

Shared Memory Architectures

Software coherence
Fences
Intro to hardware coherence

Discuss paper on
RP3

6.173
Fall 2010
L09

Agarwal


- 1 -

Today's Outline

- RP3 discussion
- How do caches change things
- Shared memory programming with caches
- Software coherence
- The meaning of shared memory
- Hardware cache coherence

- 2 -

Recall Shared Memory Jacobi



Lots of repeat accesses of data
Lots of communication over the network


And very energy inefficient

- 3 -

Pure Shared Memory

32-bit energy costs in 40nm

DRAM read: $\sim 1000\text{pJ}$



Add: $\sim 1\text{pJ}$

Register read: $\sim 1\text{pJ}$


Ideas?

Caches!

Cache read (small L1): $\sim 10\text{pJ}$

- 4 -

Ultracomputer Built as RP3



This is a
butterfly
network -
variant
of Omega
network

- Indirect network - Omega network (details later in course)
- Shared memory machine with caches
- Each memory module placed physically close to processors
- Communication/synchronization through shared memory
- Hardware routing of memory requests
- SPMD FORTRAN programming (single program multiple data)
- No latency hiding - wait for memory request
- More complex hardware
- I could not find a picture of the RP3

**What were the
big ideas?**

5

Trivia Question

Cosmic Cube - Chuck Seitz (Prof. Caltech)

RP3 - Greg Pfister (IBM, Yorktown, NY)


Ultracomputer - Allan Gottlieb (Prof. NYU)

These inventors have this in common:

- (a) None finished their Bachelor's degrees
- (b) They were all Geminis
- (c) None of them is retired even today
- (d) They are all Yankees fans
- (e) None of the above

6

Caches - the good, the bad and the ugly



Average instruction time, no caching

$$T_{ins} = 1 + lT \quad T = 100 \text{ cycles}$$

With caching

$$T_{ins} = 1 + lmT$$


$P(\text{load/st})$
~0.3

$P(\text{miss})$
~0.05

More energy efficient too

- 7 -

Caches - the good, the bad and the ugly Jacobi Example with Double Buffering




Caches exploit spatial locality here (fetch cache line)
Temporal locality too (discuss)
Network traffic dramatically reduced! Lower energy too

Finish row. As before, use barrier to start next iteration

But is the barrier enough?

- 8 -

Caches - the good, the bad and the ugly Jacobi Example



Cache coherence problem! What do we do?

- 9 -

Maintaining coherence in manycores Major approaches

- User-software managed coherence
 - RP3
 - Beehive
- System-software managed coherence
- Hardware managed coherence (next week)

- 10 -

User-software managed coherence in manycores

Typically yields weak coherence
i.e. Coherence at sync points (or fence pts)

E.g.: When using locks for shared object accesses

Code:

```


shared var
    foo1
    foo2
    foo3
    foo4
    foo_LOCK

GET_foo_LOCK
/* MUNGE WITH foos */
foo1 =
X = foo2
foo3 = .
.
.
RELEASE_foo_LOCK
    
```

How do you make this work?

- 11 -

User Software Coherence



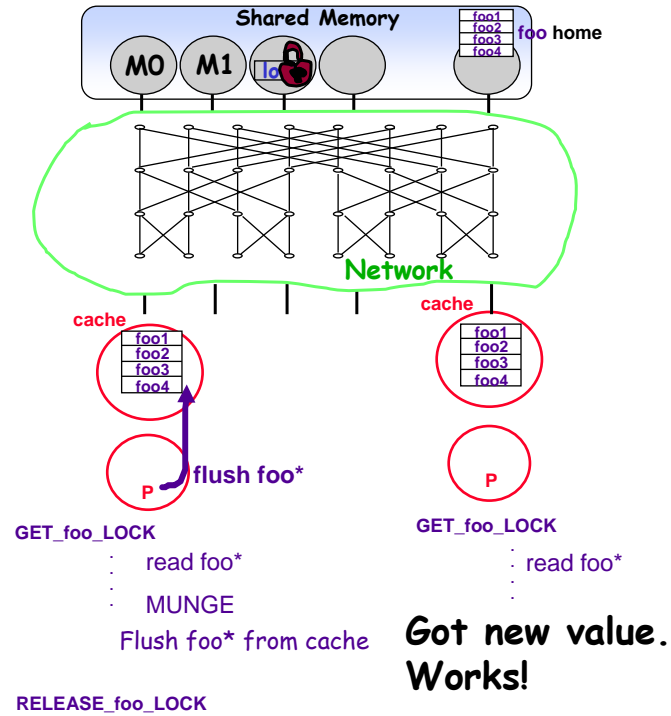
Issues

...the ugly

- Need special processor instructions for flush (e.g., beehive)
- Also need a "memory fence" (wait till all flushed local values are reflected in global store)... next
- Can you cache the lock?
- Must be conservative; when in doubt, flush
 - Lose some locality
- But, can exploit application characteristics to allow some inconsistency
 - e.g. TSP - bound does not have to be accurate
 - Chaotic relaxation

- 12 -

Good Fences Make Good Neighbors




- Need special processor instructions for flush (e.g., beehive)

- 13 -

Good Fences Make Good Neighbors

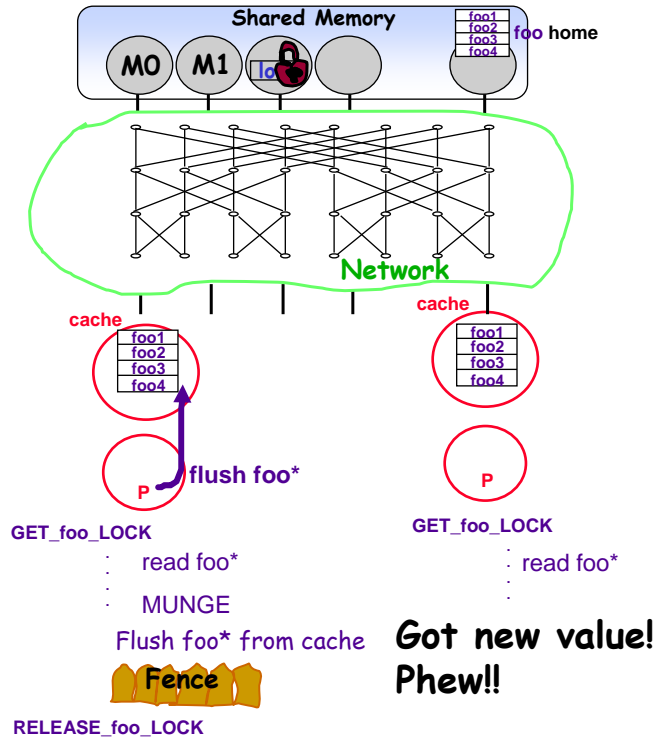
But what about this?



- 14 -

Good Fences Make Good Neighbors

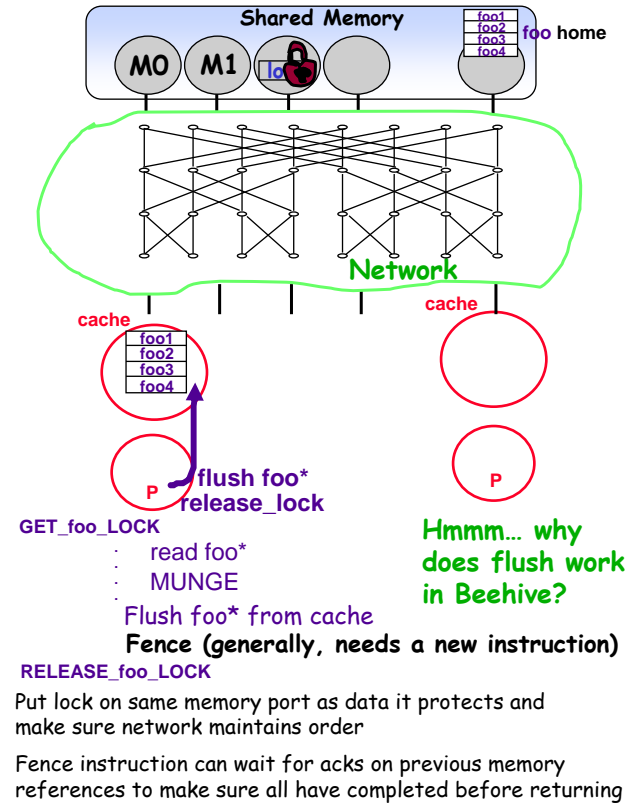
Use a fence



How to implement a fence?

- 15 -

Good Fences Make Good Neighbors How to implement memory fences




- 16 -

Foundations

What is the meaning of shared memory when you have multiple access ports into global memory?

What if you have caches?




Sequential consistency: Final state (of memory) is as if all RDs and WRTs were executed in some fixed serial order (per processor order also maintained) → Lamport

[This notion borrows from similar notions of sequential consistency in transaction processing systems.]

Foundations


A hardware designers physical perspective of sequential consistency



Key: Using fence to wait until flush is done is the key mechanism that guarantees sequential consistency

We will revisit this in more detail in a couple of weeks


One other cache nasty to watch out for



Flush foo* from cache, wait till done

Does it always work?

One other cache nasty to watch out for



Flush yyy from cache,
wait till done

Flush foo* from cache,
wait till done

Correct final value: foo1 yyy

Wrong final value: foo1 xxx

Problem called "False Sharing"


Leads to bugs with sw coherence

Leads to poor perf. with hw coherence

Solutions?

Pad shared data structures so multiple shared items do not fall into same cache line

Hardware Cache Coherence Snooping Caches



- Works for small multicores (mem off chip)
- Broadcast address on shared write
- Everyone listens (snoops) on bus/ring to see if any of their own addresses match
- How do you know when to broadcast, invalidate
 - State associated with each cache line
 - Key benefit: no global state in main mem

- 21 -

Summary of New Multicore Instructions

- Send message
- Receive message
- Synchronization
 - Barrier
 - Test and set
 - F&A and relatives (e.g., F&Op, CmpXch)
- Flush cache line
- Memory fence

- 22 -