

CS 537 - Handout Oct 29

- From last class:
 - IO time = seek + rotational + transfer
 - Estimating IO time
 - Seq vs random
 - Cache
 - Disk scheduling (SJF - NBF, SSTF, SPTF, Elevator)
- RAID
 - Redundant array of inexpensive disks
 - Big idea: use multiple disk in clever ways
 - Want to increase:
 - Performance (how fast?)
 - Capacity (how much data?)
 - Reliability (how many disks can fail without data loss?)
 - Transparent
 - Fault model: fail-stop
 - RAID-0: striping
 - Chunk size trade-offs
 - RAID-1: mirroring
 - RAID-4: parity
 - RAID-5: rotated parity
 - Main workloads: Seq read, seq write, rand read, rand write
 - Perf Analysis:
 - One logical operation -> how many physical operations?
 - Single request: latency
 - Many request: throughput

Workload/RAID	RAID 0	RAID 1	RAID 4	RAID 5
Seq Read	N.S	N/2. S	(N-1)S	(N-1)S
Seq Write	N.S	N/2. S	(N-1)S	(N-1)S
Rand Read	N.R	N.R	(N-1)R	N.R
Rand Write	N.R	N/2.R	R/2	N/4. R

Problems/Questions:

1. What are the advantages/disadvantages of RAID having the same interface as a block device?
2. Given a workload, how would you go about determining the chunk size for RAID 0?
3. If all you cared about was performance, which RAID level would you use and why?
4. Why is RAID 0 the upper limit on performance?
5. List the performance for all RAID levels for a single read/write seq/rand request. Why is this different from steady-state levels?
6. Why is sequential write performance on RAID 4 $(N-1)S$ and not NS ? Aren't all the disks in use?
7. A question direct from the book: for RAID 4, how many disks would need to be in the system so that the additive method performs fewer I/Os than the subtractive method; what is the cross-over point?
8. For RAID 4, what is the small write problem?
9. For which RAID levels can you increase random write performance by adding more disks?
10. How does RAID 5 solve the small update problem?
11. Why is random read performance for RAID 5 NR instead of $(N-1)R$ like RAID4? Remember we still have to store parity.
12. Random write performance for RAID 5 is $NR/4$ because each logical request turns into four requests. Since this is the case in RAID 4 as well, why isn't the random write performance $R/4$ (instead of $R/2$)?
13. Given RAID-5 is like RAID 4 but better, why would someone ever use RAID-4?
14. If you write to two disks in parallel, the seek time (on average) is higher than the seek time on a single disk. Why?

Check out the RAID simulator online (link at the end of the book chapter).