11 Feb 2025 - DBMS-II - Week 06

Introduction to Memory Management

→ Hard disks ⇝ obsolete

→ Pen drives

→ RAM

→ Memory technology: what technology is used to store 1 bit of information?

- von neumann architecture

memory ⟷ CPU
read
disk

why RAM/cache?

→ spatial & temporal locality

→ DRAM cell

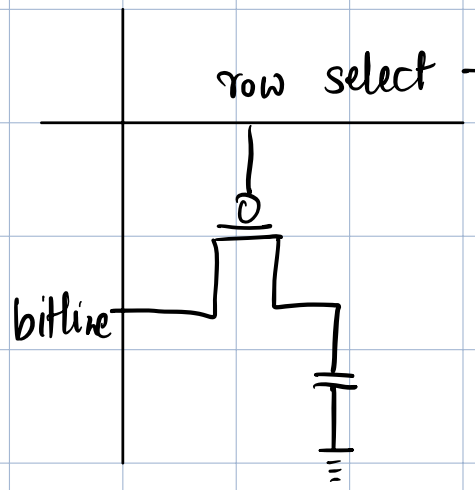

→ switch, tests charge in bitline
→ threshold
→ sense amplifier.

DRAM → store information in the form of charge.

→ Channels → DIMMs
→ North chip

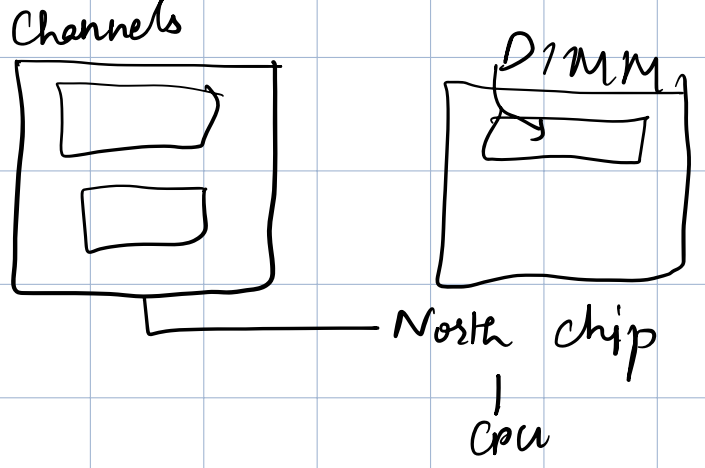

DIMM → → 2 ranks (front and back)

↓

chips

simplified design

↓

usually not used:

each chip has multiple banks

↓

each bank is a 2D array of DRAM cell

→ Row buffer (SRAM)

→ sense amplifier

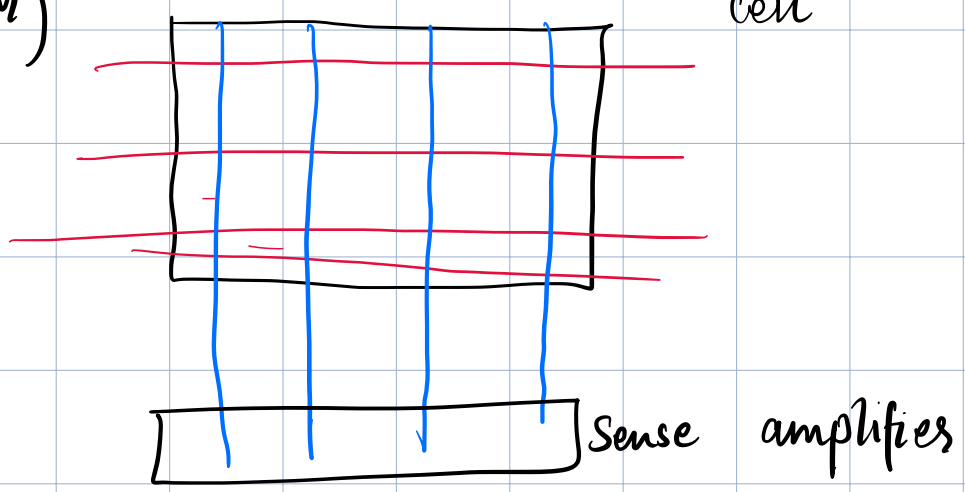

→ DRAM used in memory

→ information is stored in terms of charge.

→ needs periodic refreshment

↳ research in reduction is active

cache

main mem

flash

magnetic disk

optical disk (CD/DVD)

magnetic tapes ⇝ cassettes

Disadvantage of SRAM and DRAM

SRAM

- → density ↓
- → 6 transistors
- → high leakage
- → volatile memory

DRAM

- → destructive read operation
- → requires periodic refresh to restore data
- → High leakage power consumption as compared to other main mem tech.

→ Other mem tech. → being explored.

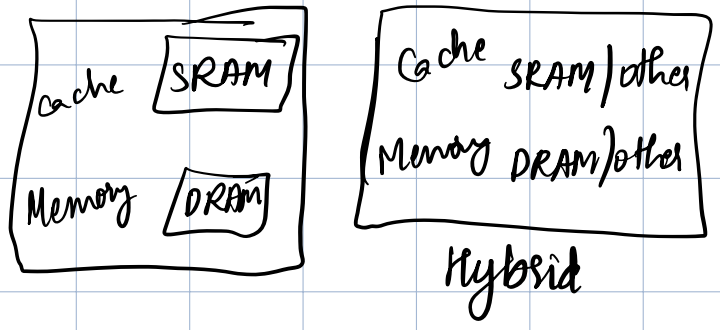

- → eDRAM (embedded) → no need of discharge
- → Magneto Resistive RAM
- → Spin-Transfer Torque RAM
- → Phase-change memory (PCM)

Non-volatile Memory (NVM)

→ lifelong (~100 years)

→ age and endurance of DRAM ↑

→ NVM : writes are expensive.

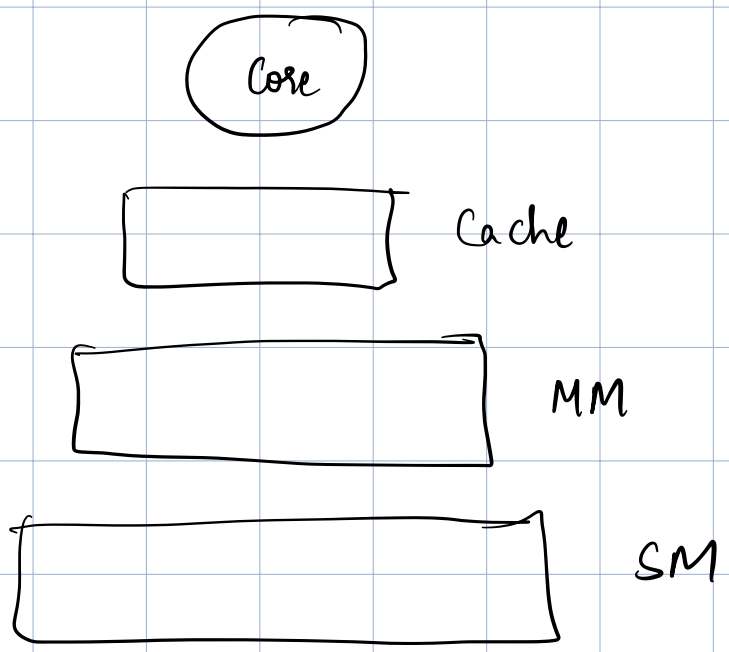

PCM

↓

heat metal

↓

resistance change

→ technologies : how 1 bit is stored

→ CS ~~works~~ on a high level

→ EE

2025-02-15

Storage and File Structure

10.2.3 → read once

→ block, page (not cache block but page)

→ every relation is stored in a file

multiple rows form a block.

each row is called a record

→ secondary memory to main memory

↓

always a page is fetched

→ RAID : redundant array of independent disk

↓

invented during disk times

→ modified versions today

Main requirements for database storage in disk:

① reliability: disk failure, power problems, etc should not affect my data.

→ solved using redundancy / mirroring

→ cost ↑ disk size ↓

200 GB available → effective 100 GB

→ in last chapter, we removed redundancies within the database

but maintaining only 1 copy of redundant-free data is not reliable.

② Performance

→ parallelism : reading multiple stuff.

→ data stripping

100 R

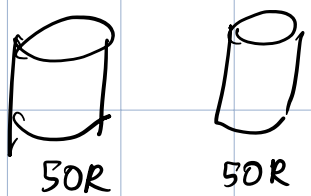

→ transfer rate can be less (both requests on same DB)
but bandwidth ↑

→ record interleaved stripping

Reliability and performance of Storage

Replication

- Mirroring
- Error correcting codes (ECC)
  - Hamming codes
  - Parity bits
    - → cannot detect 2 or even changes
    - → cannot detect where exactly error happened → use Hamming codes } → in networks

Data stripping

- → Bit-interleaved
- → Block-interleaved
  good for writing
  - → to read a block you need to read all disks: not efficient
  - B1 B2 ... P2 ...
  - → parallelism

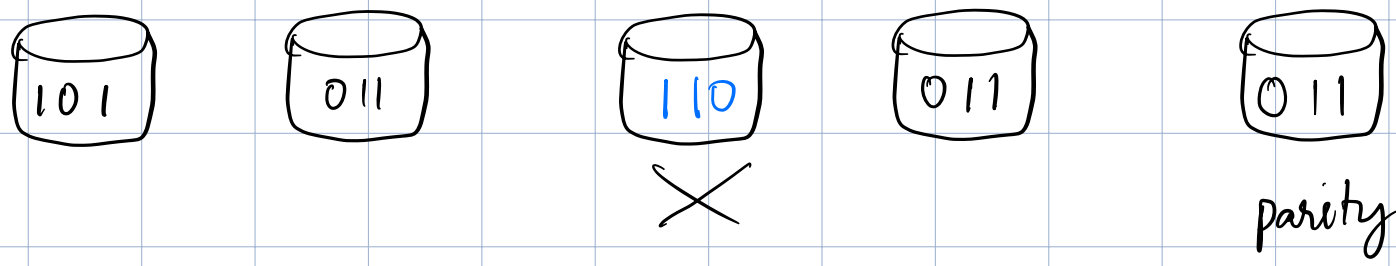

→ Assumption: only one disk fails at a time

→ Difference b/w this and networking:

you know which DB malfunctioned and you can recreate it easily

Main goals of parallelism: ① load-balance small accesses (block accesses) so that throughput of such accesses ↑

② parallelise large accesses so that the response time is ↓

RAID

Level 0 to level 10

Level 0 : non redundant

performance good