The Memory Hierarchy
CS 740

24 January 2017

Topics

- Memories
- Cache design

L3 Reading

Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers

Famous paper; won "Test of Time Award"

Norm Jouppi

- DEC, Hewlett-Packard, Google (TPUs)
- Winner of Eckert-Mauchly Award
- "The Nobel Prize of Computer Architecture"

What is the key idea of the paper?

How are ideas presented and evaluated?

Computer System

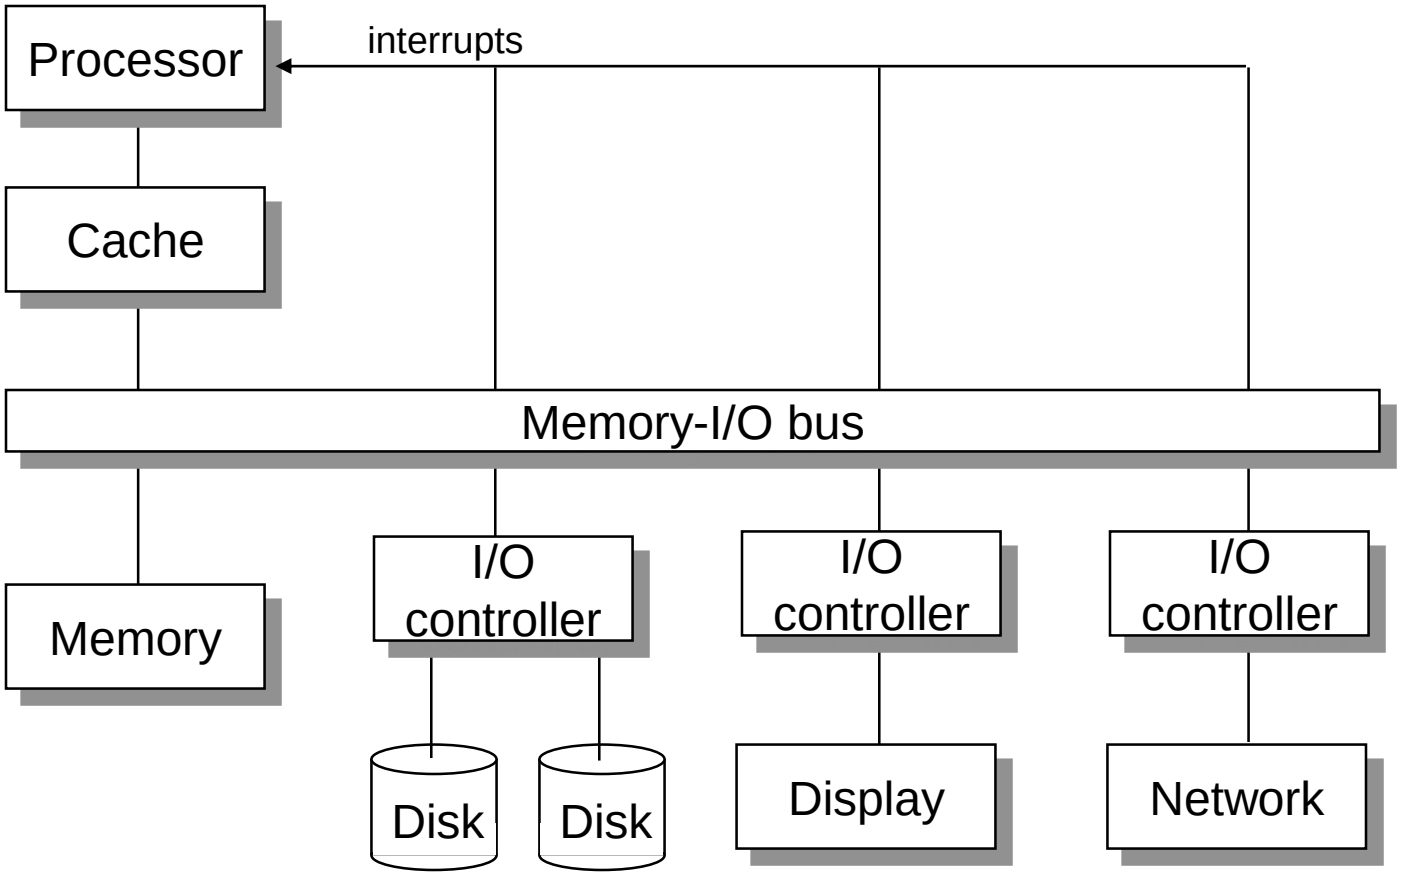

The Tradeoff

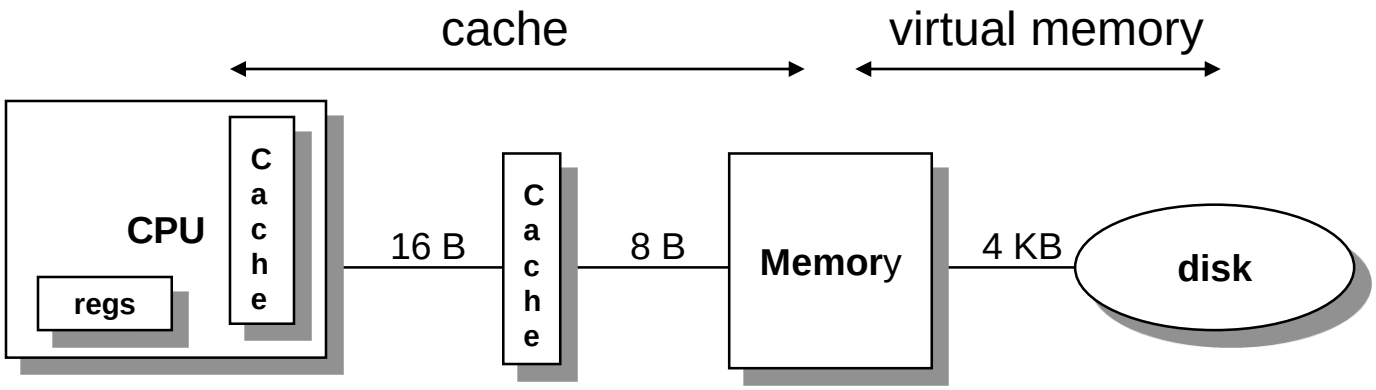

	register reference	L1-cache reference	L2-cache reference	memory reference	disk memory reference
size:	608 B	128k B	512kB -- 4MB	128 MB	27GB
speed:	1.4 ns	4.2 ns	16.8 ns	112 ns	9 ms
$/Mbyte:			$90/MB	$2-6/MB	$0.01/MB
block size:	4 B	4 B	16 B	4-8 KB	

larger, slower, cheaper

(Numbers are for a 21264 at 700MHz)

The Tradeoff

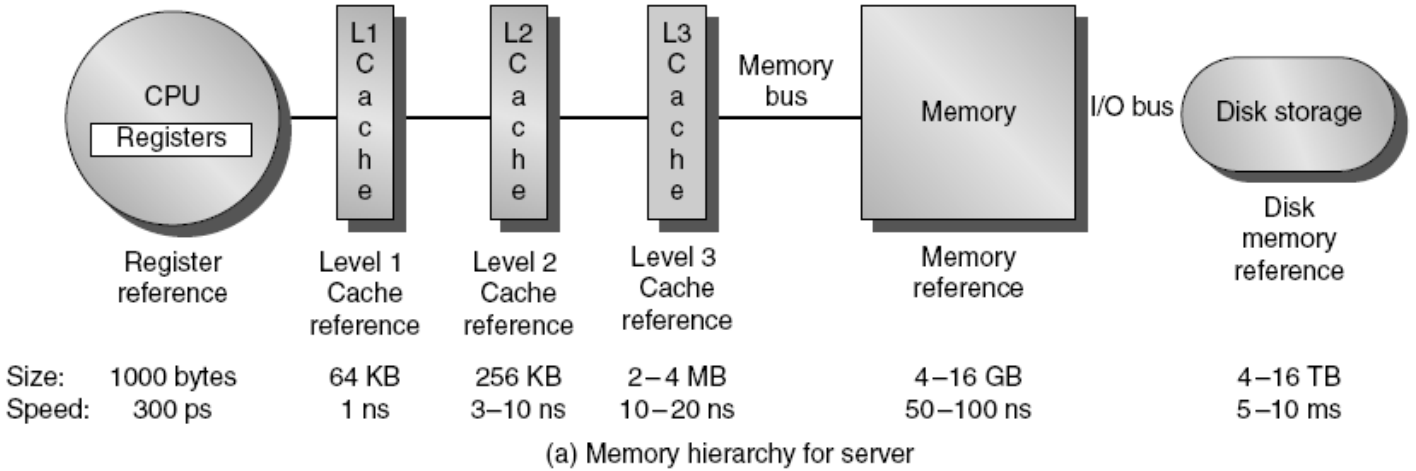

(a) Memory hierarchy for server

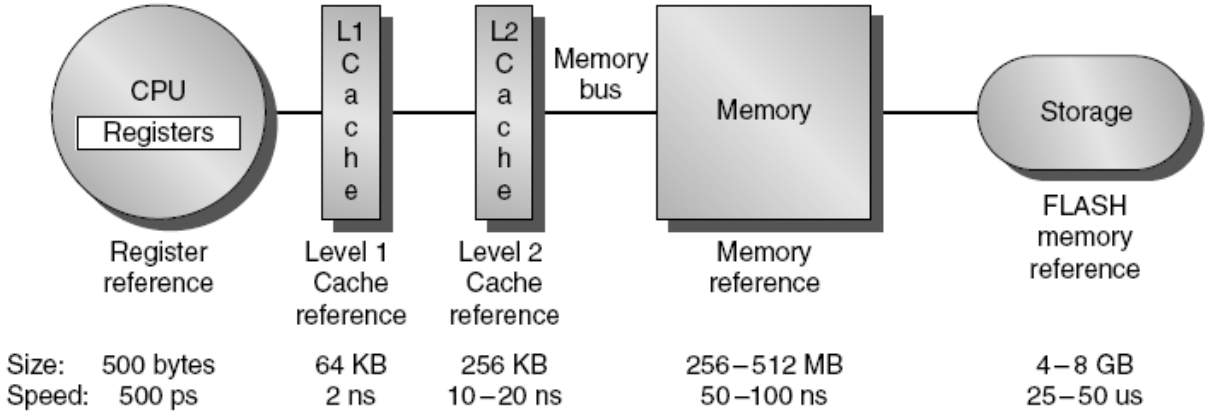

(b) Memory hierarchy for a personal mobile device

The Performance Gap

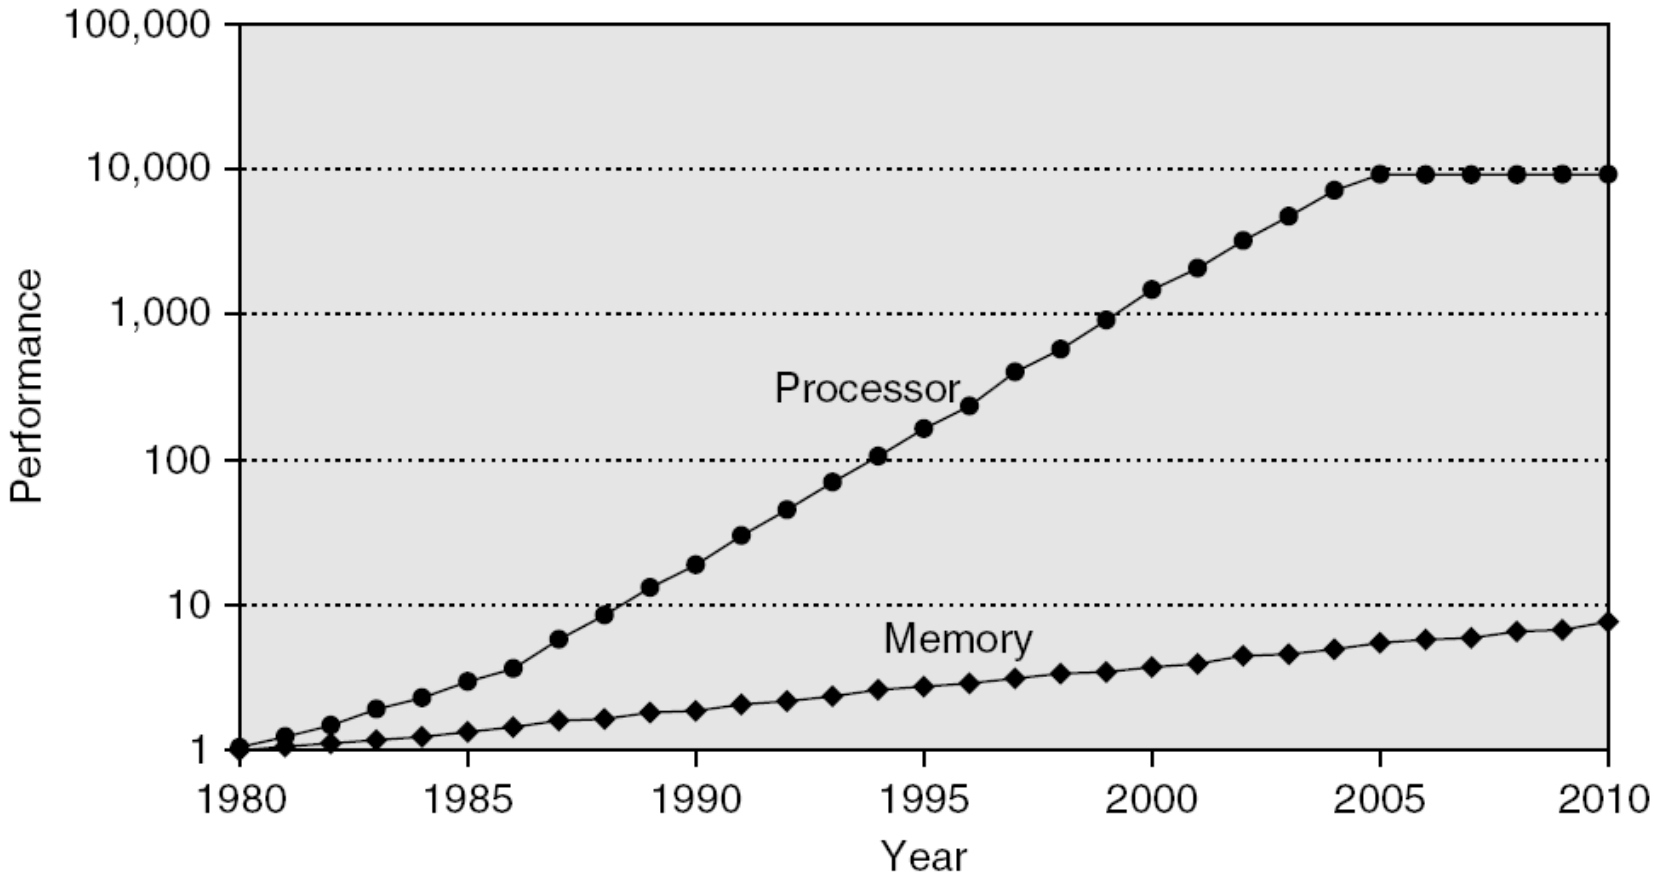

Physical Size Affects Latency

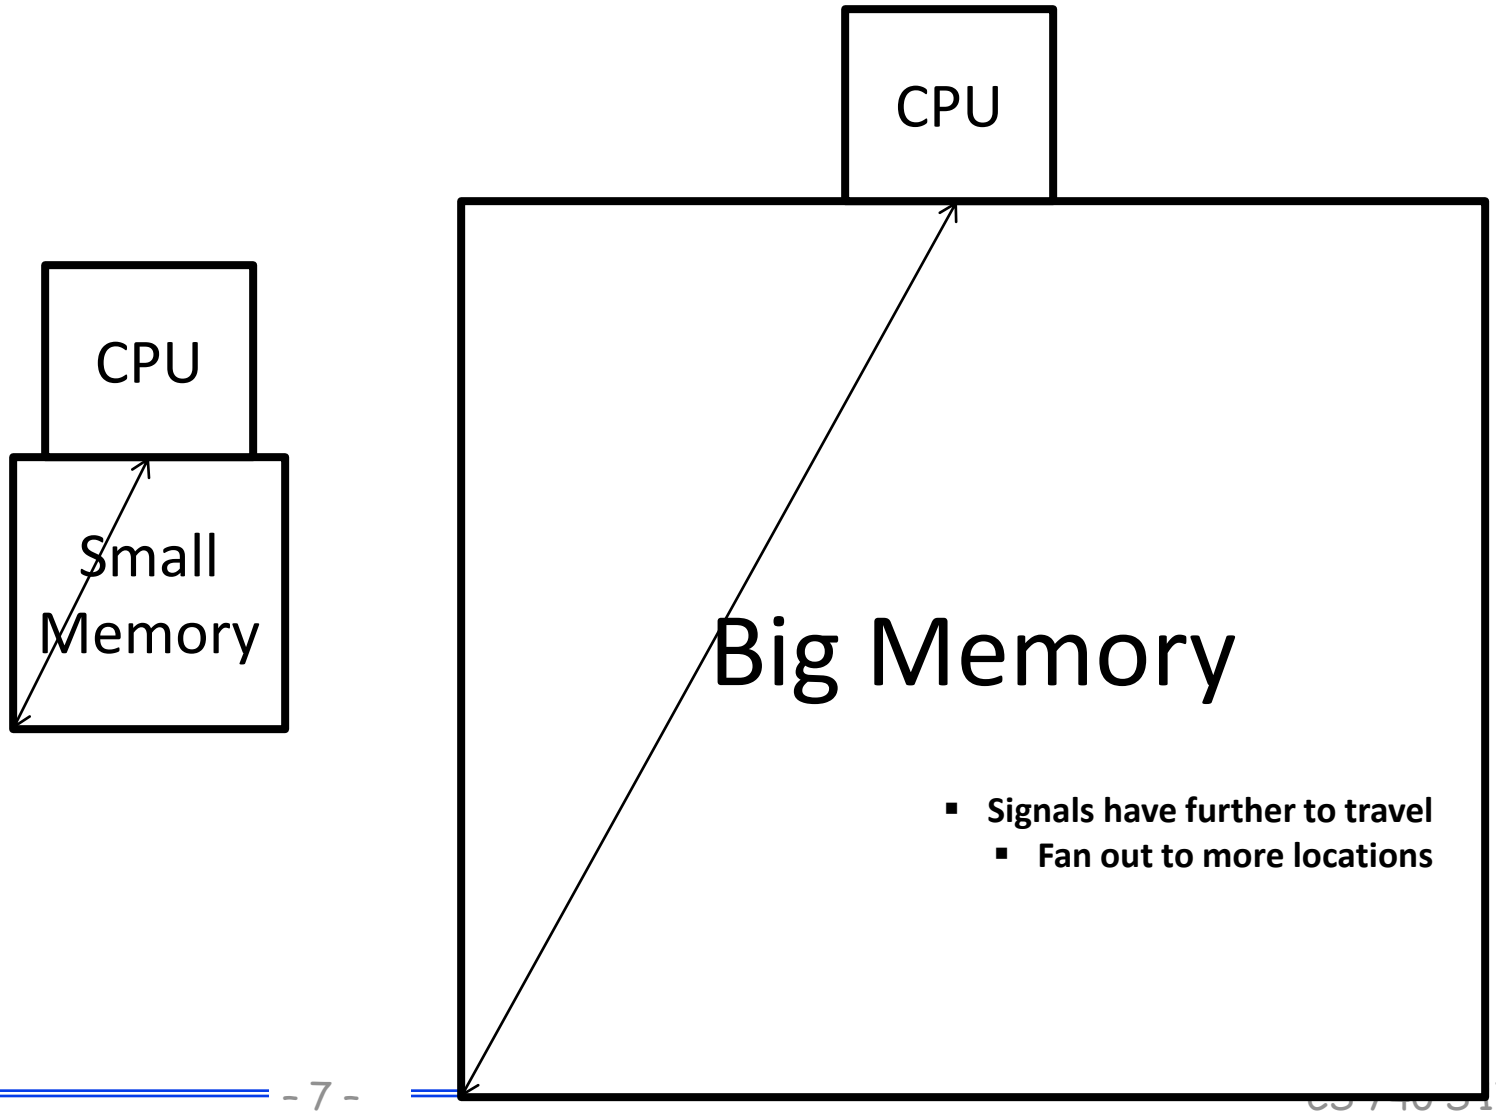

- Signals have further to travel
  - Fan out to more locations

Why is bigger slower?

- Physics slows us down
- Racing the speed of light: (3.0x10^8m/s)
  - take recent Intel chip (Haswell-E 8C)
  - how far can I go in a clock cycle @ 3 GHz?
    (3.0x10^8 m/s) / (3x10^9 cycles/s) = 0.1m/cycle
  - for comparison: Haswell-E 8C is about 19mm = .019m across
  - ➔ speed of light doesn't directly limit speed, but its in ballpark
- Capacitance:
  - long wires have more capacitance
  - either more powerful (bigger) transistors required, or slower
  - signal propagation speed proportional to capacitance
  - going "off chip" has an order of magnitude more capacitance

MEMORY

One-Transistor Dynamic RAM

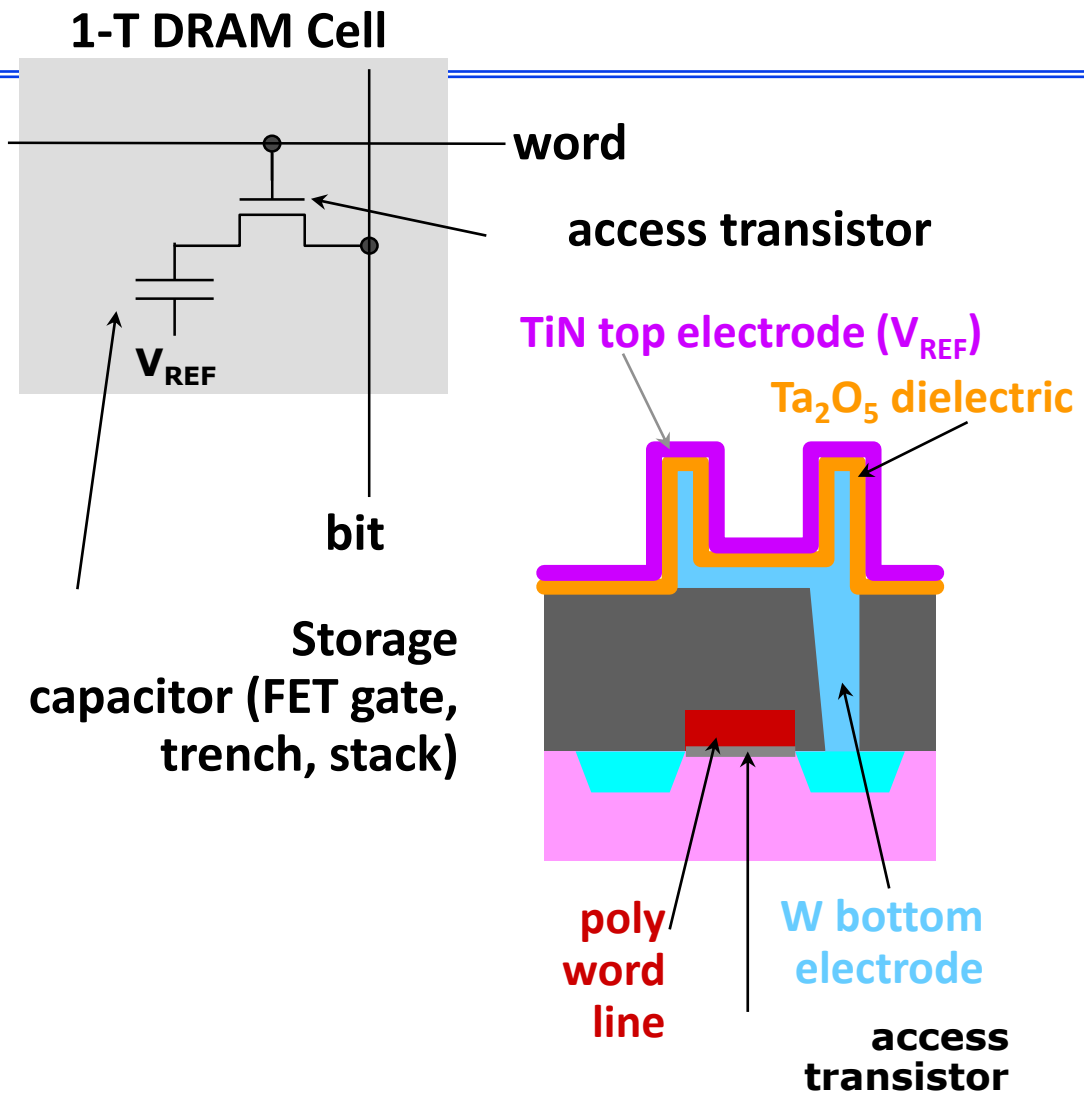

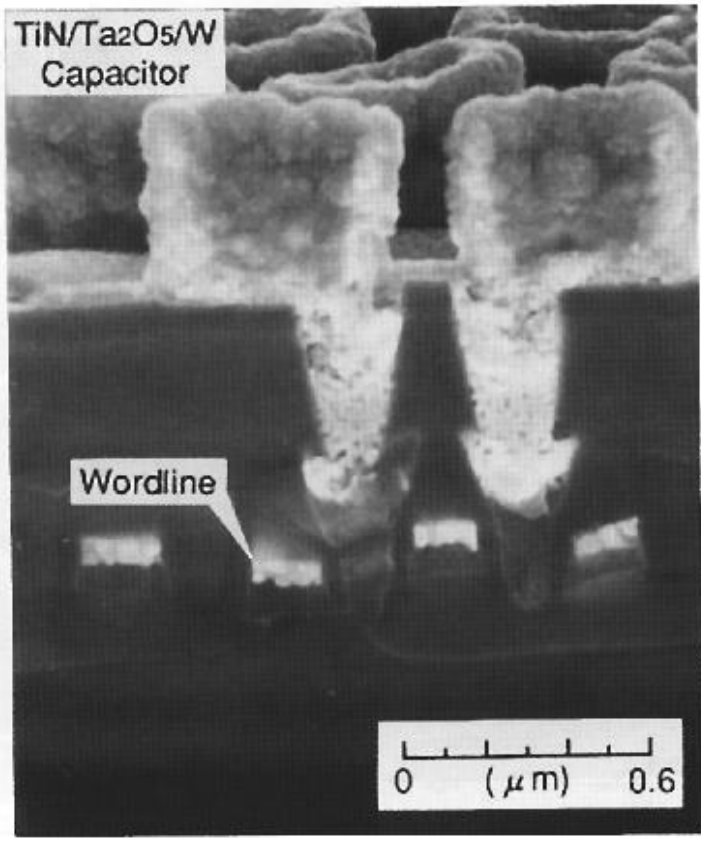

Modern DRAM Cell Structure

[Samsung, sub-70nm DRAM, 2004]

DRAM Conceptual Architecture

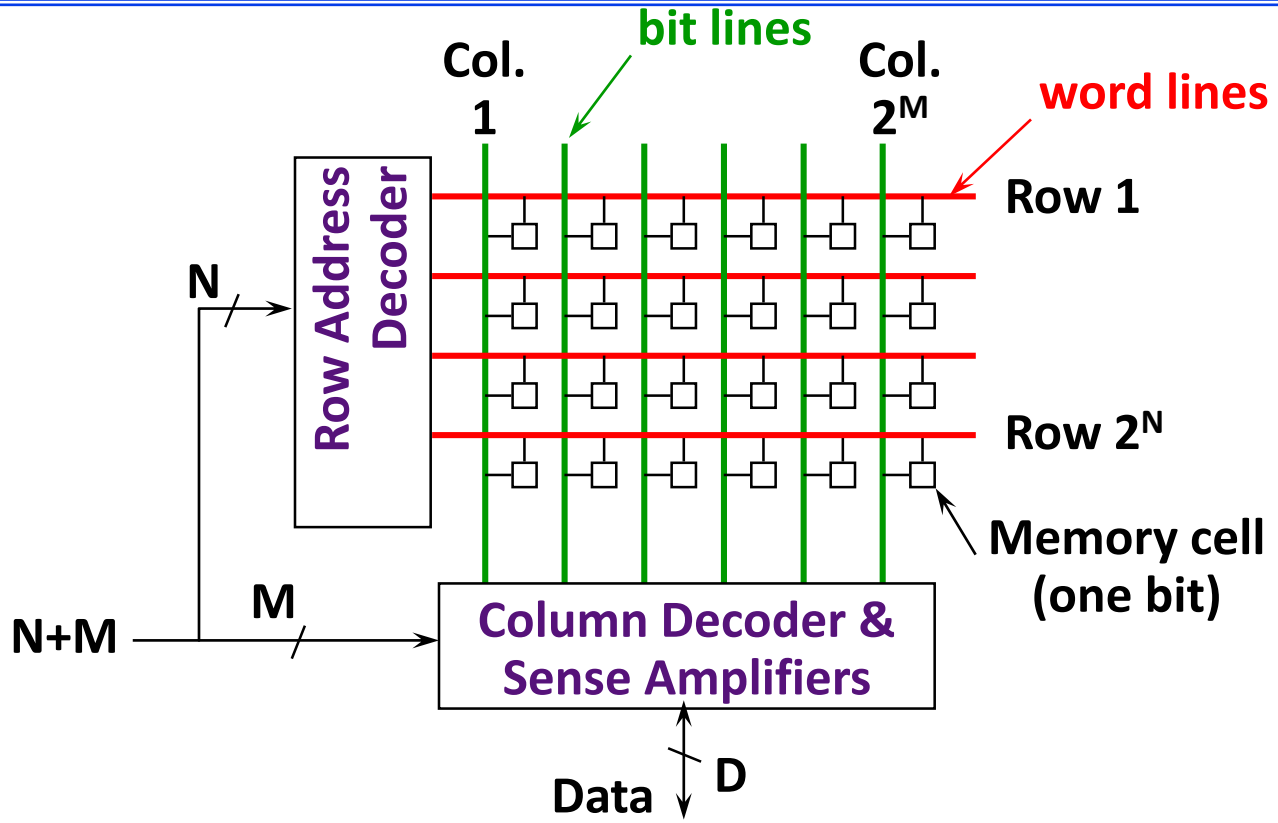

- Bits stored in 2-dimensional arrays on chip
- Modern chips have around 4-8 logical banks on each chip
  - each logical bank physically implemented as many smaller arrays

DRAM Physical Layout

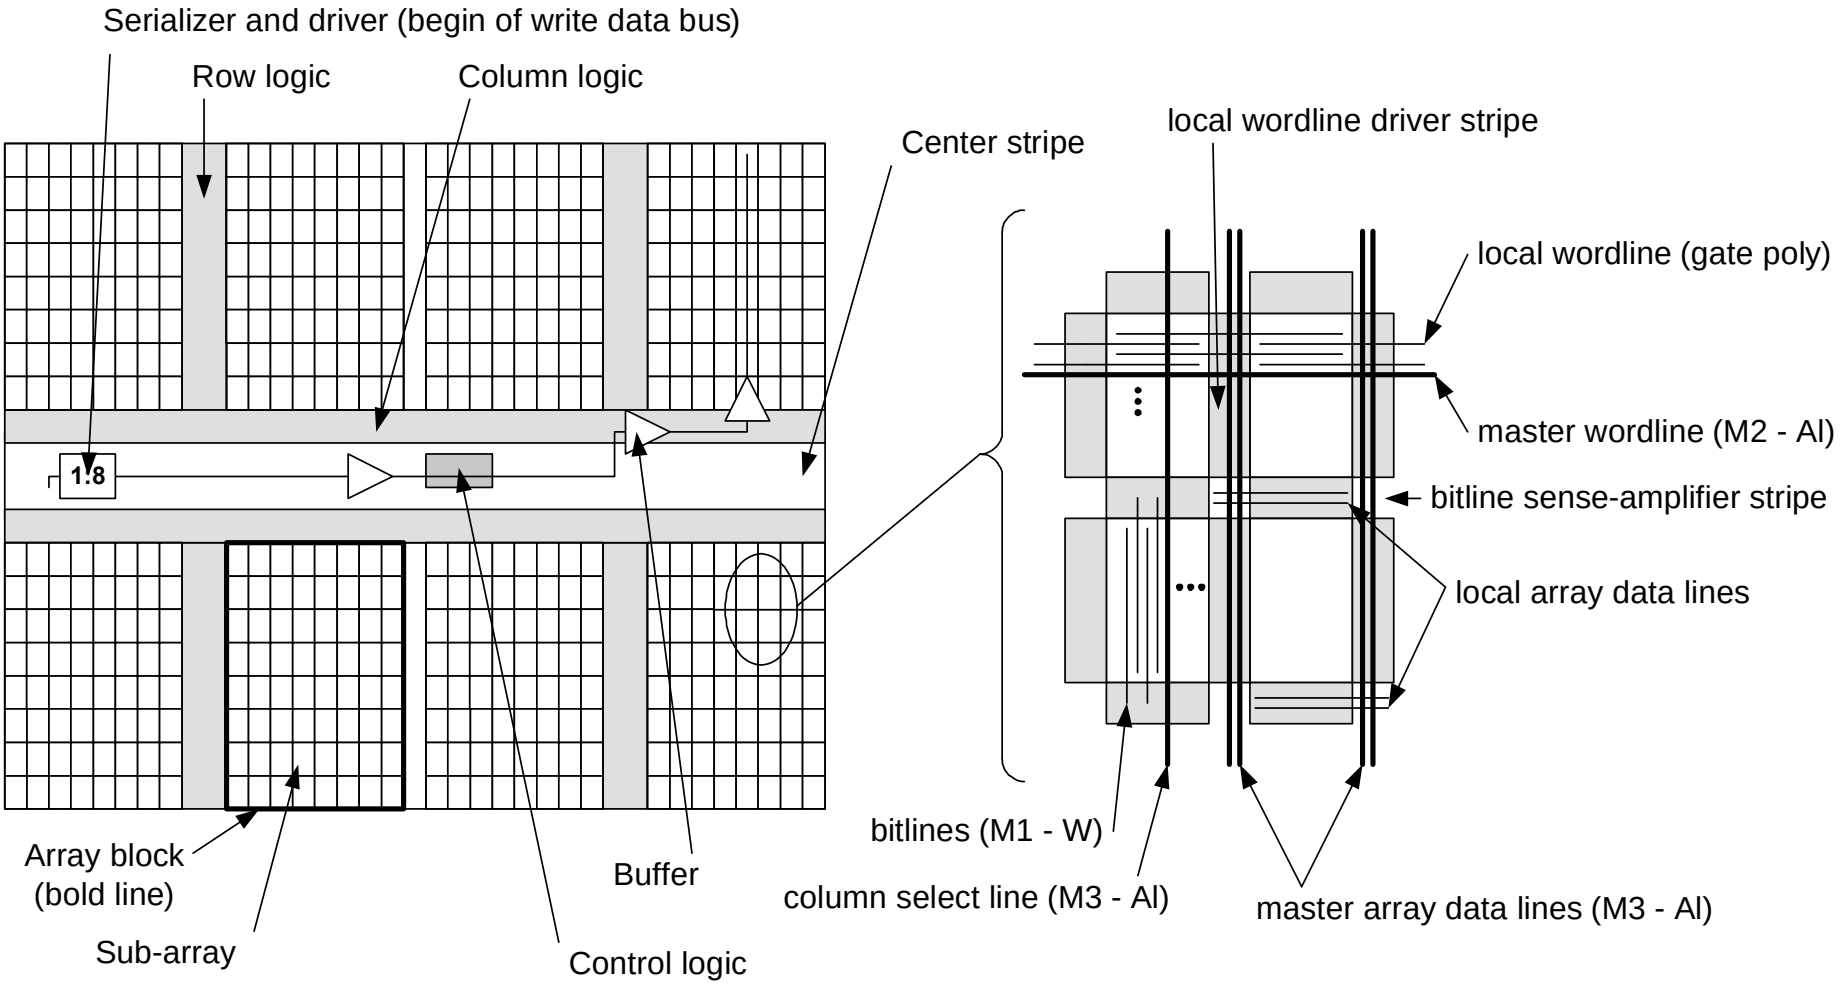

Figure 1. Physical floorplan of a DRAM. A DRAM actually contains a very large number of small DRAMs called sub-arrays.

[Vogelsang, MICRO-2010]

DRAM Operation

Three steps in read/write access to a given bank

- Precharge
- Row access (RAS)
- Column access (CAS)

Each step has a latency of around 10ns in modern DRAMs

Various DRAM standards (DDR, RDRAM) have different ways of encoding the signals for transmission to the DRAM, but all share same core architecture

DRAM Operation

Three steps in read/write access to a given bank

- Precharge
  - charges bit lines to known value, required before next row access
- Row access (RAS)
- Column access (CAS)

Each step has a latency of around 10ns

DRAM Operation

Three steps in read/write access to a given bank

- Precharge
- Row access (RAS)
  - decode row address, enable addressed row (often multiple Kb in row)
  - bitlines share charge with storage cell
  - small change in voltage detected by sense amplifiers which latch whole row of bits
  - sense amplifiers drive bitlines full rail to recharge storage cells
- Column access (CAS)

Each step has a latency of around 10ns

DRAM Operation

Three steps in read/write access to a given bank

- Precharge
- Row access (RAS)

- Column access (CAS)
  - decode column address to select small number of sense amplifier latches (4, 8, 16, or 32 bits depending on DRAM package)
  - on read, send latched bits out to chip pins
  - on write, change sense amplifier latches which then charge storage cells to required value
  - can perform multiple column accesses on same row without another row access (burst mode)

Each step has a latency of around 10ns

Static RAM

Different varieties based on # transistors

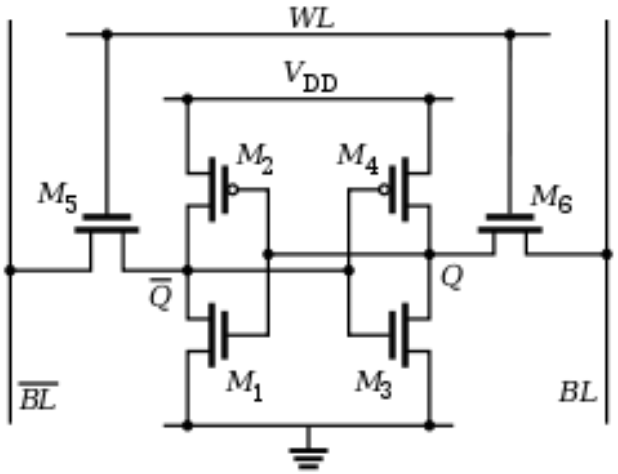

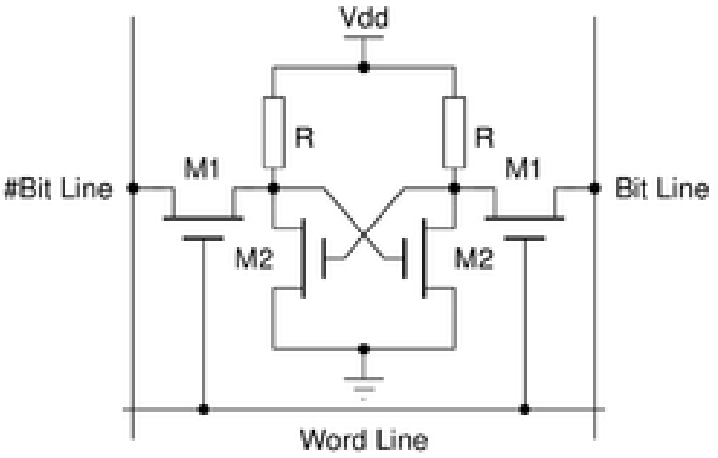

Fewer transistors ➔ more bits / mm^2, but harder to manufacture

Standby: M5 & M6 disconnected, M1-M4 make self-reinforcing inverters

Read: connect M5 & M6, sense + amplify signal on bitlines

Write: connect M5 & M6, bias bitlines to desired value

Memory Parameters

Density

- bits / mm^2, higher for DRAM than SRAM

Latency

- Time from initiation to completion of one memory read (e.g., in nanoseconds, or in CPU or DRAM clock cycles)

Bandwidth

- Rate at which requests can be processed (accesses/sec, or GB/s)

Occupancy

- Time that a memory bank is busy with one request (esp. writes)

Energy

Performance can vary significantly for reads vs. writes, or address, or access history

SRAM vs DRAM

SRAM is simpler

- non-destructive reads

SRAM is faster

DRAM is denser

SRAM used for on-chip caches, register file

DRAM used for main memory

- Often with a different manufacturing process
- That's why single chips with main memory + logic are rare

Processor-DRAM Gap (latency)

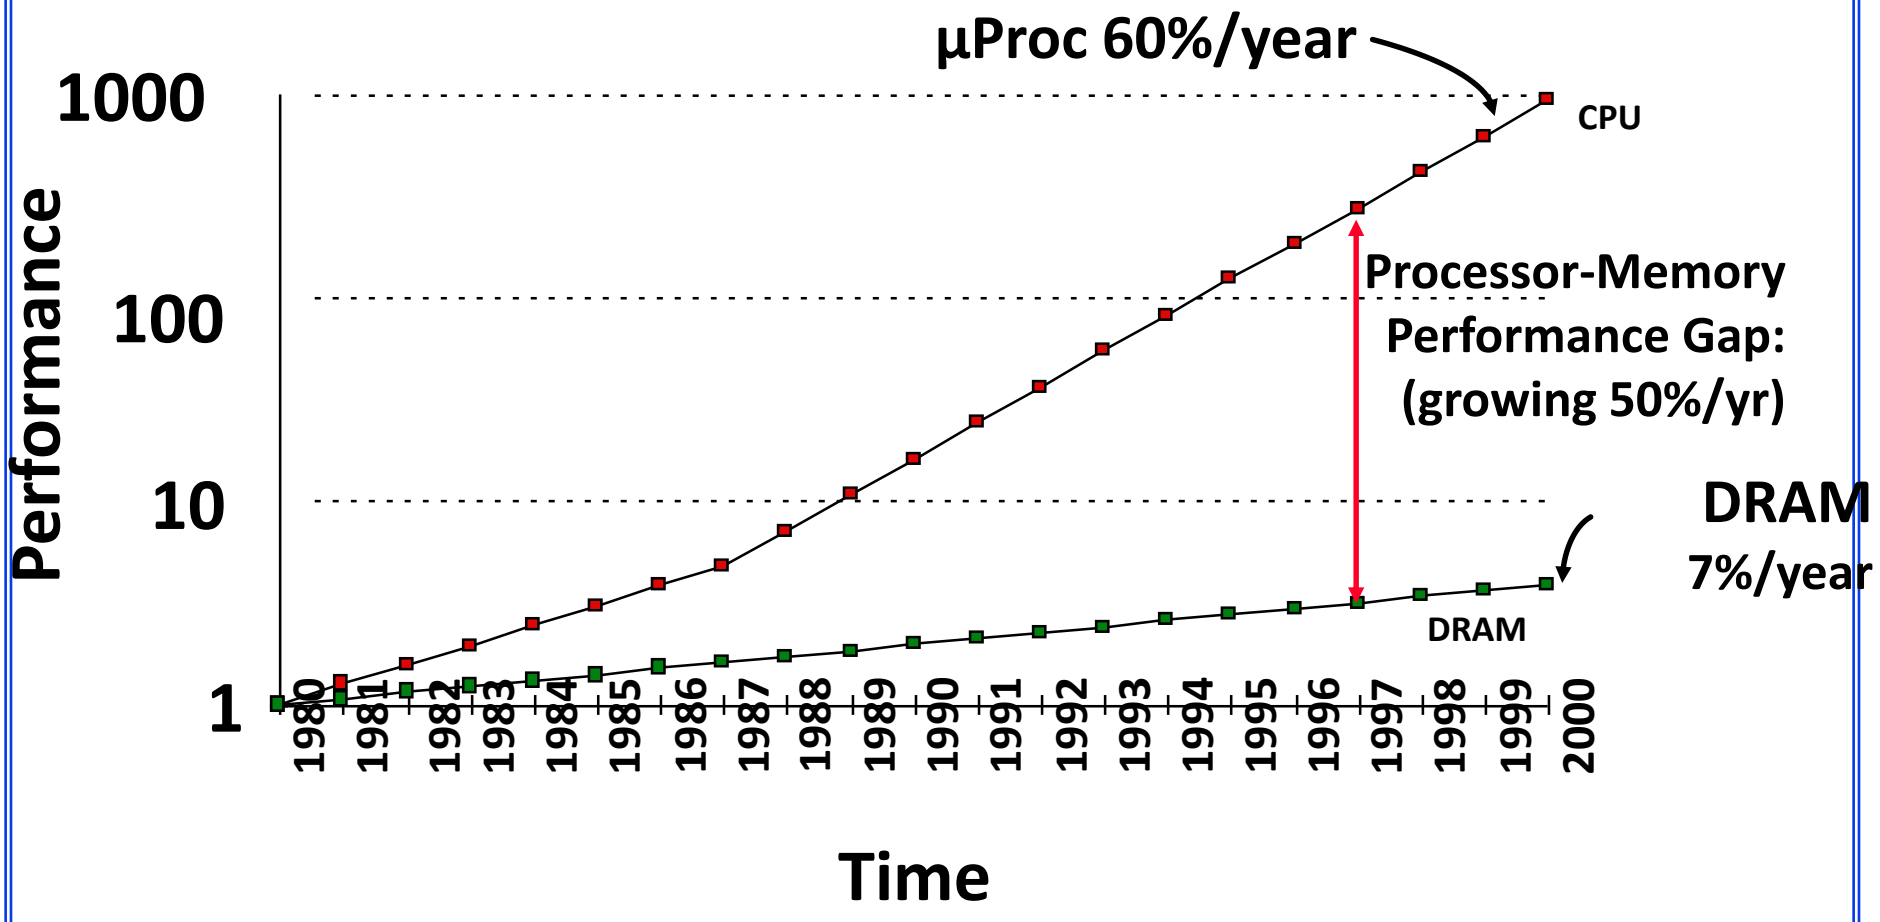

Four-issue 3GHz superscalar accessing 100ns DRAM could execute 1,200 instructions during time for one memory access!

MEMORY HIERARCHY

Two predictable properties of memory references:

Temporal Locality: If a location is referenced it is likely to be referenced again in the near future.

Spatial Locality: If a location is referenced it is likely that locations near it will be referenced in the near future.

Memory Reference Patterns

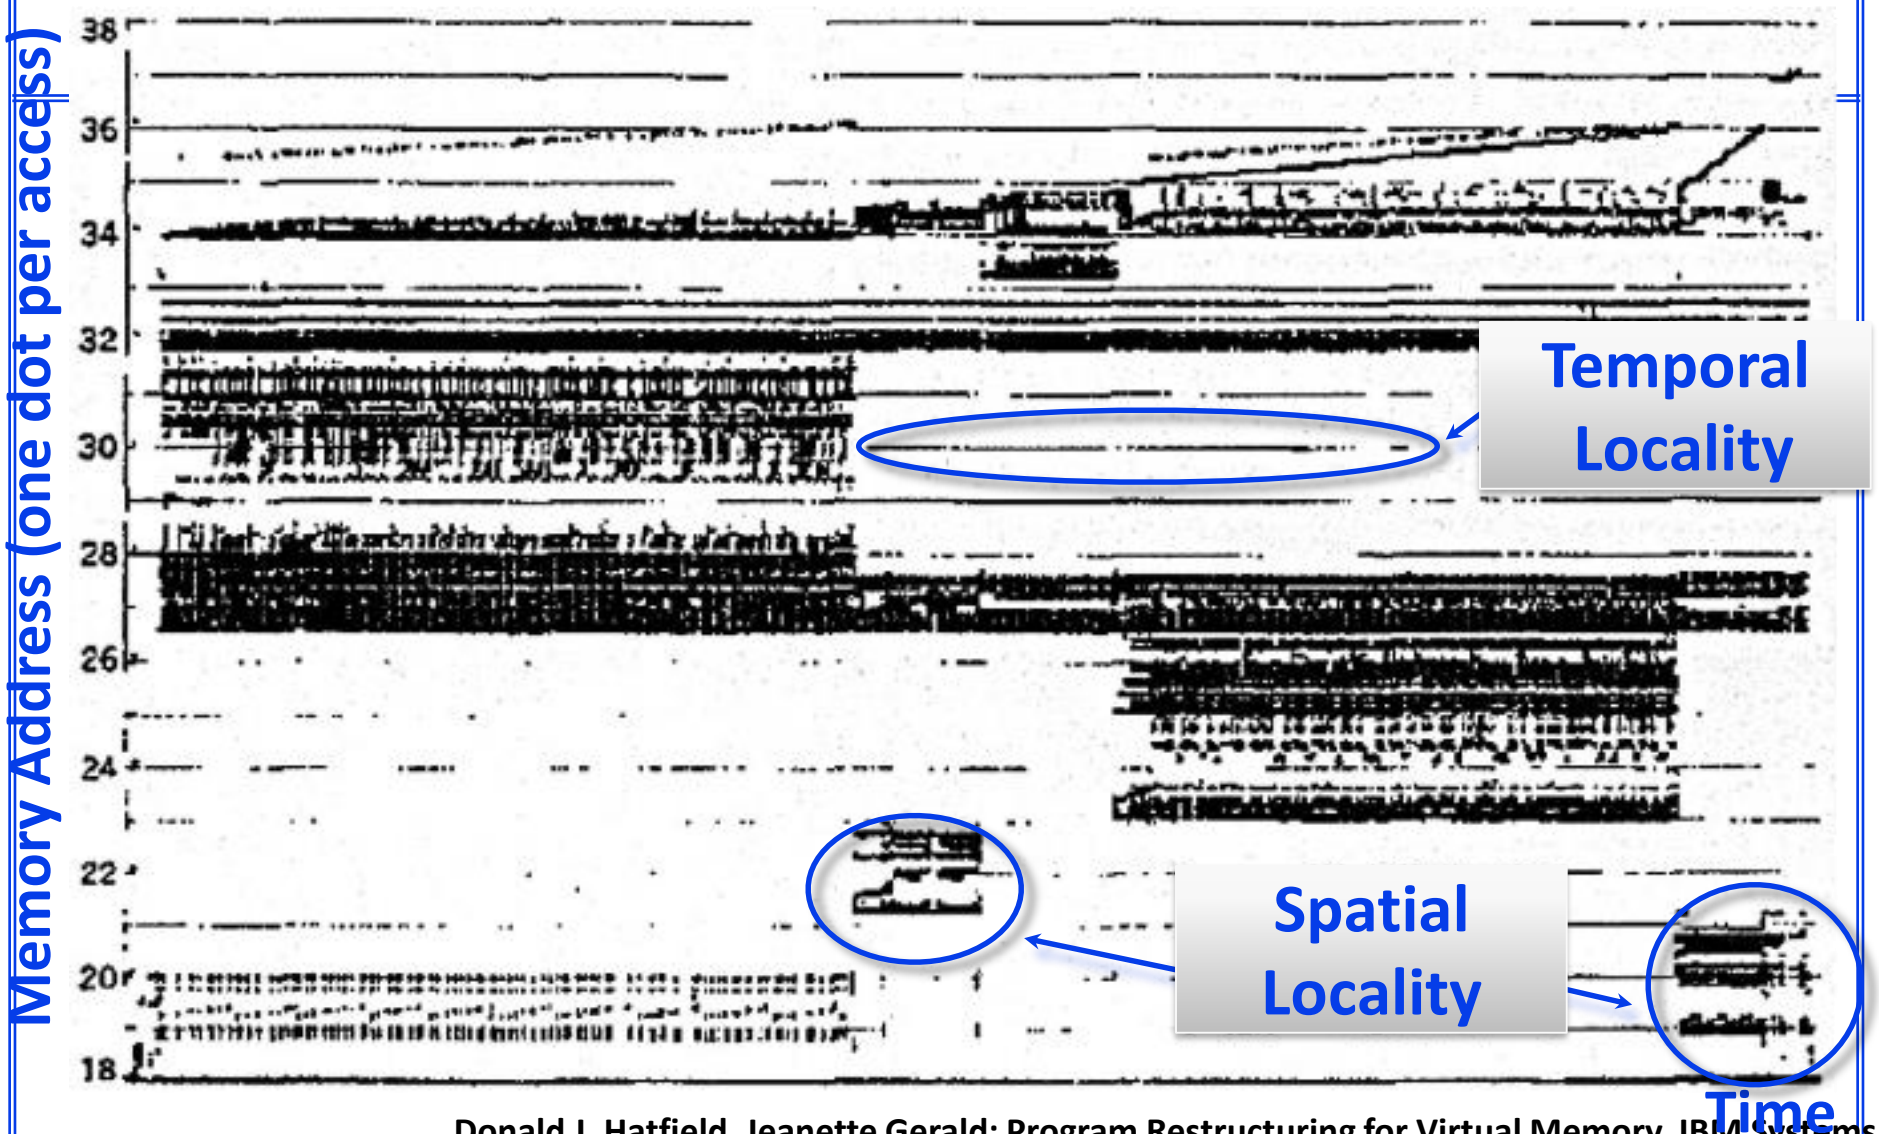

Donald J. Hatfield, Jeanette Gerald: Program Restructuring for Virtual Memory. IBM Systems Journal 10(3): 168-192 (1971)

Memory Hierarchy

Implement memories of different sizes to serve different latency / latency / bandwidth tradeoffs

Keep frequently accessed data in small memories & large datasets in large memories

Provides illusion of a large & fast memory

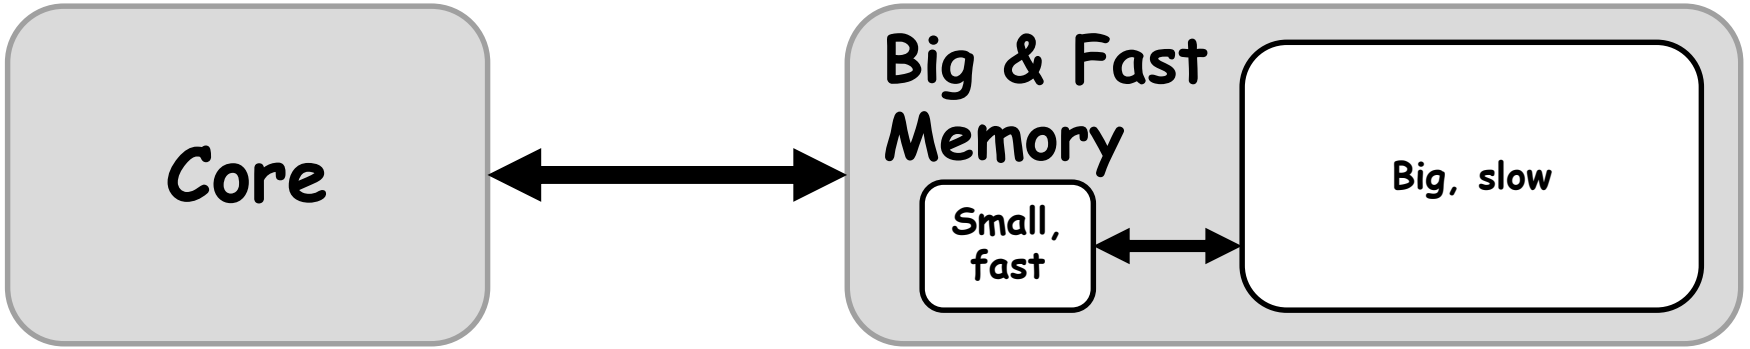

Cache vs. Memory

How to manage the hierarchy?

As memory (aka "scratchpads"): software must be aware of different memories and use them well

- In theory: most efficient
- In practice: inconvenient and difficult (eg, PS3)

As cache: transparent to software; hardware moves data between levels of memory hierarchy

- In theory: overheads and performance loss
- In practice: convenient and h/w does a good job

Management of Memory Hierarchy

Small/fast storage, e.g., registers

- Address usually specified in instruction
- Generally implemented directly as a register file
  - *but hardware might do things behind software's back, e.g., stack management, register renaming*

Larger/slower storage, e.g., main memory

- Address usually computed from values in register
- Generally implemented as a hardware-managed cache hierarchy (hardware decides what is kept in fast memory)
  - *but software may provide "hints", e.g., don't cache or prefetch*

Alpha 21164 Chip Photo

Microprocessor Report 9/12/94

Caches:

- L1 data
- L1 instruction
- L2 unified
- + L3 off-chip

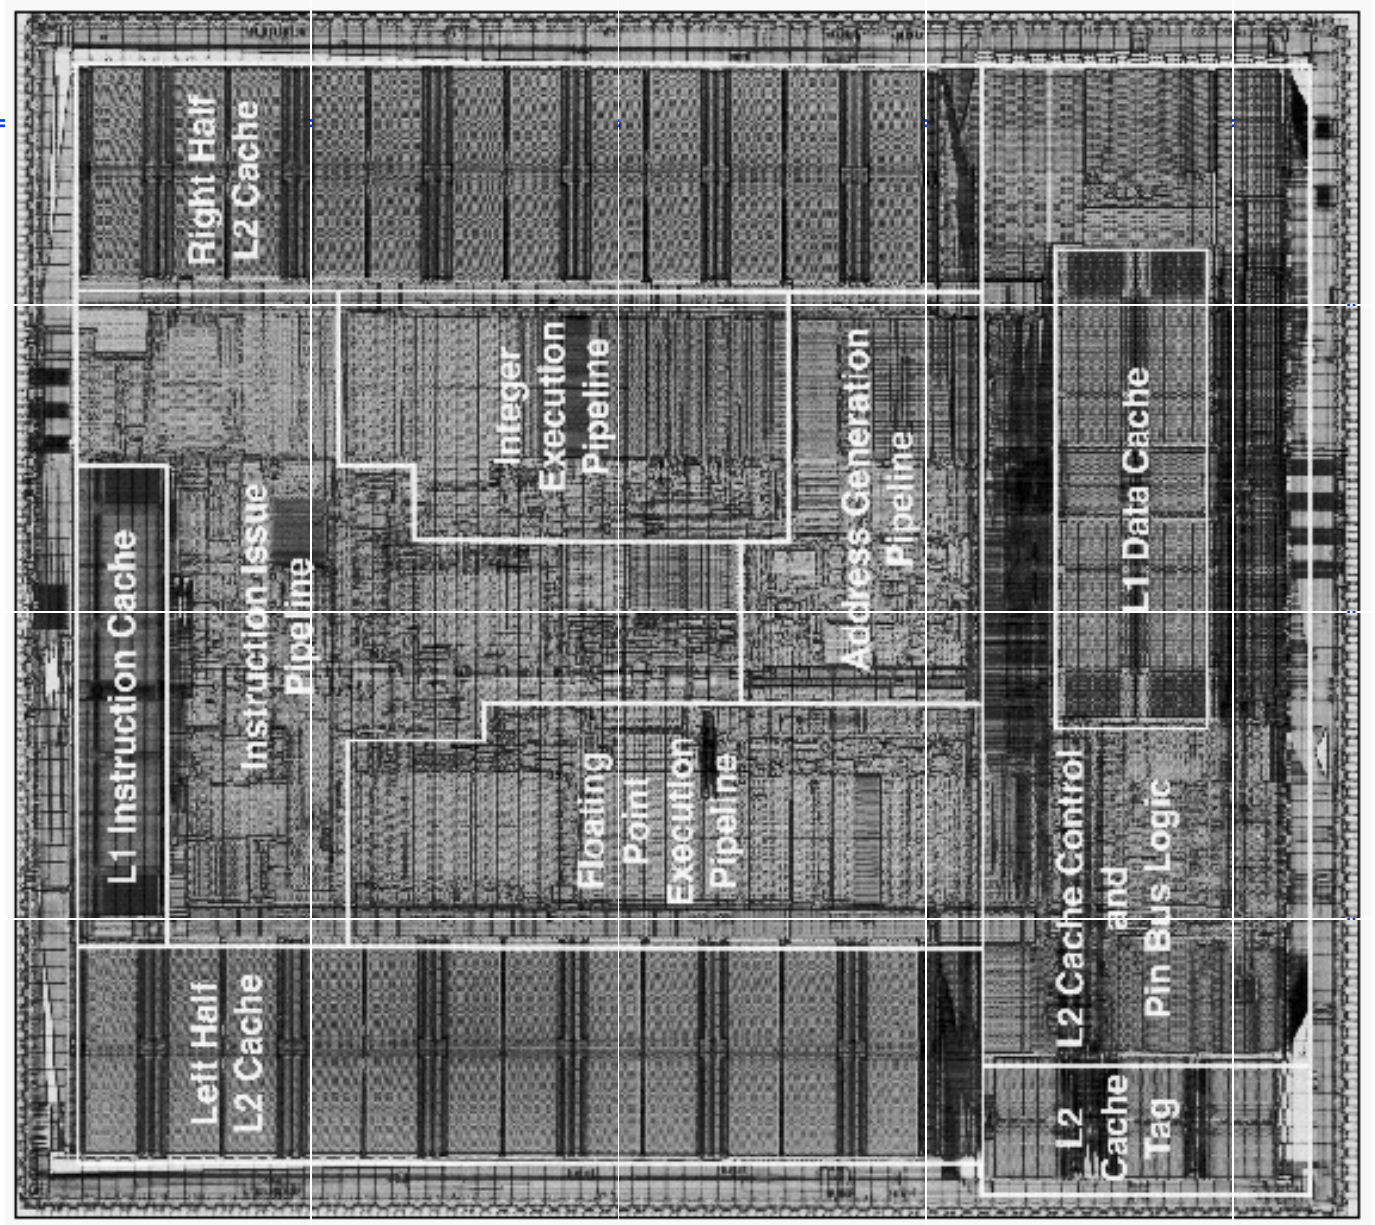

Alpha 21164 Chip Caches

Caches:

- L1 data
- L1 instruction
- L2 unified
- + L3 off-chip

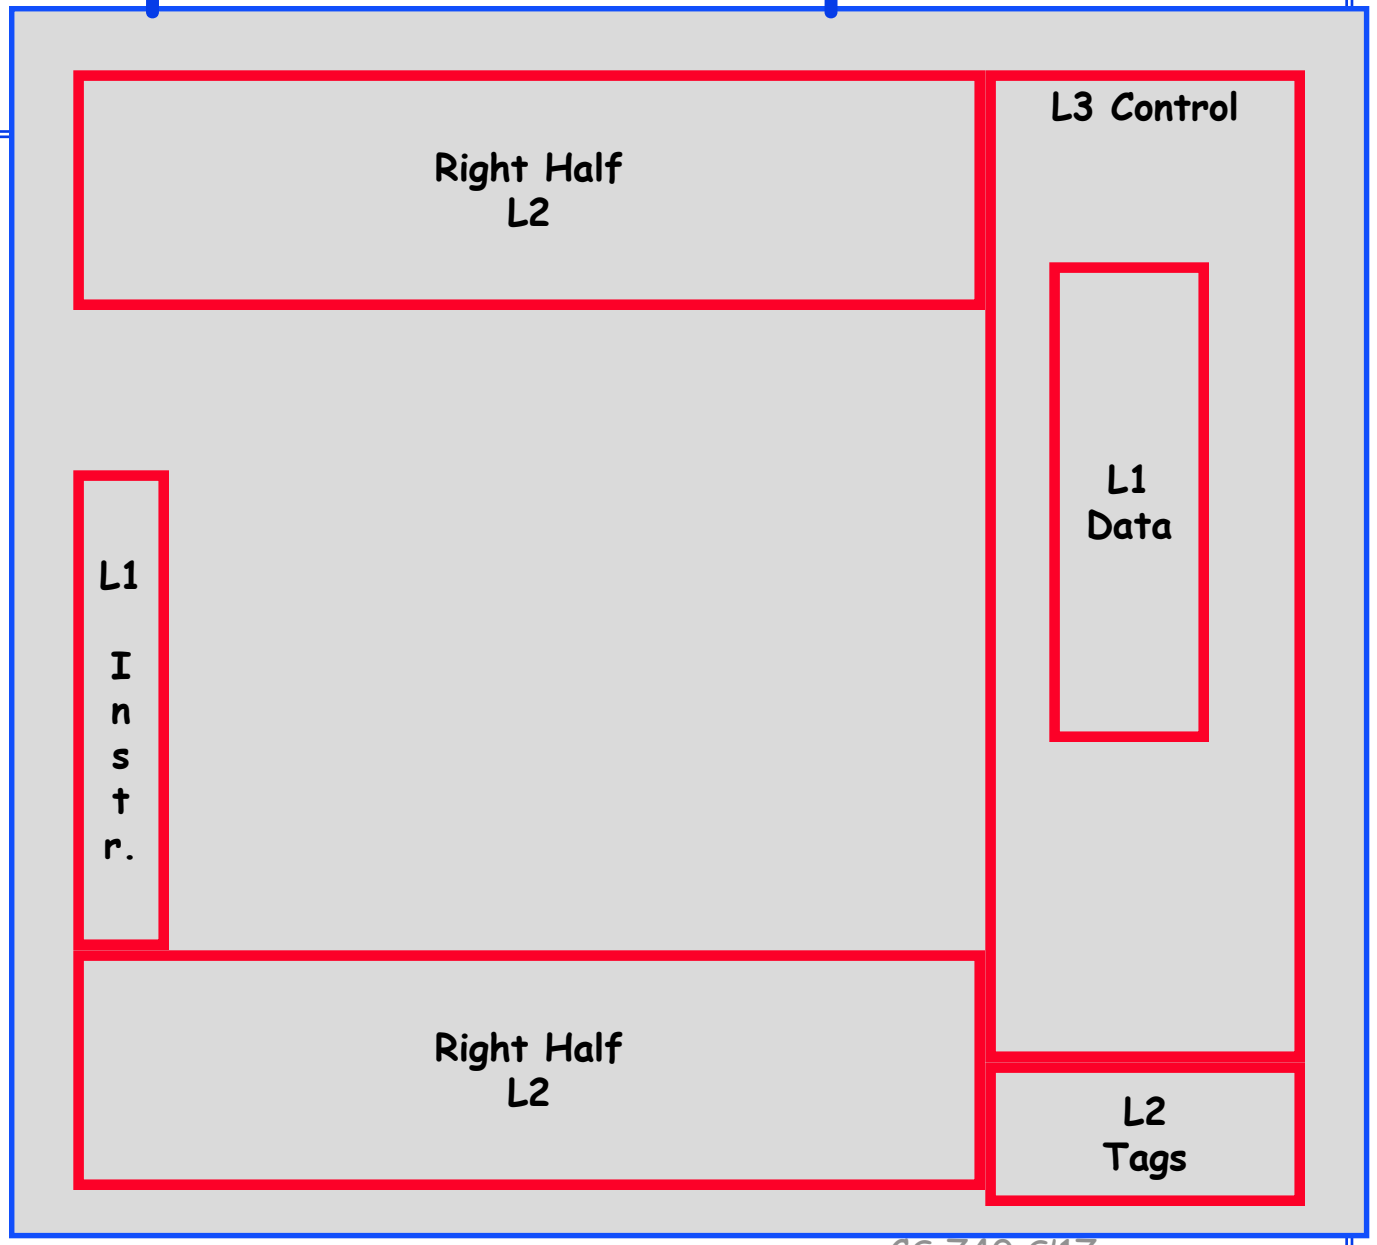

How important are caches?

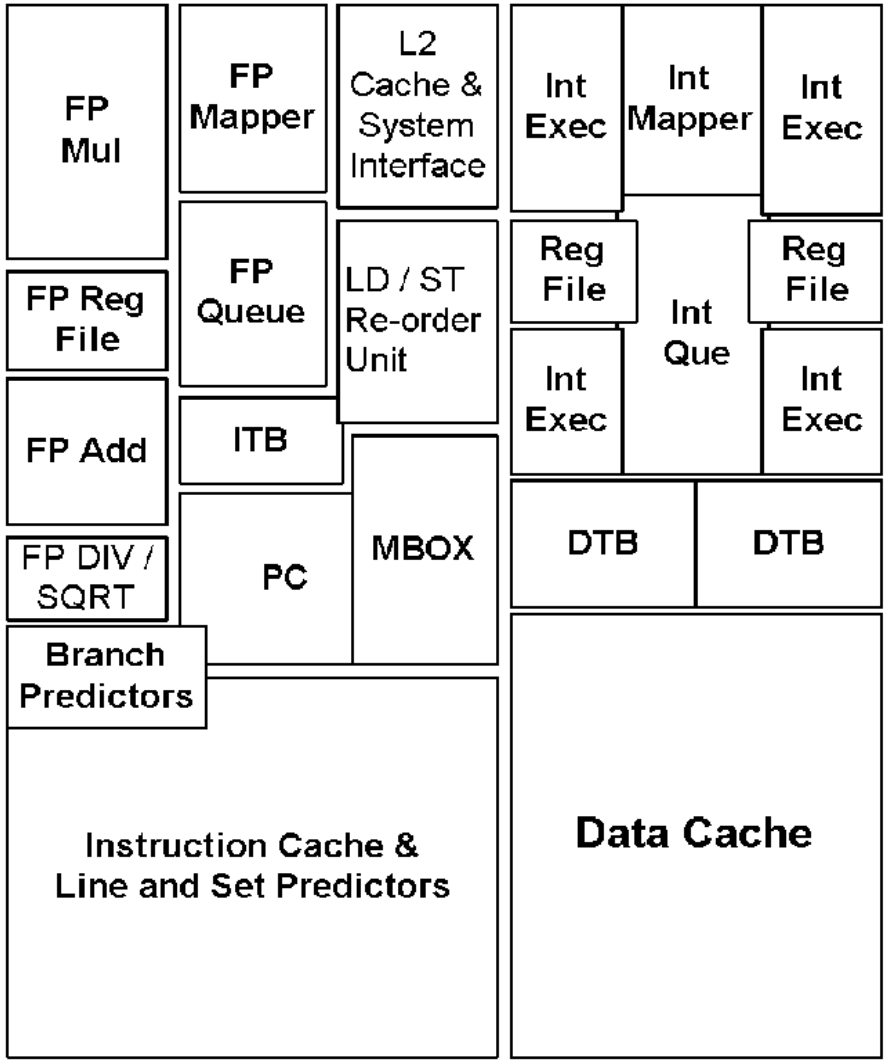

- 21264 Floorplan
- Register files in middle of execution units
- 64k instr cache
- 64k data cache
- Caches take up a large fraction of the die
  - ~30-50% in recent chips

(Figure from Jim Keller, Compaq Corp.)

Multi-Level Caches

Options: ***separate*** data and instruction caches, or a ***unified*** cache

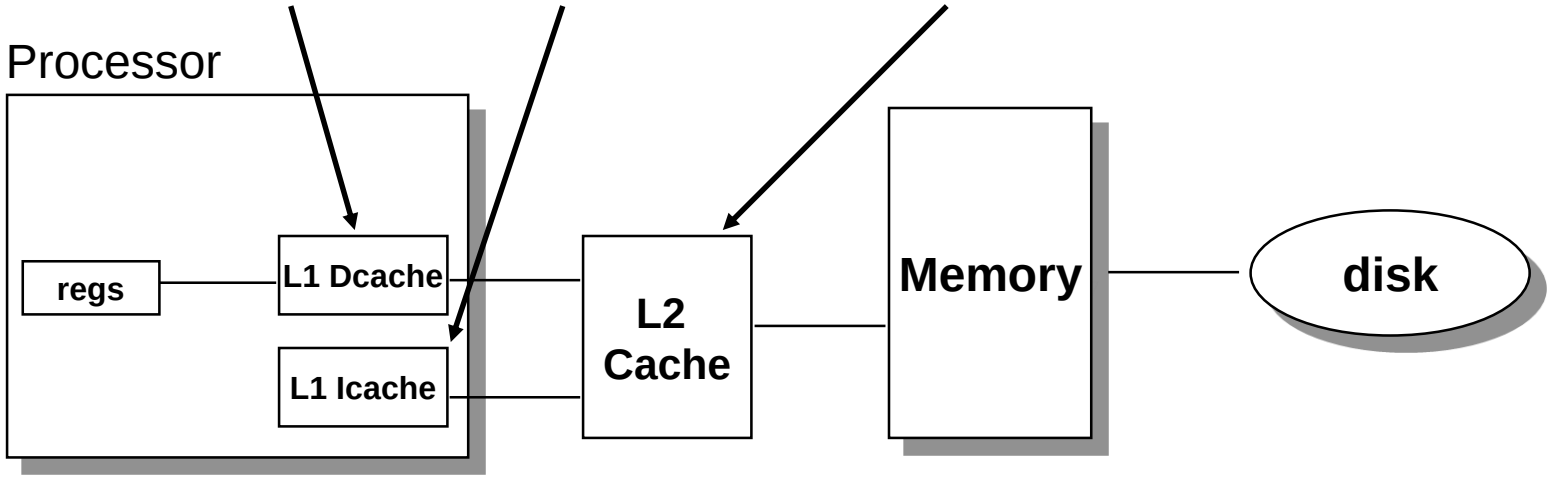

How does this affect self modifying code?

Instructions vs Data

Where to store instructions & data?

Harvard architecture: In early machines, instructions were hard-wired (switchboards) or punchcards, while data was kept in memory

Princeton/von Neumann architecture: Instructions and data are both in memory—"instructions are data"

Modern architecture: von Neumann, but…split instruction/data caches; protection bits prevent execution of data; different optimizations, etc.

Caches Exploit Locality

Temporal locality:

- Hardware decides what to keep in cache
- *Replacement/eviction policy* evicts a victim upon a cache miss
- Least-recently used (LRU) most common

Spatial locality:

- Cache stores multiple, neighboring words per *block*
- *Prefetchers* speculate about next accesses and fetch them into cache

Locality of Reference

Principle of Locality:

- Programs tend to reuse data and instructions near those they have used recently.
- *Temporal locality:* recently referenced items are likely to be referenced in the near future.
- *Spatial locality:* items with nearby addresses tend to be referenced close together in time.

```
sum = 0;
for (i = 0; i < n; i++)
    sum += a[i];
*v = sum;
```

Locality in Example:

- Data
  - Reference array elements in succession (spatial)
  - sum variable (temporal, allocated to register)
- Instructions
  - Reference instructions in sequence (spatial)
  - Cycle through loop repeatedly (temporal)

Caching: The Basic Idea

Main Memory

- **Stores words**
  A–Z in example

Cache

- **Stores subset of words**
  e.g., 4 in example
- **Organized in lines**
  - Multiple words to…
  - Exploit spatial locality
  - Amortize overheads

Access

- **Processor requests address from cache, which handles misses itself**

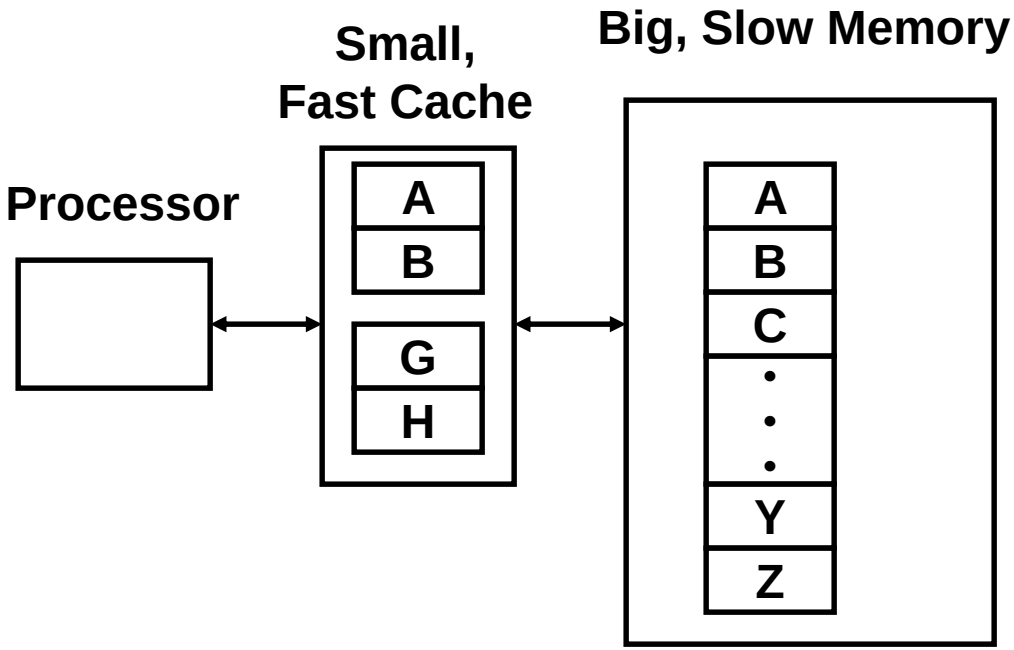

Memory Hierarchy Basics

$$\frac{\text{Misses}}{\text{Instruction}} = \frac{\text{Miss rate} \times \text{Memory accesses}}{\text{Instruction count}} = \text{Miss rate} \times \frac{\text{Memory accesses}}{\text{Instruction}}$$

$$\text{Average memory access time} = \text{Hit time} + \text{Miss rate} \times \text{Miss penalty}$$

Three ways to improve memory performance:

1. **Reduce hit time**
2. **Reduce miss rate**
3. **Reduce miss penalty**

There's a tension between these

Memory Hierarchy Basics

$$\frac{\text{Misses}}{\text{Instruction}} = \frac{\text{Miss rate} \times \text{Memory accesses}}{\text{Instruction count}} = \text{Miss rate} \times \frac{\text{Memory accesses}}{\text{Instruction}}$$

$$\text{Average memory access time} = \text{Hit time} + \text{Miss rate} \times \text{Miss penalty}$$

Note that speculative and multithreaded processors may execute other instructions during a miss

- **Reduces performance impact of misses**
- **Memory-level parallelism (MLP) overlaps miss latency**

Accessing Data in Memory Hierarchy

- **Between any two levels, memory is divided into *lines* (aka "*blocks*")**
- **Data moves between levels on demand, in line-sized chunks**
- **Invisible to application programmer**
  - Hardware responsible for cache operation
- **Upper-level lines a subset of lower-level lines**

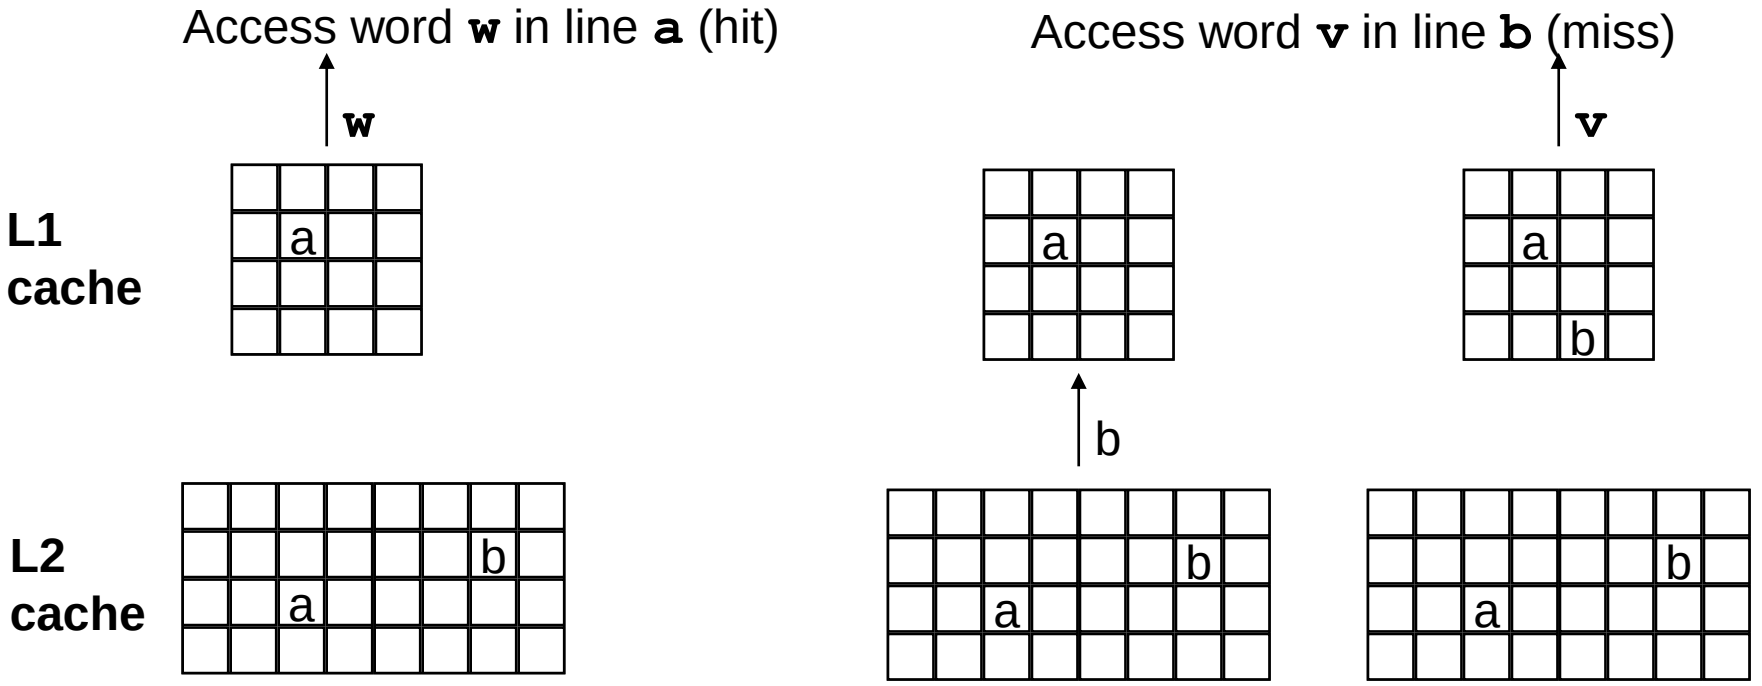

Impact of Cache Size

Increasing cache size

- Effect on cache area (tags + data)?
- Effect on hit time?
- Effect on miss rate?
- Effect on miss penalty?

Design Issues for Caches

Key Questions:

- **Where should a line be placed in the cache? (line placement)**
- **How is a line found in the cache? (line identification)**
- **Which line should be replaced on a miss? (line replacement)**
- **What happens on a write? (write strategy)**

Constraints:

- **Design must be simple**
  - Hardware realization
  - All decision making within nanosecond time scale
- **Want to optimize performance for "typical" programs**
  - Do extensive benchmarking and simulations
  - Many subtle engineering tradeoffs

Fully Associative Cache

Mapping of Memory Lines

- Cache consists of single set holding A=S lines
- Given memory line can map to any line in set
- Only practical for small caches
- Useful for analysis and simulation
- Common in software caches

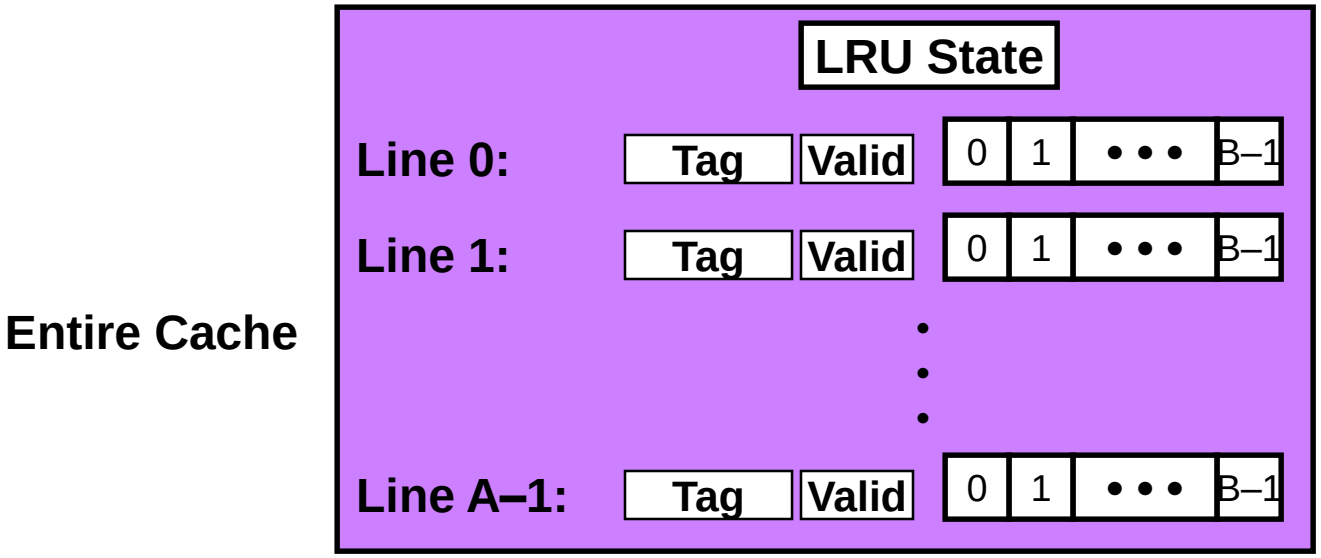

Fully Associative Cache Tag Matching

Identifying Line

- Must check all of the tags for match
- Must have *Valid = 1* for this line

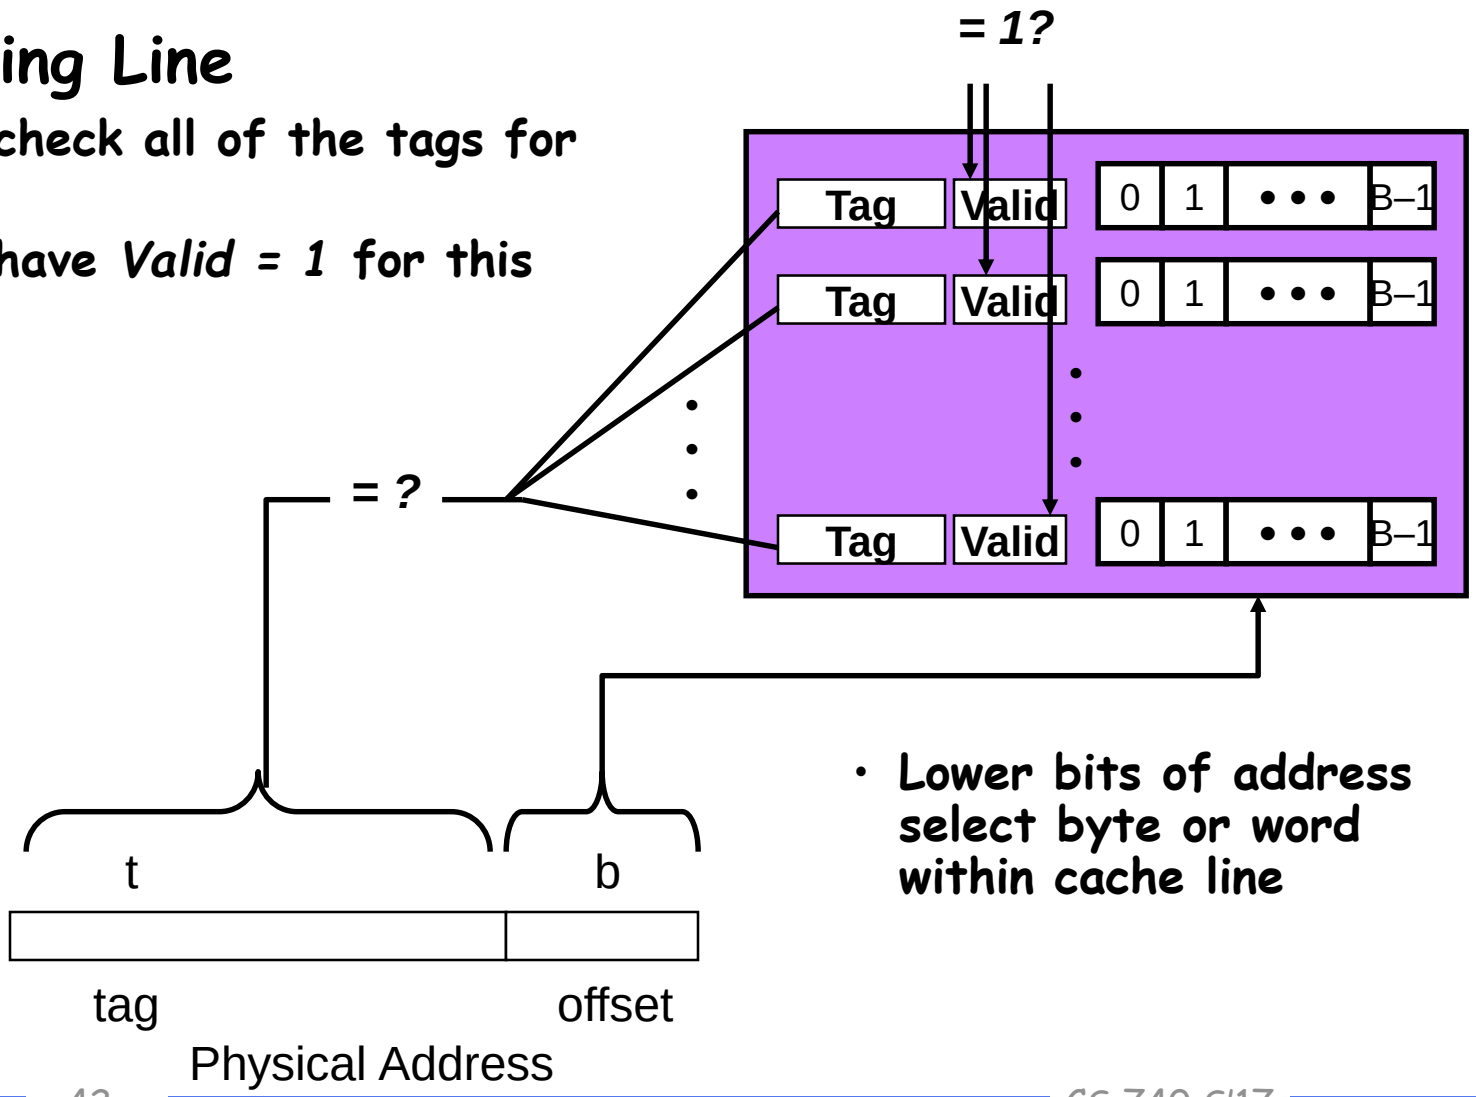

- Lower bits of address select byte or word within cache line

Direct-Mapped Caches

Simplest Design

- **Each memory line has a unique cache location**

Parameters

- **Line (aka block) size B = 2^b**
  - Number of bytes in each line
  - Typically 2X–8X word size
- **Number of sets S = 2^s**
  - Number of lines cache can hold
- **Total Cache Size = B*S = 2^{b+s}**

Physical Address

- **Address used to reference main memory**
- ***n* bits to reference N = 2^n total bytes**
- **Partition into fields**
  - *Offset:* Lower *b* bits indicate which byte within line
  - *Set:* Next *s* bits indicate how to locate line within cache
  - *Tag:* Identifies this line when in cache

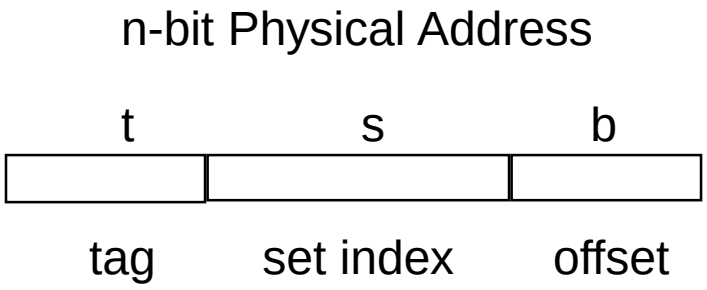

Indexing into Direct-Mapped Cache

- Use set index bits to select cache set

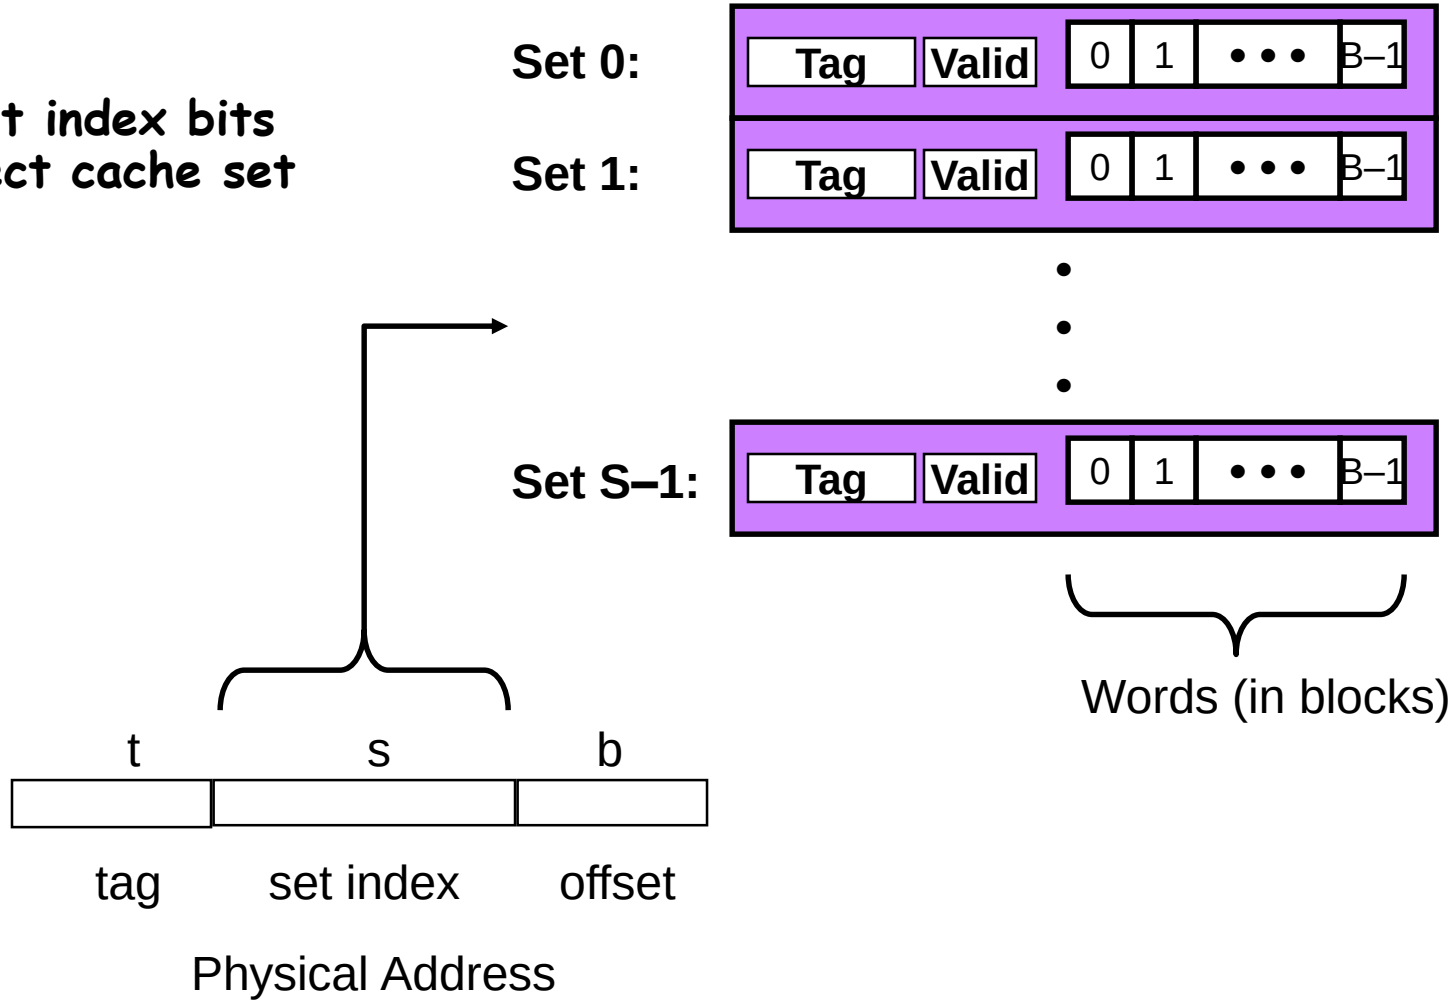

Direct-Mapped Cache Tag Matching

Identifying Line

- Must have tag match high order bits of address
- Must have *Valid = 1*

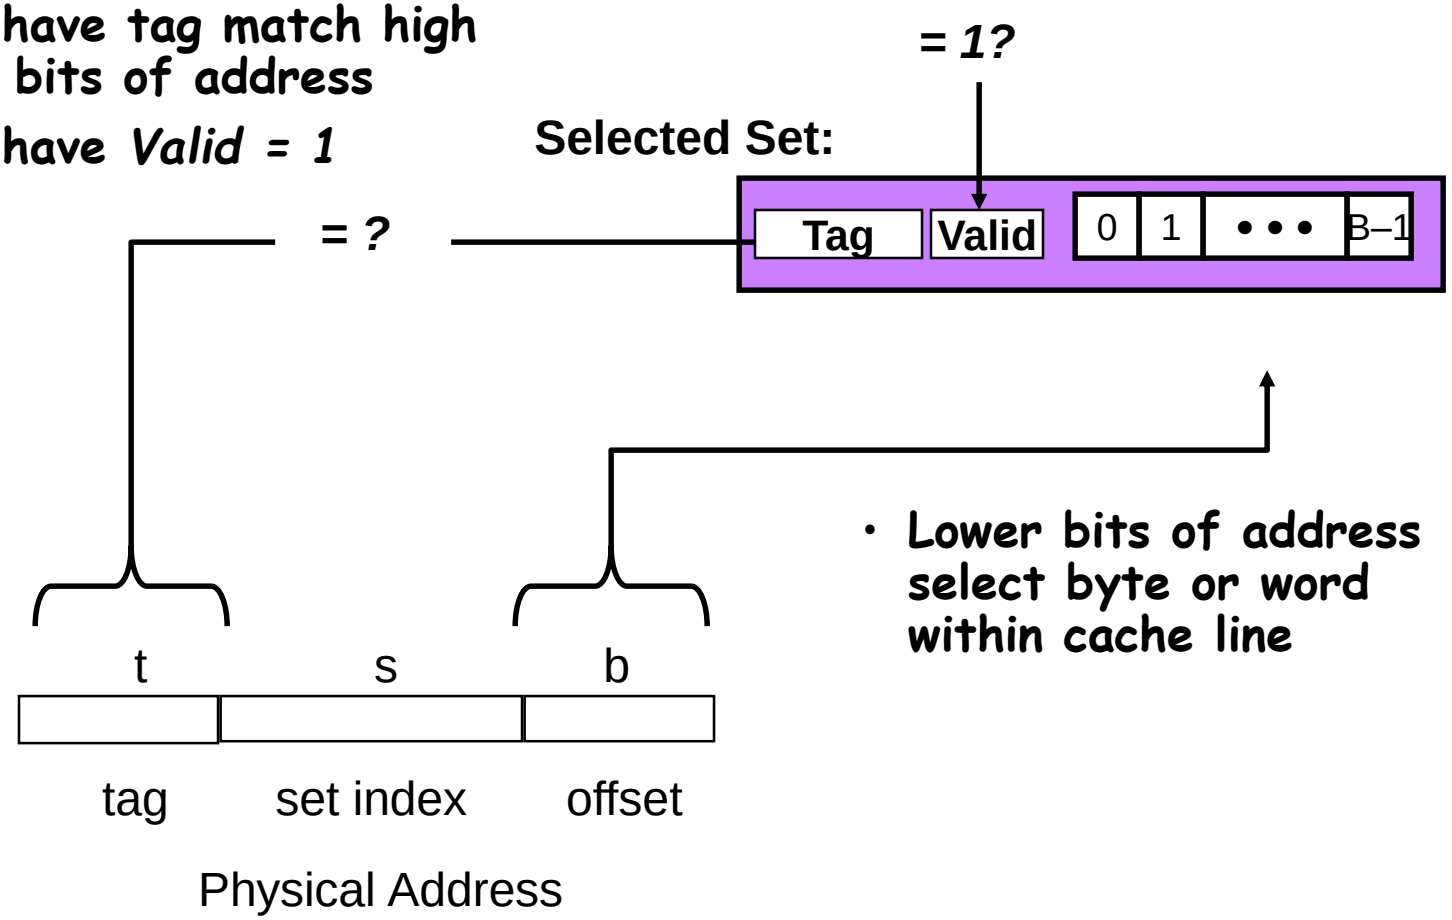

- Lower bits of address select byte or word within cache line

Properties of Direct Mapped Caches

Strength

- Minimal control hardware overhead
- Simple design
- (Relatively) easy to make fast

Weakness

- Vulnerable to thrashing
- Two heavily used lines have same cache index
- Repeatedly evict one to make room for other

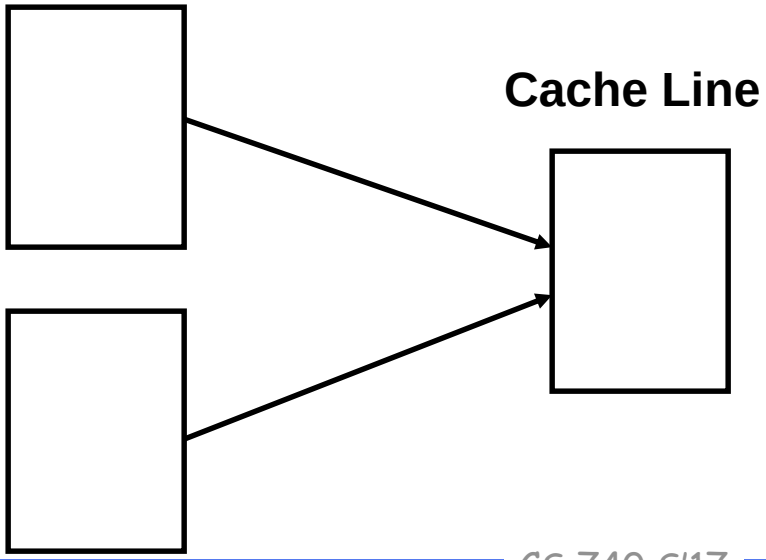

Dot Product Example

```
float dot_prod(float x[1024], y[1024])
{
  float sum = 0.0;
  int i;
  for (i = 0; i < 1024; i++)
    sum += x[i]*y[i];
  return sum;
}
```

Machine

- DECStation 5000
- MIPS Processor with 64KB direct-mapped cache, 16 B line size

Performance

- Good case: 24 cycles / element
- Bad case: 66 cycles / element

Thrashing Example

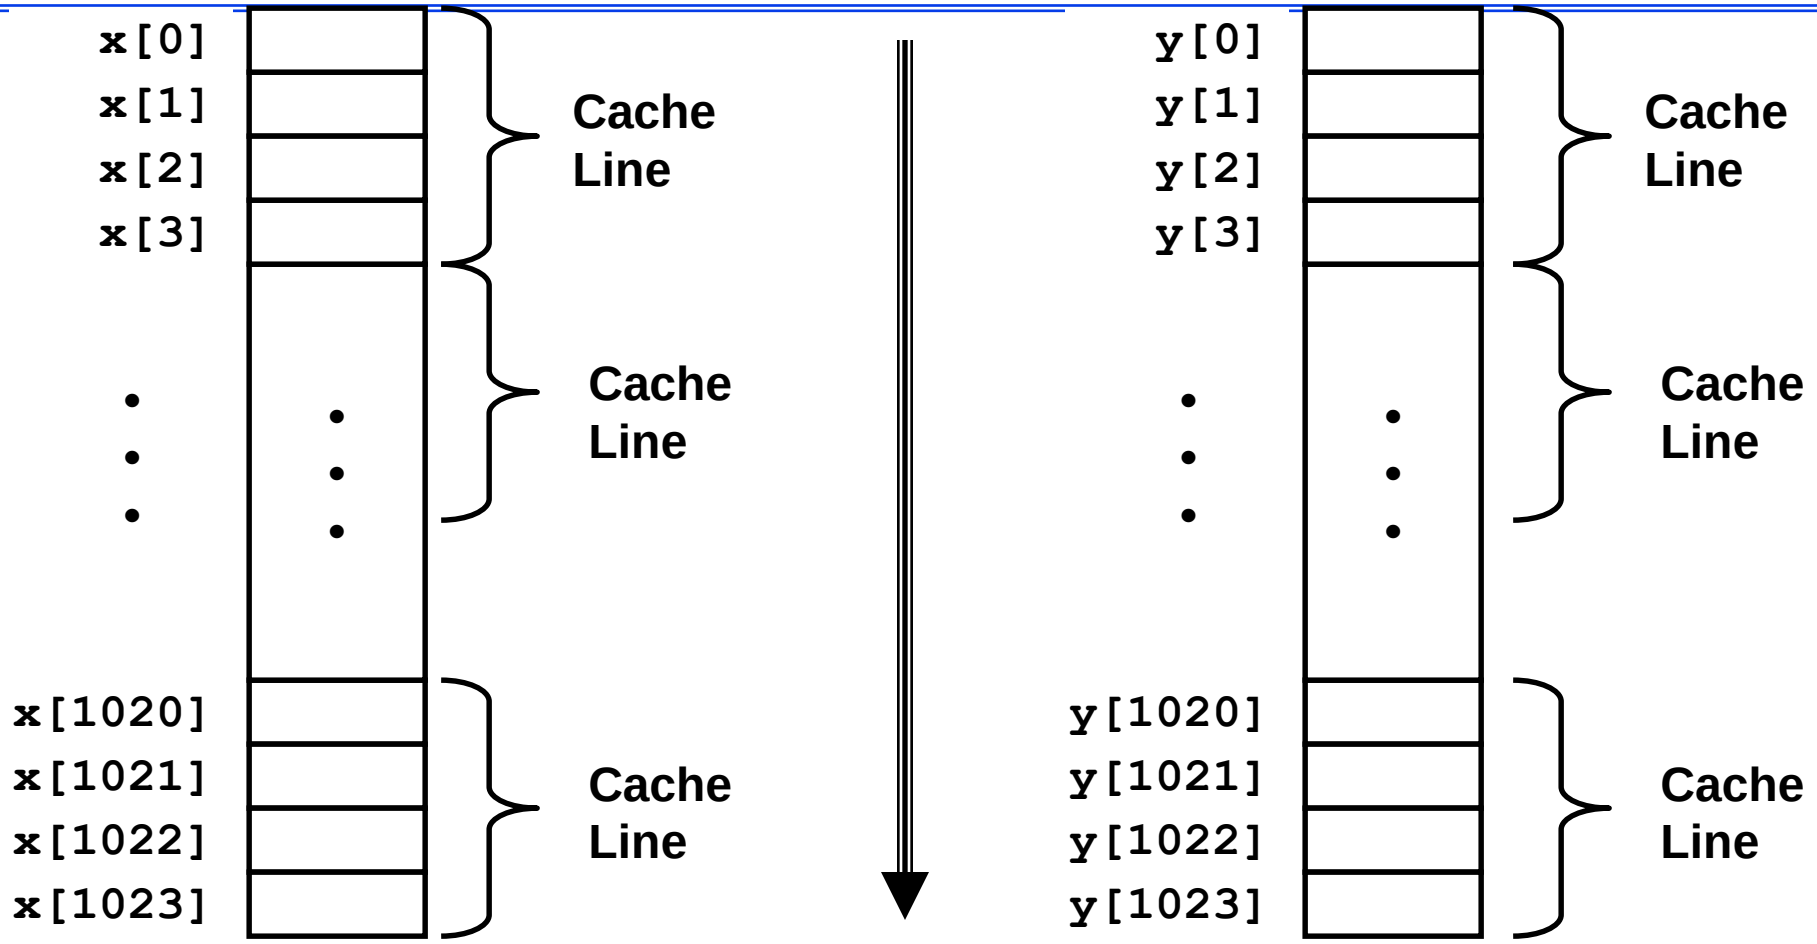

- Access one element from each array per iteration

Thrashing Example: Good Case

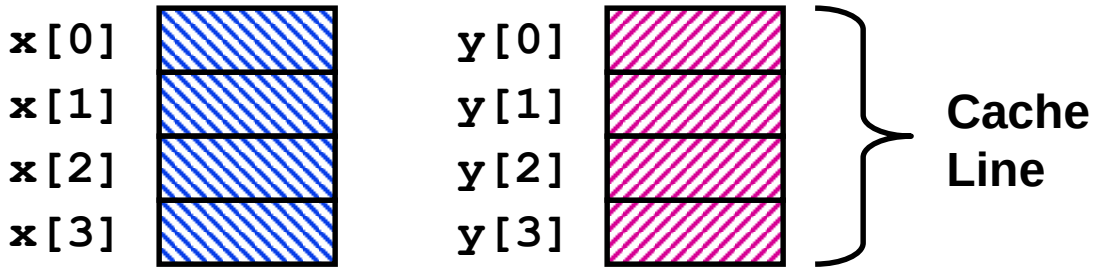

Access Sequence

- **Read x[0]**
  - x[0], x[1], x[2], x[3] loaded
- **Read y[0]**
  - y[0], y[1], y[2], y[3] loaded
- **Read x[1]**
  - Hit
- **Read y[1]**
  - Hit
- **. . .**
- **2 misses / 8 reads**

Analysis

- **x[i] and y[i] map to different cache lines**
- **Miss rate = 25%**
  - Two memory accesses / iteration
  - On every 4th iteration have two misses

Timing

- **10 cycle loop time**
- **28 cycles / cache miss**
- **Average time / iteration =**
  10 + 0.25 * 2 * 28

Thrashing Example: Bad Case

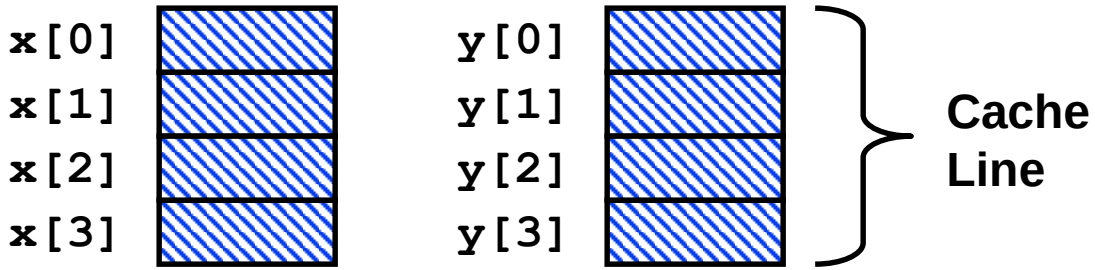

Access Pattern

- **Read x[0]**
  - x[0], x[1], x[2], x[3] loaded
- **Read y[0]**
  - y[0], y[1], y[2], y[3] loaded
- **Read x[1]**
  - x[0], x[1], x[2], x[3] loaded
- **Read y[1]**
  - y[0], y[1], y[2], y[3] loaded
- • •
- **8 misses / 8 reads**

Analysis

- **x[i] and y[i] map to same cache lines**
- **Miss rate = 100%**
  - Two memory accesses / iteration
  - On *every* iteration have two misses

Timing

- **10 cycle loop time**
- **28 cycles / cache miss**
- **Average time / iteration =**
  10 + 1.0 * 2 * 28

Miss Types: The 3 Cs

Compulsory/Cold-start Misses – address not seen previously; difficult to avoid (not impossible!)

- Compulsory misses = misses @ infinite size

Capacity Misses – cache not big enough; larger cache size

- Capacity misses = fully associative misses – compulsory misses

Conflict/Collision Misses – poor block placement evicts useful blocks

- Conflict misses = actual misses – capacity misses

Impact of Block Size

Increasing block size

- Effect on cache area (tags + data)?
- Effect on hit time?
- Effect on miss rate?
- Effect on miss penalty?

Set Associative Cache

Mapping of Memory Lines

- Each set can hold A lines (usually A=2-8 for L1, A=8-32 for L3)
- Given memory line can map to any entry within its given set

- Tradeoffs
  - Fewer conflict misses
  - Forced by virtual memory
  - Longer access latency
  - More complex to implement

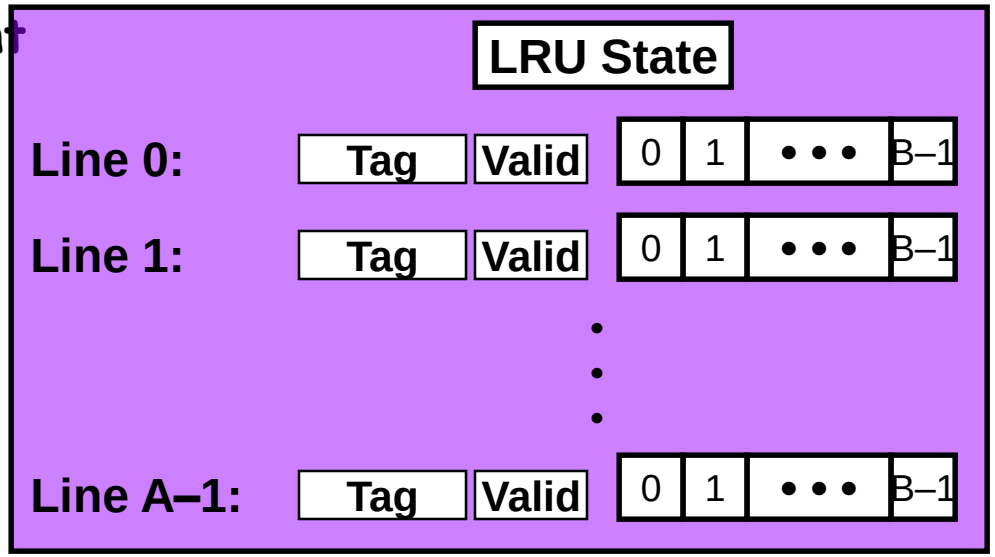

Indexing into 2-Way Associative Cache

- Use middle s bits to select from among $S = 2^s$ sets

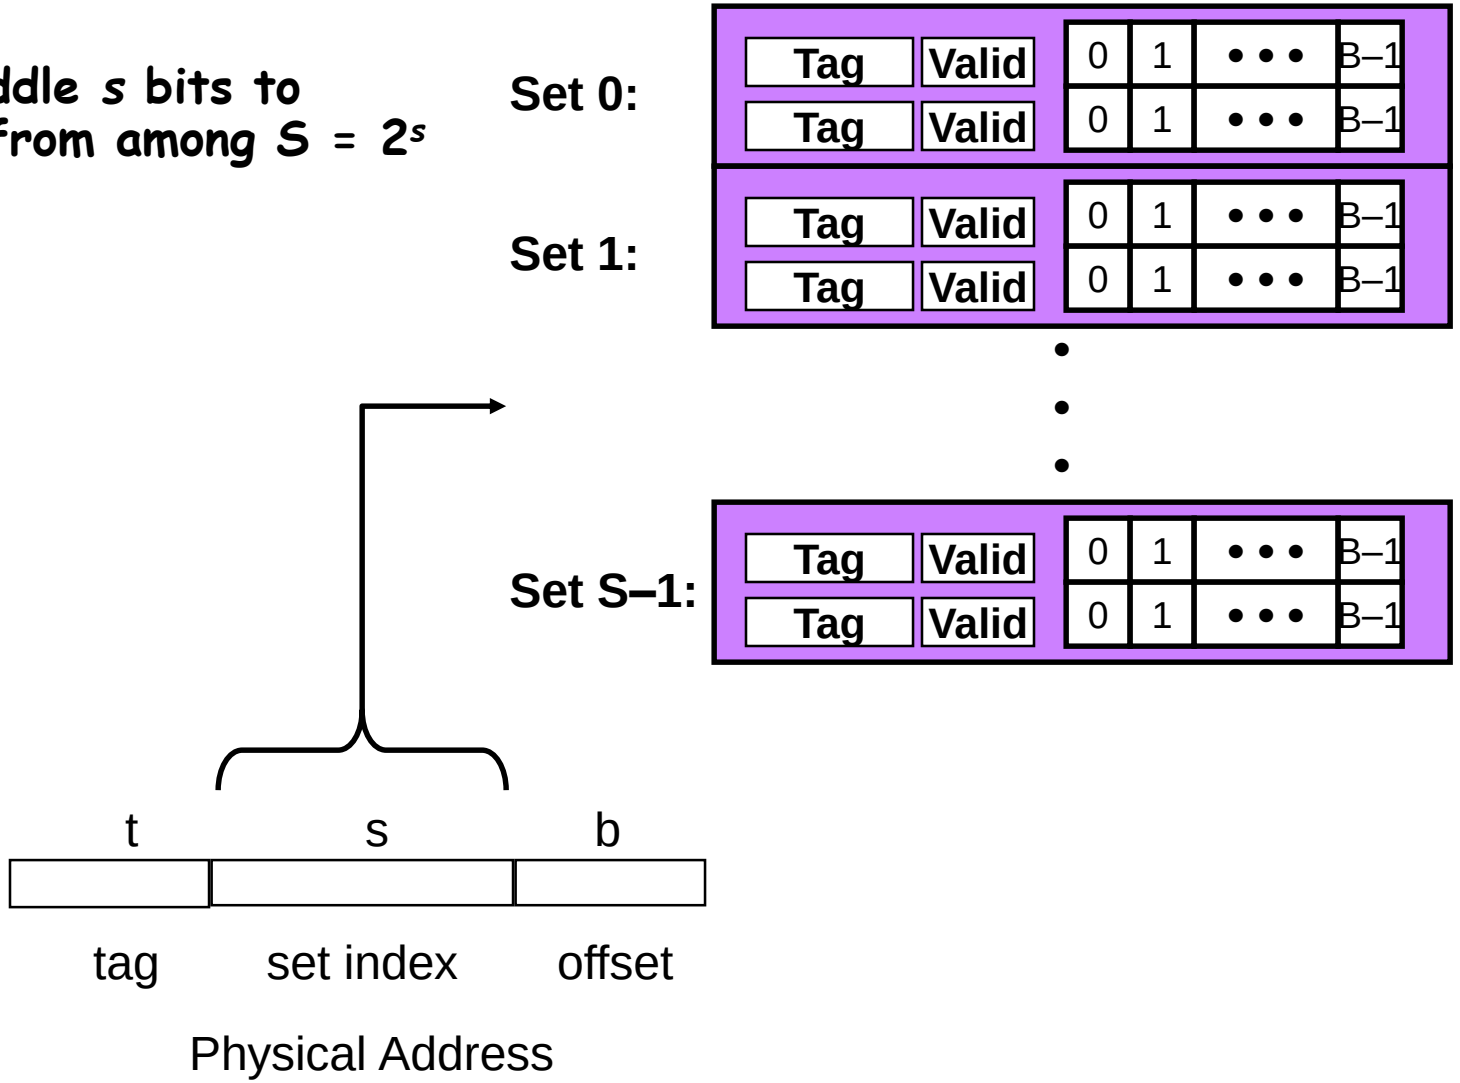

Associative Cache Tag Matching

Identifying Line

- Must have one of the tags match high order bits of address
- Must have *Valid = 1* for this line

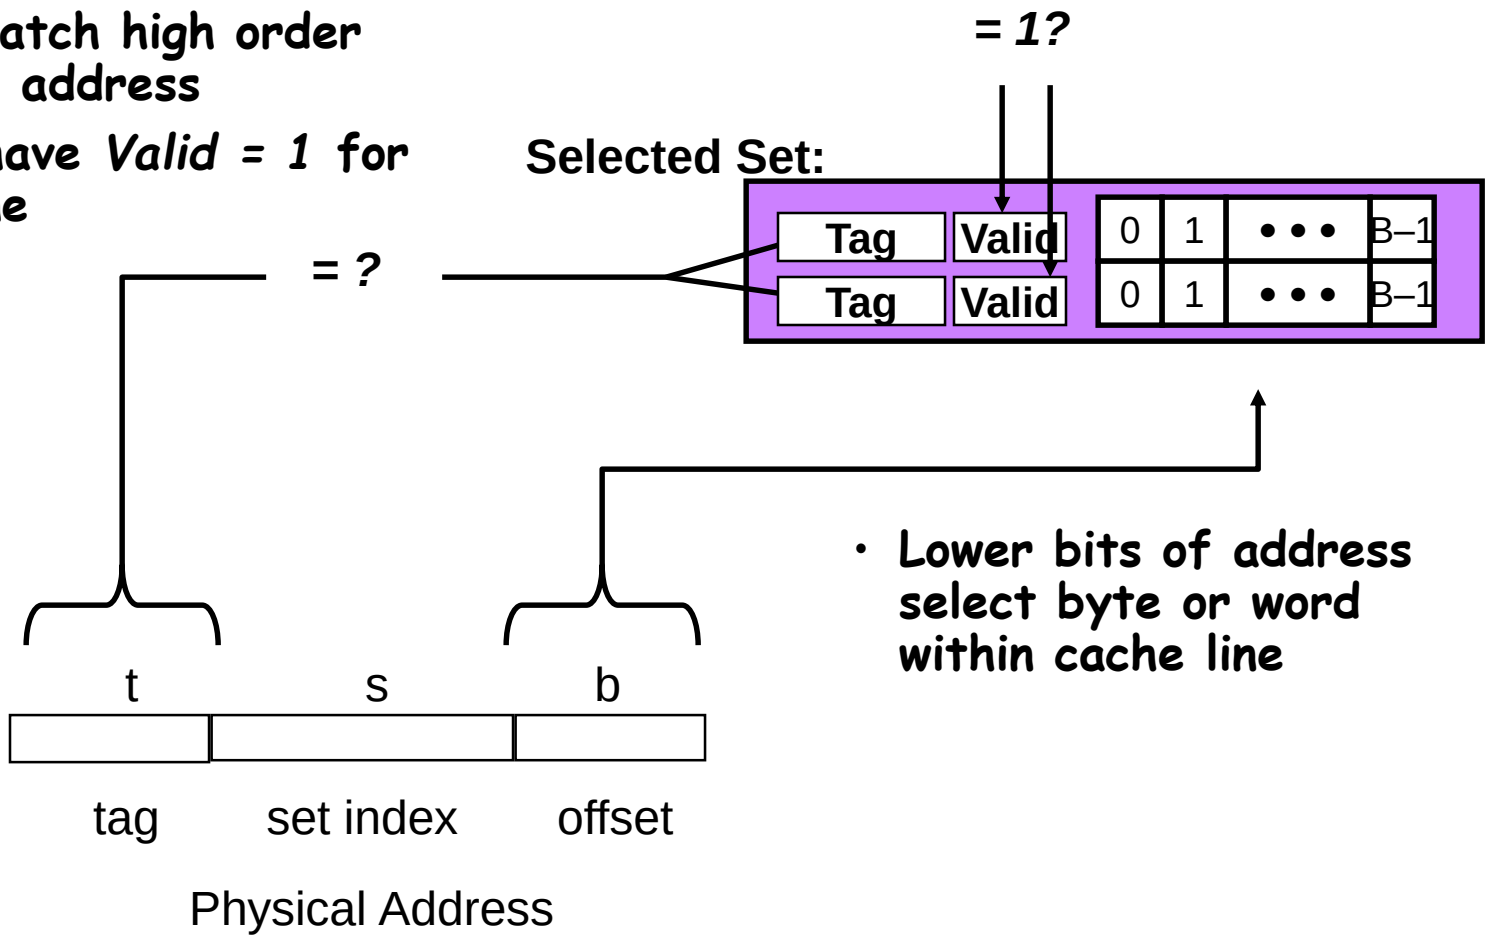

- Lower bits of address select byte or word within cache line

Two-Way Set Associative Cache Implementation

- Set index selects a set from the cache
- The two tags in the set are compared in parallel
- Data is selected based on the tag result

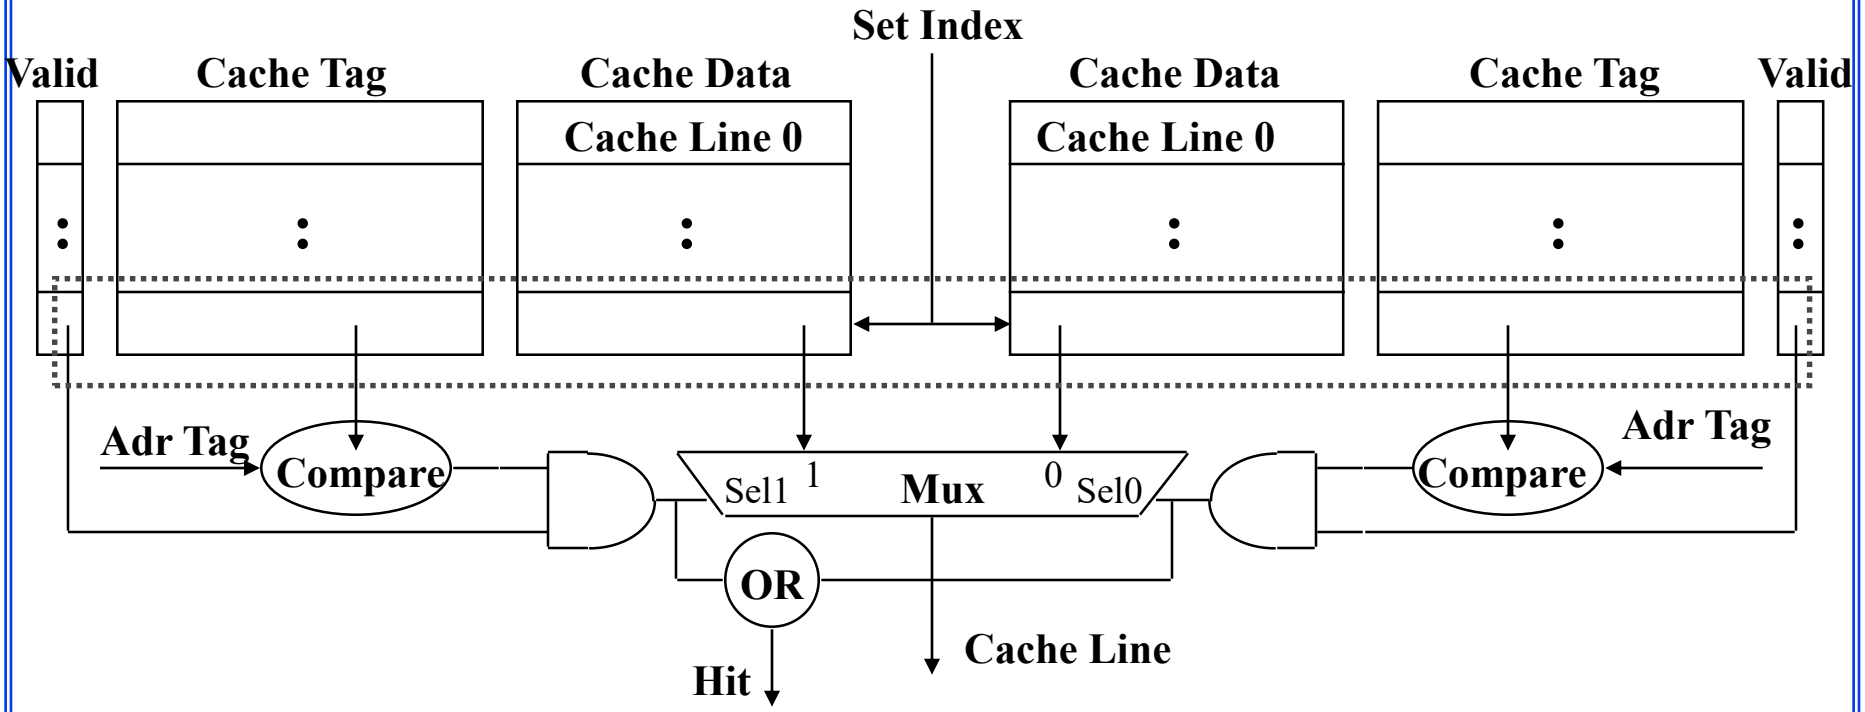

Impact of Associativity

Increasing associativity
(eg, direct-mapped ➔ set associative ➔ fully associative)

Effect on cache area (tags+data)?

Hit time?

Miss rate?

Miss Penalty?

Replacement Algorithms

Optimal algorithm: (Belady/MIN/OPT)

- replace the block that is next referenced furthest in the future
- **but** must know the future (can't be implemented)
- tricky to prove optimality; only optimal under "vanilla" cache designs

Usage based algorithms:

- **Least-recently used (LRU)**
  - replace the block that has been referenced least recently
  - hard to implement – but not really!
- **Least-frequently used (LFU)**
  - replace the block that has been referenced the fewest times
  - even harder to implement ("true" LFU—track blocks not in cache?)
- **Many approximations: CLOCK, tree-based pseudo-LRU, etc**

Non-usage based algorithms:

- **First-in First-out (FIFO)**
  - Weird pathologies (eg, hit rate degrades at larger cache size)
- **Random (RAND)**

Replacement Algorithms

Implementation

- FIFO: Keep per-set counter, replace block at counter offset + increment
- Random: Like FIFO, but a global counter instead
- Naïve LRU: encode ordering within set (n log n bits) + state machine
- Simple LRU: track time in # accesses, each candidate stores timestamp it was last accessed
  - Tradeoff?
  - Efficiency vs complexity
  - Coarsened ages (eg, high bits of timestamp) save space with ~no performance loss

Replacement Algorithms

Old & still active research area

- Fix pathologies in, eg, LRU
  - E.g.: ???
- Managing shared caches ("thread-aware" variants, cache partitioning)
  - Throughput vs fairness
- How to handle objects of different size?
- What observable factors correlate with reuse behavior?
  - Was referenced previously
  - PC of referencing instruction
  - Etc.
- What is best replacement metric if you don't know the future?

Miss Types: The 3 Cs++

Compulsory misses - unchanged

Capacity Misses – cache not big enough

- Capacity misses = fully associative misses with optimal replacement – compulsory misses

Replacement misses: those due to sub-optimal replacement decisions

- Replacement misses = fully associative misses – capacity misses

Conflict/Collision Misses – poor block placement

- Conflict misses = actual misses – replacement misses

Impact of Replacement Policy

Improving replacement policy
(eg, random → LRU)

Effect on cache area (tags+data)?

Hit time?

Miss rate?

Miss Penalty?

Write Policy

- What happens when processor writes to the cache?
- Should memory be updated as well?

Write Through:

- Store by processor updates cache *and* memory
- Memory always consistent with cache
- Never need to store from cache to memory
- ~2X more loads than stores

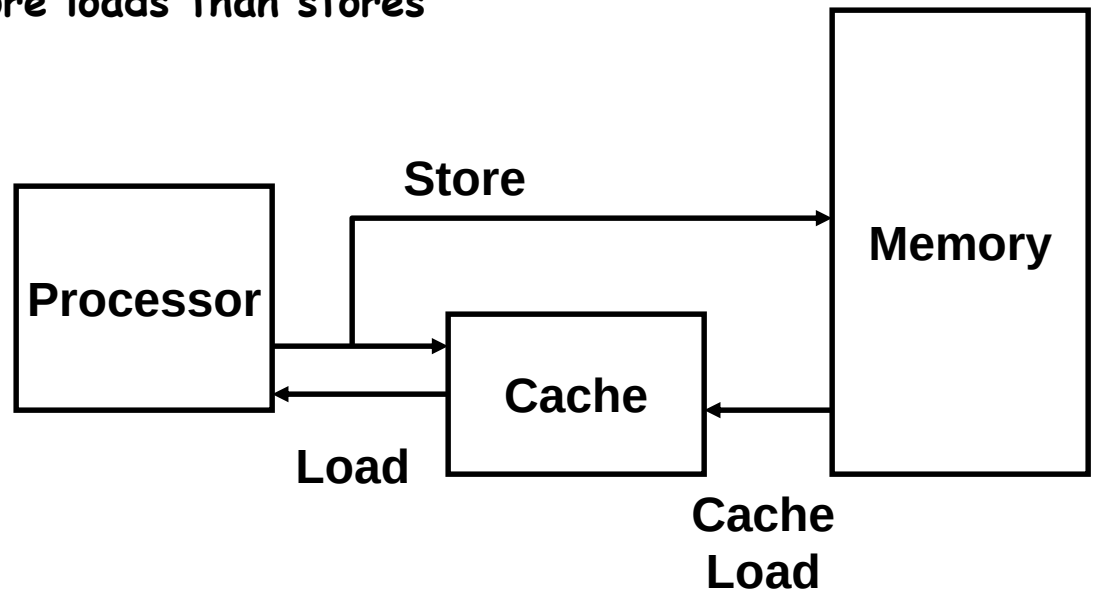

Write Policy (Cont.)

Write Back:

- **Store by processor only updates cache line**
- **Modified line written to memory only when it is evicted**
  - Requires "dirty bit" for each line
    - » Set when line in cache is modified
    - » Indicates that line in memory is stale
- **Memory not always consistent with cache**

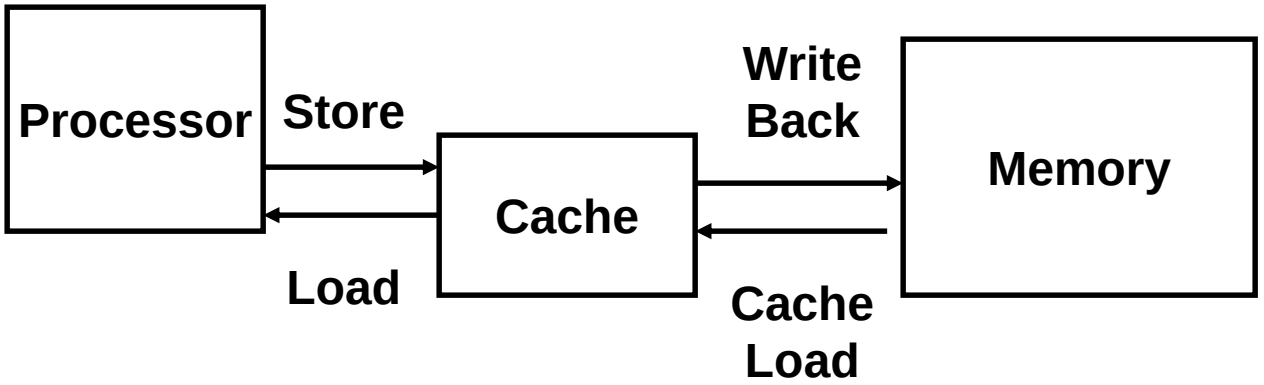

Write Buffering

Write Buffer

- Common optimization for all caches
- Overlaps memory updates with processor execution
- Read operation must check write buffer for matching address

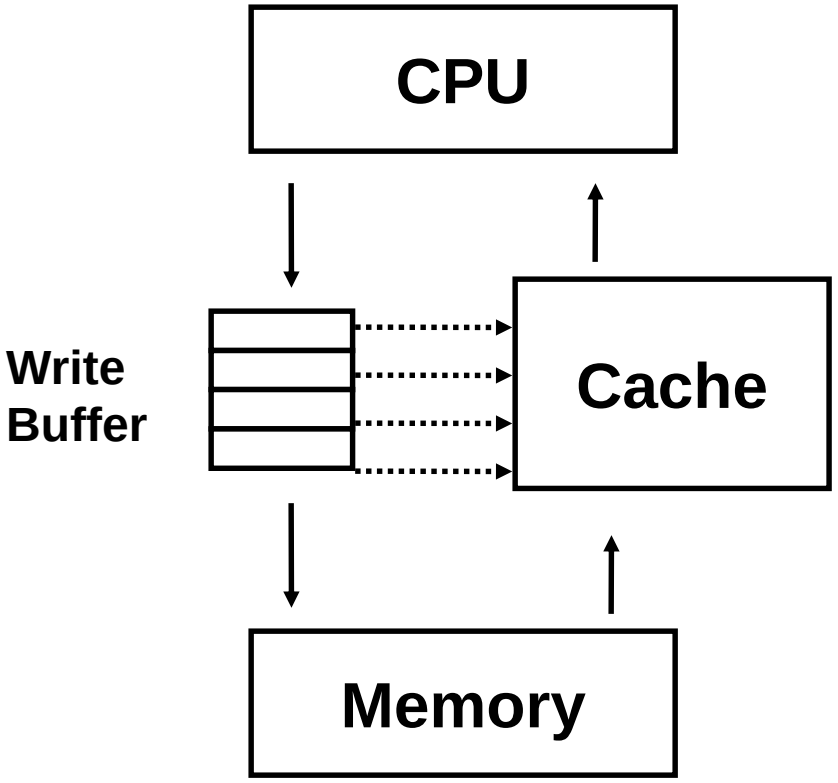

Allocation Strategies

- On a write miss, is the block loaded from memory into the cache?

Write Allocate:

- Block is loaded into cache on a write miss.
- Usually used with write back
- Otherwise, write-back requires read-modify-write to replace word within block

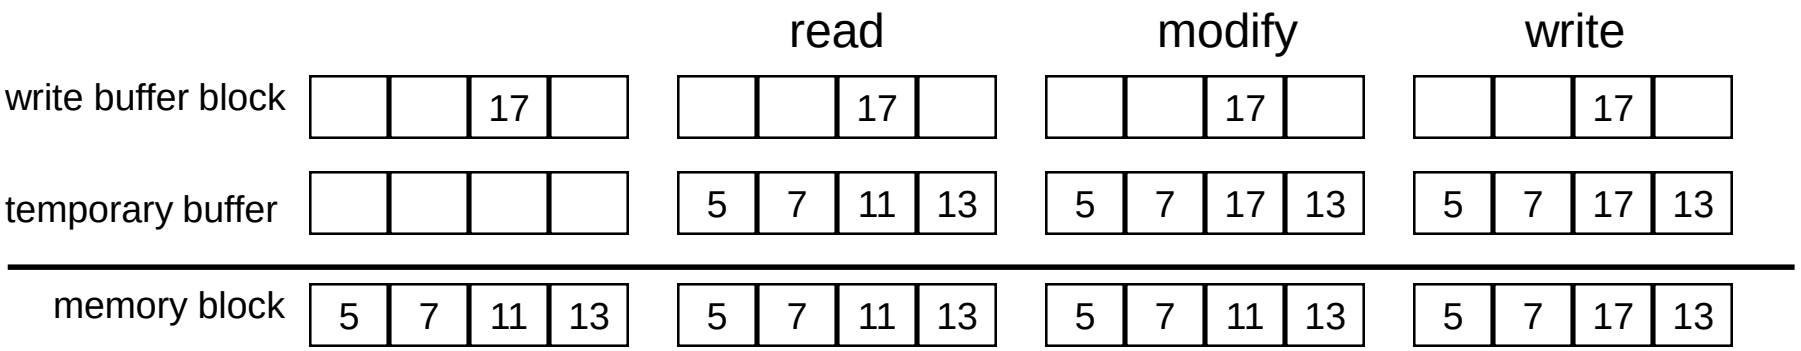

- But if you've gone to the trouble of reading the entire block, why not load it in cache?

Allocation Strategies (Cont.)

- **On a write miss, is the block loaded from memory into the cache?**

No-Write Allocate (Write Around):

- **Block is not loaded into cache on a write miss**
- **Usually used with write through**
  - Memory system directly handles word-level writes

Impact of Write Policy

Writeback vs write-through

Effect on cache area (tags+data)?

Hit time?

Miss rate?

Miss Penalty?

EXAMPLE: MATRIX MULTIPLY

Interactions Between Program & Cache

Major Cache Effects to Consider

- **Total cache size**
  - Try to keep heavily used data in highest level cache
- **Block size (sometimes referred to "line size")**
  - Exploit spatial locality

Variable `sum` held in register

```
/* ijk */
for (i=0; i<n; i++)  {
  for (j=0; j<n; j++) {
    sum = 0.0;
    for (k=0; k<n; k++)
      sum += a[i][k] * b[k][j];
    c[i][j] = sum;
  }
}
```

Example Application

- **Multiply n X n matrices**
- **$O(n^3)$ total operations**
- **Accesses**
  - n reads per source element
  - n values summed per destination
    - » But may be able to hold in register

Matmult Performance (Alpha 21164)

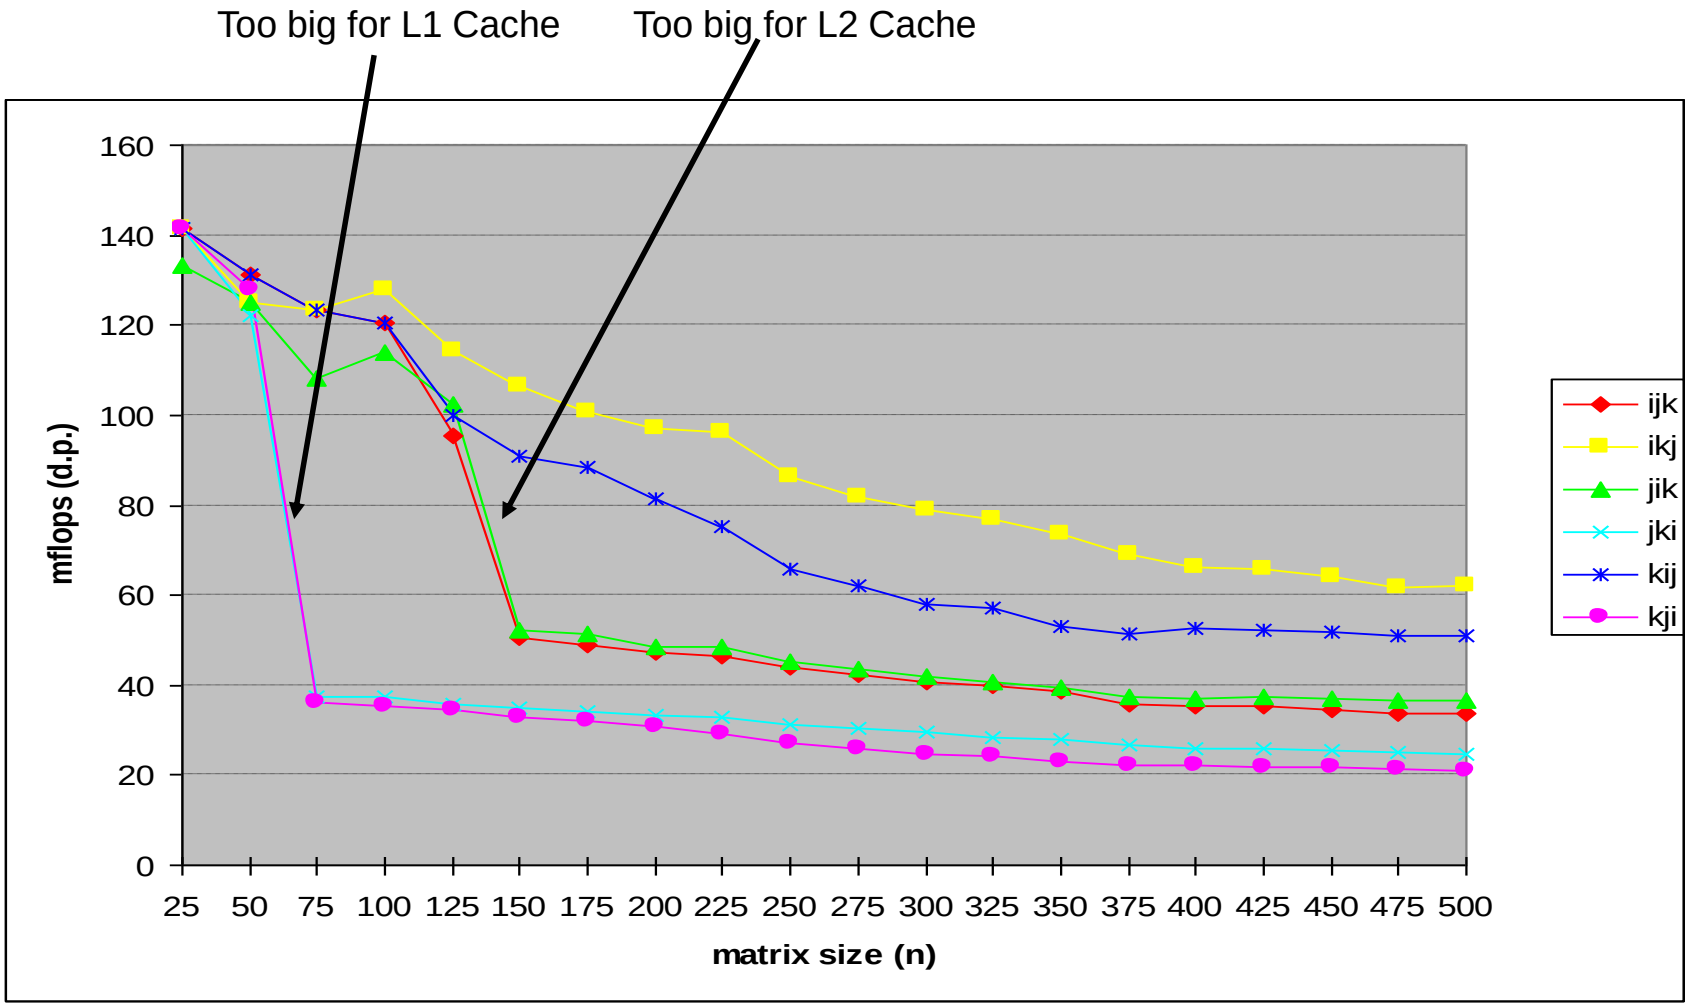

Block Matrix Multiplication

Example n=8, B = 4:

$$\begin{bmatrix} A11 & A12 \\ A21 & A22 \end{bmatrix} \times \begin{bmatrix} B11 & B12 \\ B21 & B22 \end{bmatrix} = \begin{bmatrix} C11 & C12 \\ C21 & C22 \end{bmatrix}$$

Key idea: Sub-blocks (i.e., A_{ij}) can be treated just like scalars.

$$C_{11} = A_{11}B_{11} + A_{12}B_{21} \qquad C_{12} = A_{11}B_{12} + A_{12}B_{22}$$

$$C_{21} = A_{21}B_{11} + A_{22}B_{21} \qquad C_{22} = A_{21}B_{12} + A_{22}B_{22}$$

Blocked Matrix Multiply (bijk)

```
for (jj=0; jj<n; jj+=bsize) {
  for (i=0; i<n; i++)
    for (j=jj; j < min(jj+bsize,n); j++)
      c[i][j] = 0.0;
  for (kk=0; kk<n; kk+=bsize) {
    for (i=0; i<n; i++) {
      for (j=jj; j < min(jj+bsize,n); j++) {
        sum = 0.0
        for (k=kk; k < min(kk+bsize,n); k++) {
          sum += a[i][k] * b[k][j];
        }
        c[i][j] += sum;
      }
    }
  }
}
```

Blocked Matrix Multiply Analysis

- Innermost loop pair multiplies 1 X bsize sliver of A times bsize X bsize block of B and accumulates into 1 X bsize sliver of C
- Loop over i steps through n row slivers of A & C, using same B

```
for (i=0; i<n; i++) {
    for (j=jj; j < min(jj+bsize,n); j++) {
      sum = 0.0
      for (k=kk; k < min(kk+bsize,n); k++) {
        sum += a[i][k] * b[k][j];
      }
      c[i][j] += sum;
}
```

Innermost Loop Pair

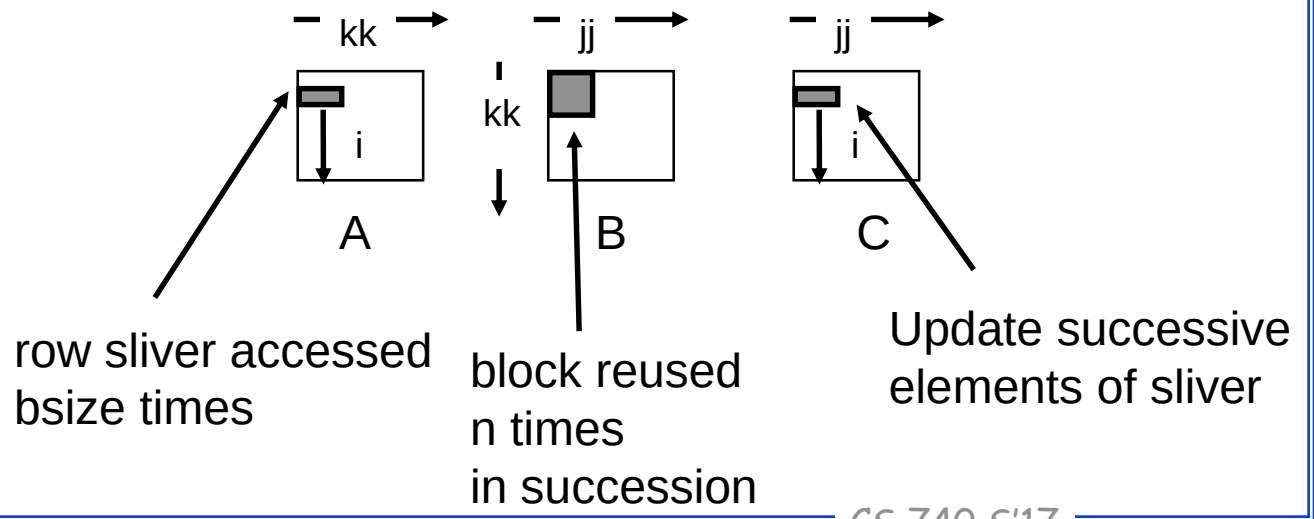

Blocked matmult perf (Alpha 21164)

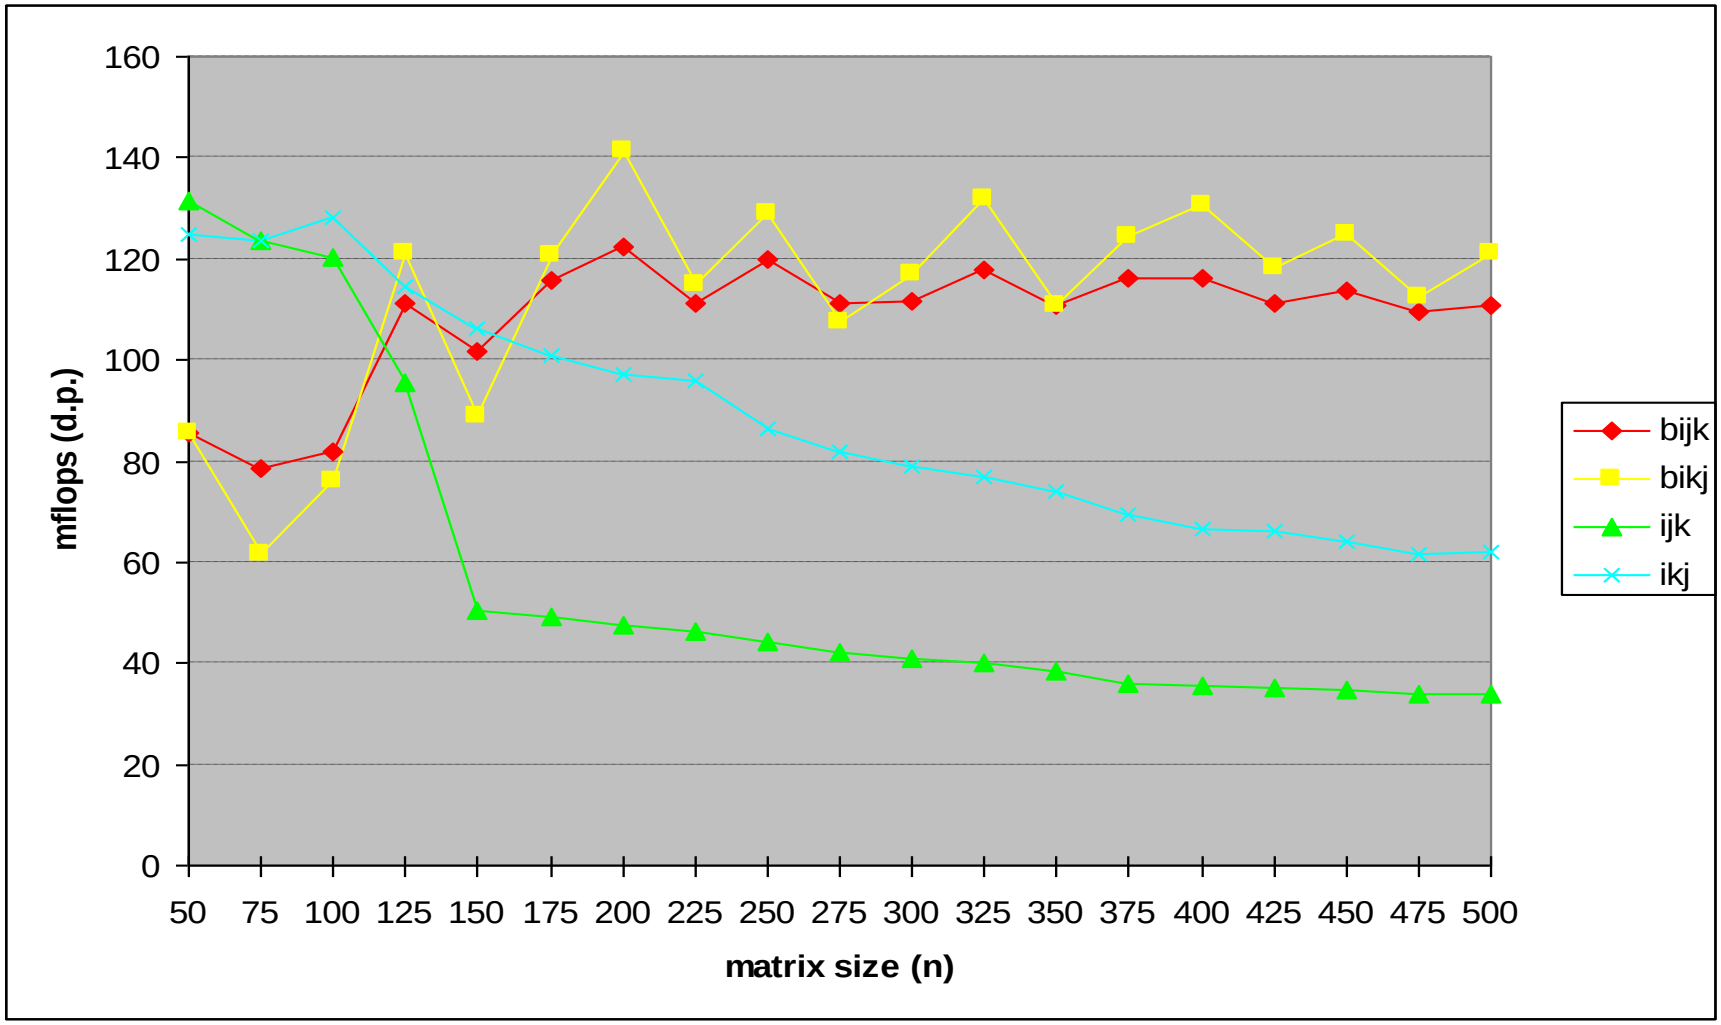

ASSOCIATIVITY ABSTRACTION

What is associativity?

Simple answer: number of replacement candidates

More associativity ➔ better hit rates

- 1-way < 2-way < 3-way < … < fully associative

What is associativity?

Simple answer: number of replacement candidates

Victim caches

- Candidates include recently evicted blocks
- Does 1-way + 1-entry victim cache == 2-way?

- 1-way < 1-way + 1-entry victim cache < 2-way

What is associativity?

Simple answer: number of replacement candidates

Hashing

- Hash address to compute set
- Reduce conflict misses
- Add latency + tag size + complexity

- 1-way < 1-way hashed < 2-way ??????
- 8-way < 8-way hashed ??????

What is associativity?

Simple answer: number of replacement candidates

Skew-associative caches [Seznec, ISCA'93]

- Use different hash function for each *way*
- Mixes candidates across sets for diff addresses

- 2-way < 2-way hash < 2-way skew < 3-way ?????

What is associativity?

Simple answer: number of replacement candidates

Associativity can be thought as a distribution of victims' eviction priority [Sanchez, MICRO'10]

- Distribution answers two questions: *Among all cached blocks, how much did I want to evict the victim? (y-axis) How likely was that? (x-axis)*
- Fully associative always evicts the highest rank
- Random sampling converges toward fully associative with larger samplers
- Can plot associativity distribution (eg, through simulation) for different cache organizations

Memory Hierarchy Summary

Gap between memory + compute is growing

Processors often spend most of their time + energy waiting for memory, not doing useful work

Hierarchy and *locality* are the key ideas to scale memory performance

Most systems use caches, which introduce many parameters to the design with many tradeoffs

- E.g., associativity—hit rate vs hit latency

Bandwidth Matching

Challenge

- CPU works with short cycle times
- DRAM (relatively) long cycle times
- *How can we provide enough bandwidth between processor & memory?*

Effect of Caching

- Caching greatly reduces amount of traffic to main memory
- But, sometimes need to move large amounts of data from memory into cache

Trends

- Need for high bandwidth much greater for multimedia applications
  - Repeated operations on image data
- Recent generation machines greatly improve on predecessors

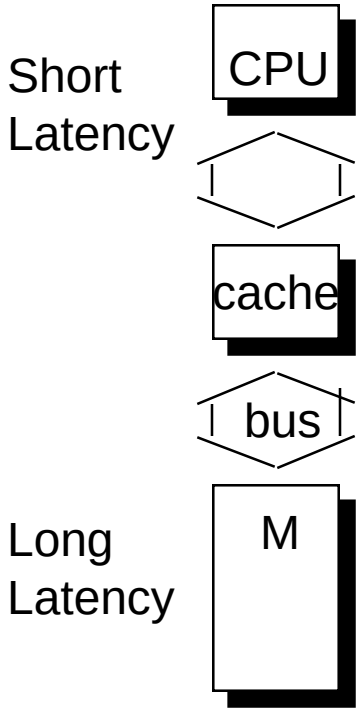

High Bandwidth Memory Systems

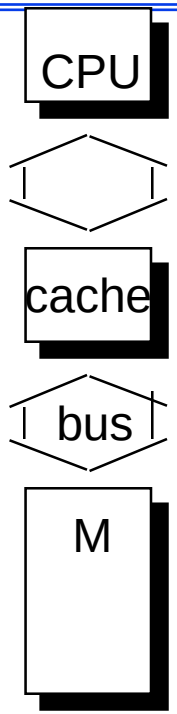

Solution 1
High BW DRAM

Example:
Page Mode DRAM
RAMbus

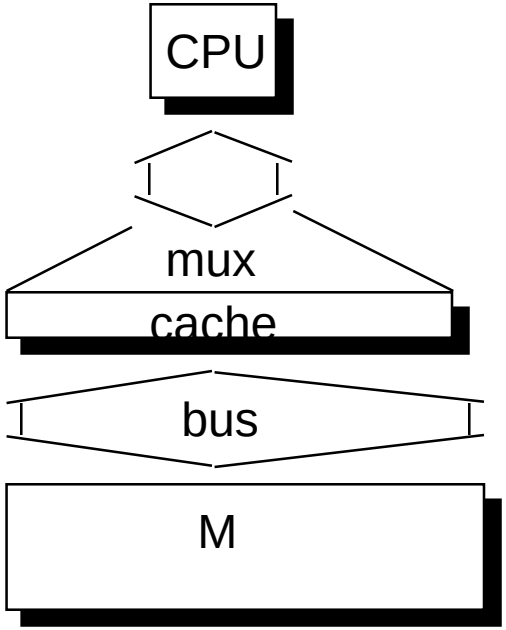

Solution 2
Wide path between memory & cache

Example: Alpha AXP 21064
256 bit wide bus, L2 cache,
and memory.

Cache Performance Metrics

Miss Rate

- **fraction of memory references not found in cache (misses/references)**
- **Typical numbers:**

  3-10% for L1

  can be quite small (e.g., < 1%) for L2, depending on size, etc.

Hit Time

- **time to deliver a line in the cache to the processor (includes time to determine whether the line is in the cache)**
- **Typical numbers:**

  1-3 clock cycles for L1

  3-12 clock cycles for L2

Miss Penalty

- **additional time required because of a miss**
  - Typically 25-100 cycles for main memory

Qualitative Cache Performance Model

Miss Types

- **Compulsory ("Cold Start") Misses**
- **Capacity Misses**
- **Replacement Misses**
  - Not universally included in classification
- **Conflict Misses**
- **Coherence Misses**
  - Block invalidated by multiprocessor cache coherence mechanism

Hit Types

- **Reuse hit**
  - Accessing same word that previously accessed
- **Line hit**
  - Accessing word spatially near previously accessed word