18-447: Computer Architecture Lecture 22: Memory Hierarchy and Caches

Prof. Onur Mutlu
Carnegie Mellon University
Spring 2013, 3/27/2013

Reminder: Homework 5

- Due April 1
- Topics: Vector processing, VLIW, Virtual memory, Caching

Reminder: Lab Assignment 5

- Lab Assignment 5
 - Due Friday, April 5
 - Modeling caches and branch prediction at the microarchitectural level (cycle level) in C
 - Extra credit: Cache design optimization
 - Size, block size, associativity
 - Replacement and insertion policies
 - Cache indexing policies
 - Anything else you would like
 - TAs will go over the baseline simulator in lab sessions

Readings for Today and Next Lecture

- Memory Hierarchy and Caches
- Cache chapters from P&H: 5.1-5.3
- Memory/cache chapters from Hamacher+: 8.1-8.7
- An early cache paper by Maurice Wilkes
 - Wilkes, "Slave Memories and Dynamic Storage Allocation,"
 IEEE Trans. On Electronic Computers, 1965.

Today

The memory hierarchy

Caches

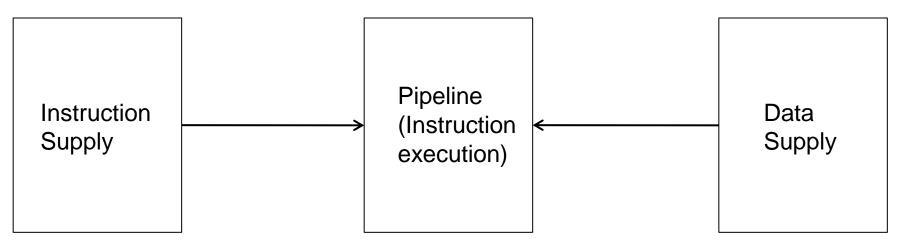
Also Today: Oracle Tech Day

- "What Keeps Us Up at Night"
- March 27, 2013 (Today!)
- 4:30 6:30pm
- Singleton Room, Roberts Engineering Hall
- Hardware and software engineers must balance customers' wants and needs with an eye towards the future and make the right choices and tradeoffs today for how technologies will unfold in the next five years.
- ... even in industry, a solution may arise only after many sleepless nights.

Also Today: Oracle Tech Day

- "What Keeps Us Up at Night"
- March 27, 2013 (Today!)
- 4:30 6:30pm
- Singleton Room, Roberts Engineering Hall
- Please join Oracle senior engineers and engineering managers from Database, Exadata, TimesTen, Solaris, and Oracle Virtualization and SPARC Architecture teams to see what problems keep them up at night.
- There will be a one-hour presentation, followed by a one-hour reception where Oracle engineers are available for ample discussion and commiseration.

Idealism



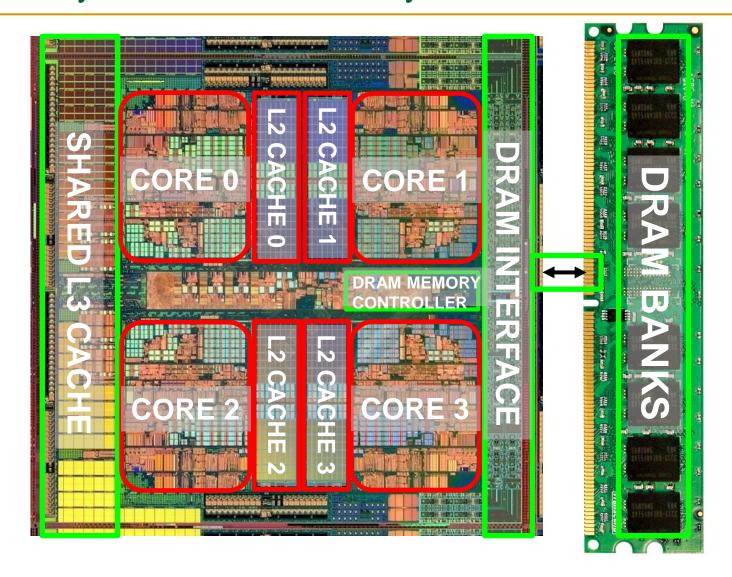
- Zero-cycle latency
- Infinite capacity
- Zero cost
- Perfect control flow

- No pipeline stalls
- -Perfect data flow (reg/memory dependencies)
- Zero-cycle interconnect (operand communication)
- Enough functional units
- Zero latency compute

- Zero-cycle latency
- Infinite capacity
- Infinite bandwidth
- Zero cost

The Memory Hierarchy

Memory in a Modern System



Ideal Memory

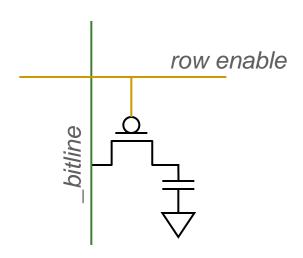
- Zero access time (latency)
- Infinite capacity
- Zero cost
- Infinite bandwidth (to support multiple accesses in parallel)

The Problem

- Ideal memory's requirements oppose each other
- Bigger is slower
 - □ Bigger → Takes longer to determine the location
- Faster is more expensive
 - Memory technology: SRAM vs. DRAM
- Higher bandwidth is more expensive
 - Need more banks, more ports, higher frequency, or faster technology

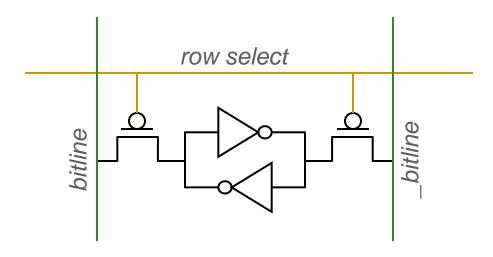
Memory Technology: DRAM

- Dynamic random access memory
- Capacitor charge state indicates stored value
 - Whether the capacitor is charged or discharged indicates storage of 1 or 0
 - 1 capacitor
 - 1 access transistor
- Capacitor leaks through the RC path
 - DRAM cell loses charge over time
 - DRAM cell needs to be refreshed

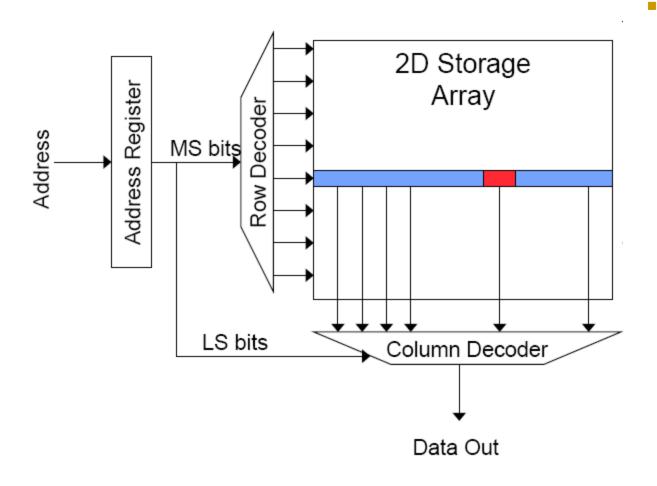


Memory Technology: SRAM

- Static random access memory
- Two cross coupled inverters store a single bit
 - Feedback path enables the stored value to persist in the "cell"
 - 4 transistors for storage
 - 2 transistors for access

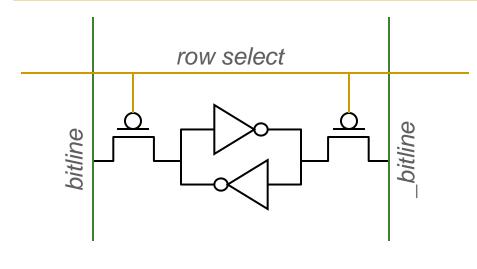


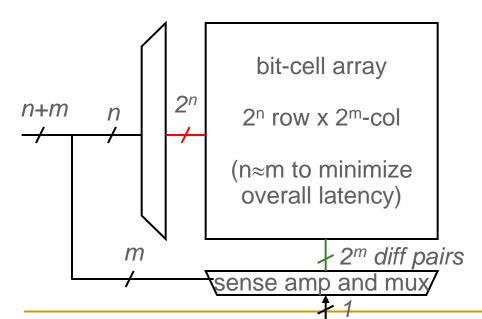
Memory Bank Organization and Operation



- Read access sequence:
 - Decode row address
 drive word-lines
 - 2. Selected bits drive bit-lines
 - Entire row read
 - 3. Amplify row data
 - 4. Decode column address & select subset of row
 - Send to output
 - 5. Precharge bit-lines
 - For next access

SRAM (Static Random Access Memory)





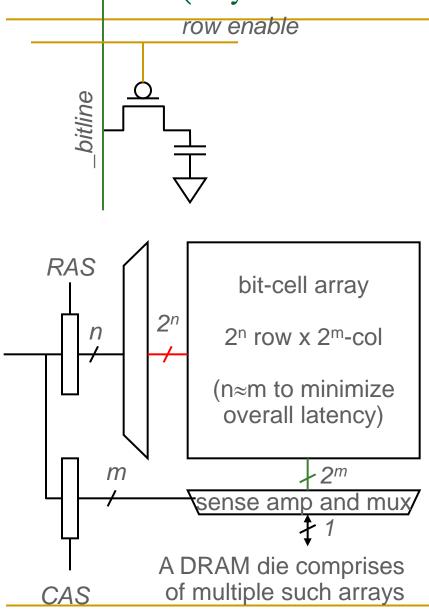
Read Sequence

- 1. address decode
- 2. drive row select
- 3. selected bit-cells drive bitlines (entire row is read together)
- 4. differential sensing and column select (data is ready)
- 5. precharge all bitlines (for next read or write)

Access latency dominated by steps 2 and 3 Cycling time dominated by steps 2, 3 and 5

- step 2 proportional to 2^m
- step 3 and 5 proportional to 2ⁿ

DRAM (Dynamic Random Access Memory)



Bits stored as charges on node capacitance (non-restorative)

- bit cell loses charge when read
- bit cell loses charge over time

Read Sequence

- 1~3 same as SRAM
- 4. a "flip-flopping" sense amp amplifies and regenerates the bitline, data bit is mux' ed out
- 5. precharge all bitlines

Destructive reads

Charge loss over time

Refresh: A DRAM controller must periodically read each row within the allowed refresh time (10s of ms) such that charge is restored

DRAM vs. SRAM

DRAM

- Slower access (capacitor)
- Higher density (1T 1C cell)
- Lower cost
- Requires refresh (power, performance, circuitry)
- Manufacturing requires putting capacitor and logic together

SRAM

- Faster access (no capacitor)
- Lower density (6T cell)
- Higher cost
- No need for refresh
- Manufacturing compatible with logic process (no capacitor)

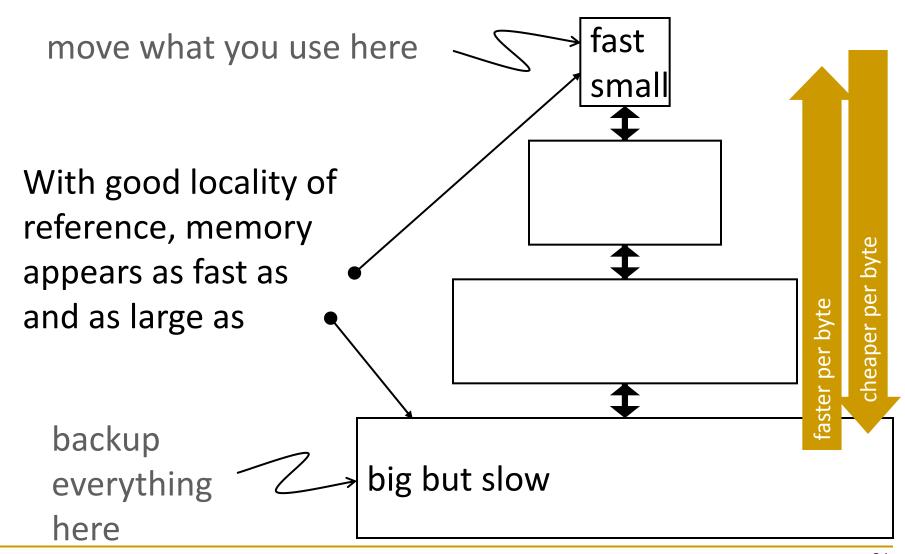
The Problem

- Bigger is slower
 - SRAM, 512 Bytes, sub-nanosec
 - □ SRAM, KByte~MByte, ~nanosec
 - □ DRAM, Gigabyte, ~50 nanosec
 - □ Hard Disk, Terabyte, ~10 millisec
- Faster is more expensive (dollars and chip area)
 - □ SRAM, < 10\$ per Megabyte</p>
 - DRAM, < 1\$ per Megabyte
 - □ Hard Disk < 1\$ per Gigabyte
 - These sample values scale with time
- Other technologies have their place as well
 - Flash memory, Phase-change memory (not mature yet)

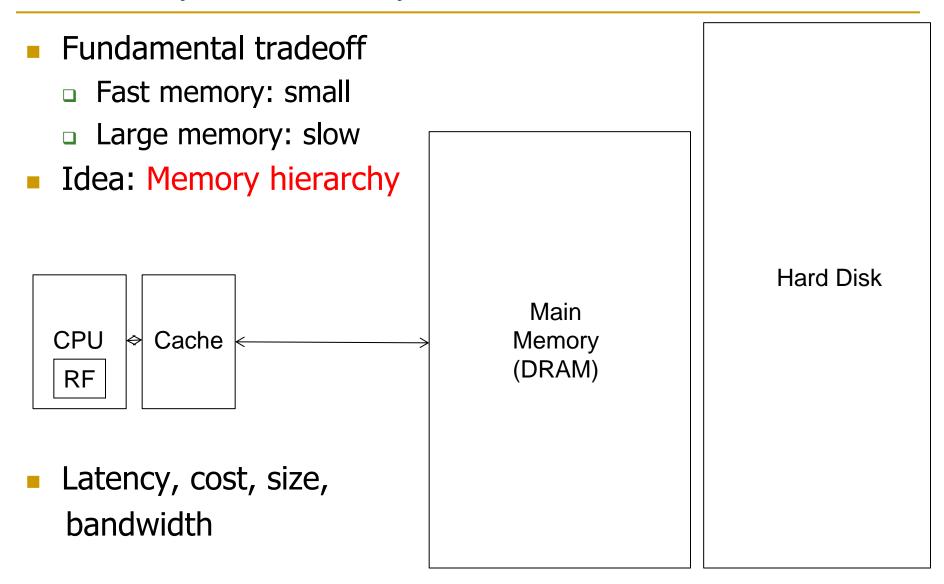
Why Memory Hierarchy?

- We want both fast and large
- But we cannot achieve both with a single level of memory
- Idea: Have multiple levels of storage (progressively bigger and slower as the levels are farther from the processor) and ensure most of the data the processor needs is kept in the fast(er) level(s)

The Memory Hierarchy



Memory Hierarchy



Locality

- One's recent past is a very good predictor of his/her near future.
- Temporal Locality: If you just did something, it is very likely that you will do the same thing again soon
 - since you are here today, there is a good chance you will be here again and again regularly
- Spatial Locality: If you did something, it is very likely you will do something similar/related (in space)
 - every time I find you in this room, you are probably sitting close to the same people

Memory Locality

- A "typical" program has a lot of locality in memory references
 - typical programs are composed of "loops"
- Temporal: A program tends to reference the same memory location many times and all within a small window of time
- Spatial: A program tends to reference a cluster of memory locations at a time
 - most notable examples:
 - 1. instruction memory references
 - 2. array/data structure references

Caching Basics: Exploit Temporal Locality

- Idea: Store recently accessed data in automatically managed fast memory (called cache)
- Anticipation: the data will be accessed again soon
- Temporal locality principle
 - Recently accessed data will be again accessed in the near future
 - This is what Maurice Wilkes had in mind:
 - Wilkes, "Slave Memories and Dynamic Storage Allocation," IEEE Trans. On Electronic Computers, 1965.
 - "The use is discussed of a fast core memory of, say 32000 words as a slave to a slower core memory of, say, one million words in such a way that in practical cases the effective access time is nearer that of the fast memory than that of the slow memory."

Caching Basics: Exploit Spatial Locality

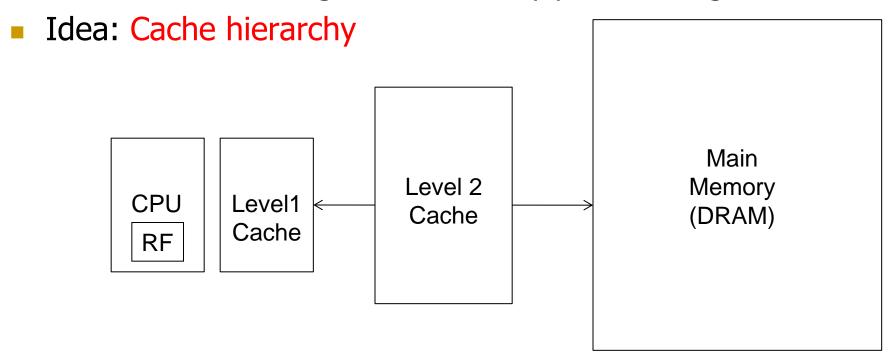
- Idea: Store addresses adjacent to the recently accessed one in automatically managed fast memory
 - Logically divide memory into equal size blocks
 - Fetch to cache the accessed block in its entirety
- Anticipation: nearby data will be accessed soon
- Spatial locality principle
 - Nearby data in memory will be accessed in the near future
 - E.g., sequential instruction access, array traversal
 - This is what IBM 360/85 implemented
 - 16 Kbyte cache with 64 byte blocks
 - Liptay, "Structural aspects of the System/360 Model 85 II: the cache," IBM Systems Journal, 1968.

The Bookshelf Analogy

- Book in your hand
- Desk
- Bookshelf
- Boxes at home
- Boxes in storage
- Recently-used books tend to stay on desk
 - Comp Arch books, books for classes you are currently taking
 - Until the desk gets full
- Adjacent books in the shelf needed around the same time
 - If I have organized/categorized my books well in the shelf

Caching in a Pipelined Design

- The cache needs to be tightly integrated into the pipeline
 - Ideally, access in 1-cycle so that dependent operations do not stall
- High frequency pipeline → Cannot make the cache large
 - But, we want a large cache AND a pipelined design



A Note on Manual vs. Automatic Management

- Manual: Programmer manages data movement across levels
 - -- too painful for programmers on substantial programs
 - "core" vs "drum" memory in the 50's
 - astill done in some embedded processors (on-chip scratch pad SRAM in lieu of a cache)
- Automatic: Hardware manages data movement across levels, transparently to the programmer
 - ++ programmer's life is easier
 - simple heuristic: keep most recently used items in cache
 - □the average programmer doesn't need to know about it
 - You don't need to know how big the cache is and how it works to write a "correct" program! (What if you want a "fast" program?)

Automatic Management in Memory Hierarchy

Wilkes, "Slave Memories and Dynamic Storage Allocation," IEEE Trans. On Electronic Computers, 1965.

Slave Memories and Dynamic Storage Allocation

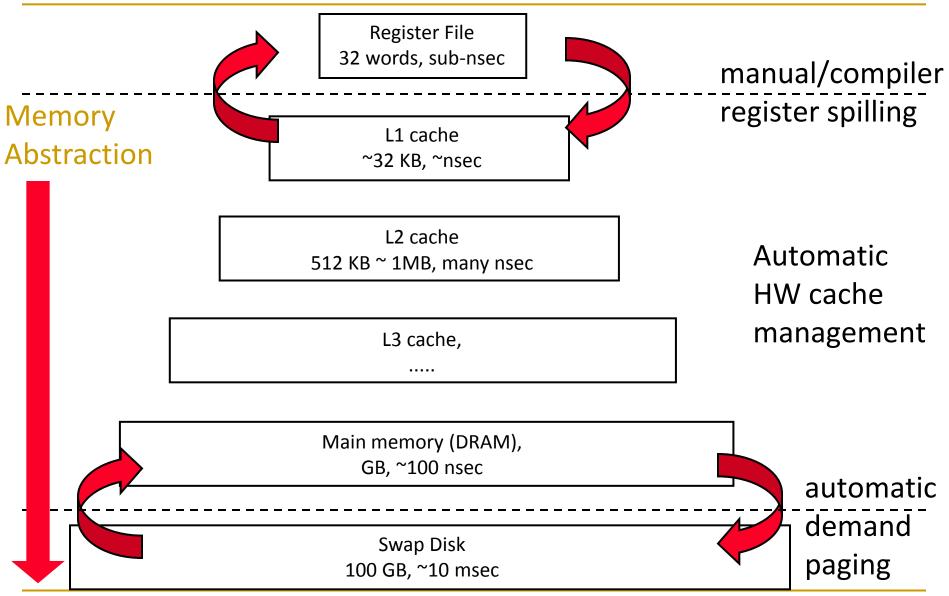
M. V. WILKES

SUMMARY

The use is discussed of a fast core memory of, say, 32 000 words as a slave to a slower core memory of, say, one million words in such a way that in practical cases the effective access time is nearer that of the fast memory than that of the slow memory.

"By a slave memory I mean one which automatically accumulates to itself words that come from a slower main memory, and keeps them available for subsequent use without it being necessary for the penalty of main memory access to be incurred again."

A Modern Memory Hierarchy



Hierarchical Latency Analysis

- For a given memory hierarchy level i it has a technology-intrinsic access time of t_i. The perceived access time T_i is longer than t_i
- Except for the outer-most hierarchy, when looking for a given address there is
 - □ a chance (hit-rate h_i) you "hit" and access time is t_i
 - a chance (miss-rate m_i) you "miss" and access time t_i +T_{i+1}
 - $h_i + m_i = 1$
- Thus

$$T_i = h_i \cdot t_i + m_i \cdot (t_i + T_{i+1})$$
$$T_i = t_i + m_i \cdot T_{i+1}$$

keep in mind, h_i and m_i are defined to be the hit-rate and miss-rate of just the references that missed at L_{i-1}

Hierarchy Design Considerations

Recursive latency equation

$$T_i = t_i + m_i \cdot T_{i+1}$$

- The goal: achieve desired T₁ within allowed cost
- $T_i \approx t_i$ is desirable
- Keep m_i low
 - increasing capacity C_i lowers m_i, but beware of increasing t_i
 - lower m_i by smarter management (replacement::anticipate what you don't need, prefetching::anticipate what you will need)
- Keep T_{i+1} low
 - faster lower hierarchies, but beware of increasing cost
 - introduce intermediate hierarchies as a compromise

Intel Pentium 4 Example

- 90nm P4, 3.6 GHz
- L1 D-cache

 - \Box $t_1 = 4$ cyc int / 9 cycle fp
- L2 D-cache
 - \Box C₂ = 1024 KB
 - = t₂ = 18 cyc int / 18 cyc fp
- Main memory
 - $t_3 = 70$ 50ns or 180 cyc
- Notice
 - best case latency is not 1
 - worst case access latencies are into 500+ cycles

if
$$m_1=0.1$$
, $m_2=0.1$
 $T_1=7.6$, $T_2=36$

if
$$m_1=0.01$$
, $m_2=0.01$
 $T_1=4.2$, $T_2=19.8$

if
$$m_1=0.05$$
, $m_2=0.01$
 $T_1=5.00$, $T_2=19.8$

if
$$m_1=0.01$$
, $m_2=0.50$
 $T_1=5.08$, $T_2=108$

Cache Basics and Operation

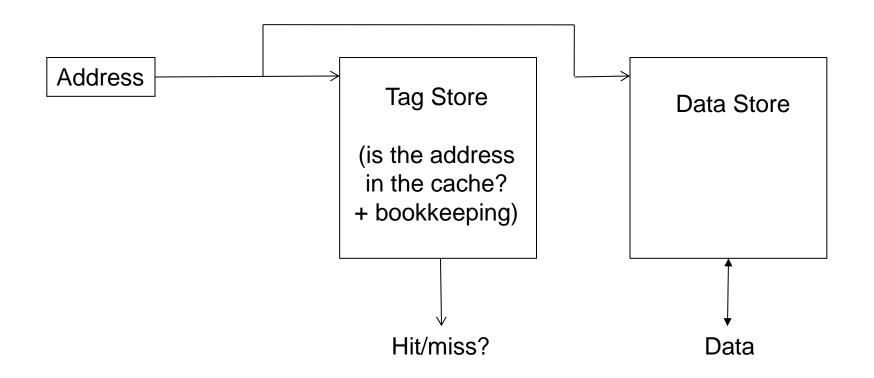
Cache

- Generically, any structure that "memoizes" frequently used results to avoid repeating the long-latency operations required to reproduce the results from scratch, e.g. a web cache
- Most commonly, an automatically-managed memory hierarchy based on SRAM
 - memoize in SRAM the most frequently accessed DRAM memory locations to avoid repeatedly paying for the DRAM access latency

Caching Basics

- Block (line): Unit of storage in the cache
 - Memory is logically divided into cache blocks that map to locations in the cache
- When data referenced
 - HIT: If in cache, use cached data instead of accessing memory
 - MISS: If not in cache, bring block into cache
 - Maybe have to kick something else out to do it
- Some important cache design decisions
 - Placement: where and how to place/find a block in cache?
 - Replacement: what data to remove to make room in cache?
 - Granularity of management: large, small, uniform blocks?
 - Write policy: what do we do about writes?
 - Instructions/data: Do we treat them separately?

Cache Abstraction and Metrics



- Cache hit rate = (# hits) / (# hits + # misses) = (# hits) / (# accesses)
- Average memory access time (AMAT)= (hit-rate * hit-latency) + (miss-rate * miss-latency)
- Aside: Can reducing AMAT reduce performance?

Blocks and Addressing the Cache

- Memory is logically divided into cache blocks
- Each block maps to a location in the cache, determined by the index bits in the address
 - used to index into the tag and data stores

tag index byte in block

2b 3 bits 3 bits

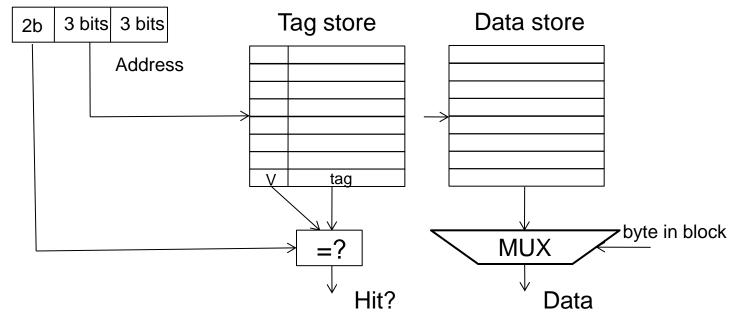
8-bit address

- Cache access: index into the tag and data stores with index bits in address, check valid bit in tag store, compare tag bits in address with the stored tag in tag store
- If a block is in the cache (cache hit), the tag store should have the tag of the block stored in the index of the block

Direct-Mapped Cache: Placement and Access

- Assume byte-addressable memory:
 256 bytes, 8-byte blocks → 32 blocks
- Assume cache: 64 bytes, 8 blocks
 - Direct-mapped: A block can go to only one location

tag index byte in block



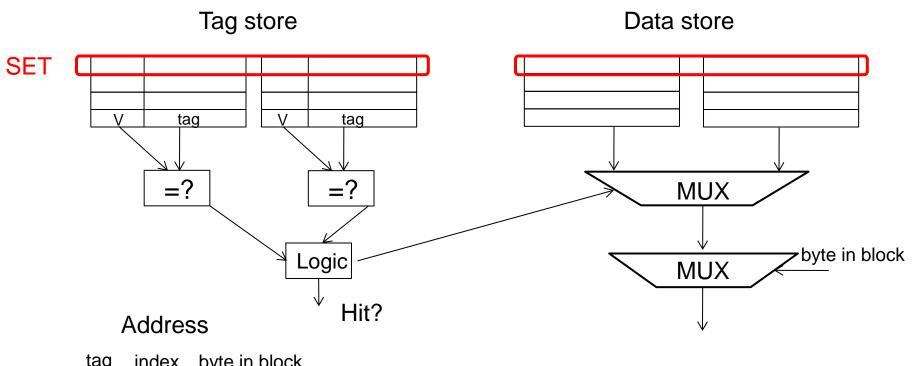
- Addresses with same index contend for the same location
 - Cause conflict misses

Direct-Mapped Caches

- Direct-mapped cache: Two blocks in memory that map to the same index in the cache cannot be present in the cache at the same time
 - □ One index → one entry
- Can lead to 0% hit rate if more than one block accessed in an interleaved manner map to the same index
 - Assume addresses A and B have the same index bits but different tag bits
 - \Box A, B, A, B, A, B, ... \rightarrow conflict in the cache index
 - All accesses are conflict misses

Set Associativity

- Addresses 0 and 8 always conflict in direct mapped cache
- Instead of having one column of 8, have 2 columns of 4 blocks



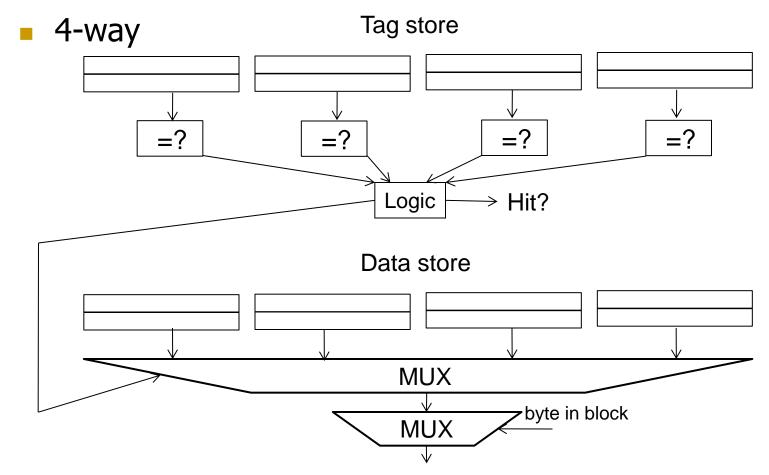
tag index byte in block

3b 2 bits 3 bits

Associative memory within the set

- -- More complex, slower access, larger tag store
- + Accommodates conflicts better (fewer conflict misses)

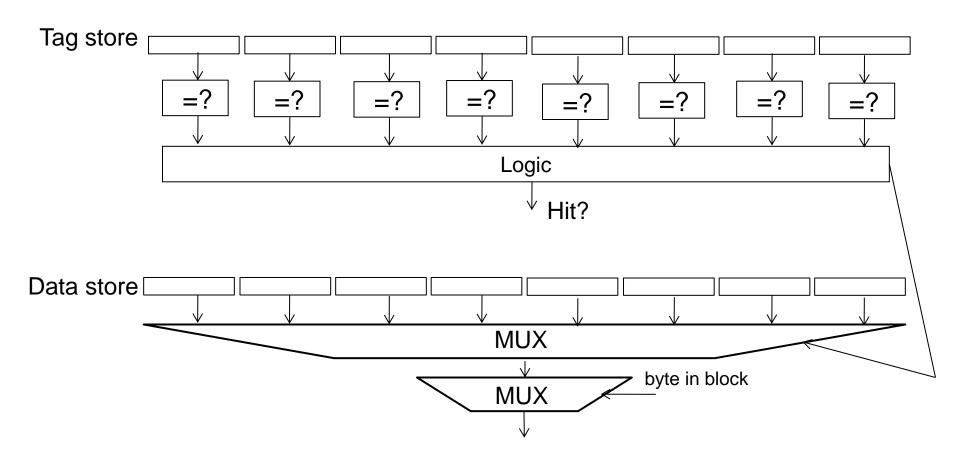
Higher Associativity



- -- More tag comparators and wider data mux; larger tags
- + Likelihood of conflict misses even lower

Full Associativity

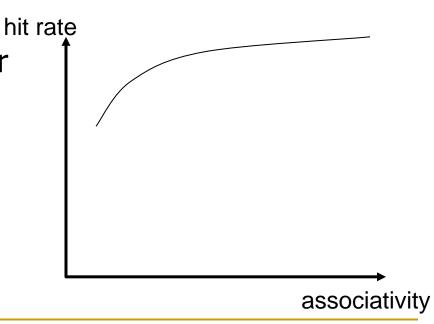
- Fully associative cache
 - A block can be placed in any cache location



Associativity (and Tradeoffs)

- How many blocks can map to the same index (or set)?
- Higher associativity
 - ++ Higher hit rate
 - -- Slower cache access time (hit latency and data access latency)
 - -- More expensive hardware (more comparators)

Diminishing returns from higher associativity



Set-Associative Caches (I)

- Diminishing returns in hit rate from higher associativity
- Longer access time with higher associativity
- Which block in the set to replace on a cache miss?
 - Any invalid block first
 - If all are valid, consult the replacement policy
 - Random
 - FIFO
 - Least recently used (how to implement?)
 - Not most recently used
 - Least frequently used?
 - Least costly to re-fetch?
 - Why would memory accesses have different cost?
 - Hybrid replacement policies
 - Optimal replacement policy?

Implementing LRU

- Idea: Evict the least recently accessed block
- Problem: Need to keep track of access ordering of blocks
- Question: 2-way set associative cache:
 - What do you need to implement LRU?
- Question: 4-way set associative cache:
 - How many different orderings possible for the 4 blocks in the set?
 - How many bits needed to encode the LRU order of a block?
 - What is the logic needed to determine the LRU victim?

Approximations of LRU

 Most modern processors do not implement "true LRU" in highly-associative caches

Why?

- True LRU is complex
- LRU is an approximation to predict locality anyway (i.e., not the best possible replacement policy)

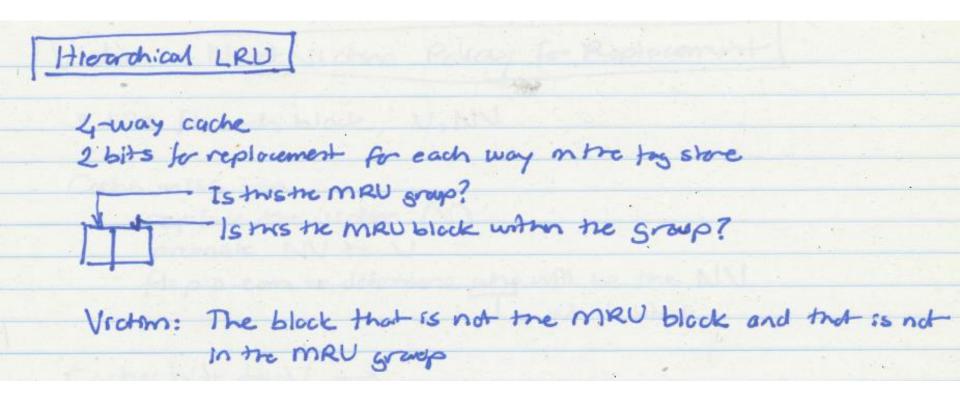
Examples:

- Not MRU (not most recently used)
- Hierarchical LRU: divide the 4-way set into 2-way "groups", track the MRU group and the MRU way in each group
- Victim-NextVictim Replacement: Only keep track of the victim and the next victim

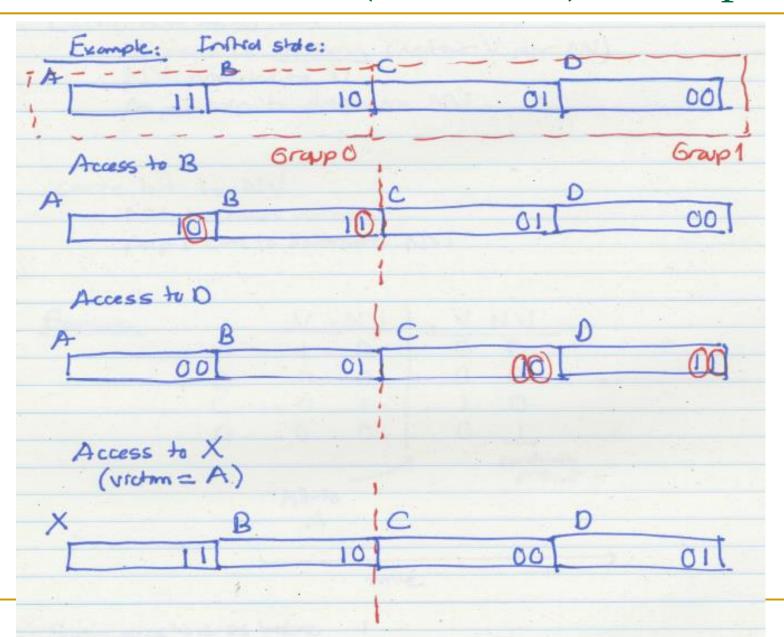
Hierarchical LRU (not MRU)

- Divide a set into multiple groups
- Keep track of the MRU group
- Keep track of the MRU block in each group
- On replacement, select victim as:
 - A not-MRU block in one of the not-MRU groups

Hierarchical LRU (not MRU) Example



Hierarchical LRU (not MRU) Example



Hierarchical LRU (not MRU): Questions

- 8-way cache
- 2 4-way groups
- What is an access pattern that performs worse than true LRU?
- What is an access pattern that performs better than true LRU?

Victim/Next-Victim Policy

- Only 2 blocks' status tracked in each set:
 - victim (V), next victim (NV)
 - □ all other blocks denoted as (O) Ordinary block
- On a cache miss
 - Replace V
 - Promote NV to V
 - Randomly pick an O block as NV
- On a cache hit to V
 - Promote NV to V
 - Randomly pick an O block as NV
 - Turn V to O

Victim/Next-Victim Policy (II)

- On a cache hit to NV
 - Randomly pick an O block as NV
 - Turn NV to O
- On a cache hit to O
 - Do nothing

Victim/Next-Victim Example

