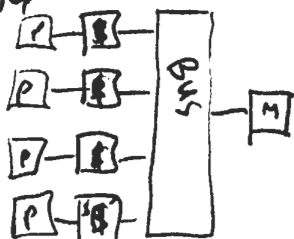


Today we analyze an algorithm that uses an MESI-like protocol.

Setup:



- Caches
- Processors and memory can see all messages
- Only one (for M) can send a message at a time
- Processors (P) talk to caches (C) which talk to bus.

Locality of reference: ~~Spatial locality of a program in~~

Spatial locality is the property that if a processor reads ^{just} location v , it is likely to operate on $v+1$ or other locations near v .

Programs exhibit spatial locality \Rightarrow transmit blocks of size p .

Here

Definition side bar
 $p = \text{block size}$

if $b=1$, fetch costs 2 cycles
send address \rightarrow
fetch word \leftarrow

if b bytes, it takes $b+1$ cycles, saving 2x for reads, b for writes.

Blocks stored at >1 processor sockets. Reduces read cost

To update a replicated block requires communication

- Either invalidate for other copies
 - or update them
- } which is right?

Two common strategies

Exclusive Write: Invalidate everything else

Cache-rot : Update everything else

Worst-case behavior of exclusive-write: there is a bad case



read \Rightarrow total cost is $k \cdot (n-1) \cdot p$

optimal strategy update all (cache rot)
strategy for this sequence is cache rot

cost = $k + np$ all read \leftarrow
update $k + \dots$

$$\lim_{k \rightarrow \infty} \frac{k \cdot (n-1) \cdot p \cdot p}{k + n \cdot p} = \frac{\text{for large } k}{(n-1) \cdot p}$$

(Review:

$$\begin{aligned} \lim_{a \rightarrow \infty} \frac{a \cdot b}{a+c} &= \lim_{a \rightarrow \infty} \frac{(a+c)b - cb}{a+c} \\ &= \lim_{a \rightarrow \infty} \frac{(a+c)b}{a+c} - \lim_{a \rightarrow \infty} \frac{cb}{a+c} \\ &= b - 0 \\ &= b \end{aligned}$$

~~Best~~

Best (worst?) use calculus for packet:

Po units

Pi units

repeat w time

Po units

Total cost $\geq w$

Optimal is ~~exclude wire, cost of DP +~~

drop after writing, for cost of DP

$\lim_{w \rightarrow \infty} \frac{w}{\partial p}$ is unbounded.

Strategy: ~~Competitive algorithm~~

Defn: an on-line algorithm A is c -competitive if,

$\exists d$ ~~cost(A) on a sequence~~

\forall sequences of inputs

optimal $\leq c \cdot \text{cost}(A) + d$ ~~cost(A)~~ \leq OPTIMUM DIFFERENCE MAXIMUM

for fixed sequence.

$c + d$ constant or independent of sequence

Defn: Strongly c -competitive \Rightarrow no better c .

Another algorithm: Goodman's algorithm

On first write, update everything else

On second write, invalidate everything else

Analysis: ? But it is fishy to do something once.
Why not twice?

~~STRATEGY: COMPETITIVE ALGORITHM~~

~~Def'n: An ~~etc~~ online algorithm, is c-competitive (or competitive) iff~~

~~$\exists \forall$ sequences of inputs~~

~~$$\text{cost}(A) \leq c \cdot \text{cost}(\text{OPTIMAL OFFLINE ALG.})$$~~

~~Def'n: An algorithm is an online algorithm if ~~it must make its~~
its behavior depends only on the inputs~~

~~Def'n: An Algo~~

Consider an algorithm that takes a sequence of inputs - produces a sequence of outputs

Def'n Such an algo is on-line if it can produce the i th output ~~with~~,
having seen only the first i inputs, otherwise offline.

~~Intuition: It is offline if it looks at all inputs to determine~~
~~offline. "looks into the future."~~

INTUITION OFFLINE ALGORITHMS LOOK INTO THE FUTURE.

STRATEGY: COMPETITIVE ALGORITHMS

DEF'N. An ON-LINE ALGORITHM A , is c -COMPETITIVE (or
"COMPETITIVE") IFF $\exists d \forall$ sequences of inputs

$$\text{cost}(A) \leq c \cdot \text{cost}(\text{OPTIMAL OFFLINE ALGORITHM}) + d$$

note: c, d independent of sequence.

Def'n: STRONGLY c -competitive if no ~~a~~ better c can be found.

Example of a competitive algorithm: (Due to Long Rudolph)

Skis, ski rent. skis cost \$10 to rent, \$100 to buy.

What should I do?

~~At~~

- Buy skis? \Rightarrow could save once & spend 10x too much.

- Rent skis? \Rightarrow could rent on final day.

2-competitive algorithm:

Rent ~~until~~ ~~to~~ times ~~to~~ 10 times, then buy.

Analysis: if I ski ≤ 10 times, the optimal alg is rent, which I did.

if I ski > 10 times, the optimal alg is buy.

~~Cost of (10x100) = 1000~~

I spent \$200 instead of \$100, a factor of 2 off.

COMPETITIVE SNOOP CACHING, by Karlin, Manasse, Rudolph, Sleator
(No online copy to be found)

Examples using kinds of caches.

~~to~~

I'll focus on DIRECT-MAPPED SNOOP CACHES

- Direct mapped. i.e. \exists a function h mapping block #s to a single slot in each cache.

- caches snoop on variables.

That is processor i updates location v

~~block~~ in its cache, & broadcasts

|

On read, if not in cache must fetch a whole cache line of p bytes

On write, if we communicate, we send the new

~~if dirty~~ \rightarrow

value for v over the bus. This costs 1 cycle.

Definition sidebar

$[v]$ = block # of address

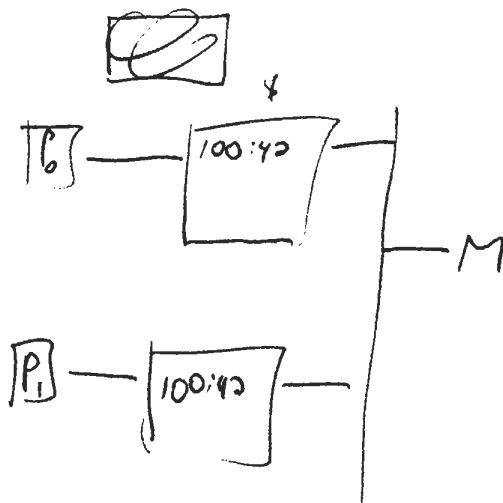
e.g. $[v] = \lfloor \frac{v}{p} \rfloor$ (v/p in C)

$h([v])$ = cache slot #

e.g. $h(B) = B \% \text{mod}$ (Code size)

$h(B) = B \% \text{CACHESIZE}$ in C

Only freedom: when should a processor drop a block from its cache



P_0 writes to 100 over + over.

P_1 may decide to drop the block so P_0 can have it

The only way for data to get into a cache is if the processor reads or writes it.

Once it is there, it must remain up-to-date.

Algorithm: Direct-Mapped-Snoopy-CACHING (DSC)

~~Ass~~
A data structure

$w_i[B]$ an array, i is processor #
 B is ~~block #~~ block #

Invariant:

$w_i[B] = 0$ IFF Block B is \neq cache i .

C a booking keeping integer

Algorithm

TO READ v on processor i :

let $B := [v]$

if $w_i[B] = 0$ then GETBLOCK(i, B)

else

~~$w_i[B]++$~~ (actually can be incremented to any value up to

$w_i[B] := \min(p, w_i[B] + 1)$ (or ANY value in between)

provide B from cache.

To write v from processor i :

let $B := [v]$

if ~~w~~ $w_i(B) = 0$ then GETBLOCK(i, B)

else $w_i(B) := P$

if $\exists j$ s.t. $j \neq i$ and $w_j(B) \neq 0$ then

~~if $i < j$ in same order cache~~

~~up~~

BROADCAST to update on the bus.

$C++$;

$w_j(B) --$;

if $(w_i(B) = 0)$ then drop block B from cache j .

else

update local copy (now B is dirty)

Procedure GETBLOCK(i, B):

if $\exists B'$ s.t. $h(B) = h(B')$ and ~~w~~ $w_i(B) \neq 0$

~~if B collides with B' . Drop B .~~

if B' is dirty then

WRITE B' block to mem

$C++P$;

~~w~~

$w_i(B') = 0$

Drop (i, B')

Fetch block B into cache

$C++P$;

~~w~~ $w_i(B) := P$

Two choices: increment on bus > 1
 choice of j dirty write

Theorem: Algorithm DSC is ~~strongly~~ 2-competitive.

That is, for any sequence of reads + writes, & any
(online or offline) algorithm A

$$\text{COST}_{\text{DSC}}(\sigma) \leq 2 \text{COST}_A(\sigma) + k$$

k depends only on the initial code state, and $k=0$ if all caches are initially empty

We use a potential function Φ

$\Phi(t)$ is a function that depends on the code states of
DSC + A after processing the first t steps of σ .

$$\Phi(t) = \sum_{(i,B) \in S_A} (w_i(B) - 2p) + \sum_{(i,B) \notin S_A} -w_i(B)$$

where S_A is the set of pairs (i,B) s.t. A has block B in cache i at step t .
note $\Phi(t) \leq 0$

We'll prove inductively

$$\text{COST}_{\text{DSC}}(t) - 2 \text{COST}_A(t) \leq \Phi(t) - \Phi(0)$$

for $t=0$, both sides are 0.

The inductive step is to show

$$\Delta \text{COST}_{\text{DSC}} - 2 \Delta \text{COST}_A \leq \Delta \Phi$$

Strategy: given σ construct code sequence τ

τ is ~~the~~ constructed from $\sigma = s$

for each request in σ

Strategy: Break up operations into

Fetchblock(i, B)

Block B added to code i ; cost = p

Drop(i, B)

Block B is dropped from i . (B not dirty)
cost = 0

WRITEBACK(i, B)

Block B is made clean. B is broadcast to
cost = p

Supply(i, v)

Variable v is supplied to proc i by its code.
block $[v]$ must be in code. cost = 0

Update local(i, v)

Variable v is update in code i .
 $[v]$ must be unique for i .
 $[v]$ becomes dirty. cost = 0

Update Global(i, v)

Variable v is updated in code i + broadcast.
 $[v]$ must be clean and in code i .
cost = 1

~~Also~~

Algorithm ~~des~~ performs, for a

Any algorithm must perform those ops; for a read, end with supply(i, v)
for a write, end with update.

consider τ a sequence of these subops.

$\tau =$ subops due to A on ~~the~~ first request in σ , not included to supply or update
subops due to d on first request

~~seq~~

A: Fetch block(); Drop(); supply(); ...

dsc: fetch block(); supply();

Idea: line these supplies + updates up so we have a new seq

$\gamma =$ (A fetch block
~~A~~ drop
 dsc fetch block
 dsc + A supply block
 ...)

} put A actions in first
 } then dsc actions
 } then the common actions

Now define a potential function: depends on cache states of dsc + A after processing the first t steps of γ

$$\Phi(A) = \sum_{(i,B) \in SA} w_i(B)$$

$$\mathcal{J}(t) = \sum_{\substack{(i,B) \\ A \text{ has } B \\ \text{in cache } i \\ \text{at step } t}} (w_i(B) - 2p) + \sum_{\substack{(i,B) \\ A \text{ doesn't} \\ \text{have } B \text{ in} \\ \text{cache } i \\ \text{at step } t}} -w_i(B)$$

note: $\Phi(t) \leq 0$

Prove by induction

$$\text{cost}_{dsc}(t) - 2 \text{cost}_A(t) \leq \Phi(t) - \Phi(0)$$

Theorem will follow, as $K = \Phi(0)$ since $\Phi(t) \leq 0$

BASE CASE: $t=0$, both sides 0

11

INDUCTIVE CASE:

SHOW

$$\Delta \text{cost}_{\text{disc}} - 2 \Delta \text{cost}_A \leq \Delta \Phi$$

CASE ANALYSIS:

if step i is "A does $\text{fetchBlock}(i, B)$ "

$$\Delta C_A = p$$

must show $\Delta \Phi \geq -2p$

Before this code ~~(L)~~ - Block B not in cache i in A .

Afterwards, it is in.

Reverse

$$\begin{aligned} \Delta \Phi &= w_i(B) - 2p - (-w_i(B)) \\ &= 2w_i(B) - 2p \\ &\geq -2p. \end{aligned}$$

Step i is "A drops (i, B) "

$$\Delta C_A = 0$$

must show $\Delta \Phi \geq 0$

Before B in cache i , after B not in cache i .

$$\text{check is } 2p - 2w_i(B) \geq 0$$

"A writes $\text{loc}(i, B)$ "

$$\Delta C_A = p$$

must show $\Delta \Phi \geq -2p$

but $\Delta \Phi = 0$.

"disc Feasible (i, p)

$$\Delta_{disc} = p$$

must show $\Delta \phi \geq p$

$w_i(B)$ changes from 0 to p.

Because we did A's first, A must have B in cell i now.

$$\text{so } \Delta \phi = p$$

⋮

"both supply (i, v) to i

$$\text{Cost to both} = 0$$

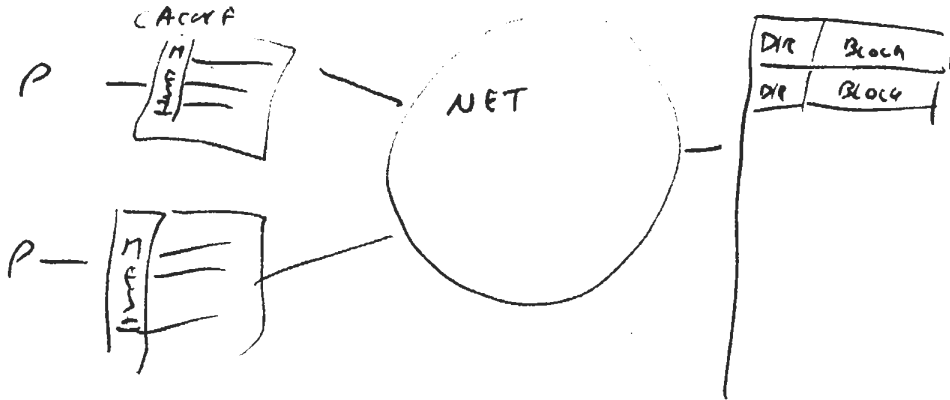
$$\Delta w_i(B) \geq 0 \quad \text{since it's a top}$$

$$B \text{ in } i \text{ so } \Delta \phi \geq 0.$$

DIRECTORY-BASED CACHE

Problem with MESI: Doesn't scale

Solution: Add info to my ~~code~~ block in memory



Idea: leave info at memory to identify who needs to be updated when cache must be invalidated or updated

Examples:

- A

Nobody has it

- B

"Proc S has it"

- C

Proc 3, 5, 9 have S

To obtain E access for proc 2

case A: Fetch from mem, set M to "2 has it in M"

case B: send message to proc S
proc S sends back to mem,
then case A

case C: send messages to 3, 5, 9 to invalidate. (in parallel)
~~then send~~ + rd data to 2
3, 5, 9 send to 2
2 then proceeds to update.

NO ANALYSIS!