

Chapter 4

Memory System

π Contents

- › Computer memory system
- › Cache memory
- › Internal memory
- › External memory

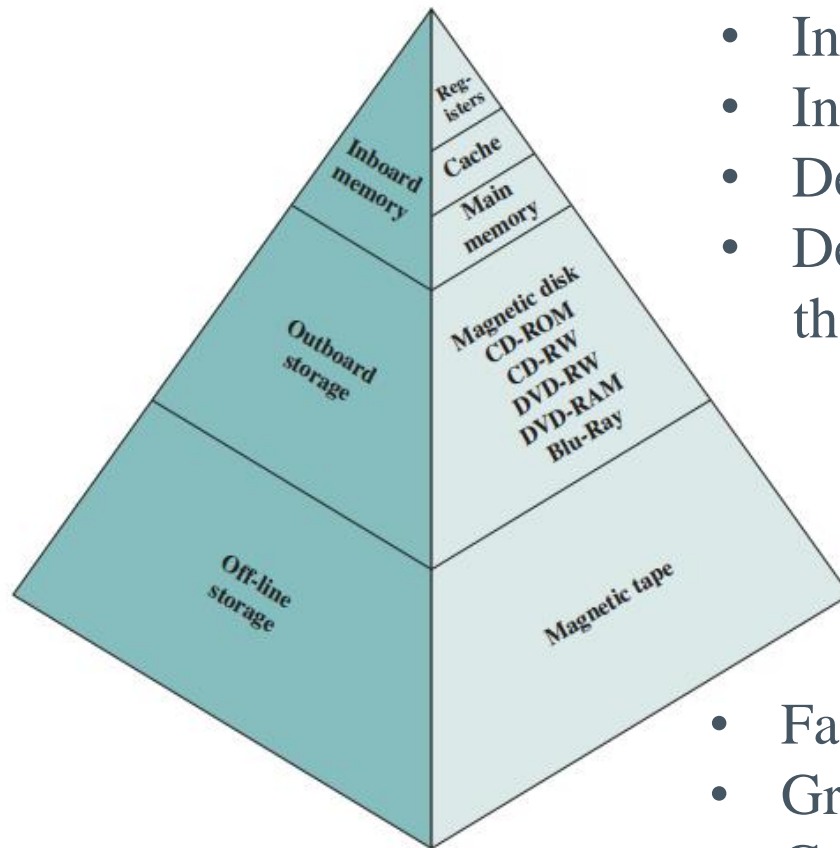


Computer memory system

› Key characteristics:

Location	Performance
Internal (e.g., processor registers, cache, main memory)	Access time
External (e.g., optical disks, magnetic disks, tapes)	Cycle time
	Transfer rate
Capacity	Physical Type
Number of words	Semiconductor
Number of bytes	Magnetic
	Optical
Unit of Transfer	Magneto-optical
Word	Physical Characteristics
Block	Volatile/nonvolatile
	Erasable/nonerasable
Access Method	Organization
Sequential	Memory modules
Direct	
Random	
Associative	

Memory Hierarchy



As one goes down the hierarchy:

- Increasing capacity
- Increasing access time
- Decreasing cost per bit
- Decreasing frequency of access of the memory by the processor

- Faster access time, greater cost per bit
- Greater capacity, smaller cost per bit
- Greater capacity, slower access time

π Contents


- › Computer memory system
- › Cache memory
- › Internal memory
- › External memory



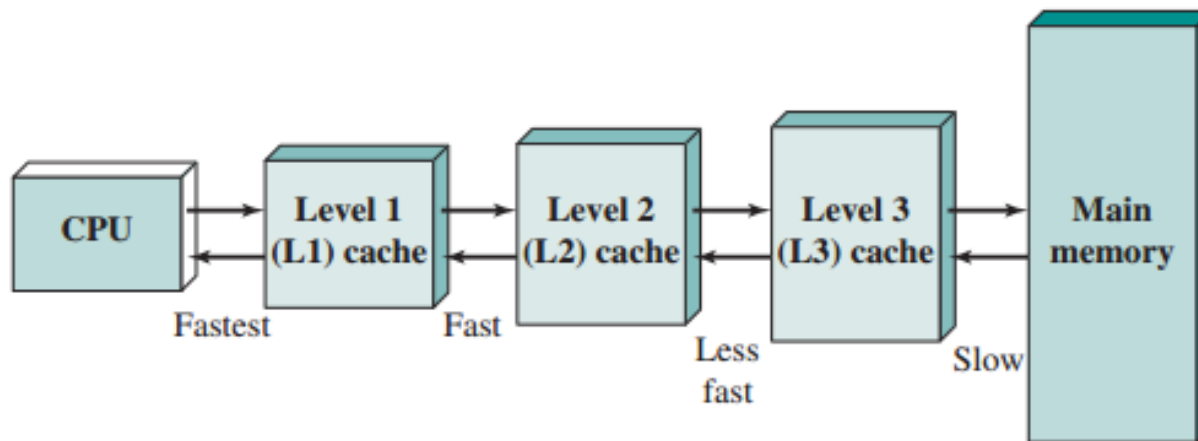
Cache memory principles

- › A cache is a high speed memory that is located between CPU and main memory and keeps a copy of the frequently used data
- › When the CPU wants a data value from memory, it first looks in the cache
- › If the data is in the cache, it uses that data
- › If the data is not in the cache, it copies a **line** of data from memory to the cache and gives the CPU what it wants
- › **Cache hit:** When the processor wants information at a given address and that information is already in the cache
- › **Cache miss:** If the processor wants information that is not already in the cache

π Cache memory principles



(a) Single cache



(b) Three-level cache organization

π Cache line

- › The CPU usually accesses a **word** of memory at a time
- › When memory is copied from the RAM to the cache, it usually copies a **line** of data
- › A line is usually much larger than a word, often 4 to 16 words

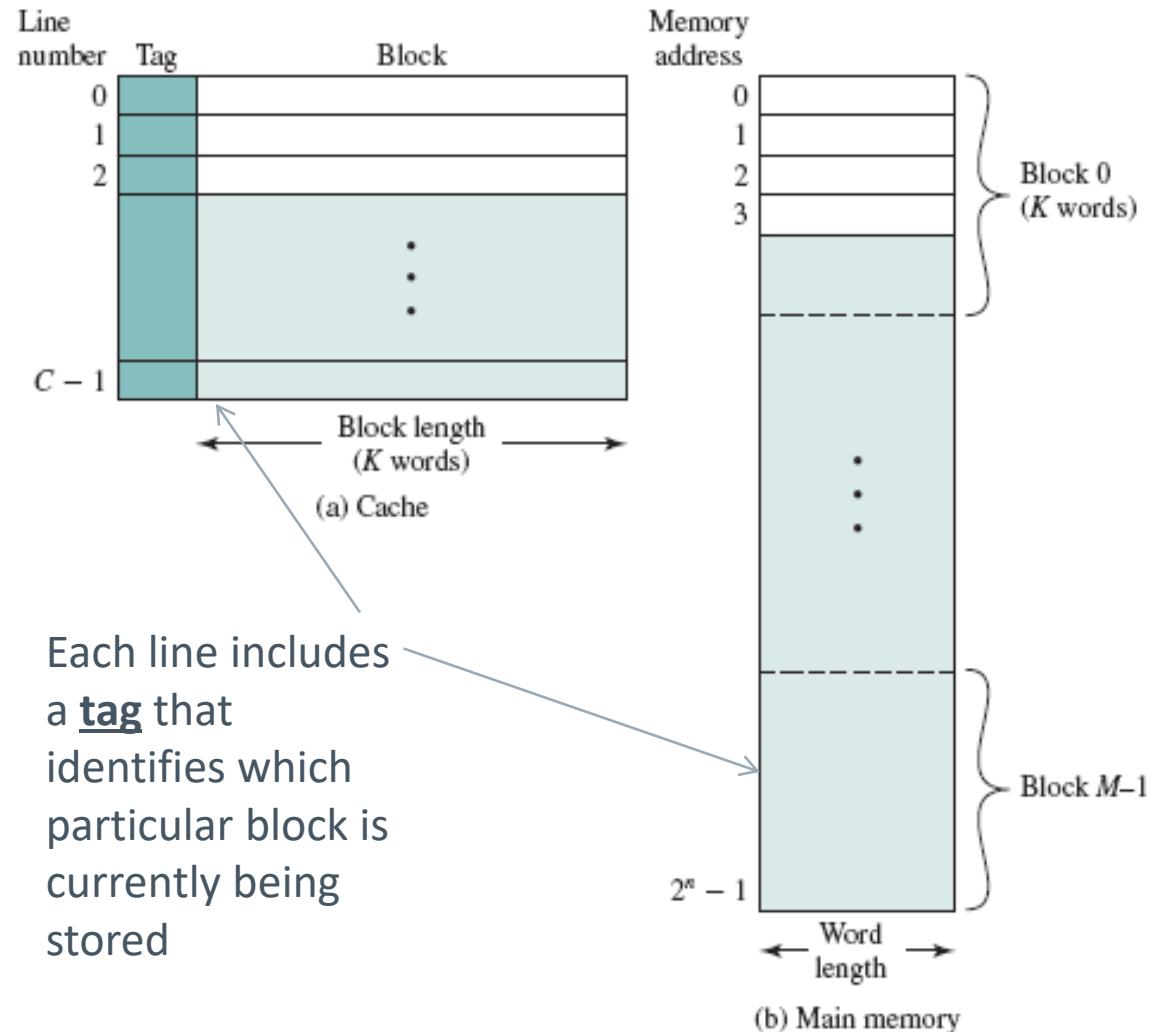
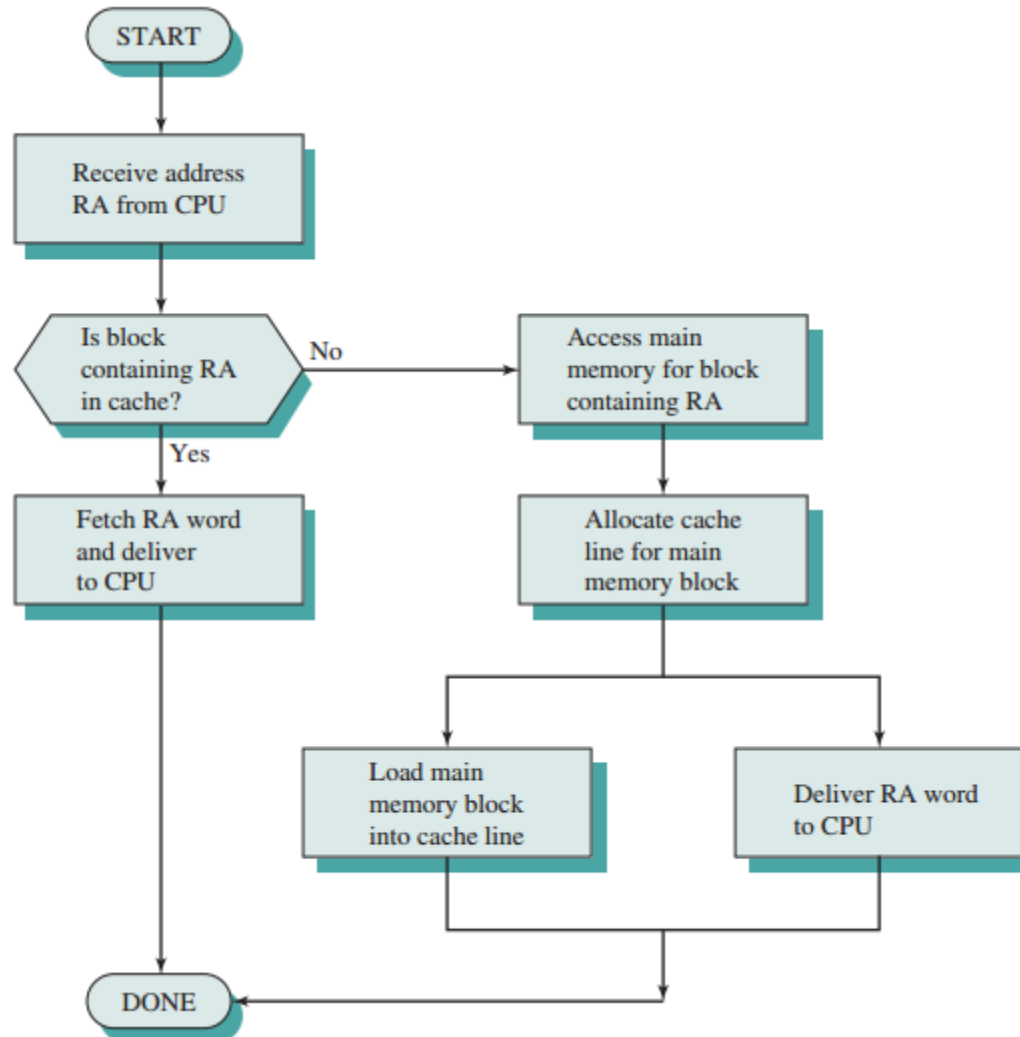


Figure 4.4 Cache/Main Memory Structure

π Cache read operation



Elements of cache design

Cache Addresses

Logical

Physical

Cache Size**Mapping Function**

Direct

Associative

Set associative

Replacement Algorithm

Least recently used (LRU)

First in first out (FIFO)

Least frequently used (LFU)

Random

Write Policy

Write through

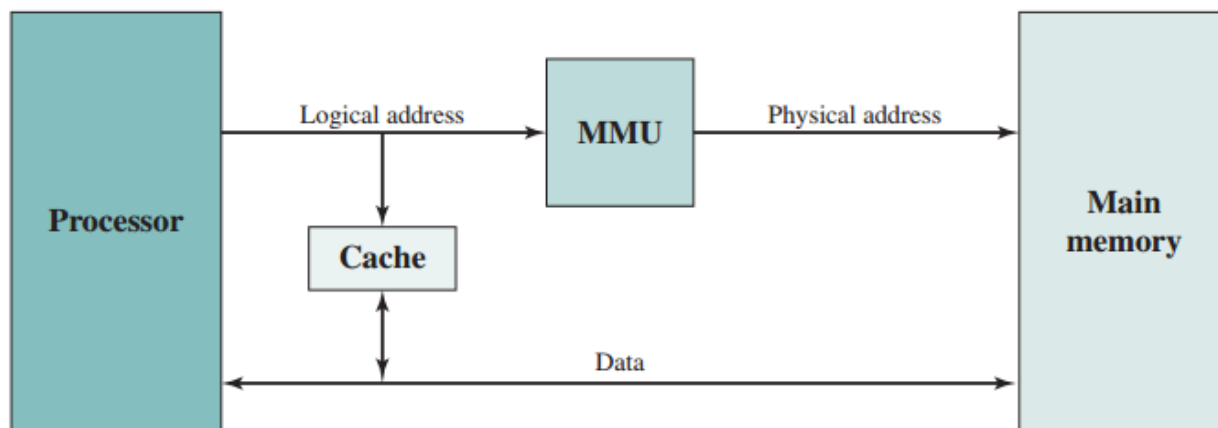
Write back

Line Size**Number of Caches**

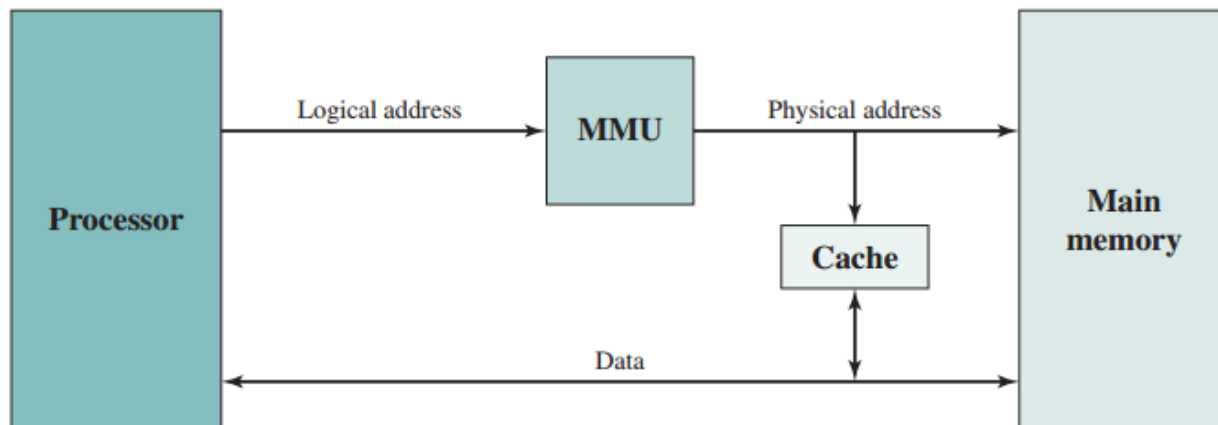
Single or two level

Unified or split

π Cache addresses



(a) Logical cache



(b) Physical cache



Cache size


- › Bigger is better
- › Cache size is limited by:
 - Space on the processor chip
 - Manufacturing limitations
 - Cost

π Mapping

- › Algorithm for mapping main memory blocks into cache lines
- › Mapping techniques:
 - Direct: each address has a specific place in the cache
 - Fully associative: search the entire cache for an address
 - Set associative: each address can be in any of a small set of cache locations

Direct mapping

- › The block j in main memory will be loaded to the line i of the cache: $i = j \text{ mode } m$



(a) Direct mapping

π Direct mapping

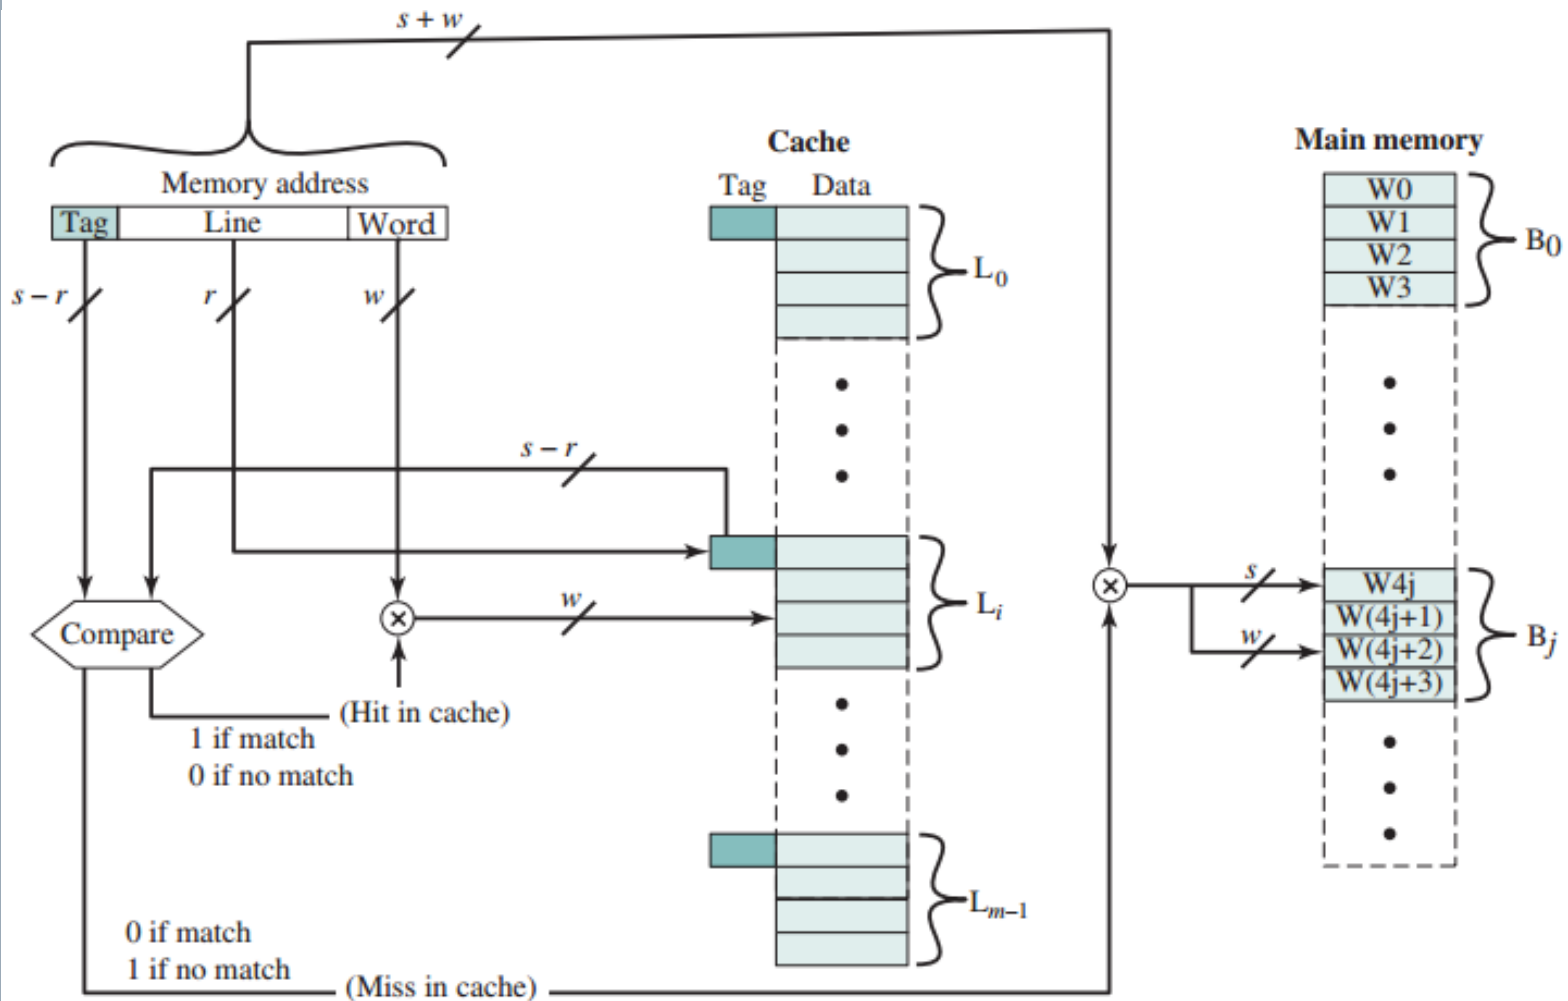



Figure 4.9 Direct-Mapping Cache Organization

π Fully associative mapping

- › A block in main memory can be load to any line of the cache



(b) Associative mapping

π Fully associative mapping

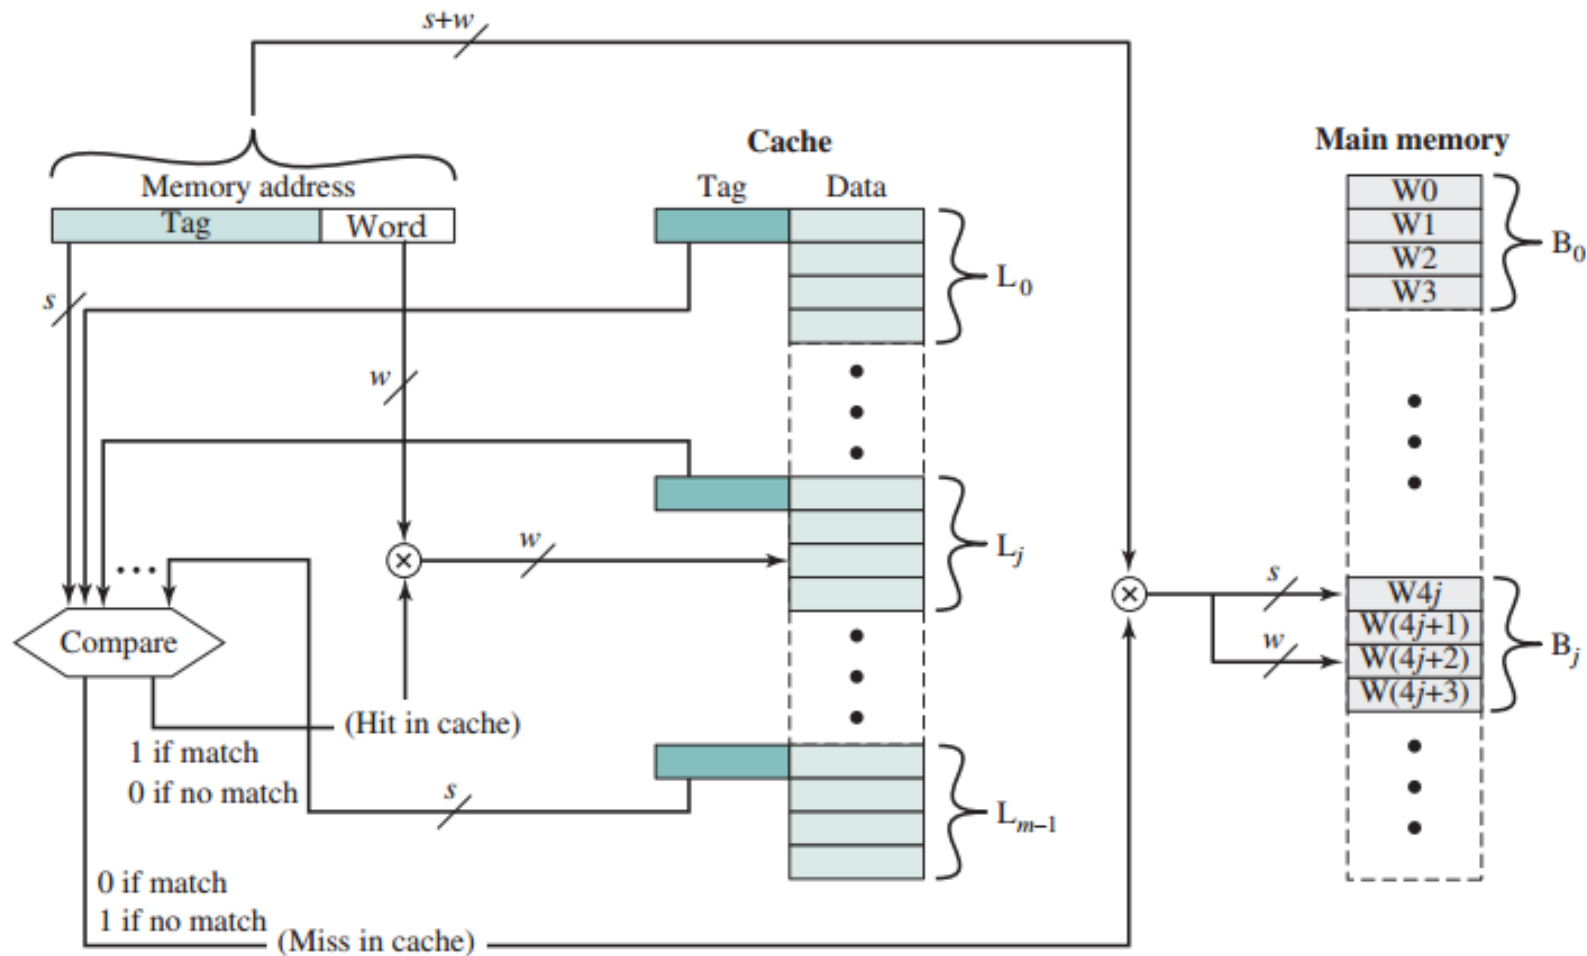


Figure 4.11 Fully Associative Cache Organization

π Set associative mapping

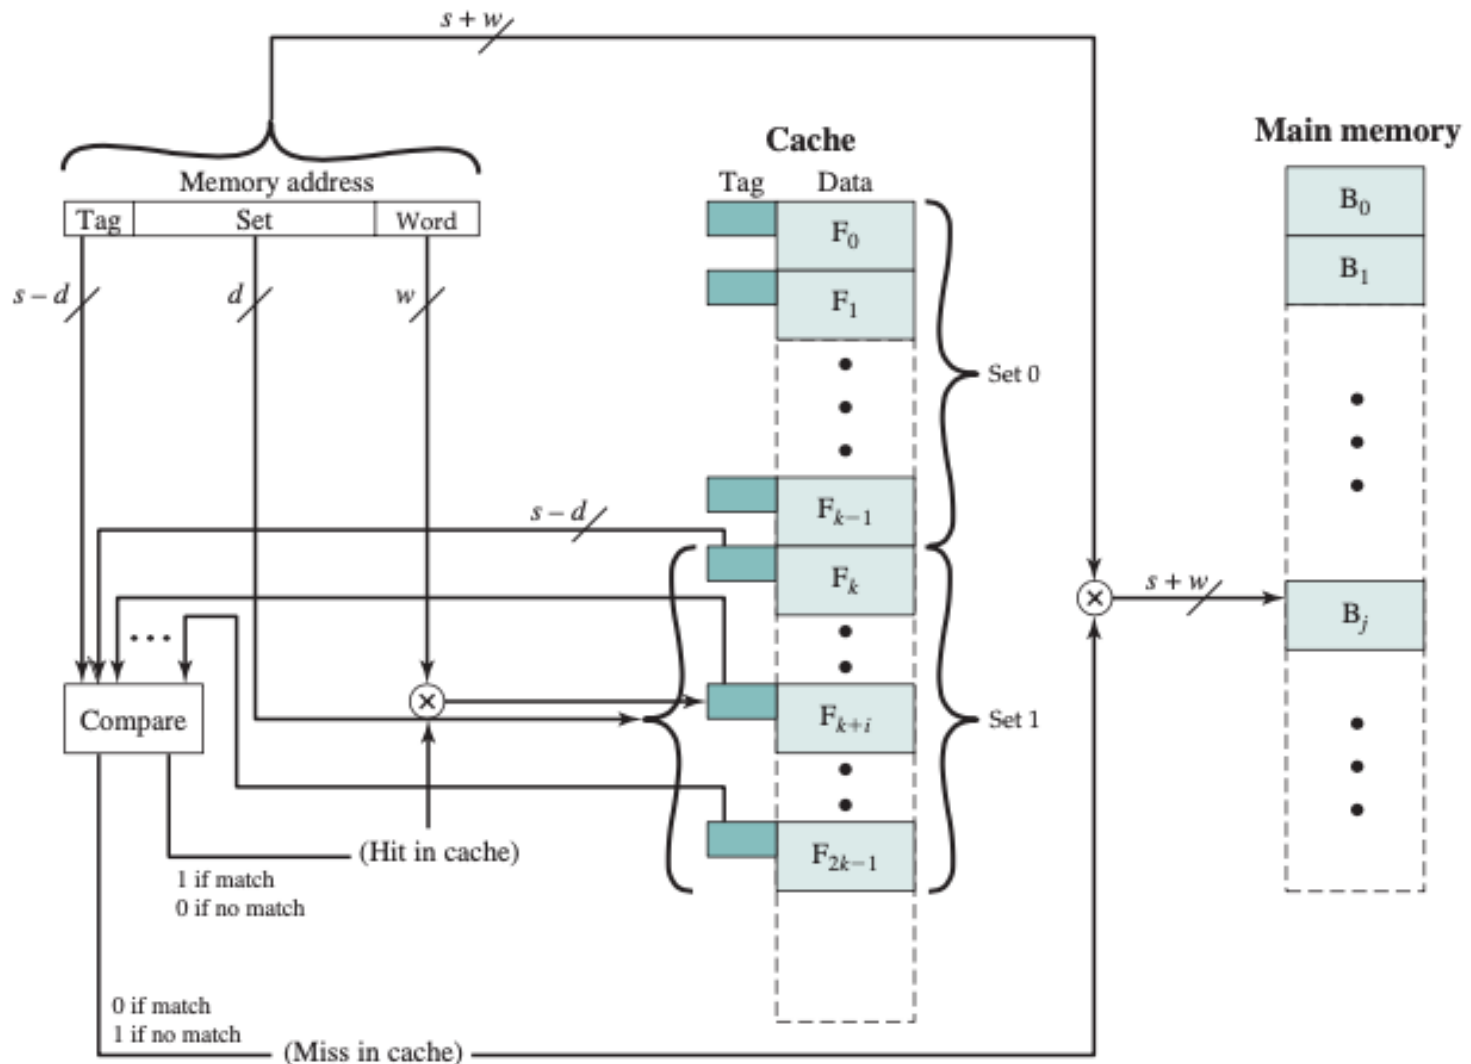


Figure 4.14 K-Way Set Associative Cache Organization

π

Set-Associative

- 4.1** A set-associative cache consists of 64 lines, or slots, divided into four-line sets. Main memory contains 4K blocks of 128 words each. Show the format of main memory addresses.
- 4.2** A two-way set-associative cache has lines of 16 bytes and a total size of 8 Kbytes. The 64-Mbyte main memory is byte addressable. Show the format of main memory addresses.

4.1 A set-associative cache consists of 64 lines, or slots, divided into four-line sets. Main memory contains 4K blocks of 128 words each. Show the format of main memory addresses.

π

Set-Associative

- › 4.3.a Direct (Tag, Line, Word)
- › 4.3.c Set Ass (Tag, Set, Word)
- › Chỉ tính cho địa chỉ BBBBBB



4.3 For the hexadecimal main memory addresses 111111, 666666, BBBBBB, show the following information, in hexadecimal format:

- a. Tag, Line, and Word values for a direct-mapped cache, using the format of Figure 4.10
- b. Tag and Word values for an associative cache, using the format of Figure 4.12
- c. Tag, Set, and Word values for a two-way set-associative cache, using the format of Figure 4.15

- a. Specify how the bits B0, B1, and B2 are set and then describe in words how they are used in the replacement algorithm depicted in Figure 4.20.

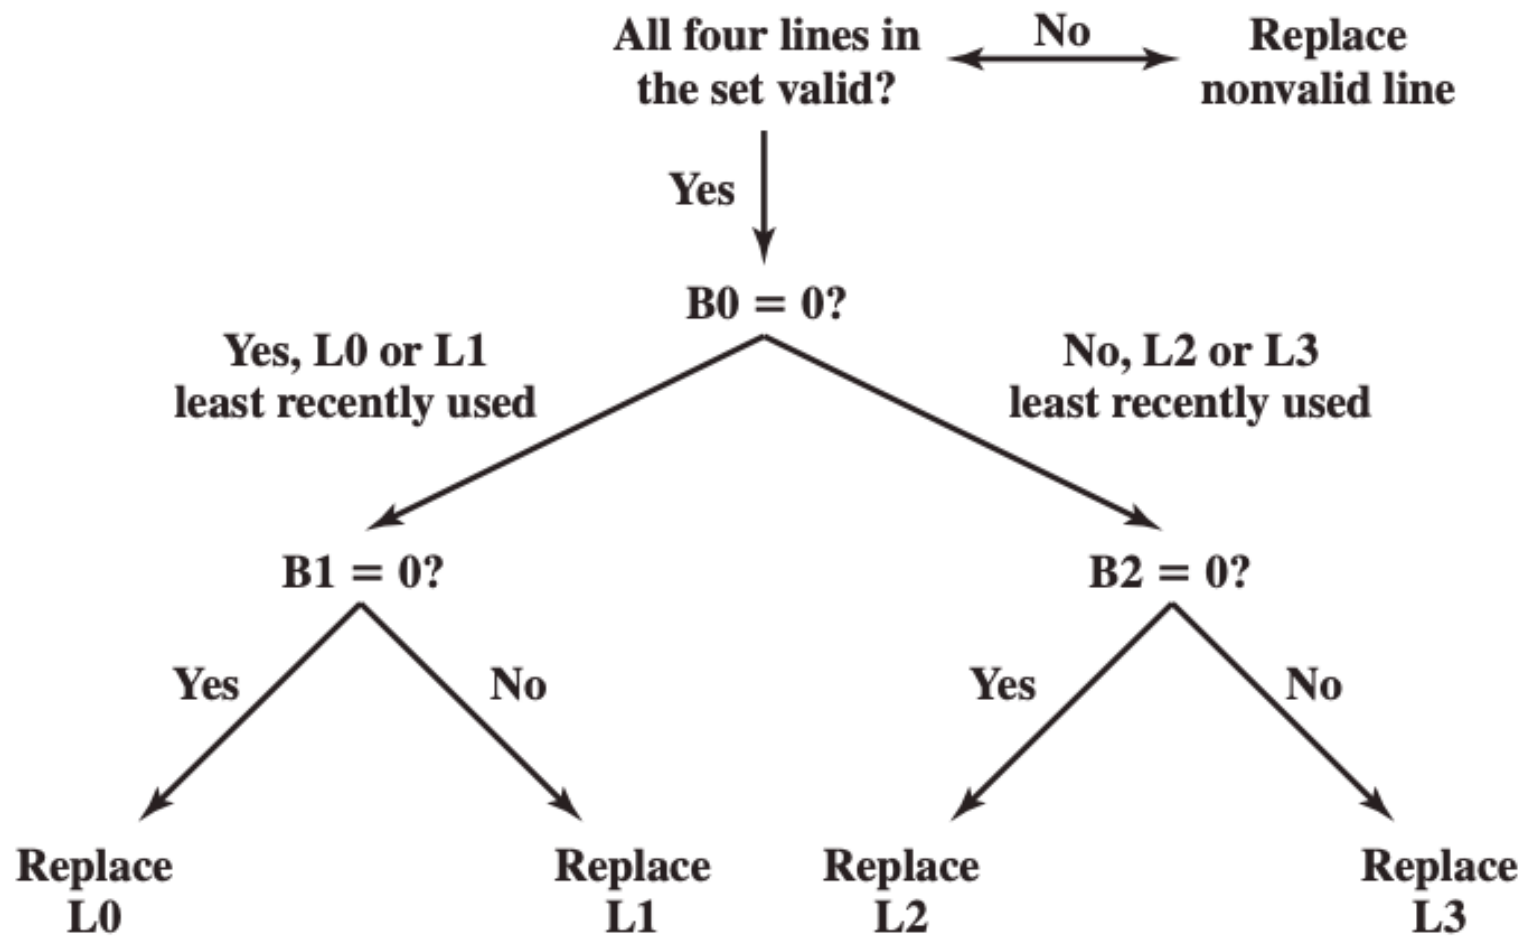


Figure 4.20 Intel 80486 On-Chip Cache Replacement Strategy

- 4.8** Consider a machine with a byte addressable main memory of 2^{16} bytes and block size of 8 bytes. Assume that a direct mapped cache consisting of 32 lines is used with this machine.
- How is a 16-bit memory address divided into tag, line number, and byte number?
 - Into what line would bytes with each of the following addresses be stored?
0001 0001 0001 1011
1100 0011 0011 0100
1101 0000 0001 1101
1010 1010 1010 1010
 - Suppose the byte with address 0001 1010 0001 1010 is stored in the cache. What are the addresses of the other bytes stored along with it?
 - How many total bytes of memory can be stored in the cache?
 - Why is the tag also stored in the cache?

π Replacement policy

- › When a cache miss occurs, data is copied into some location in cache
- › With Set Associative or Fully Associative mapping, the system must decide where to put the data and what values will be replaced
- › Cache performance is greatly affected by properly choosing data that is unlikely to be referenced again
- › Replacement algorithms:
 - First In First Out (FIFO)
 - Least Recently Used (LRU)
 - Least Frequently Used (LFU)
 - Random



Write policy

- › When the processor changes a data item, it is changed in the cache. If this line of cache is replaced, main memory needs to be updated
- › Write through: cache and main memory are updated at the same time whenever data is changed
- › Write back: main memory is updated when the cache line is replaced

π Line size

- › A cache line is the amount of data copied into the cache at one time; 16 to 64 bytes is common
- › When there is a cache miss, a line of memory is copied into the cache
- › The bigger the line size, the fewer lines in the cache
- › Big lines will copy more nearby addresses
- › The relationship between block size and hit ratio depends on the locality characteristics of a particular program



Number of caches

- › Multi-level cache:
 - Most systems have two or three cache levels
 - Higher level caches are usually larger
- › Unified vs. split caches
 - Unified cache: a single cache used to store references to both data and instructions
 - Split caches: the cache is split into two, one dedicated to instructions and one dedicated to data




Contents

- › Computer memory system
- › Cache memory
- › Internal memory
- › External memory

Semiconductor Main Memory

- › Basic element of a semiconductor memory is the memory cell.
- › Cell properties:
 - They exhibit two stable (or semistable) states, which can be used to represent binary 1 and 0.
 - They are capable of being written into (at least once), to set the state.
 - They are capable of being read to sense the state
- Memory cell operation:



π Semiconductor Memory Types

Memory Type	Category	Erase	Write Mechanism	Volatility
Random-access memory (RAM)	Read-write memory	Electrically, byte-level	Electrically	Volatile
Read-only memory (ROM)	Read-only memory	Not possible	Masks	Nonvolatile
Programmable ROM (PROM)			Electrically	
Erasable PROM (EPROM)	UV light, chip-level			
Electrically Erasable PROM (EEPROM)	Electrically, byte-level			
Flash memory	Electrically, block-level			
	Read-mostly memory			

π DRAM and SRAM

DRAM (Dynamic RAM)	SRAM (Static RAM)
Made with cells that store data as charge on capacitors	Digital device that uses the same logic elements used in the processor
Presence or absence of charge in a capacitor is interpreted as a binary 1 or 0	Binary values are stored using traditional flip-flop logic gate configurations
Require periodic charge refreshing to maintain data storage	Hold its data as long as power is supplied to it
Smaller, cheaper but slower	Faster but bigger and more costly
Often used for main memory	Often used for cache

Memory Organization

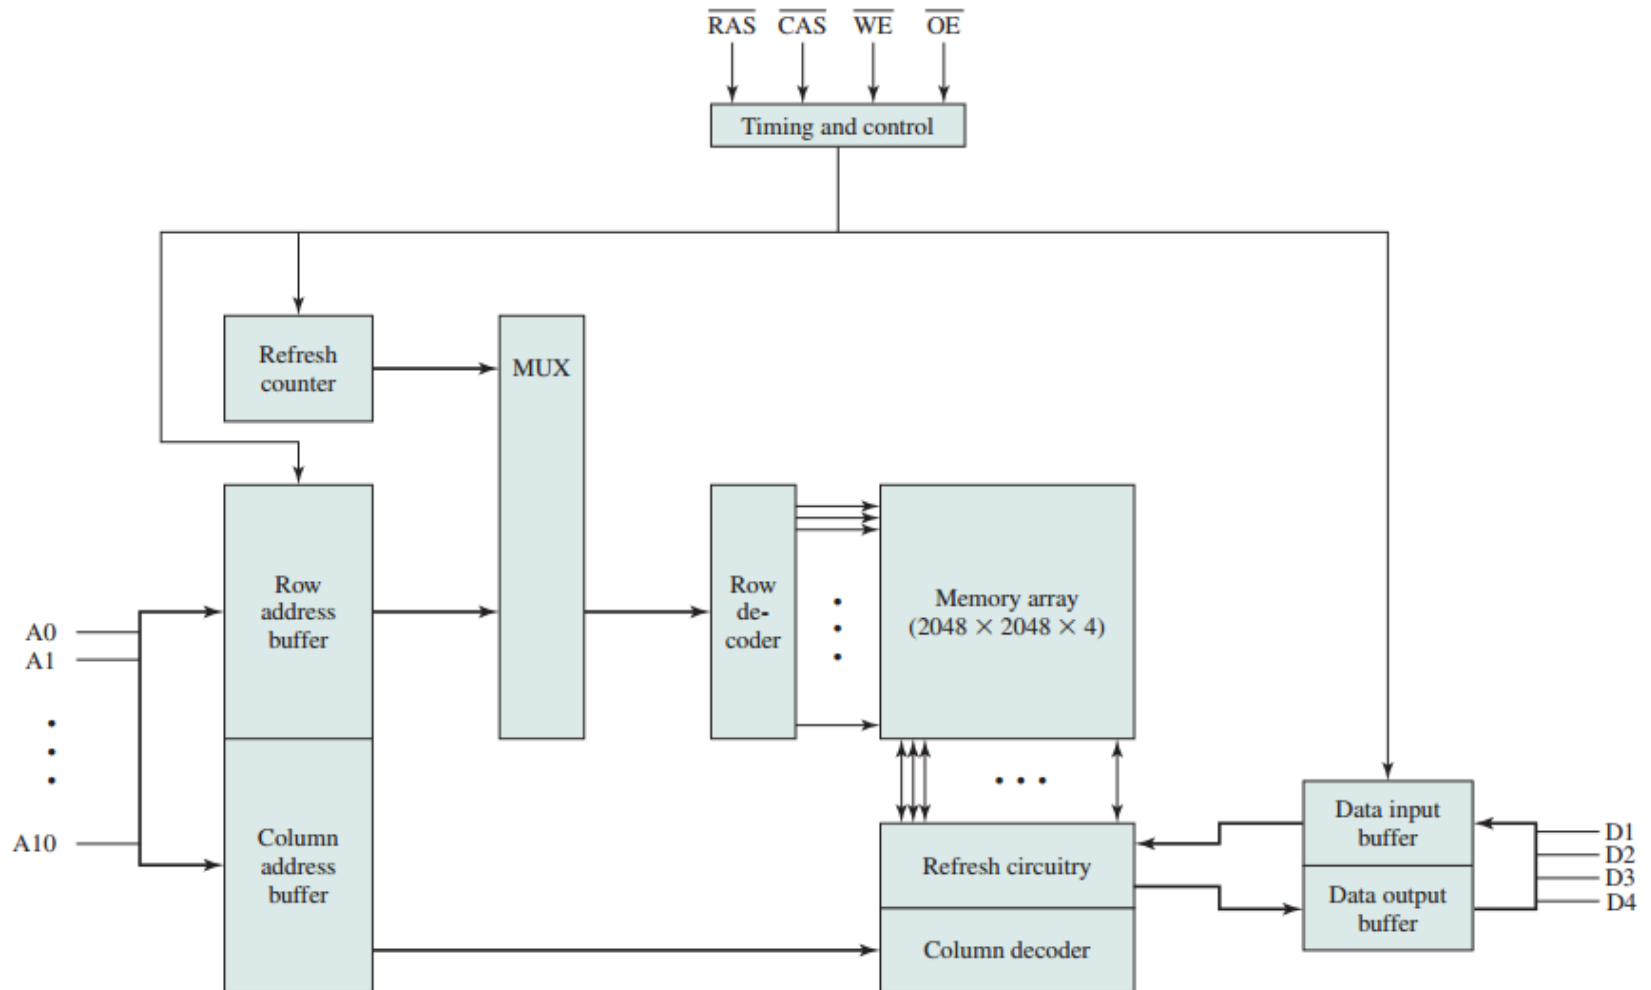


Figure 5.3 Typical 16 Megabit DRAM (4M x 4)

π Memory Organization

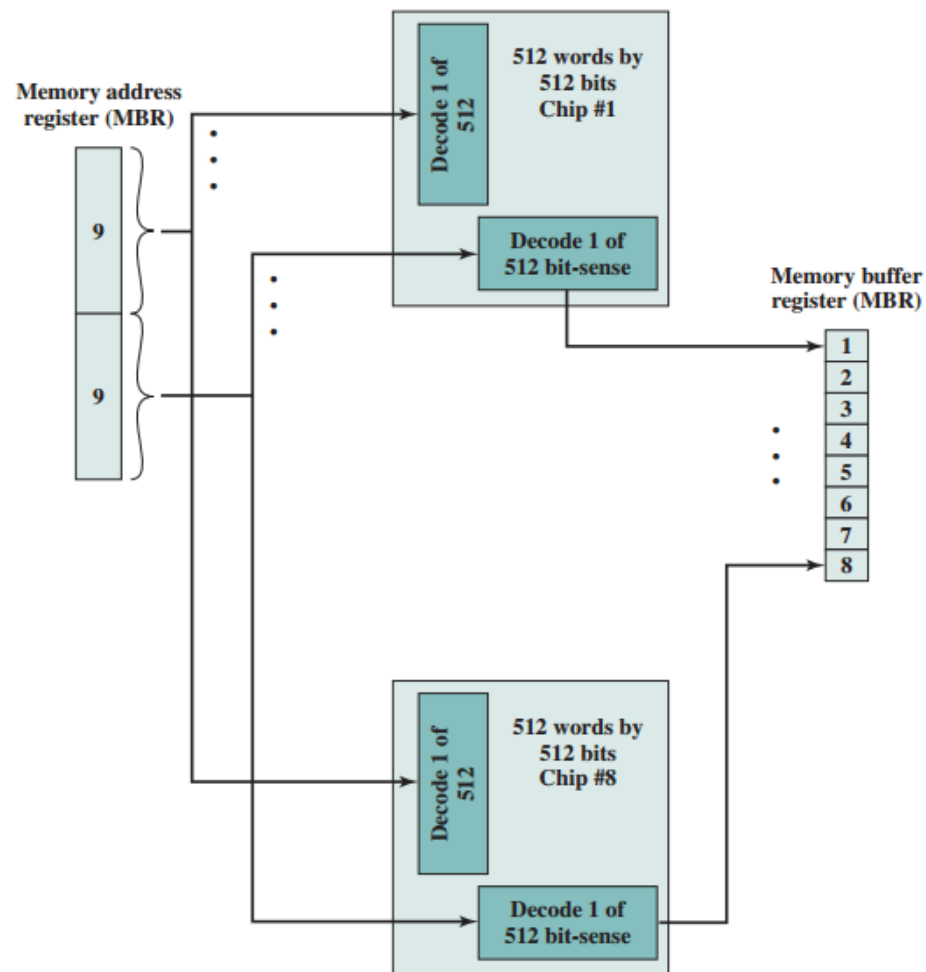


Figure 5.5 256-KByte Memory Organization

Advanced DRAM Organization

- › SDRAM (Synchronous DRAM): exchanges data with the processor synchronized to an external clock signal and running at the full speed of the processor/memory bus without imposing wait states.
- › RDRAM (Rambus DRAM): delivers address and control information using an asynchronous block-oriented protocol.
- › DDR DRAM (Double Data Rate DRAM): can send data twice per clock cycle, once on the rising edge of the clock pulse and once on the falling edge.

π Contents

- › Computer memory system
- › Cache memory
- › Internal memory
- › External memory
 - Cấu trúc đĩa cứng
 - Đặc tính kỹ thuật (tốc độ truyền, quay đọc, công nghệ sản xuất,...)
 - Các cơ chế sao lưu RAID 0 đến RAID 6
 - Các công cụ, tiện ích đọc thông tin phần cứng, kiểm tra,...

Magnetic Disk

- › A disk is a circular platter constructed of nonmagnetic material, called the substrate, coated with a magnetizable material
 - Traditionally the substrate has been an aluminium or aluminium alloy material
 - Recently glass substrates have been introduced
- › Benefits of the glass substrate:
 - Improvement in the uniformity of the magnetic film surface to increase disk reliability
 - A significant reduction in overall surface defects to help reduce read-write errors
 - Ability to support lower fly heights
 - Better stiffness to reduce disk dynamics
 - Greater ability to withstand(anti) shock and damage

Magnetic Read and Write Mechanisms

- › Data are recorded on and later retrieved from the disk via a conducting coil named the head
 - On many systems, there are two heads, a read head and a write head.
 - During a read or write operation, the head is stationary while the platter rotates beneath it.
- › Write mechanism: electricity flowing through a coil produces a magnetic field.
 - Electric pulses are sent to the write head, and the resulting magnetic patterns are recorded on the surface below, with different patterns for positive and negative currents.
- › Read mechanism: a magnetic field moving relative to a coil produces an electrical current in the coil.
 - When the surface of the disk passes under the head, it generates a current of the same polarity as the one already recorded.

π Data Organization and Formatting

- › Data on the platter is organized in a concentric set of rings, called tracks.
 - Each track is the same width as the head.
 - There are thousands of tracks per surface.
- › Adjacent tracks are separated by inter-track gaps which prevent or minimize errors
- › Data are transferred to and from the disk in sectors.
 - There are hundreds of sectors per track, either fixed or variable length.
 - Adjacent sectors are separated by intra-track (inter-sector) gaps.




Figure 6.2 Disk Data Layout

π Physical Characteristics

Table 6.1 Physical Characteristics of Disk Systems

Head Motion	Platters
Fixed head (one per track)	Single platter
Movable head (one per surface)	Multiple platter
Disk Portability	Head Mechanism
Nonremovable disk	Contact (floppy)
Removable disk	Fixed gap
	Aerodynamic gap (Winchester)
Sides	
Single sided	
Double sided	



Disk Performance Parameters

- › Seek time: time it takes to position the head at the track on a movable-head system
- › Rotational delay (rotational latency): time it takes for the beginning of the sector to reach the head
- › Access time : the time it takes to get into position to read or write; sum of the seek time and the rotational delay
- › Transfer time: time required for the data transfer portion of the operation

π RAID

- › Redundant Array of Independent Disks/ Redundant Array of Inexpensive Disks
- › Consist of 7 levels
- › Levels do not imply a hierarchical relationship but designate different design architectures that share three common characteristics:
 - Set of physical disk drives viewed by the operating system as a single logical drive
 - Data are distributed across the physical drives of an array in a scheme known as striping
 - Redundant disk capacity is used to store parity information, which guarantees data recoverability in case of a disk failure


π RAID Levels

Table 6.3 RAID Levels


Category	Level	Description	Disks Required	Data Availability	Large I/O Data Transfer Capacity	Small I/O Request Rate
Striping	0	Nonredundant	N	Lower than single disk	Very high	Very high for both read and write
Mirroring	1	Mirrored	$2N$	Higher than RAID 2, 3, 4, or 5; lower than RAID 6	Higher than single disk for read; similar to single disk for write	Up to twice that of a single disk for read; similar to single disk for write
Parallel access	2	Redundant via Hamming code	$N + m$	Much higher than single disk; comparable to RAID 3, 4, or 5	Highest of all listed alternatives	Approximately twice that of a single disk
	3	Bit-interleaved parity	$N + 1$	Much higher than single disk; comparable to RAID 2, 4, or 5	Highest of all listed alternatives	Approximately twice that of a single disk
Independent access	4	Block-interleaved parity	$N + 1$	Much higher than single disk; comparable to RAID 2, 3, or 5	Similar to RAID 0 for read; significantly lower than single disk for write	Similar to RAID 0 for read; significantly lower than single disk for write
	5	Block-interleaved distributed parity	$N + 1$	Much higher than single disk; comparable to RAID 2, 3, or 4	Similar to RAID 0 for read; lower than single disk for write	Similar to RAID 0 for read; generally lower than single disk for write
	6	Block-interleaved dual distributed parity	$N + 2$	Highest of all listed alternatives	Similar to RAID 0 for read; lower than RAID 5 for write	Similar to RAID 0 for read; significantly lower than RAID 5 for write

Note: N = number of data disks; m proportional to $\log N$


π RAID Levels



(a) RAID 0 (non-redundant)




(b) RAID 1 (mirrored)




(c) RAID 2 (redundancy through Hamming code)

Figure 6.8 RAID Levels (page 1 of 2)


π RAID Levels




(d) RAID 3 (bit-interleaved parity)



(e) RAID 4 (block-level parity)



(f) RAID 5 (block-level distributed parity)



(g) RAID 6 (dual redundancy)

Figure 6.8 RAID Levels (page 2 of 2)



Solid State Drives (SSD)

- › A SSD is a memory device made with solid state components that can be used as a replacement to a hard disk drive (HDD)
- › Flash-memory-based SSD
 - › NOR flash memory
 - Reads and writes in bytes
 - Used to store cell phone operating system code and on Windows computers for the BIOS program that runs at start-up
 - › NAND flash memory
 - Reads and writes in small blocks
 - Used in USB flash drives, memory cards, and in SSDs

π SSD compared to HDD

- › SSDs have the following advantages over HDDs:
 - High-performance input/output operations per second (IOPS): Significantly increases performance I/O subsystems.
 - Durability: Less susceptible to physical shock and vibration.
 - Longer lifespan: SSDs are not susceptible to mechanical wear.
 - Lower power consumption: SSDs use as little as 2.1 watts of power per drive, considerably less than comparable-size HDDs.
 - Quieter and cooler running capabilities: Less floor space required, lower energy costs, and a greener enterprise.
 - Lower access times and latency rates: Over 10 times faster than the spinning disks in an HDD.
- › Currently, HDDs enjoy a cost per bit advantage and a capacity advantage, but these differences are shrinking.

SSD Organization

- › Interface to host system
- › Controller: Provides SSD device level interfacing and firmware execution.
- › Addressing: Logic that performs the selection function across the flash memory components.
- › Data buffer/cache: High speed RAM memory components used for speed matching and to increased data throughput.
- › Error correction: Logic for error detection and correction.
- › Flash memory components: Individual NAND flash chips.

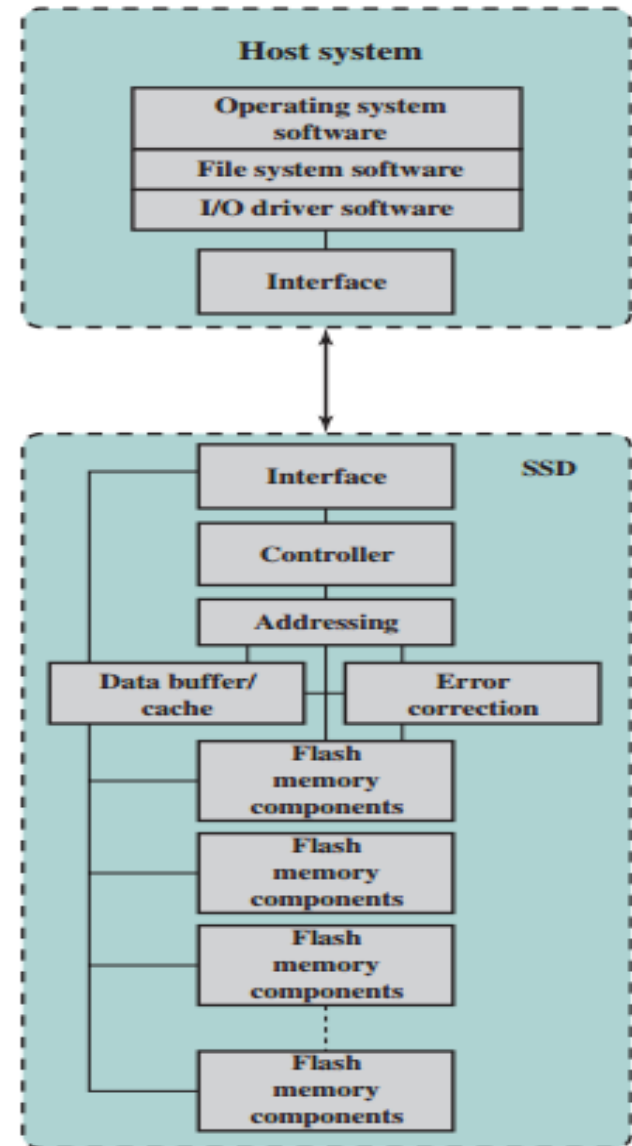


Figure 6.11 Solid State Drive Architecture

Practical Issues

- › SDD performance has a tendency to slow down as the device is used
 - Files are stored on disk as a set of pages, typically 4 KB in length; however, flash memory is accessed in blocks, with a typically block size of 512 KB
 - The entire block must be read from the flash memory and placed in a RAM buffer. Then the appropriate page in the RAM buffer is updated.
 - Before the block can be written to flash memory, the entire block of flash memory must be erased.
 - The entire block from the buffer is now written back to the flash memory.
- › Flash memory becomes unusable after a certain number of writes, typical limit is 100,000 write
- › Most flash devices estimate their own remaining lifetimes so systems can anticipate failure and take preemptive action

π Optical memory

- › CD (Compact Disk): A nonerasable disk that stores digitized audio information. The standard system uses 12-cm disks and can record more than 60 minutes of uninterrupted playing time.
- › CD-ROM (Compact Disk Read-Only Memory): A nonerasable disk used for storing computer data. The standard system uses 12-cm disks and can hold more than 650 Mbytes.
- › CD-R (CD Recordable): Similar to a CD-ROM. The user can write to the disk only once.
- › CD-RW (CD Rewritable): Similar to a CD-ROM. The user can erase and rewrite to the disk multiple times.

π Optical memory

- › DVD (Digital Versatile Disk): A technology for producing digitized, compressed representation of video information, as well as large volumes of other digital data. Both 8 and 12 cm diameters are used, with a double-sided capacity of up to 17 Gbytes. The basic DVD is read-only (DVD-ROM).
- › DVD-R (DVD Recordable): Similar to a DVD-ROM. The user can write to the disk only once. Only one-sided disks can be used.
- › DVD-RW (DVD Rewritable): Similar to a DVD-ROM. The user can erase and rewrite to the disk multiple times. Only one-sided disks can be used.
- › Blu-ray DVD (High-definition video disk.): Provides considerably greater data storage density than DVD, using a 405-nm (blue-violet) laser. A single layer on a single side can store 25 Gbytes.