CMSC 724: RAID; 5 Minute Rule

Amol Deshpande

University of Maryland, College Park

February 22, 2007

Technology Trends

- “Laws”
  - Moore’s law (transisters/chip doubles every 18 months)
  - Joy’s law (MIPS doubles every year) . . .
  - **Amdahl’s Law:** No point in improving the performance of something that is 1% of the total cost.
- Memory, CPU, disk densities increasing very rapidly. . .
- But raw seek times improve very slowly $\rightarrow$ Random I/Os are not going to keep pace
- Still mostly true. . .

Redundant Array of *Independent* Disks

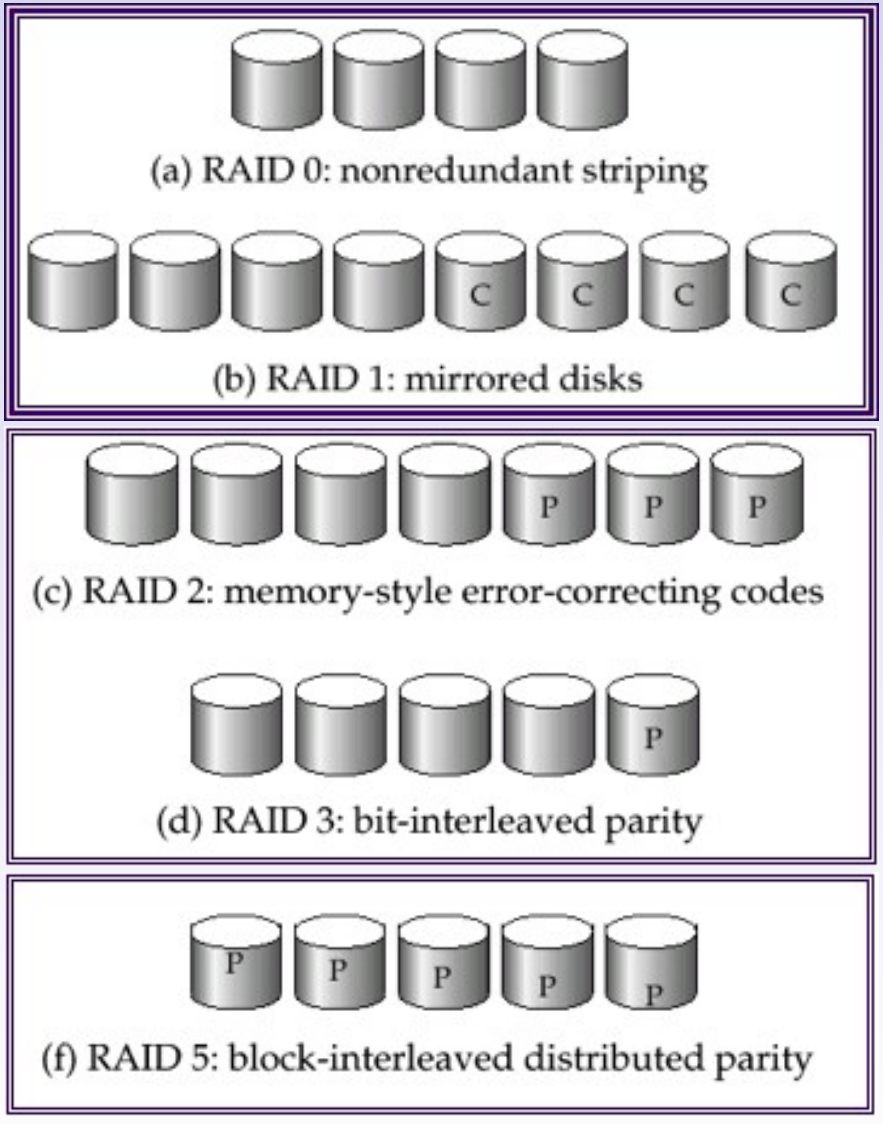

RAID

- RAID 5 $\rightarrow$ High storage efficiency, High write overheads
- RAID 1 $\rightarrow$ Lower storage efficiency, Best performance
- RAID 1 most suitable for Transaction Processing
- HP AutoRAID
  - Monitored the read/write frequencies
  - Switched between RAID 1 and RAID 5

5 Minute Rule

- Born out by observing real systems, technology trends..
- 5-minute rule: If the data is going to be accessed again in 5 minutes, keep it in memory.
  - Purely economic decision.
  - Rent on keeping disk block in memory for 5 minutes $\approx$ Cost of one disk access
  - So "cheaper" to just access it from disk if re-reference in $>$ 5 minutes.
- To be applied when the system is being designed.
  - When deciding how much memory, disk to use.
  - Typically based on a workload.
  - See original 1987 paper for an example.
- Technology trends suggest it will stay in this range.

5 Minute Rule

- Sequential data: Reduces to 1-minute. Why ?
  - Sorts take 5GB/Minute.
  - So $<$ 5GB data, sort in one pass
  - $>$ 5GB data, sort in two passes (can sort upto 100TB)
- Exchanging Memory for CPU ?
  - 5-Byte Rule (Gray, Putzolu, 1987)
  - Spend 5 bytes of main memory to save 1 instruction per second.
- Such rules of thumb are critical during real system design/deployment.