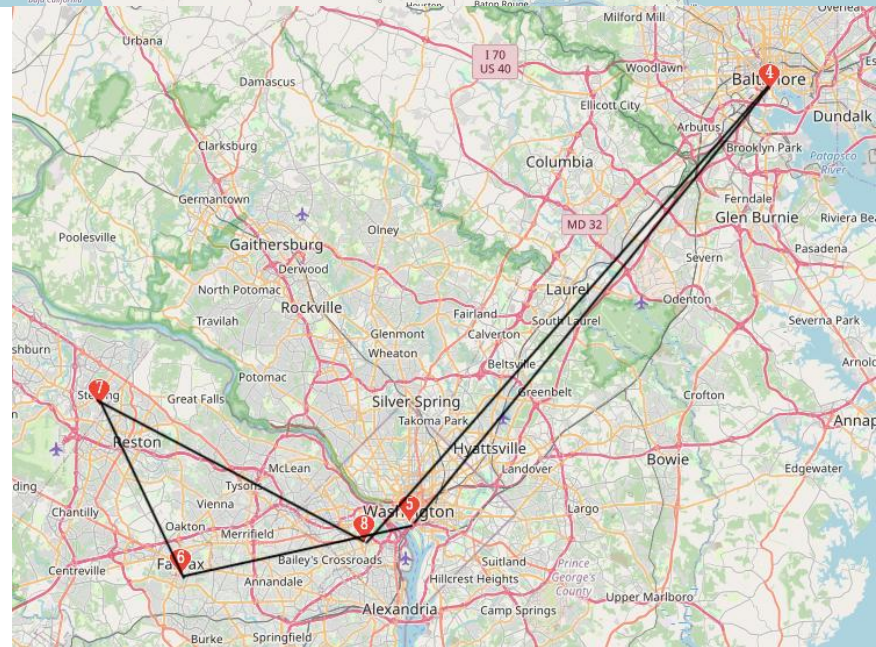
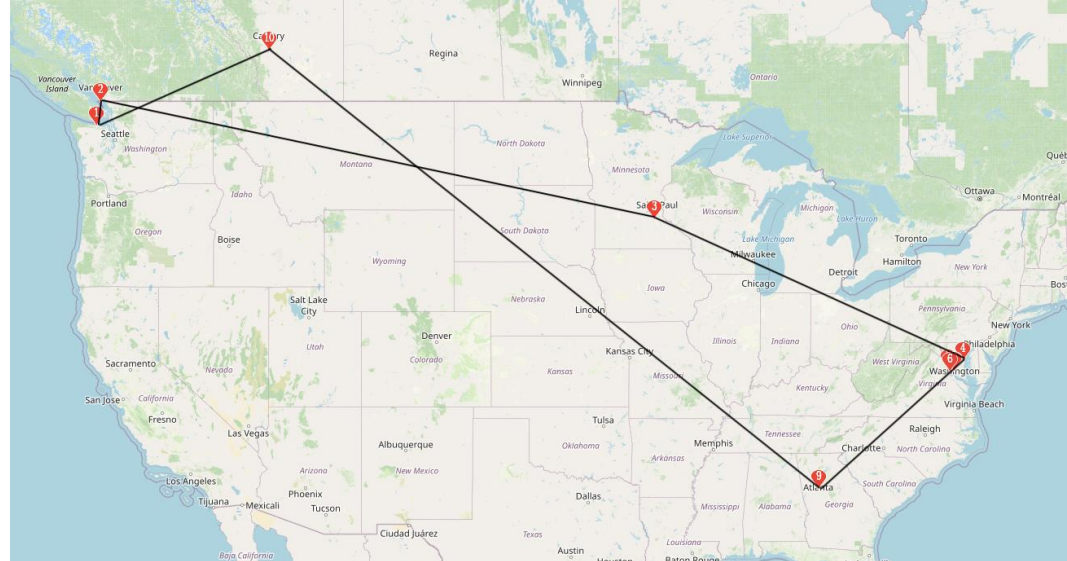





CSc 360 Operating Systems Mass Storage

Wenjun Yang
Fall 2025



CSc 360





Mass storage

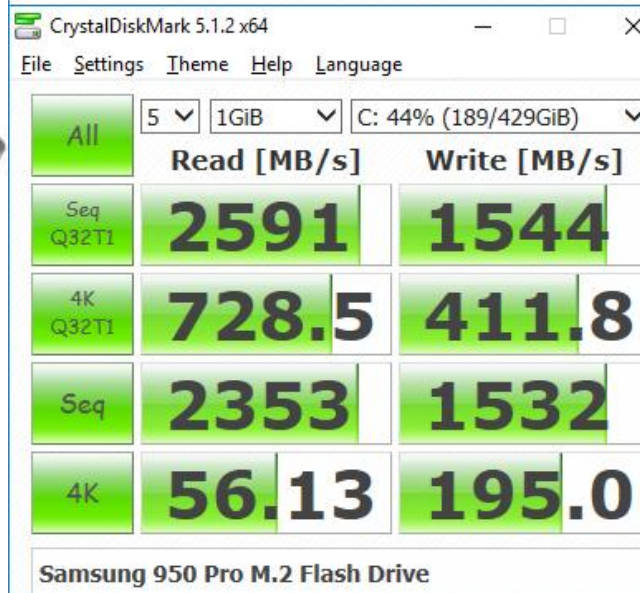




- Magnetic disk
 - performance: data rate, access time
 - interface: ATA, SATA, SCSI, FC, etc
 - address: geometry (cylinder, track, sector)
- Magnetic tape
- Network attached storage (NAS)
- Storage area network (SAN)
 - iSCSI



CSc 360

* 1tb sata2 7200rpm 4.2ms 32mb 3.5in?




Performance Summary:

Hard Drive: 103 MB/sec read, 96 MB/sec write speed
 SSD Drive: 558 MB/sec read, 532 MB/sec write speed
 M.2 Drive: 2591 MB/sec read, 1544 MB/sec write speed

In Read Performance:
SSD is 5x Faster than HDD
M.2 is 5x Faster than SSD
M.2 is 25x Faster than HDD

Magnetic disk




CSc 360

* multi-arm? roll-back arm movement?

Disk scheduling

queue = 98, 183, 37, 122, 14, 124, 65, 67
head starts at 53



CSc 360

5

- FCFS
- Example
 - total head movement
 - 640 cylinders

* fcfs in cpu scheduling? fifo in page replacement?




Shortest seek time first

queue = 98, 183, 37, 122, 14, 124, 65, 67

head starts at 53

- SSTF
 - similar to SJF
- Example
 - 236 cyl
- Starvation




* how to deal with starvation?



Scan

queue = 98, 183, 37, 122, 14, 124, 65, 67

head starts at 53 going toward 0



CSc 360

7

- AKA
 - elevator
- Example
 - 236 cyl
 - shortcut
 - 208 cyl


* going up?

Circular scan

queue = 98, 183, 37, 122, 14, 124, 65, 67

head starts at 53 going toward 199 for service

- Circular list of cylinders
- Example
 - 382 cyl!
 - fairness




* only going up?

Circular look

queue 98, 183, 37, 122, 14, 124, 65, 67

head starts at 53 going toward 199 for service




CSc 360

9

- No need to reach the end
- Choice?
 - requests
 - dynamic
 - type
 - file allocation

* what's the "disk scheduling" problem for ssd?

Disk management



- Physical formatting
 - creating physical sector structures
- Partition
- Logical formatting
 - creating an empty file system
- Boot block
- Handle bad sectors: sector sparing, slipping


Swap-space management

- Swap space
 - disk space as an extension of main memory
- Swap-space location
 - a regular big file
 - swap partition
- Swap-space management
 - text vs data segment
 - proactive vs reactive allocation


RAID




(a) RAID 0: non-redundant striping.




(b) RAID 1: mirrored disks.



(c) RAID 2: memory-style error-correcting codes.



(d) RAID 3: bit-interleaved parity.



(e) RAID 4: block-interleaved parity.



(f) RAID 5: block-interleaved distributed parity.



(g) RAID 6: P + Q redundancy.

- Performance through parallelism
 - striping
- Reliability through redundancy
 - mirroring
 - error correcting

This lecture

- Disk management
 - disk organization
 - geometry address
 - disk head scheduling
 - FCFS, SSTF, SCAN, C-SCAN, C-LOOK
 - swap-space
 - RAID

Next lecture

- I/O systems
 - OSC7Ch13