Internal and External Memory

Chapters 5 and 6

Joannah Nanjekye

July 29, 2024

Semiconductor Main Memory

The memory cell is the basic element of semiconductor memory

Properties of semiconductor memory cells:

- Exhibit two stable or semistable states (0 and 1)
- Capable of being written into, to **set** state
- Capable of being read to sense state

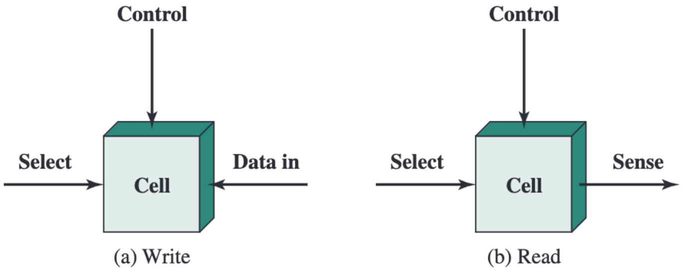

Types of Semiconductor Memory

Characteristics of RAM:

- Possible to read/write easily and rapidly
- Volatile
- Exists in two forms, static or dynamic

Memory Type	Category	Erasure	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)	Read-mostly memory	UV light, chip-level		
Electrically Erasable PROM (EEPROM)		Electrically, byte-level		
Flash memory		Electrically, block-level		

Dynamic RAM (DRAM)

- Made with cells
- Stores data as charge on capacitors
- Dynamic: because the stored charge can leak even when powered, hence requires charge refreshing
- It is an analog device, charge value (0 or 1) depends on a threshhold of the charge
- It is used in main memory and is slower
- Less expensive and simpler

DRAM Structure

- The address line is activated on read/write
- The transistor acts as a switch
- **Write:**
  - A voltage signal is applied to the bit line
  - A signal is applied to the address line
  - Allowing charge flow to the capacitor
- **Read:**
  - Address line is selected (turns the transistor on)
  - Charge from the capacitor is fed out onto a bit line
  - Then to a sense amplifier

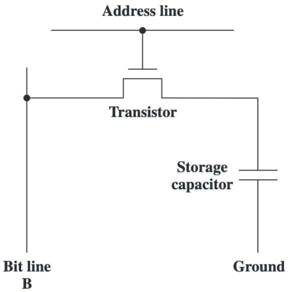

Static RAM (SRAM)

- It uses the same logic elements as a processor
- Binary values are stored using flip-flops (digital circuit)
- If powered, it will always hold data
- It is used for cache memory and is faster

SRAM Structure

- Transistors are arranged to produce a stable logic state
- **Logic State 1:**
  - C_1 high, C_2 low
  - T_1 and T_4 off, T_2 and T_3 on
- **Logic State 0:**
  - C_2 high, C_1 low
  - T_2 and T_3 off, T_1 and T_4 on
- **Write:**
  - Desired bit value is applied to line B
  - Its complement is applied to line $\overline{B}$
- **Read:**
  - A bit value is read from line B

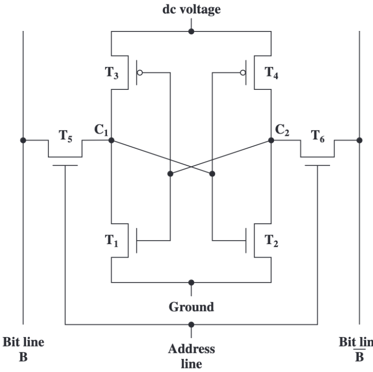

Read-only memory (ROM)

- Non-volatile
- Data is written during manufacture
- Not possible to write new raw data
- No room for error and large fixed cost
- Applications:
  - Micro-programming
  - Library subroutines
  - System programs
  - Function tables

Types of ROM

- **Programmable ROM (PROM)**
  - Written at most once
  - Can be written at a later time after fabrication
  - Requires special writing equipment
- **Read-mostly Memory**
  - **Erasable programmable read-only memory (EPROM)**
    - Must be erased before write to initial state
    - Erased by exposure to ultraviolet radiation
  - **Electrically erasable programmable read-only memory (EEPROM)**
    - Can be written into without erasing content
    - Takes longer to write than read
  - **Flash Memory**
    - Can be erased faster and supports block erasure
    - Intermediate in both cost and functionality

Memory Chip Organization

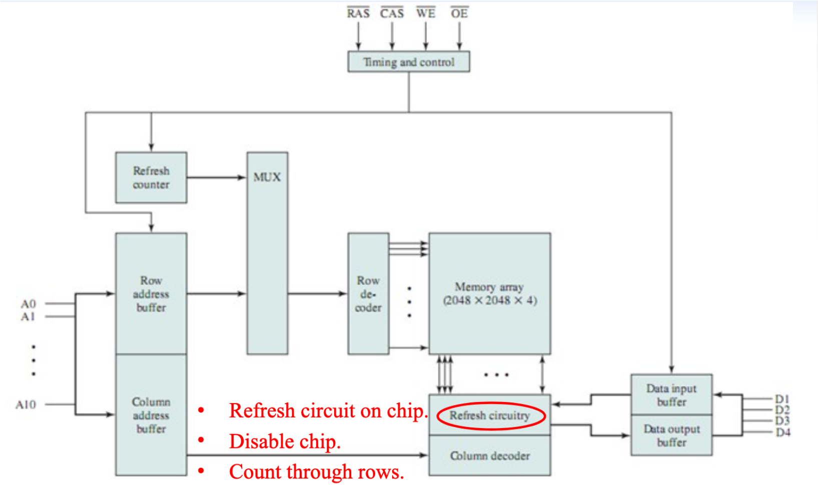

- Refresh circuit on chip.
- Disable chip.
- Count through rows.
- Read & Write back.

Chip Packaging

- 8 Mbit chip, organized as $1M \times 8$
- The address of the word being accessed is achieved in pins $(A_0 - A_{19})$
- The data to be read out, consisting of 8 lines, $D_0 - D_7$

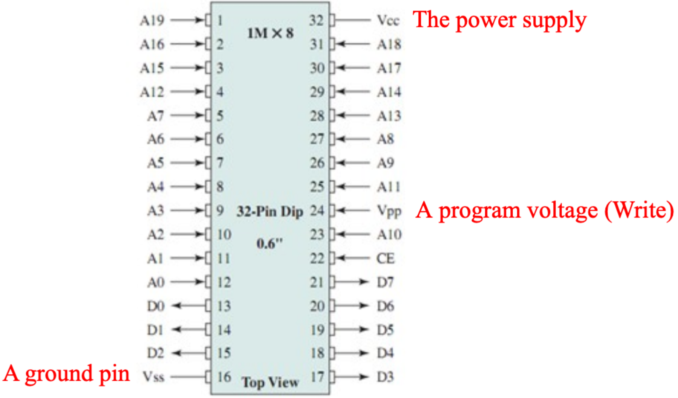

(a) 8-Mbit EPROM

Module Organization

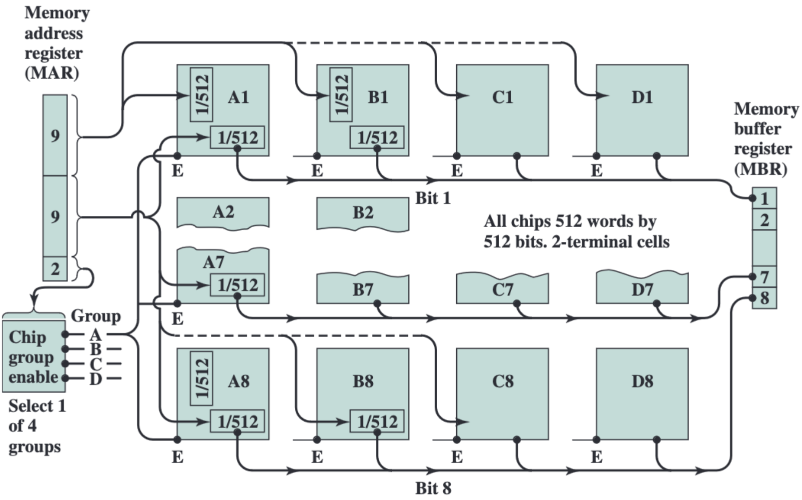

Error Detection

- **Hard Failure:**
  - Permanent physical defect
- **Soft Failure**
  - Random and non destructive
  - No permanent damage

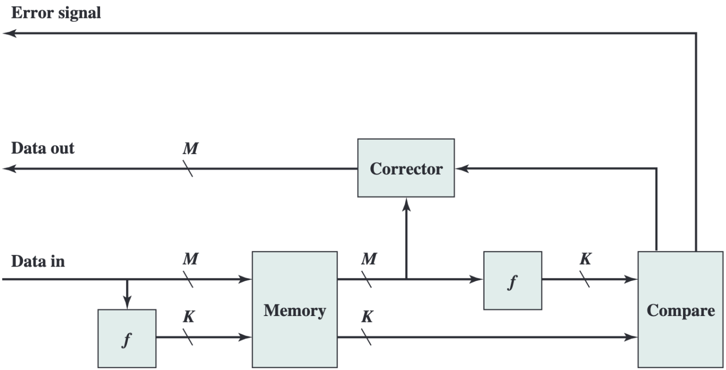

Advanced DRAM Architecture

Enhanced DRAM and Cache DRAM

- Synchronous DRAM (SDRAM)
  - Under control of the system clock
  - Processors issues instructions and DRAM responds after some clock cycles
  - Processor can safely do other tasks while the SDRAM is processing the request
- DDR
  - Has higher data rates
  - Data transfer synchronized to both the rising and falling edge of the clock
  - Uses higher data rate on the bus
  - A buffering scheme is used

External Memory

- Magnetic Disk
- Redundant Array of Independent Disks (RAID)
- Solid State Disk
- Optical Memory
- Magnetic Tapes

Magnetic Disk

Circular platter constructed with a substrate, coated with magnetic material

- Data recorded/retrieved via a conducting coil named the head
- During read/write, head is stationery, platter rotates
- **Write Mechanism:**
  - A magnetic field is produced by current through the coil
  - Pulses sent to head
- **Read (traditional) Mechanism:**
  - Magnetic field moving relative coil produces current
  - The coil is for both read and write
- **Read (contemporary) Mechanism:**
  - Requires a separate read head, close to the write head
  - Allows high frequency operation

Data organization and Formatting

- Data is organized on the platter in a concentric set of rings, called tracks
- Tracks divided into sectors
- Scan rate controlled by constant angular velocity (CAV)

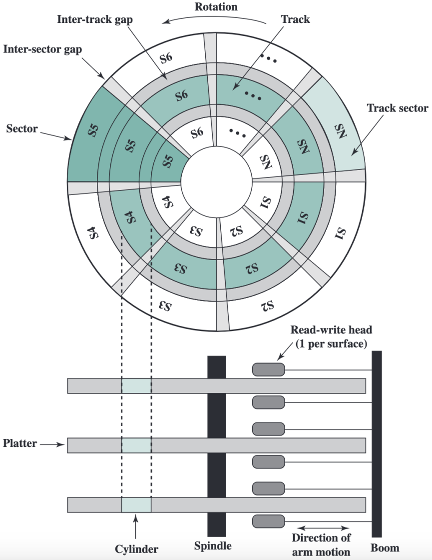

Physical Characteristics

Head Motion

- Fixed head (one per track)
- Movable head (one per surface)

Disk Portability

- Nonremovable disk
- Removable disk

Sides

- Single sided
- Double sided

Platters

- Single platter
- Multiple platter

Head Mechanism

- Contact (floppy)
- Fixed gap
- Aerodynamic gap (Winchester)

Disk Performance Parameters

- **Seek time:** The time it takes to position the head at the track
- **Rotation Delay:** The 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, seek time + rotational delay/latency
- **Transfer Time:** The time required for data transfer

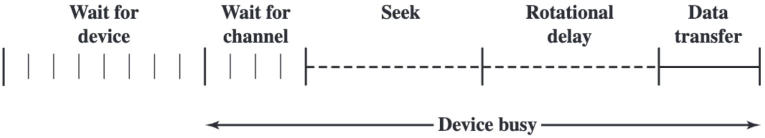

Disk Performance Parameters

The transfer time to or from the disk depends on the rotation speed of the disk in the following fashion:

$$T = \frac{b}{rN}$$

T = transfer time
b = number of bytes to be transferred
N = number of bytes on a track
r = rotation speed, in revolutions per second

The total average read or write time T_{total} can be expressed as:

$$T_a = T_s + \frac{1}{2r} + \frac{b}{rN}$$

T_s is the average seek time.

RAID

A set of physical disk drives viewed as a single logical drive by the operating system

- Redundant Array of Independent Disks
- The scheme consists of seven levels, 0 - 6
- A scheme known as **striping** distributes the data across the physical drives
- Redundant disk capacity is used to store parity information
- Parity information ensures recoverability in case of failure

RAID 0

- No redundancy
- Data is striped across the disks
- The striping is round robin
- How to improve performance:
  - High transfer capacity (transfer rate)
  - Application must must make I/O requests that efficiently use the disk
  - Effective load balancing (response time)

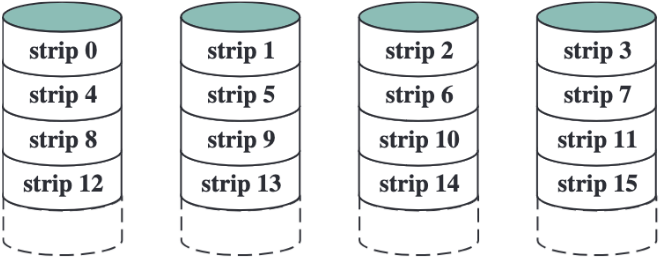

RAID 1

- Redundancy is by the expedient of duplicating all the data
- Each logical strip is mapped to two separate physical disks
- A request is serviced by either disk
- Read either, write to both
- Recovery is simple
- Expensive (twice the disk space)

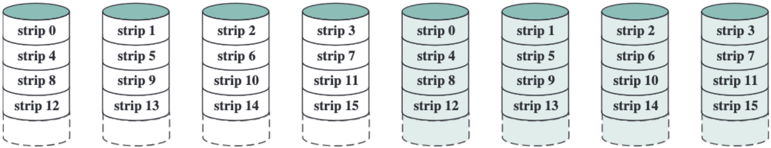

RAID 2

- Use the parallel access technique[1]
- All member disks participate in the execution of every I/O request
- Error correction calculated across respective bits on disks
- Effective when disk errors occur
- Hamming code is used
- RAID 2 is overkill and is not implemented

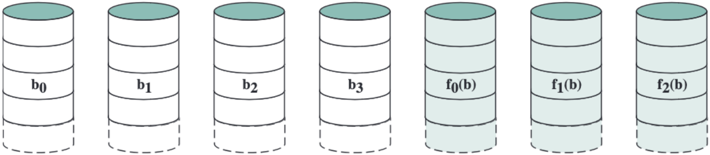

[1] The book says RAID 2 uses *log*$_n$, n data disks but the example is inconsistent

RAID 3

- Similar to RAID 2
- But requires one redundant disks, not log_n, n data disks
- A simple parity bit is computed for the set of individual bits in the same position
- Recovery is by reconstructing surviving data and parity information
- All data is available as *reduced mode* using **XOR** calculation
- High Transfer rates

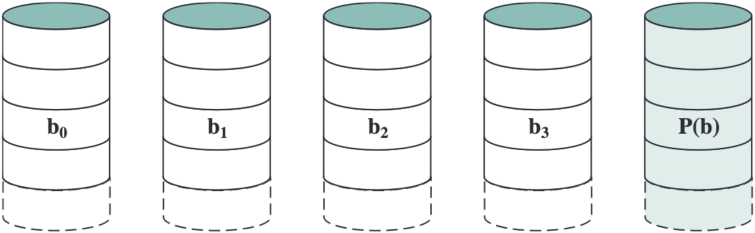

RAID 4

- Independent access technique
- Stripes are large
- a bit-by-bit parity strip is calculate for disks
- Parity stored on parity disk
- Small writes are expensive due to updates on data and parity bits

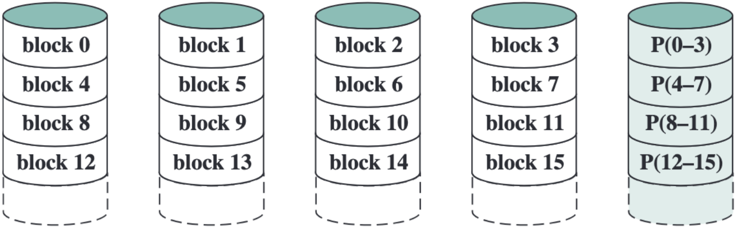

RAID 5

- Same as RAID 4
- Distributes the parity strips across all disks
- Round Robin allocation for parity stripe
- Avoids RAID 4 bottleneck at parity disk
- Used in network servers

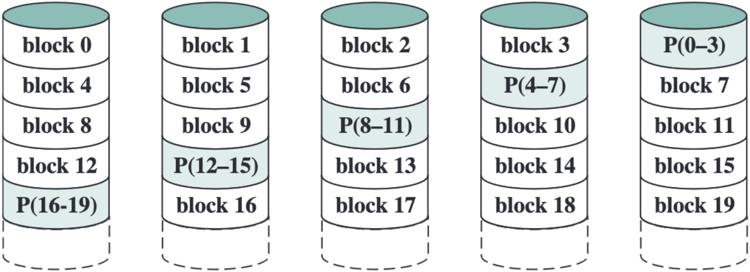

RAID 6

- Two parity calculations
- Stored in separate blocks on different disks
- N disks require N+2 disks
- High Data availability:
  - Three disks have to fail for data loss
  - Significant write penalty

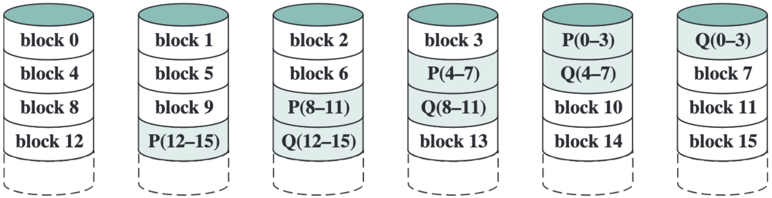

Summary

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

Solid State Drives (SSD)

A memory device made with solid state components that can be used as a replacement to hard disk drive

- Uses a type of semiconductor referred to as flash memory
- Types of flash memory
  - NOR: the basis unit of access is a bit
  - NAND: the basic unit of access is 16 or 32 bits

Advantages of SSD over HDD

- Durability
- Longer lifespan
- Lower power consumption
- High-performance input/output operations per second (IOPS)
- Quieter and cooler running capabilities
- Lower access times and latency rates

Magnetic Tape

- Serial access as in a magnetic disk
- Slow and very cheap
- Backup and archive
- Linear Tape-Open (LTO) Tape Drive:
  - Developed in 1990s
  - Open source alternative to proprietary tape systems

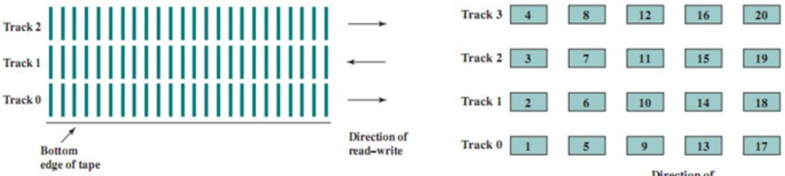

(a) Serpentine reading and writing

LTO Tape Drives Evolution

	LTO-1	LTO-2	LTO-3	LTO-4	LTO-5	LTO-6	LTO-7	LTO-8
Release date	2000	2003	2005	2007	2010	2012	TBA	TBA
Compressed capacity	200 GB	400 GB	800 GB	1600 GB	3.2 TB	8 TB	16 TB	32 TB
Compressed transfer rate	40 MB/s	80 MB/s	160 MB/s	240 MB/s	280 MB/s	400 MB/s	788 MB/s	1.18 GB/s
Linear density (bits/mm)	4880	7398	9638	13,250	15,142	15,143		
Tape tracks	384	512	704	896	1280	2176		
Tape length (m)	609	609	680	820	846	846		
Tape width (cm)	1.27	1.27	1.27	1.27	1.27	1.27		
Write elements	8	8	16	16	16	16		
WORM?	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Encryption Capable?	No	No	No	Yes	Yes	Yes	Yes	Yes
Partitioning?	No	No	No	No	Yes	Yes	Yes	Yes

Home Work

- Read about Optical Memory
- Attempt Problems 6.3, 6.5 and 6.9

Sources Acknowledgement

- https://slideplayer.com/slide/10852334/