



SNS COLLEGE OF ENGINEERING

Kurumbapalayam (Po), Coimbatore – 641 107

An Autonomous Institution

Accredited by NBA – AICTE and Accredited by NAAC – UGC with 'A' Grade
Approved by AICTE, New Delhi & Affiliated to Anna University, Chennai

DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

**COURSE NAME : 23CST203 - DATABASE MANAGEMENT
SYSTEMS**

II YEAR / IV SEMESTER

**Unit 2- Physical Storage and MongoDB
Topic 2 : RAID**

**23CST207 DBMS/ K.KARTHIKEYAN/AP-
CSE/SNSCE**



The Motivation for RAID



- Computing speeds double every 3 years
- Disk speeds can't keep up
- Data needs higher MTBF than any component in system
- IO Performance and Availability Issues!

23CST207 DBMS/ K.KARTHIKEYAN/AP-
CSE/SNSCE



RAID to the Rescue!



- PERFORMANCE
 - ◆ Parallelism
 - ◆ Load Balancing
- AVAILABILITY
 - ◆ Redundancy: Mirroring, or Striping with Parity
- FLEXIBILITY
 - ◆ Selectable Performance/Availability/Cost



What is RAID Technology?



- Redundant Array of Independent Disks (a.k.a. “Disk Array”)
 - ◆ *Multiple* drives, *single* host disk unit
- Provides opportunity to increase:
 - ◆ Performance via Parallelism
 - ◆ Data Availability via Redundancy
- Cheap cost via commodity disks



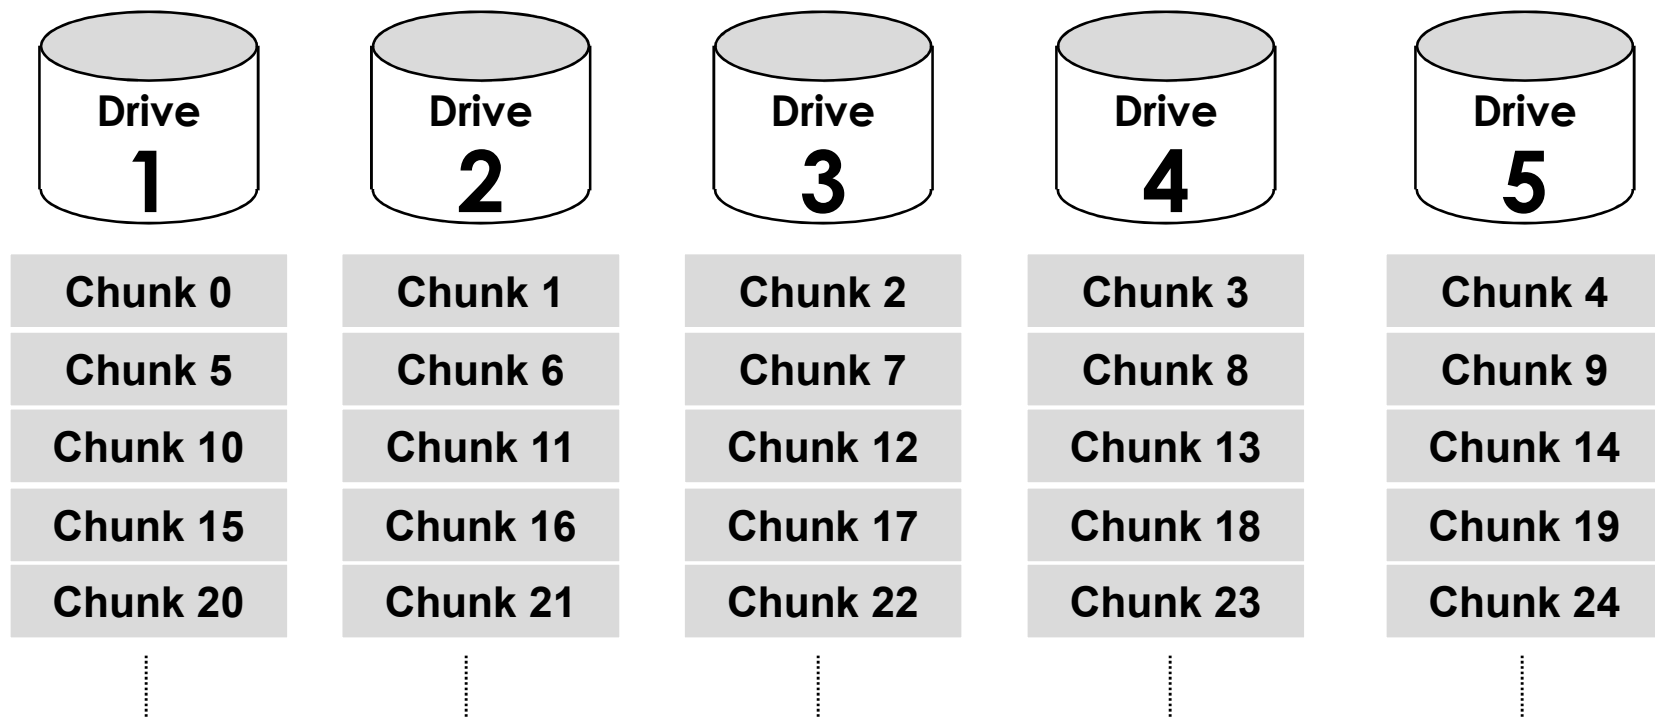
Performance: Disk Striping



- Chunk size tuneable for BW vs Thruput tradeoffs

◆ Large Chunk → High **Throughput** (IO/sec)

◆ Small Chunk → High **Bandwidth** (MB/sec)







Availability: Redundancy



Single Host Unit



Mirroring

Same data written to both disks.



Striping with Parity

Data Stream, OR IO Stream, can be multiplexed across multiple disks, depending on BW vs Thruput. Parity data is also stored on disk.




> Add XOR'd parity for increased availability.



Parity Redundancy

- Parity = XOR of data from every disk in the RAID unit

DATA



- Any single disk's data can be recovered by XOR'ing the data of the surviving disks.



RAID Levels



- Many to choose from
- Each offers unique tradeoffs
 - ◆ Performance
 - ◆ Availability
 - ◆ Costs
- We offer levels 0, 1, 3, 5, 10



RAID 0

Disk Striping with No Redundancy

- High Performance; Low Availability
- Data Striped on Multiple Disks
- Multi-threaded Access

Data Stream, OR IO Stream, can be Striped across multiple disks (BW vs Thruput).

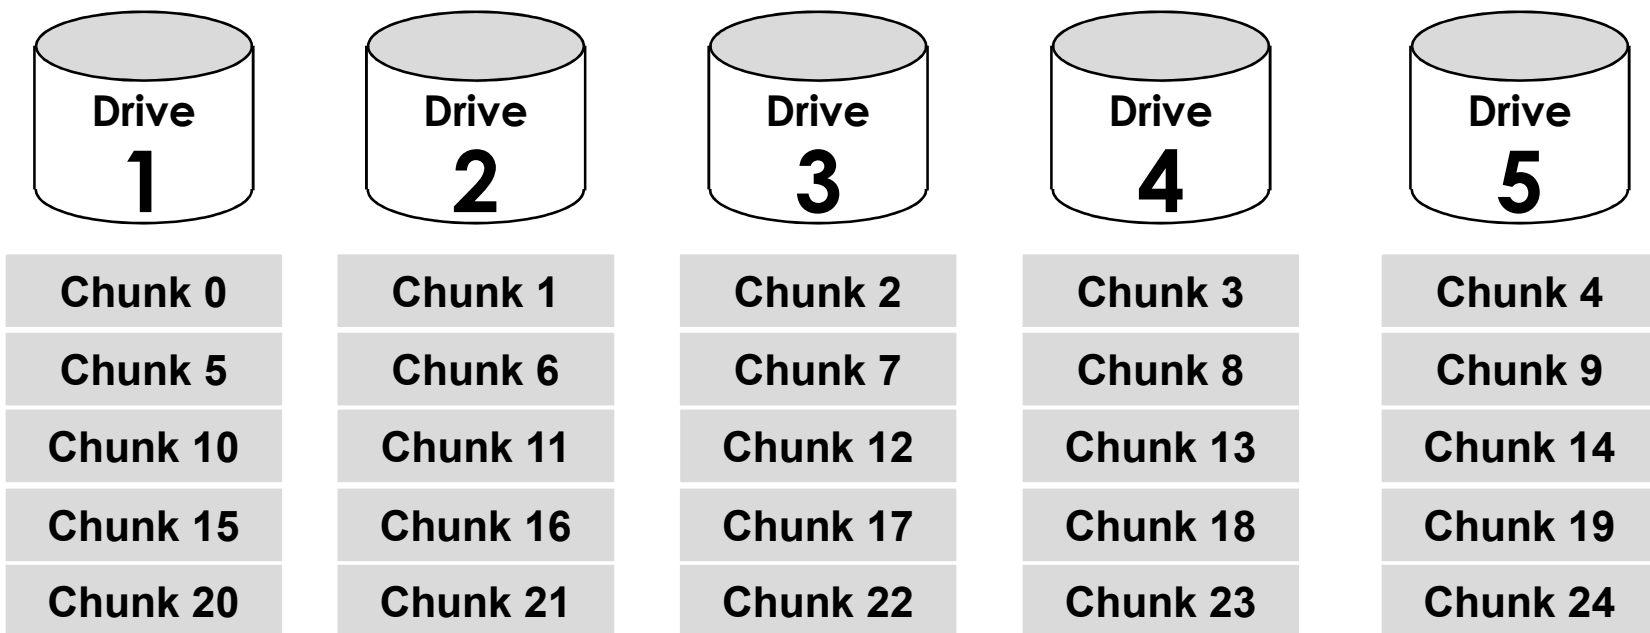


23CST207 DBMS/ K.KARTHIKEYAN/AP-
CSE/SNSCE



RAID 0 Striping

- *Chunk size tuneable for BW or Thruput*
- *No redundancy*






RAID 1


Disk Mirroring

- Single-disk Performance; Expensive Availability
- Data 100% duplicated across both spindles.
- Single-threaded access

SAME data written to BOTH disks -- no segmenting.



Single Host Unit







RAID 1/0

Striped Mirrors

- Highest Performance; Most Expensive Availability
- Multi-threaded Access

Data Stream, OR IO Stream, can be Striped across multiple disks (BW vs Thruput).






RAID 3


Disk Striping with dedicated parity drive

- High BW Performance; Cheap Availability
- Sector-granular data striping
- Single-threaded Access

Data Stream is Striped across N-1 disks for high bandwidth.



XOR Parity Data

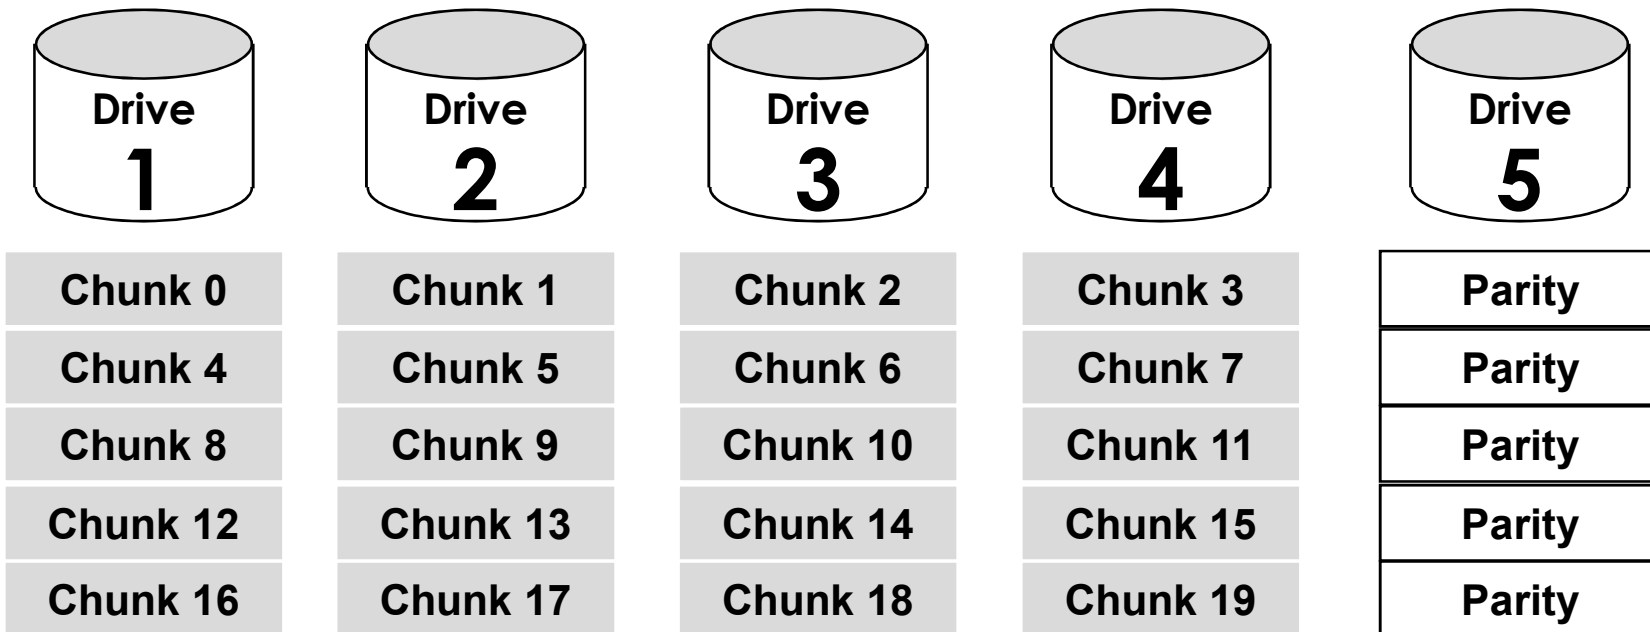


23CST207 DBMS/ K.KARTHIKEYAN/AP-
CSE/SNSCE



RAID 3 Striping

- **Chunk size = single sector (pure RAID 3 would be single byte)**
- **All parity data on same spindle**







RAID 5

Disk Striping with rotating parity drive

- High Read Performance, expensive Write performance; Cheap Availability
- Tuneable Stripe granularity
- Optimized for multi-thread access

IO Stream is Striped across N-1 disks for high IOs per second (thruput).



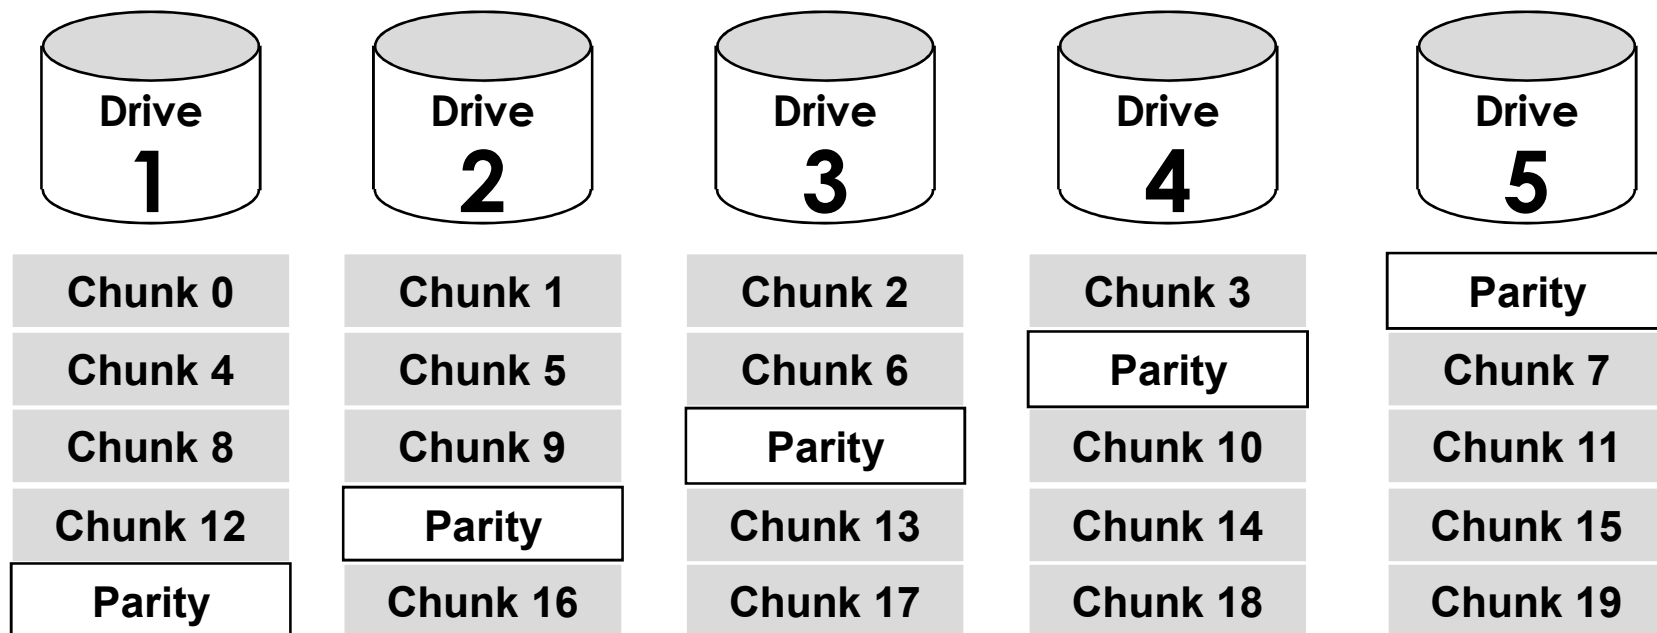
XOR Parity Data

23CST207 DBMS/ K.KARTHIKEYAN/AP-
CSE/SNSCE



RAID 5 Striping

- **Chunk size is tuned such that typical IO aligns on single disk.**
- **Parity rotates amongst disks to avoid write bottleneck**





RAID 5 - Write Operation

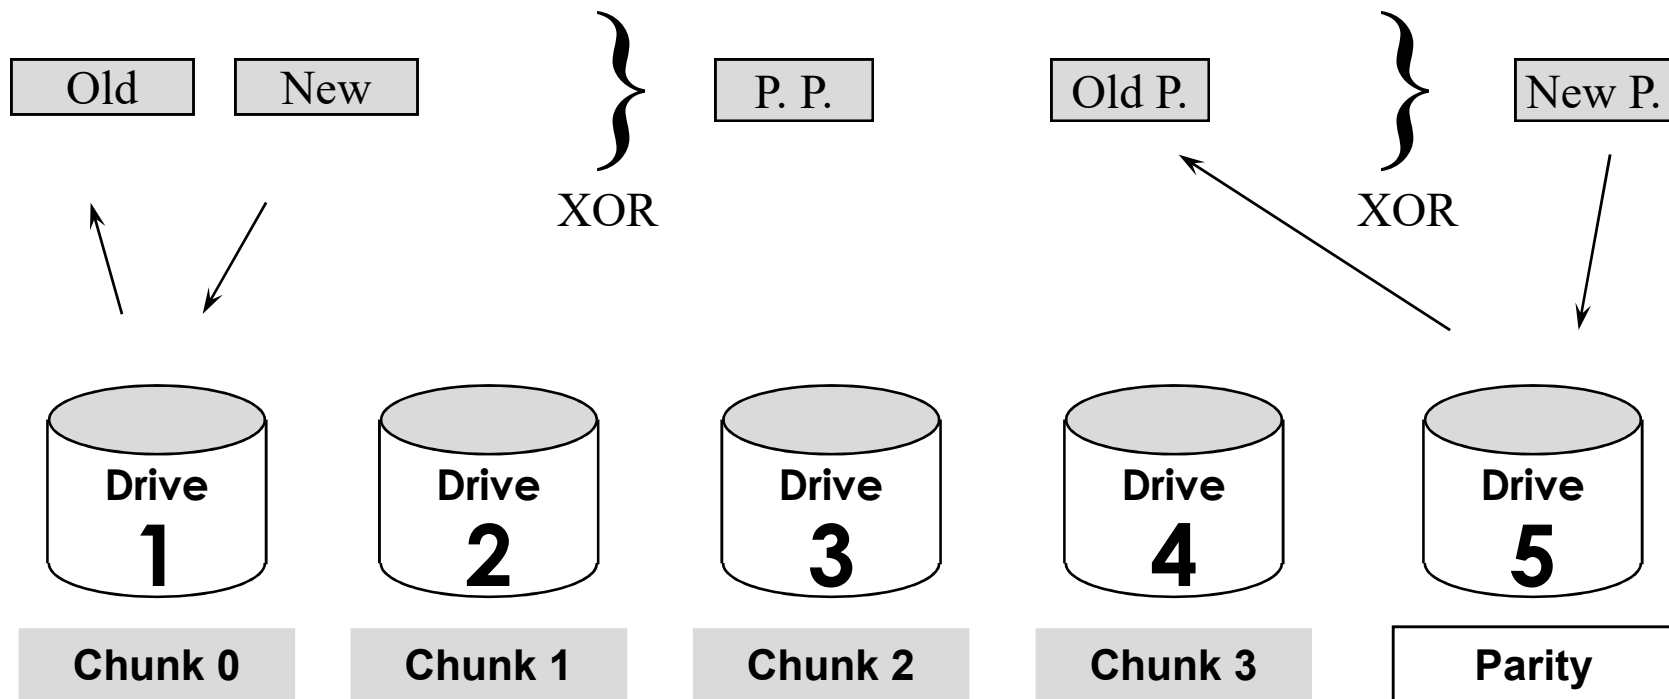
1. Read old data.

2. Write new data

3. XOR old and new data to create "Partial Product".

4. Read old parity data.

5. Xor old parity with partial product, writing out result as new parity.





RAID Level Review



- ◆ RAID 0 - Data striping, Non-redundant.
 - ◆ *High Performance, Low Availability*
- ◆ RAID 1 - Mirroring
 - ◆ *Moderate Performance, Expensive High Availability*
- ◆ RAID 1/0 - Striping and Mirroring
 - ◆ *High Performance, Expensively High Availability*
- ◆ RAID 3 - Striping, single parity disk.
 - ◆ *High **Bandwidth** Performance, Cheap Availability*
- ◆ RAID 5 - Striping, rotating parity disk.
 - ◆ *High **Thruput** Performance, Cheap Availability*

23CST207 DBMS/ K.KARTHIKEYAN/AP-
CSE/SNSCE



Summary



- Increasing performance gap between CPU and IO
- Data availability a priority
- RAID meets the IO challenge:
 - ◆ Performance via parallelism
 - ◆ Data Availability via redundancy
 - ◆ Flexibility via multiple RAID levels, each offer unique performance/availability/cost tradeoffs

23CST207 DBMS/ K.KARTHIKEYAN/AP-
CSE/SNSCE



THANK YOU