expect-old-value == memory[address])  
            memory[address] = new-value  
    }  
}
```

33

## Implementing compare-and-swap

get block into **Exclusive** or **Modified** state  
 read from memory/cache if necessary  
 invalidate other caches if necessary

compare, if value matches, do write (Modified state)

34

## Coherency

common property: **single 'responsible' cache** for  
possibly changed values

Owned, Exclusive, Modified states

responsible cache must **reply** to reads of address

variation:

when is responsibility acquired? (only on write?)  
when is it relinquished? (only on other's write?)

35