ECSE 425 Lecture 29: More Snoopy Coherence

H&P Chapter 4

Last Time

- Snoopy Coherence

Today

- Snoopy Coherence Example
- SMP Performance

Write-back State Machine-III

- State machine for CPU requests and for bus requests for each cache block

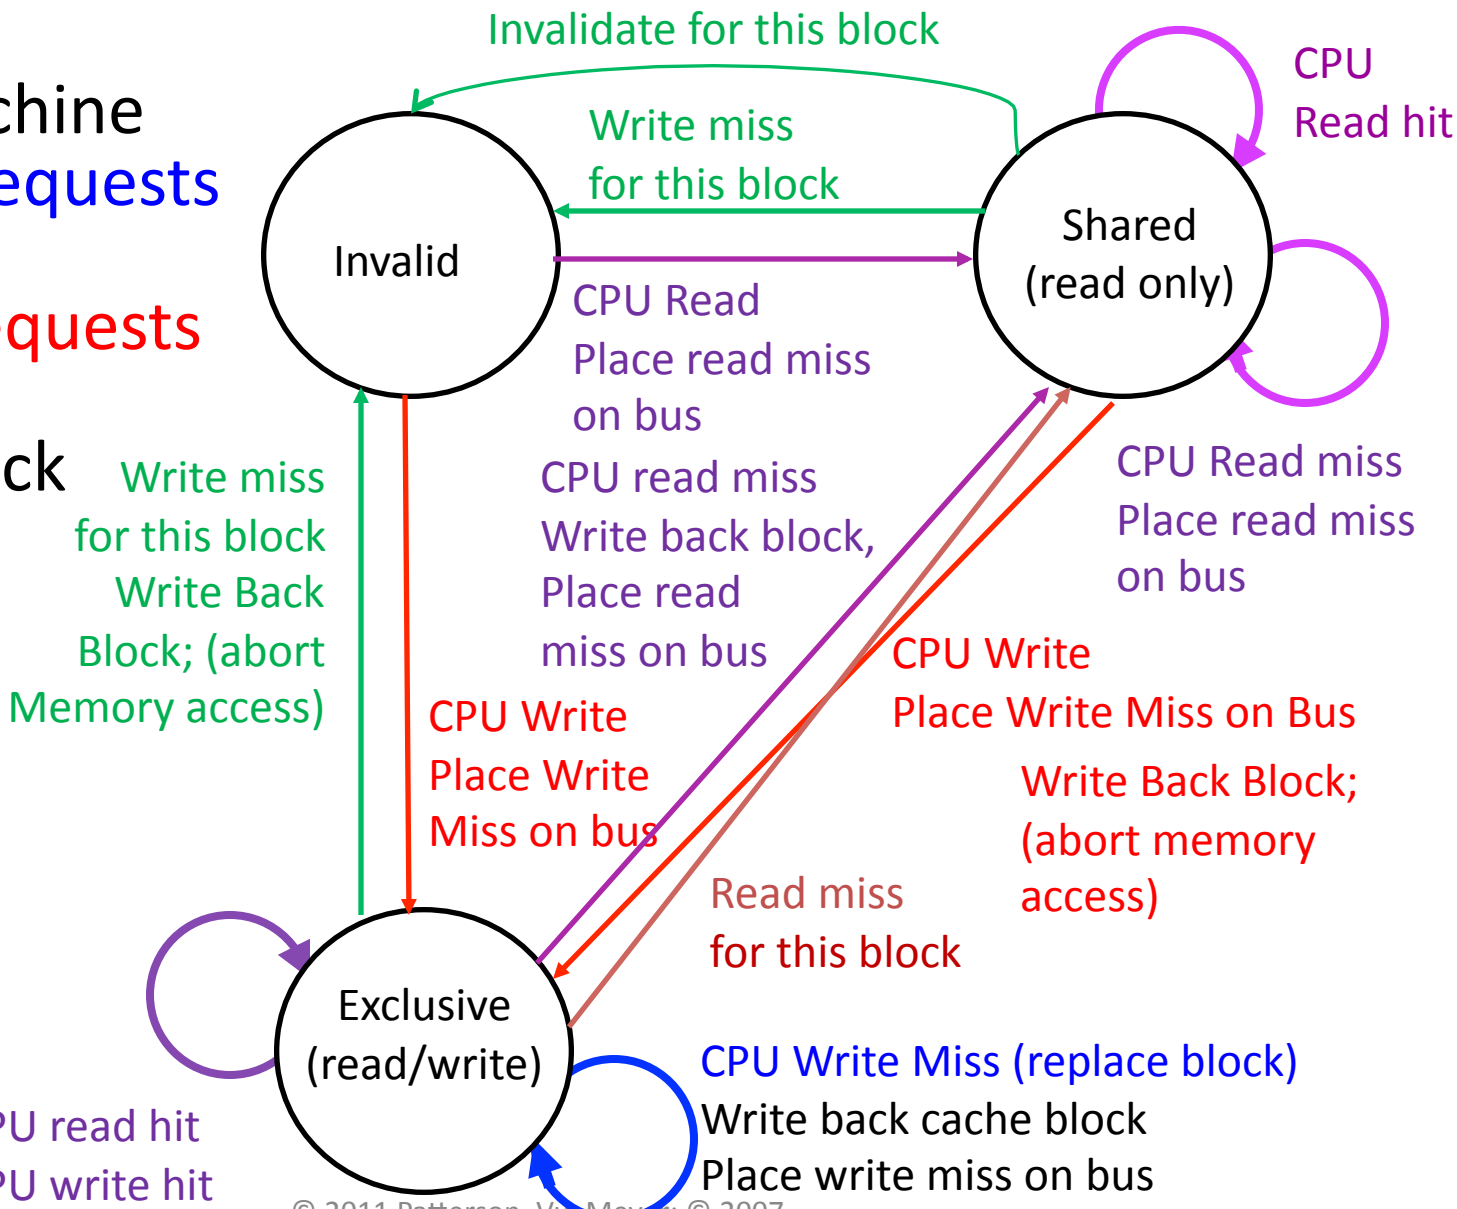

WB Invalidate Example

	P1			P2			Bus				Memory	
step	State	Addr	Value	State	Addr	Value	Action	Proc.	Addr	Value	Addr	Value
P1 Write 10 to A1												
P1: Read A1												
P2: Read A1												
P2: Write 20 to A1												
P2: Write 40 to A2												

Assumes A1 and A2 map to same cache block, initial cache state is invalid

WB Invalidate Example

	P1			P2			Bus				Memory	
step	State	Addr	Value	State	Addr	Value	Action	Proc.	Addr	Value	Addr	Value
P1 Write 10 to A1	Excl.	A1	10				WrMs	P1	A1			
P1: Read A1												
P2: Read A1												
P2: Write 20 to A1												
P2: Write 40 to A2												

Assumes A1 and A2 map to same cache block

WB Invalidate Example

	P1			P2			Bus				Memory	
step	State	Addr	Value	State	Addr	Value	Action	Proc.	Addr	Value	Addr	Value
P1 Write 10 to A1	Excl.	A1	10				WrMs	P1	A1			
P1: Read A1	Excl.	A1	10									
P2: Read A1												
P2: Write 20 to A1												
P2: Write 40 to A2												

Assumes A1 and A2 map to same cache block

WB Invalidate Example

	P1			P2			Bus				Memory	
step	State	Addr	Value	State	Addr	Value	Action	Proc.	Addr	Value	Addr	Value
P1 Write 10 to A1	Excl.	A1	10				WrMs	P1	A1			
P1: Read A1	Excl.	A1	10									
P2: Read A1				Shar.	A1		RdMs	P2	A1			
	Shar.	A1	10				WrBk	P1	A1	10	A1	10
				Shar.	A1	10	RdDa	P2	A1	10	A1	10
P2: Write 20 to A1												
P2: Write 40 to A2												

Assumes A1 and A2 map to same cache block

WB Invalidate Example

	P1			P2			Bus				Memory	
step	State	Addr	Value	State	Addr	Value	Action	Proc.	Addr	Value	Addr	Value
P1 Write 10 to A1	Excl.	A1	10				WrMs	P1	A1			
P1: Read A1	Excl.	A1	10									
P2: Read A1				Shar.	A1		RdMs	P2	A1			
	Shar.	A1	10				WrBk	P1	A1	10	A1	10
				Shar.	A1	10	RdDa	P2	A1	10	A1	10
P2: Write 20 to A1	Inv.			Excl.	A1	20	WrMs	P2	A1		A1	10
P2: Write 40 to A2												

Assumes A1 and A2 map to same cache block

WB Invalidate Example

	P1			P2			Bus				Memory	
step	State	Addr	Value	State	Addr	Value	Action	Proc.	Addr	Value	Addr	Value
P1 Write 10 to A1	Excl.	A1	10				WrMs	P1	A1			
P1: Read A1	Excl.	A1	10									
P2: Read A1				Shar.	A1		RdMs	P2	A1			
	Shar.	A1	10				WrBk	P1	A1	10	A1	10
				Shar.	A1	10	RdDa	P2	A1	10	A1	10
P2: Write 20 to A1	Inv.			Excl.	A1	20	WrMs	P2	A1		A1	10
P2: Write 40 to A2							WrMs	P2	A2		A1	10
				Excl.	A2	40	WrBk	P2	A1	20	A1	20

Assumes A1 and A2 map to same cache block,
but A1 != A2

Limitations of Snoopy SMPs

- Bus-based multiprocessor
  - Bus for coherence traffic, memory traffic
  - Becomes the bottle neck
  - Solns: Multiple buses or interconnection networks
- Single memory accommodates all CPUs
  - Multi-bank memory
- Coherence traffic can overwhelm the network and memory system

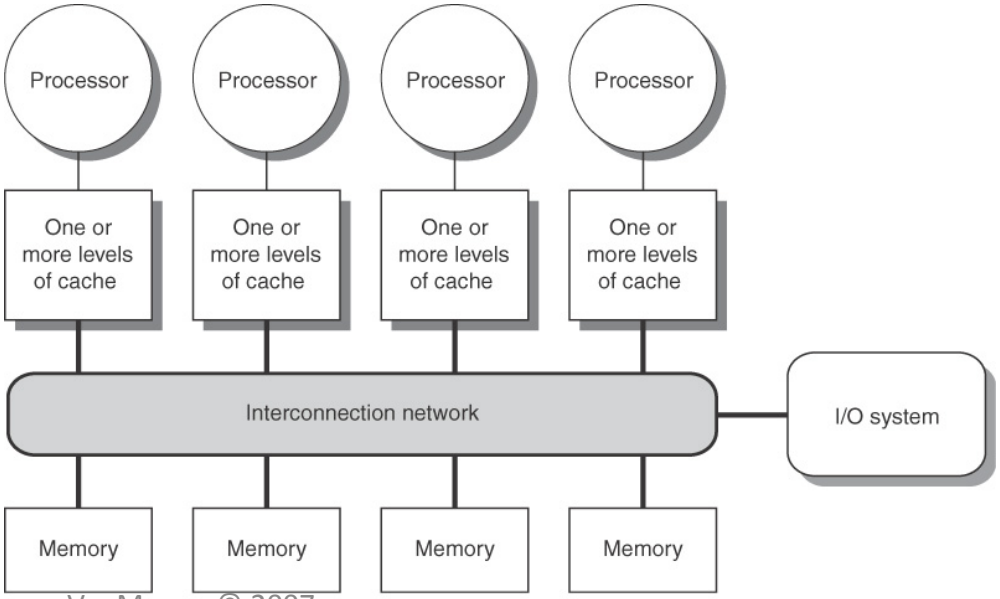

Opteron: a hybrid solution

- Memory connected directly to each chip
  - Up to four chips, or eight cores
- In addition, point-to-point connections
  - Coherent protocol by broadcast on the point-to-point connection to up to three other chips (snooping idea)
  - Invalidates complete via explicit ack (directory idea)
- Similar remote and local memory latency
  - OS can treat Opteron as a UMA MP

Performance of SMPs

- Cache performance is combination of
  - Uni-processor cache miss traffic
  - Traffic caused by communication
    - Results in invalidations and subsequent cache misses
- Cache misses:
  - Compulsory, Capacity, Conflict misses
  - 4th C: coherence miss

Coherency Misses

1. True sharing misses
   – Multiple processors read and write the same **word**
   – Miss would still occur if block size were one word
2. False sharing misses
   – Multiple processors read and write the same **block**
   – Block size matters: no miss if for one word block size
   – Occurs because only one dirty bit per block

Example: True or False Sharing misses?

- Assume x1 and x2 in same cache block.
  P1 and P2 both read x1 and x2 before.

Time	P1	P2	**True, False, Hit? Why?**
1	Write x1		True miss; invalidate x1 in P2
2		Read x2	False miss; x1 irrelevant to P2
3	Write x1		False miss; x1 irrelevant to P2
4		Write x2	True miss; occur even if block size 1
5	Read x2		True miss; x2 was written by P2

MP Performance: Commercial Workload

- Three-level cache
- Uni-processor cache misses improve
- True sharing and false sharing misses are unchanged

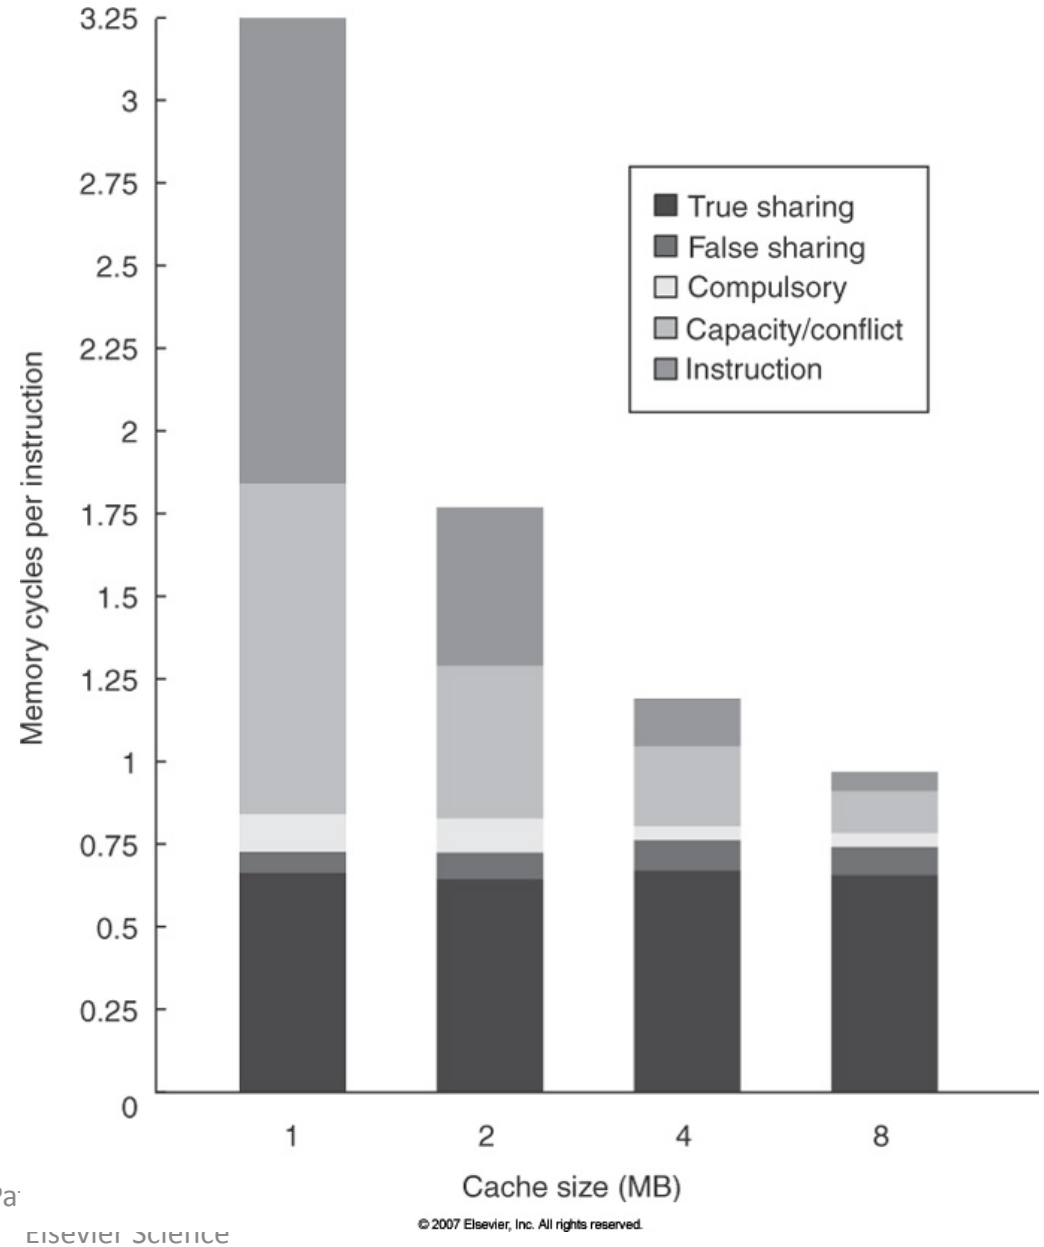

MP Performance: Commercial Workload

- 2MB cache
- True sharing, false sharing misses increase with number of CPUs going from 1 to 8 CPUs

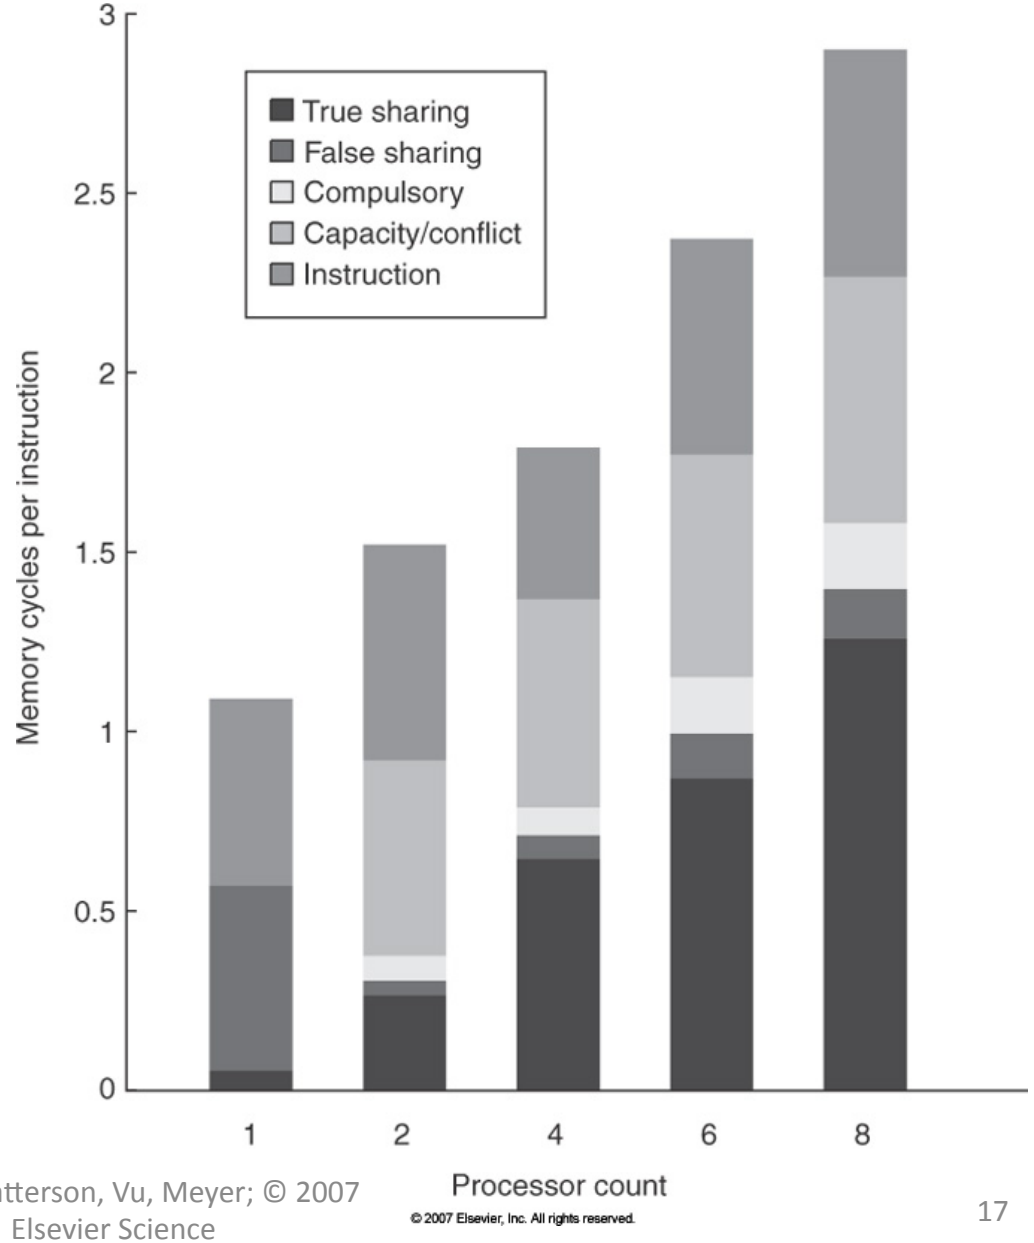

Next Time

- Directory Coherence
- Synchronization and Consistency