

# CSE 4/521

# Introduction to Operating Systems

Lecture 20 – Mass Storage Structure II

(Disk Scheduling)

Summer 2018

# Overview

---

Objective:

- To evaluate disk scheduling algorithms

- Disk Scheduling

# Recap

---

- **Overview of Mass Storage Structure**
  - Hard Disk Mechanism, Hard Disk Performance, SSD, Magnetic Tapes
- **Disk Structure**
  - Logical vs. Physical address translation transparent to the OS.

# Questions

---

1. What is **logical** and **physical** address, in terms of **Storage**? (**Easy**)
2. Why does **OS require accurate information** on how blocks are stored on disk? (**Open-ended**)
3. In what ways using SSDs as **1)a cache** and **2)as a disk-drive replacement** compare with using **only magnetic disks**. (**Open-ended**)

# Overview

---

- Disk Scheduling

# Disk Scheduling

---

- The operating system is responsible for using hardware efficiently — for the disk drives, this means having a **fast access time** and **disk bandwidth**
- **Minimize seek time**
- Seek time  $\approx$  seek distance
- **Disk bandwidth** is the **total number of bytes transferred**, divided by the total time between the first request for service and the completion of the last transfer

# Disk Scheduling

---

- There are many sources of disk I/O request
  - OS
  - System processes
  - Users processes
- OS maintains queue of requests, per disk or device
- Idle disk can immediately work on I/O request, busy disk means work must queue
  - Optimization algorithms only make sense when a queue exists

# Disk Scheduling

---

- Note that drive controllers have small buffers and can manage a queue of I/O requests (of varying “depth”)
- Several algorithms exist to schedule the servicing of disk I/O requests
- The analysis is true for one or many platters
- We illustrate scheduling algorithms with a request queue (0-199)

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

Head pointer 53


---



# Disk Scheduling: FCFS


---

Illustration shows total head movement of 640 cylinders



# Disk Scheduling: SSTF

- **Shortest Seek Time First** selects the request with the minimum seek time from the current head position
- SSTF scheduling is a form of SJF scheduling; may cause starvation of some requests
- Illustration shows total head movement of **236 cylinders**




# Disk Scheduling: SCAN

---

- The disk arm starts at one end of the disk, and moves toward the other end, servicing requests until it gets to the other end of the disk, where the head movement is reversed and servicing continues.
- **SCAN algorithm** Sometimes called the **elevator algorithm**
- Illustration shows total head movement of **236 cylinders**
- But note that if **requests are uniformly dense**, **largest density is at other end of disk** and those **wait the longest**

# Disk Scheduling: SCAN

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



# Disk Scheduling: C-SCAN


---

- Provides a **more uniform wait time than SCAN**
- The head moves from one end of the disk to the other, servicing requests as it goes
  - When it reaches the other end, however, it **immediately returns to the beginning of the disk**, without servicing any requests on the return trip
- Treats the cylinders as a circular list that wraps around from the last cylinder to the first one

# Disk Scheduling: C-SCAN

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

head starts at 53



# Disk Scheduling: C-LOOK

---


- LOOK a version of SCAN, C-LOOK a version of C-SCAN
- Arm only goes as far as the last request in each direction, then reverses direction immediately, without first going all the way to the end of the disk

# Disk Scheduling: C-LOOK

---

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

head starts at 53





# Disk Scheduling: Selecting a Disk-Scheduling Algorithm

---

- **SSTF is common** and has a natural appeal
- **SCAN and C-SCAN** perform better for systems that place a heavy load on the disk
  - Less starvation
- Performance depends on the **number** and **types of requests**
- Requests for disk service can be influenced by the **file-allocation method**
  - And metadata layout
- The disk-scheduling algorithm should be written as a separate module of the operating system, allowing it to be replaced with a different algorithm if necessary
- Either **SSTF** or **LOOK** is a reasonable choice for **the default algorithm**
- What about rotational latency?
  - Difficult for OS to calculate
- How does disk-based queueing effect OS queue ordering efforts?

# Credits for slides

Silberschatz, Galvin and Gagne