Introduction

Why memory subsystem design is important

- CPU speeds increase 55% per year
- DRAM speeds increase 3% (5%?) per year
- rate of increase is also widening

Memory Hierarchy

Levels of memory with different sizes & speeds

- close to the CPU: small, fast access
- close to memory: large, slow access

Memory hierarchies improve performance

- **caches**: demand-driven storage
- principal of **locality of reference**
  - **temporal:** a referenced word will be referenced again soon
  - **spatial:** words near a reference word will be referenced soon
- speed/size trade-off in technology

⇒ fast access for most references

First Cache: IBM 360/85 in the late '60s

Cache Review

Cache organization

- **Block:**
  - # bytes associated with 1 tag
  - usually the # bytes transferred on a memory request
- **Set:** the blocks that can be accessed with the same index bits
- **Associativity:** an implementation that indicates a specific number of blocks in a set
  - direct mapped
  - set associative
  - fully associative
- **Size:** # bytes of **data**

  How do you calculate this?

Logical Diagram of a Cache

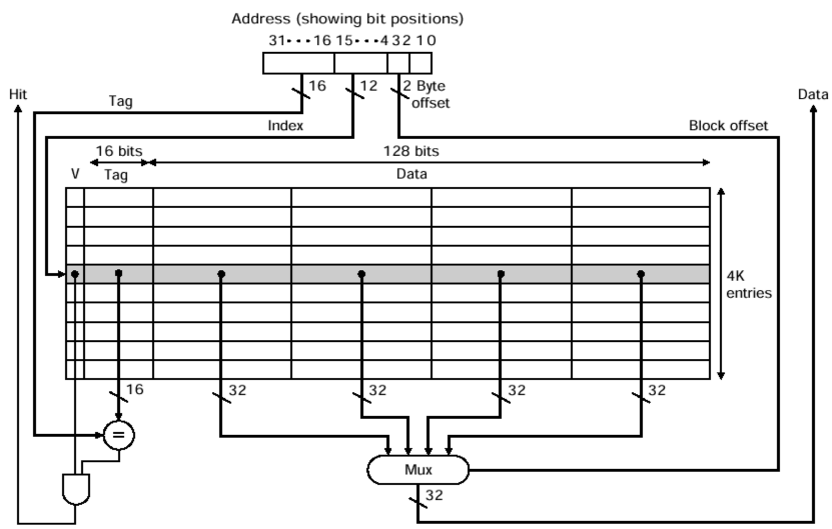

Cache Review

Accessing a cache

- number of index bits = $\log_2$(cache size / block size)
  (for a direct mapped cache)
- number of index bits = $\log_2$(cache size /(block size * associativity))
  (for a set-associative cache)

Logical Diagram of a Set-associative Cache

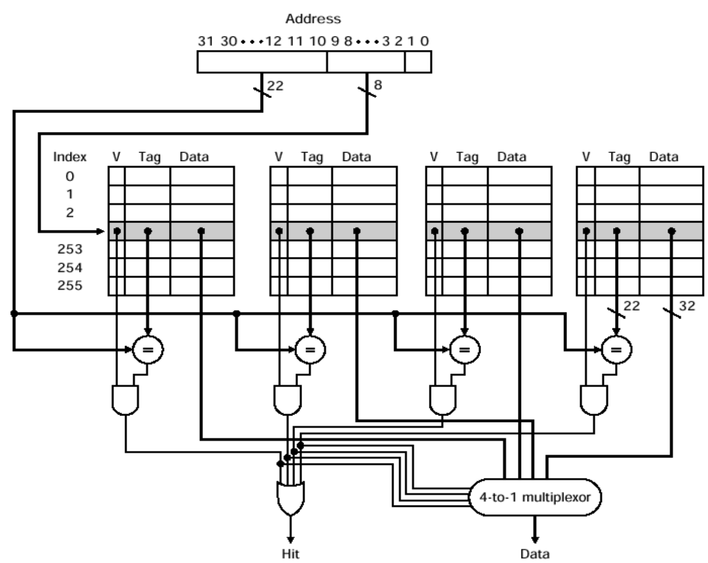

Design Tradeoffs

Cache size

the bigger the cache,

+ the higher the hit ratio
- the longer the access time

Design Tradeoffs

Block size

the bigger the block,

+ the better the spatial locality
+ less block transfer overhead/block
+ less tag overhead/entry (assuming same number of entries)
- might not access all the bytes in the block

Design Tradeoffs

Associativity

the larger the associativity,

+ the higher the hit ratio
- the larger the hardware cost (comparator/set)
- the longer the hit time (a larger MUX)
- need block replacement hardware
- increase in tag bits (if same size cache)

Associativity is more important for small caches than large

because more memory locations map to the same line

e.g., **TLBs**!

Design Tradeoffs

Memory update policy

- **write-through**
  - performance depends on the # of writes
  - write buffer decreases this
    - check on load misses
    - store compression
- **write-back**
  - performance depends on the # of dirty block replacements

    but...
  - dirty bit & logic for checking it
  - tag check before the write
  - must flush the cache before I/O
  - optimization: fetch before replace
- both use a merging write buffer

Design Tradeoffs

Cache contents

- **separate** instruction & data caches
  - separate access ⇒ double the bandwidth
  - different configurations for I & D
  - shorter access time
- **unified** cache
  - lower miss rate
  - less cache controller hardware

Review of Address Translation

Address translation:

- maps a virtual address to a physical address, using the page tables
- number of page offset bits = **page size**

Translation Lookaside Buffer (TLB):

- (often) associative cache of most recently translated virtual-to-physical page mappings
  - 64/128-entry, fully associative
  - 4-8 byte blocks
  - .5 -1 cycle hit time
  - low tens of cycles miss penalty
  - misses can be handled in software, software with hardware assists, firmware or hardware
- works because of locality of reference
- much faster than address translation using the page tables

Using a TLB

(1) Access a TLB using the virtual page number.

(2) If a **hit**,

concatenate the physical page number & the page offset bits, to form a physical address;

set the **reference bit**;

if writing, set the **dirty bit**.

(3) If a **miss**,

get the physical address from the page table;

evict a TLB entry & update dirty/reference bits in the page table;

update the TLB with the new mapping.

Design Tradeoffs

Virtual or physical addressing

Virtually-addressed caches:

- access with a virtual address (index & tag)
- do address translation on a cache miss
- \+ faster for hits because no address translation
- \+ compiler support for better data placement

Design Tradeoffs

Virtually-addressed caches:

- need to flush the cache on a context switch
  - process identification (PID) can avoid this
- synonyms
  - **“the synonym problem”**
    - if 2 processes are sharing data, two (different) virtual addresses map to the same physical address
    - 2 copies of the same data in the cache
    - on a write, only one will be updated; so the other has old data
  - a solution: **page coloring**
    - processes share segments, so all data aliases have the same low-order bits, i.e., same offset from the beginning of a segment
    - cache must be <= the segment size
      (more precisely, each set of the cache must be <= the segment size)
    - index taken from offset, tag compare on segment #

Design Tradeoffs

Virtual or physical addressing

physically addressed caches

- do address translation on every cache access
- access with a physical index & compare with physical tag
- if a straightforward implementation, hit time increases because must translate the virtual address before access the cache
  - \+ increase in hit time can be avoided if address translation is done in parallel with the cache access
    - restrict cache size so that cache index bits are in the page offset (virtual & physical bits are the same): **virtually indexed**
    - access the TLB at the same time
    - compare the physical tag from the cache to the physical address (page frame #) from the TLB: **physically tagged**
    - can increase cache size by increasing associativity, but still use page offset bits for the index
- \+ no cache flushing on a context switch
- \+ no synonym problem

Cache Hierarchies

Cache hierarchy

- different caches with different sizes & access times & purposes

\+ decrease effective memory access time:
  - many misses in the L1 cache will be satisfied by the L2 cache
  - avoid going all the way to memory

Cache Hierarchies

Level 1 cache goal: fast access
so minimize hit time (the common case)

Cache Hierarchies

Level 1 cache goal: fast access
so minimize hit time (the common case)

- small, so can access it in one (or few) CPU cycle
  (also some chip real estate concerns)
- virtually-accessed, so cache accesses can be fast without constraints on cache size
- direct mapped or set associative?
  - direct mapped: faster access
  - set associative: better hit ratio & larger cache while still accessing with virtual address
- often write-through
- separate caches for instructions & data
  - each is smaller than a unified cache, so the access time is lower
  - need the instruction/data parallel access in a pipelined processor
  - configured differently:
    - instruction cache may have larger blocks
    - instruction cache has no write-back hardware & dirty bits

Cache Hierarchies

Level 2 cache goal: keep traffic off the system bus

Cache Hierarchies

Level 2 cache goal: keep traffic off the system bus

- big cache, so it will have a high hit ratio
- physically-addressed:
  - plenty of time to do address translation
  - no flushing on a context switch
- direct mapped:
  - big direct-mapped caches have almost the same hit ratio as big set-associative caches
  - slightly less hardware cost
- unified, because its hit ratio is higher than that of two separate caches (I&D) half the size
- write-back

Review of Cache Metrics

Hit (miss) ratio = $\frac{\#hits\ (\#misses)}{\#references}$

- measures how well the cache functions
- useful for understanding cache behavior relative to the number of references
- intermediate metric

Effective access time = $HitTime + MissRatio \bullet MissPenalty$

- (rough) average time it takes to do a memory reference
- performance of the memory system, including factors that depend on the implementation
- intermediate metric

Measuring Cache Hierarchy Performance

Effective Access Time for a cache hierarchy:...

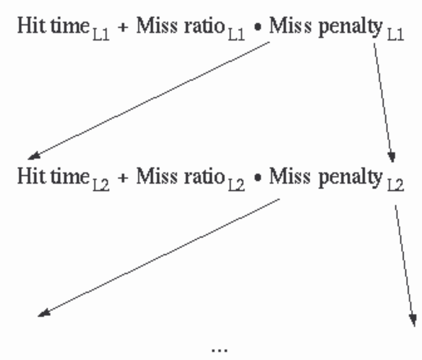

Measuring Cache Hierarchy Performance

Local Miss Ratio: $\frac{\#\text{misses}}{\#\text{accesses}}$ for that cache!

- # accesses for the L1 cache: the number of references
- # accesses for the L2 cache: the number of misses in the L1 cache

Example: 1000 references
40 L1 misses
10 L2 misses

local MR (L1):
local MR (L2):

Measuring Cache Hierarchy Performance

Global Miss Ratio:

$$\text{globalMR} = \frac{\#\,\text{misses in cache}}{\#\,\text{references generated by CPU}}$$

Example: 1000 References
40 L1 misses
10 L2 misses

global MR (L1):

global MR (L2):

Miss Classification

Usefulness is in providing insight into the causes of misses

- does not explain what caused particular, individual misses

Compulsory

- first reference misses
- decrease by increasing block size

Capacity

- due to finite size of the cache
- decrease by increasing cache size

Conflict

- too many blocks map to the same set
- decrease by increasing associativity

Coherence (invalidation)

- decrease by decreasing block size + improving processor locality