

15-410

“...What goes around comes around...”

Disks

Mar. 24, 2004

Dave Eckhardt & Bruce Maggs

Brian Railing & Steve Muckle

Contributions from

- **Eno Thereska**
- **15-213**
- **“How Stuff Works” web site**

Synchronization

Project 3 status

- Checkpoint 1 – not what we had hoped for
- Checkpoint 2 – most people caught up

Suggestions

- Three regular meeting times per week
 - Two hours or more at each meeting
 - Begin by asking questions about each other's code
 - » Requires having read code before meeting
 - » Requires “quiet time” between check-ins and meeting
- Source control
 - Frequent merges, not a single “big bang” at end
- Leave time at end for those multi-day bugs

Overview

Anatomy of a Hard Drive

Common Disk Scheduling Algorithms

Anatomy of a Hard Drive

On the outside, a hard drive looks like this



Taken from "How Hard Disks Work"
<http://computer.howstuffworks.com/hard-disk2.htm>

Anatomy of a Hard Drive

If we take the cover off,
we see that there
actually is a “hard
disk” inside



Taken from “How Hard Disks Work”
<http://computer.howstuffworks.com/hard-disk2.htm>

Anatomy of a Hard Drive

A hard drive usually contains multiple disks, called *platters*

These spin at thousands of RPM (5400, 7200, etc)



Taken from "How Hard Disks Work"
<http://computer.howstuffworks.com/hard-disk2.htm>

Anatomy of a Hard Drive

Information is written to
and read from the
platters by the
read/write heads on
the *disk arm*



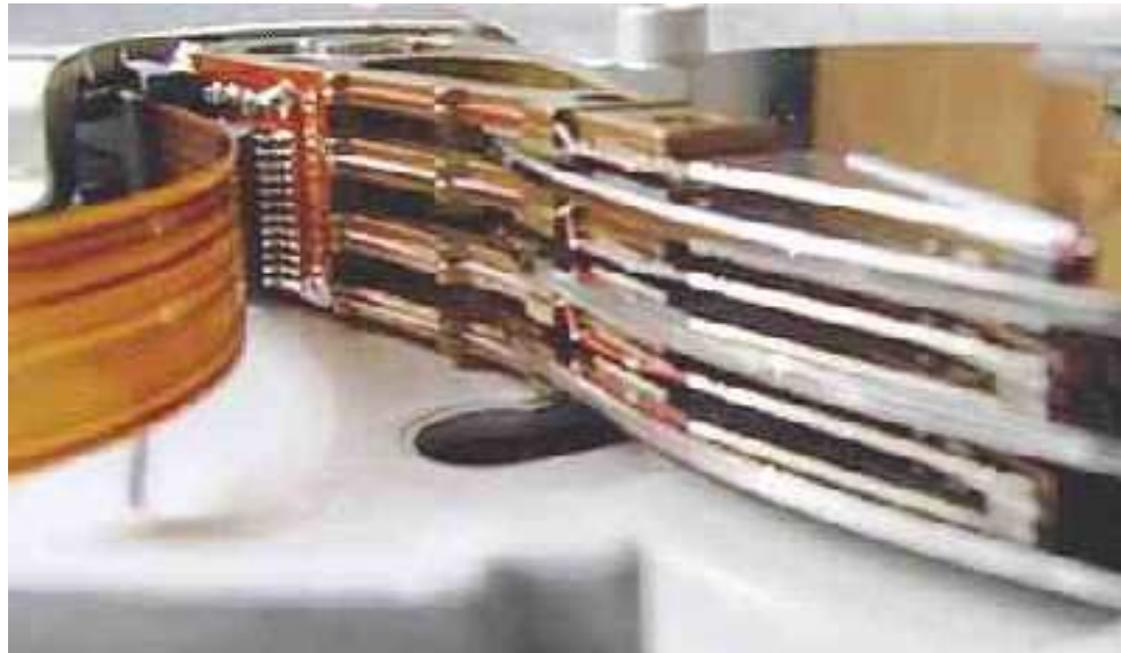
Taken from "How Hard Disks Work"
<http://computer.howstuffworks.com/hard-disk2.htm>

Anatomy of a Hard Drive

Both sides of each platter store information

Each side of a platter is called a *surface*

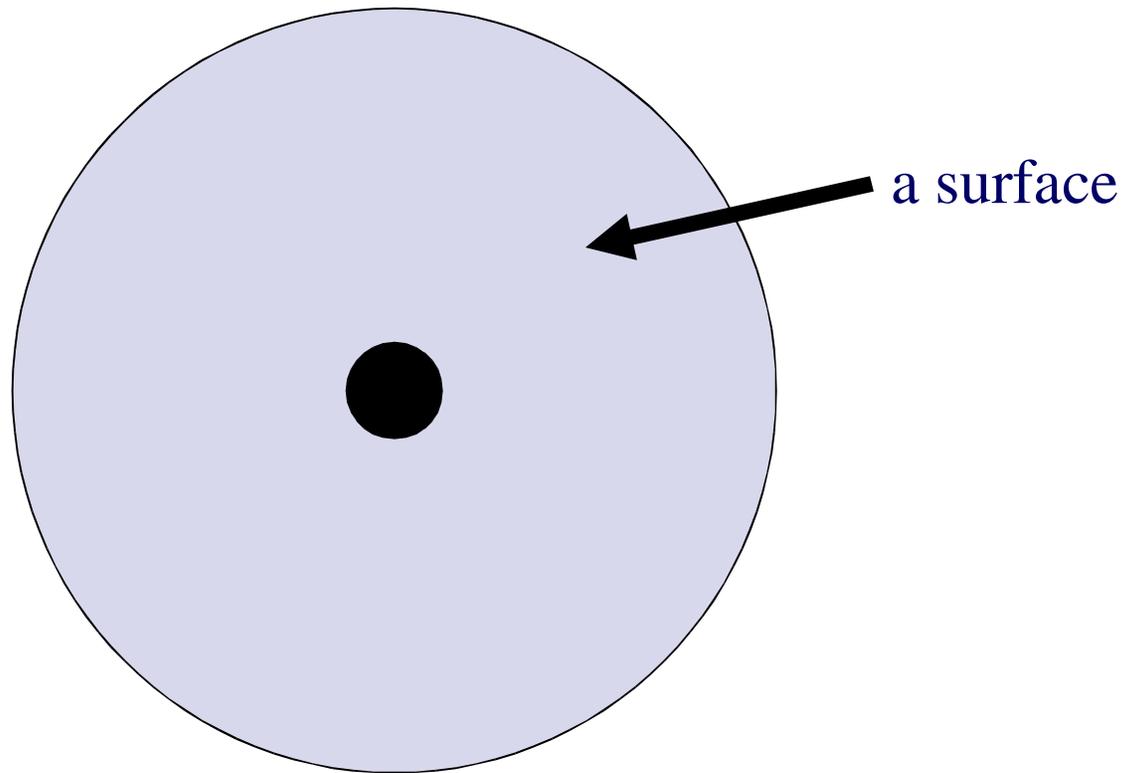
Each surface has its own read/write head



Taken from "How Hard Disks Work"
<http://computer.howstuffworks.com/hard-disk2.htm>

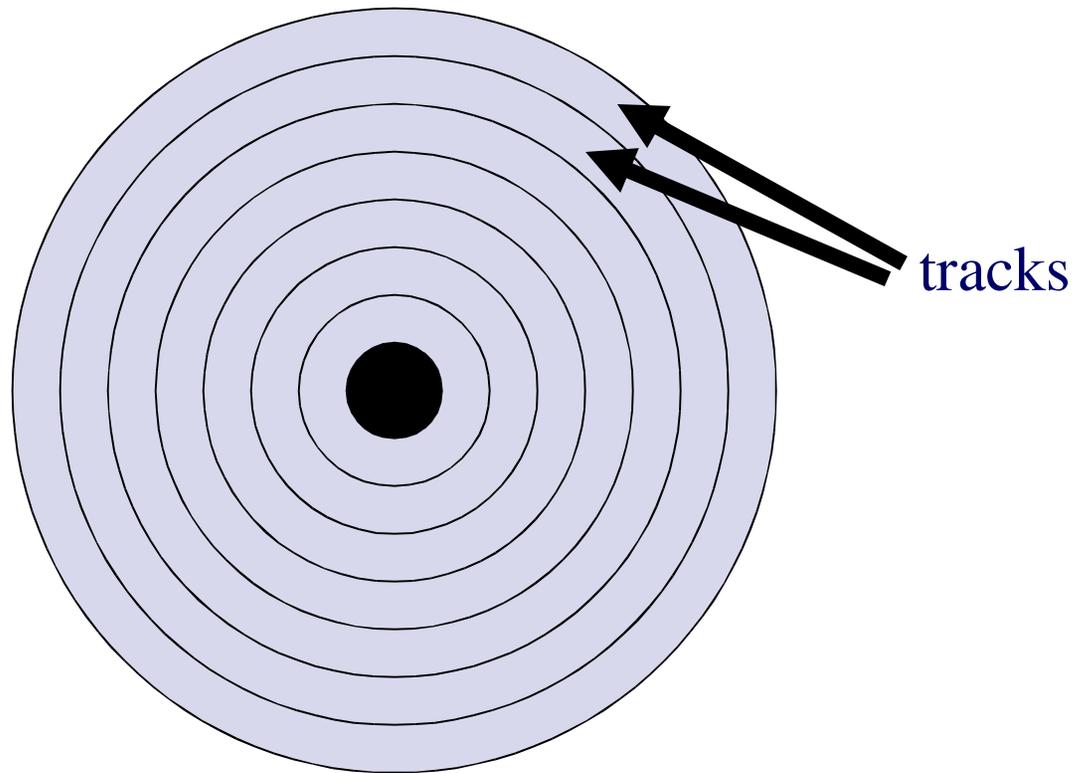
Anatomy of a Hard Drive

How are the surfaces organized?



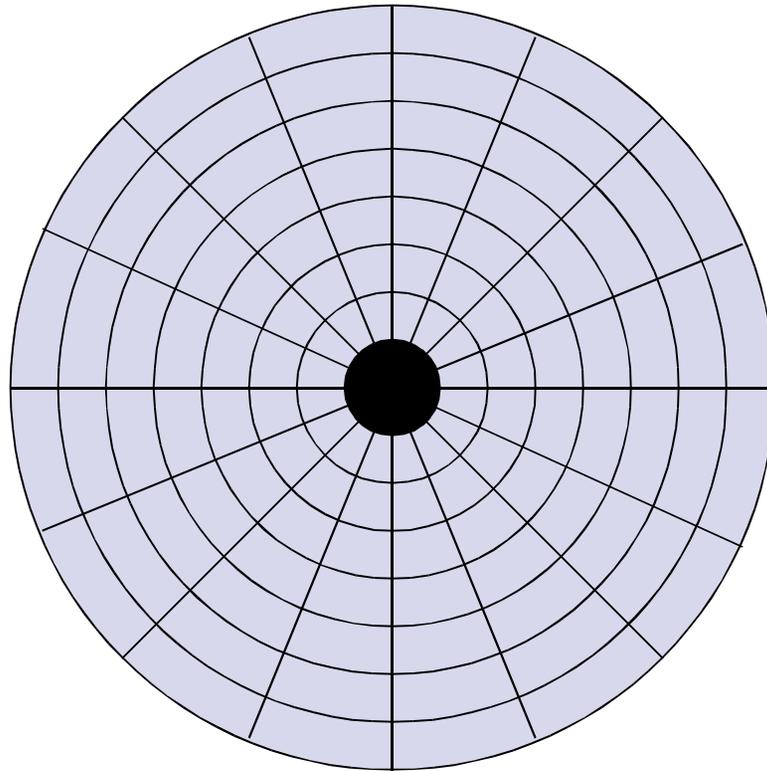
Anatomy of a Hard Drive

Each surface is divided by concentric circles, creating *tracks*



Anatomy of a Hard Drive

These tracks are further divided into *sectors*



Anatomy of a Hard Drive

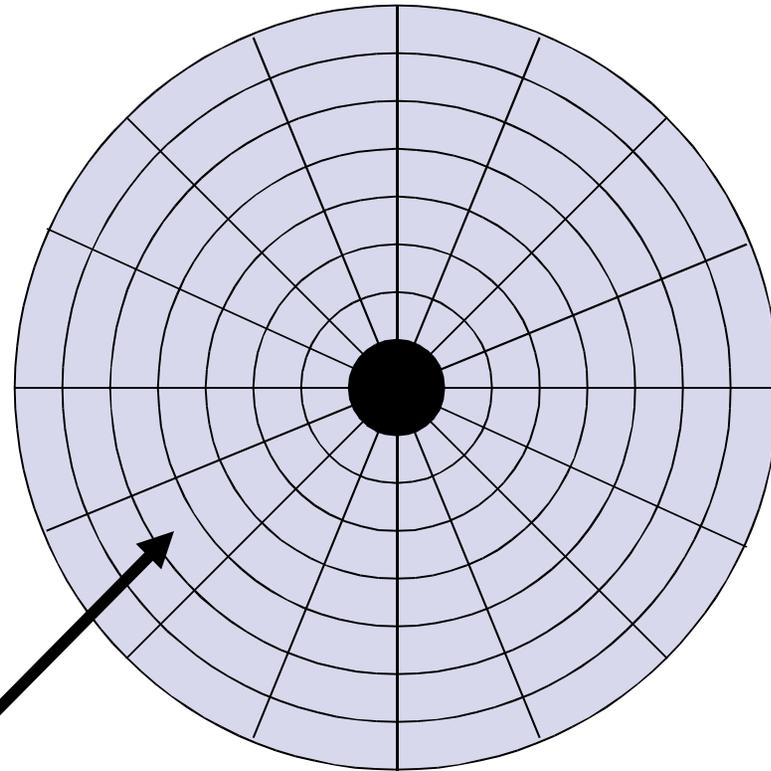
A sector is the smallest unit of data transfer to or from the disk

Most modern hard drives have 512 byte sectors

CD-ROM sectors are 2048 bytes

Gee, those outer sectors look bigger...

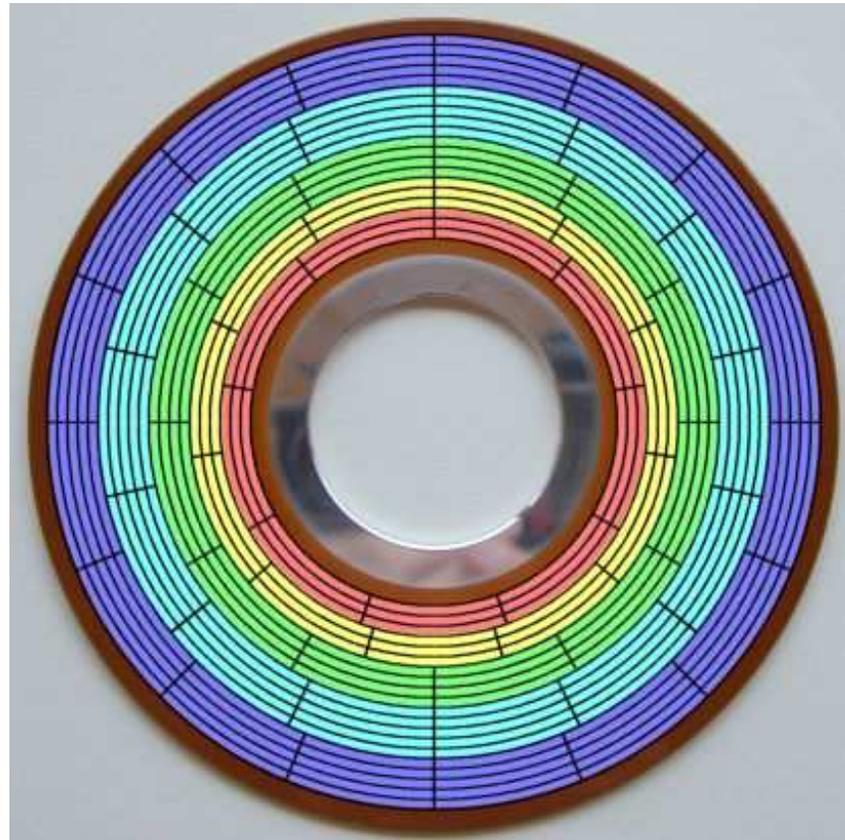
a sector



Anatomy of a Hard Drive

Does this mean that sectors on the outside of a surface are larger than those on the inside?

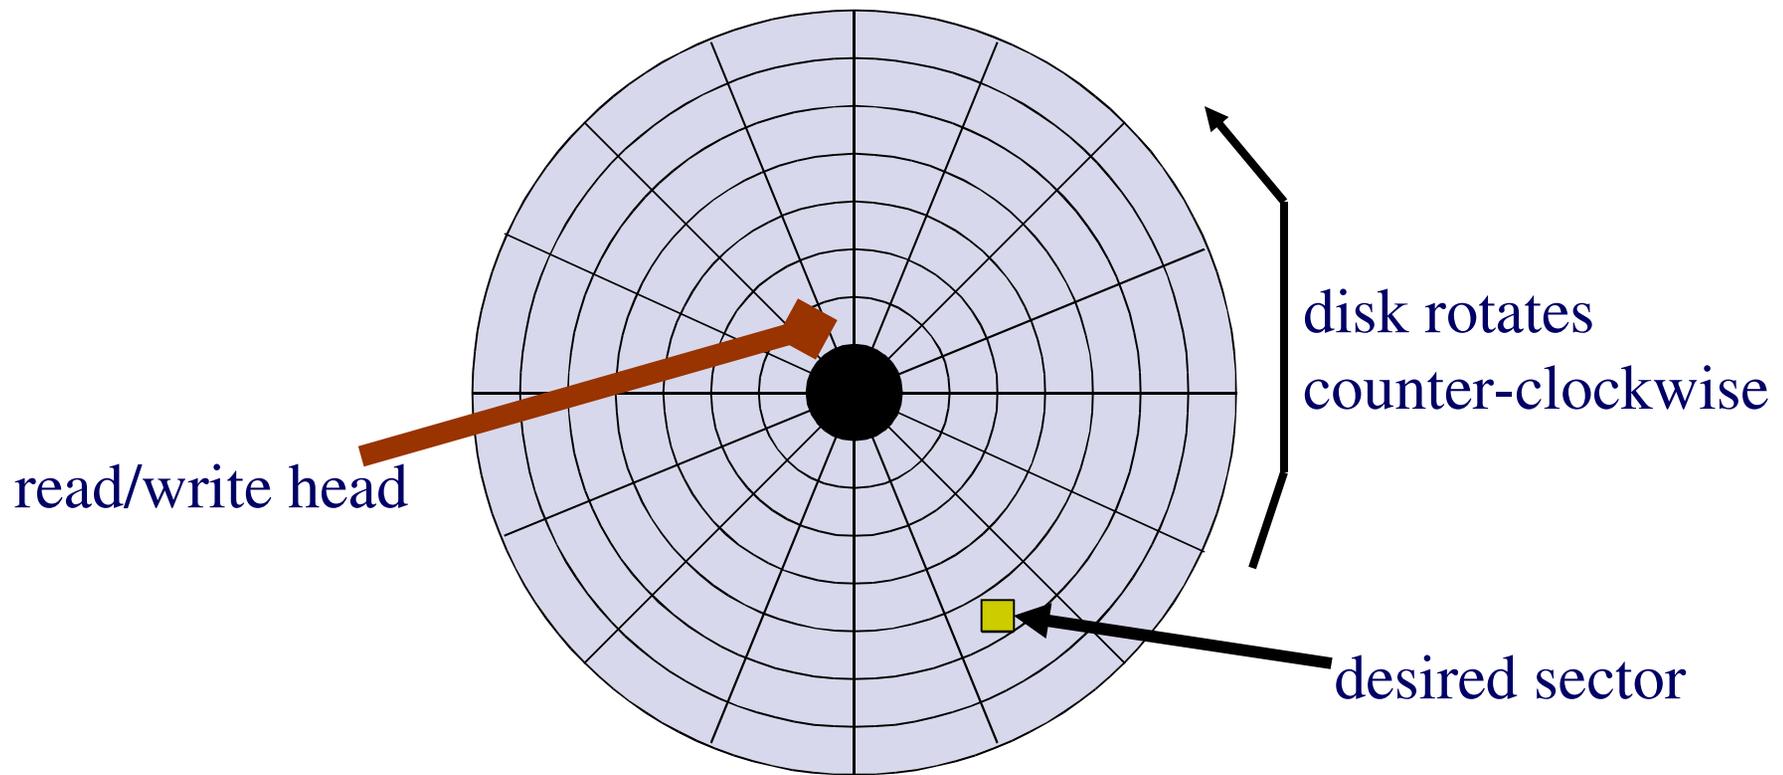
Modern hard drives fix this with *zoned bit recording*



Taken from "Reference Guide – Hard Disk Drives"
<http://www.storagereview.com/map/lm.cgi/zone>

Anatomy of a Hard Drive

Let's read in a sector from the disk



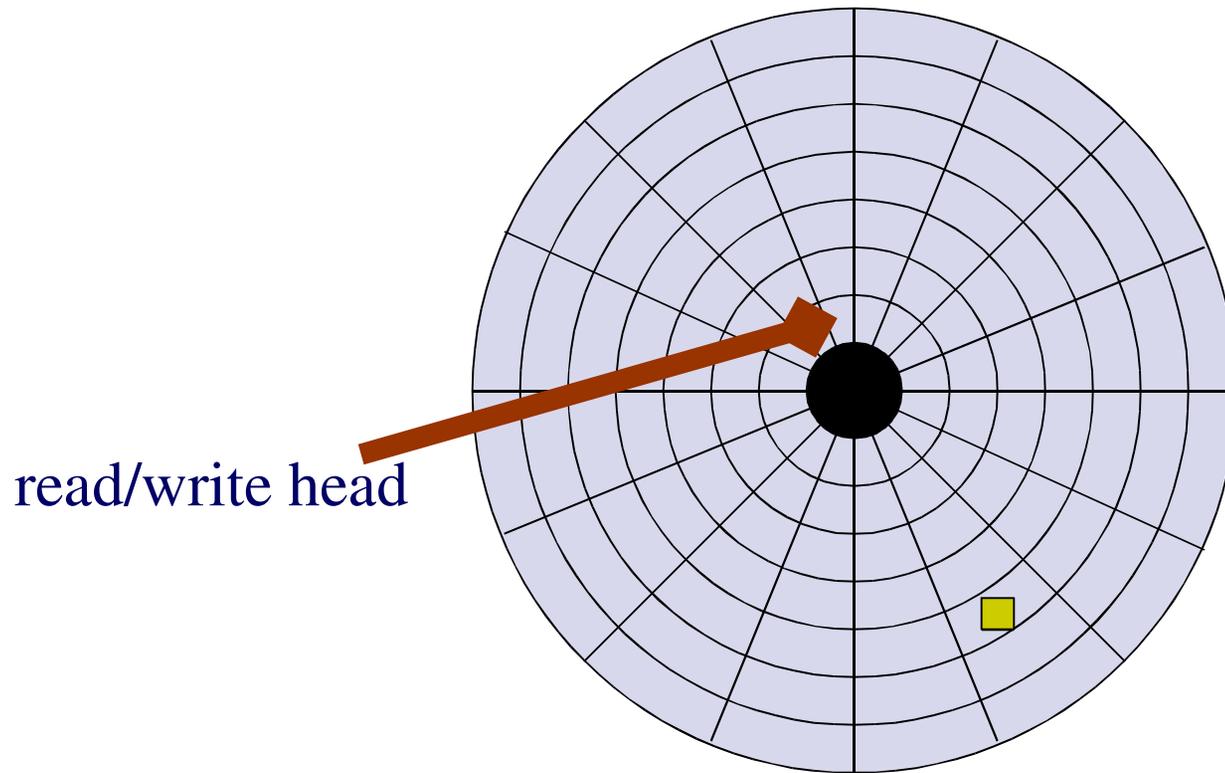
Anatomy of a Hard Drive

We need to do two things to transfer a sector

- 1. Move the read/write head to the appropriate track (seek)**
- 2. Wait until the desired sector spins around**

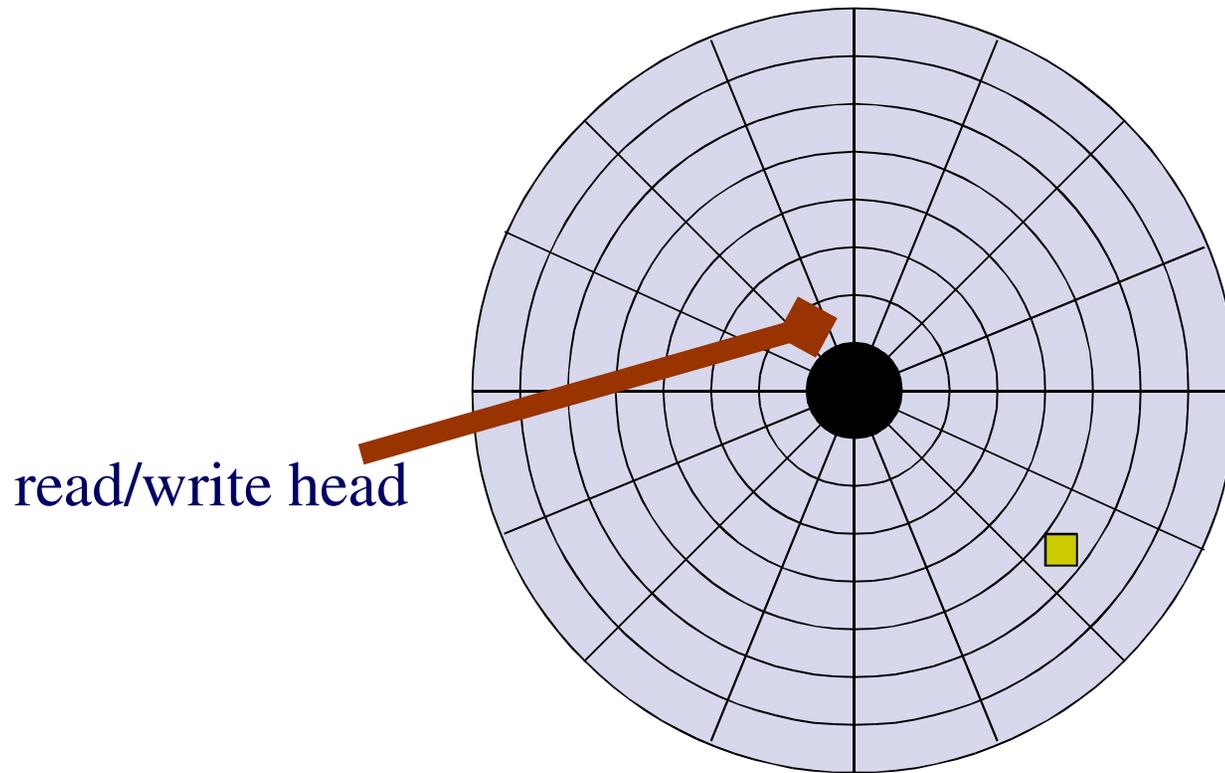
Anatomy of a Hard Drive

Let's read in a sector from the disk



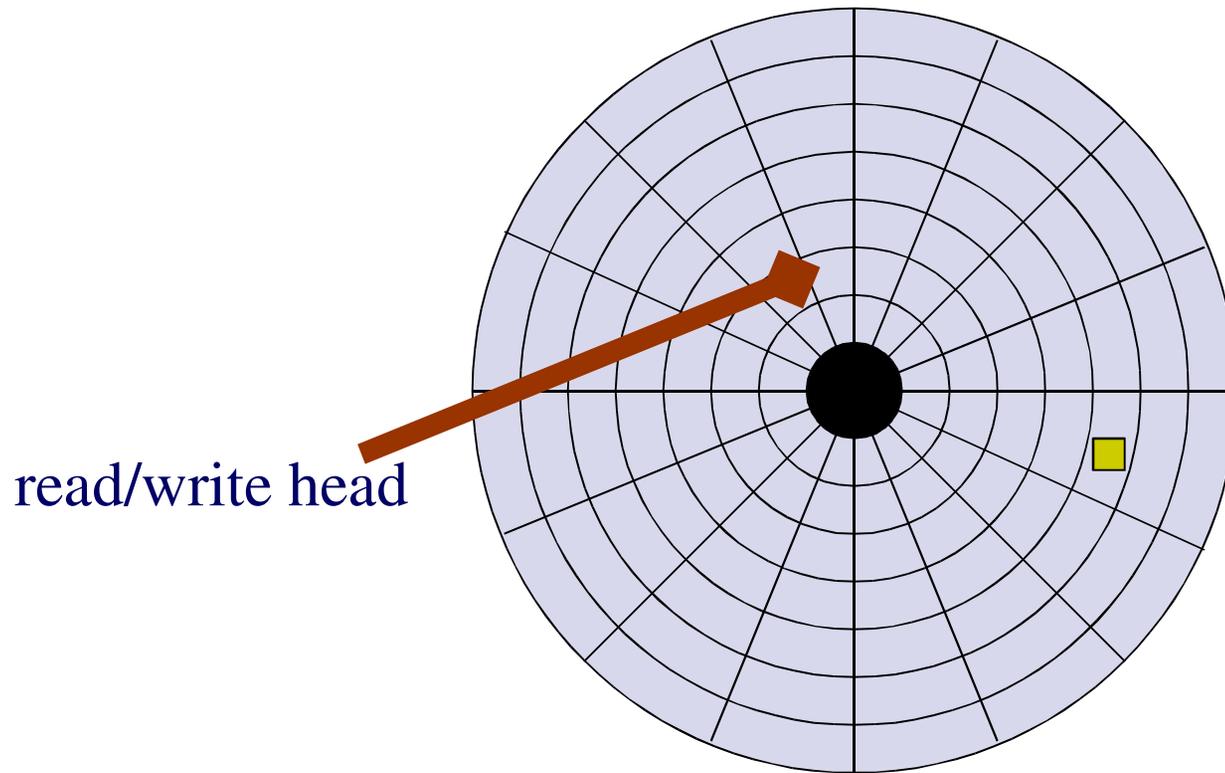
Anatomy of a Hard Drive

Let's read in a sector from the disk



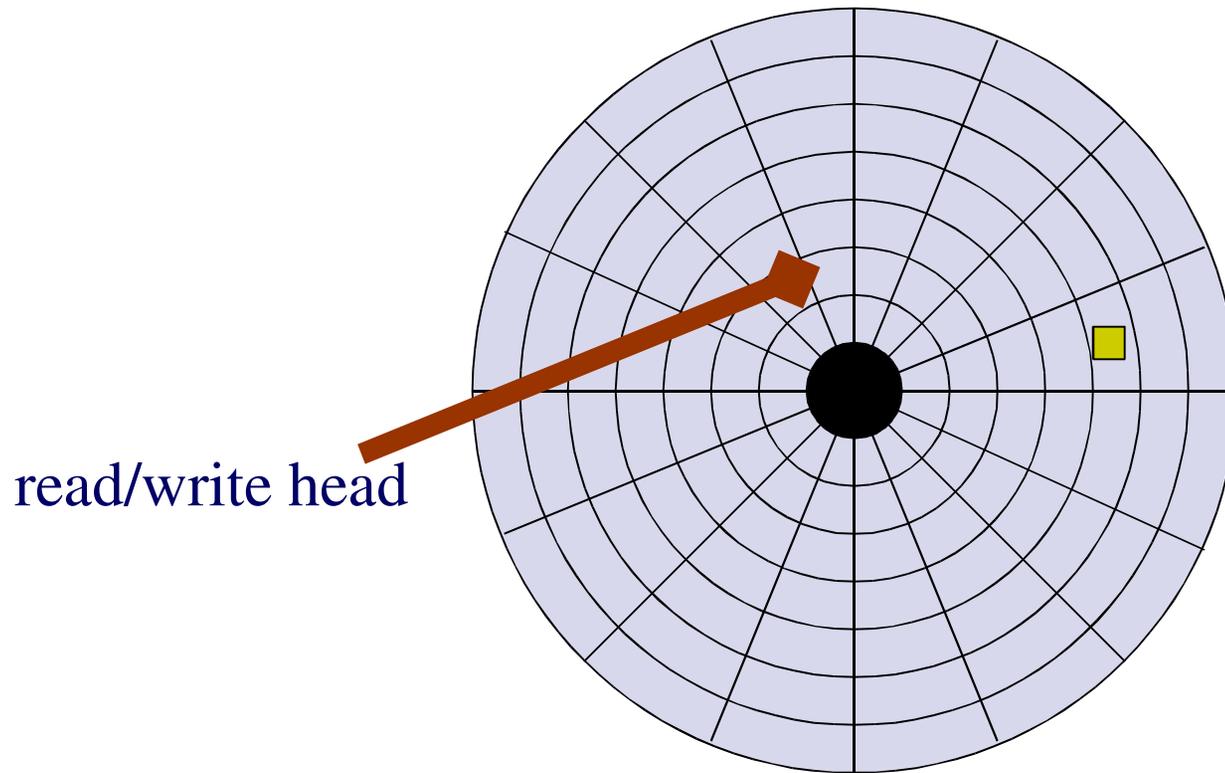
Anatomy of a Hard Drive

Let's read in a sector from the disk



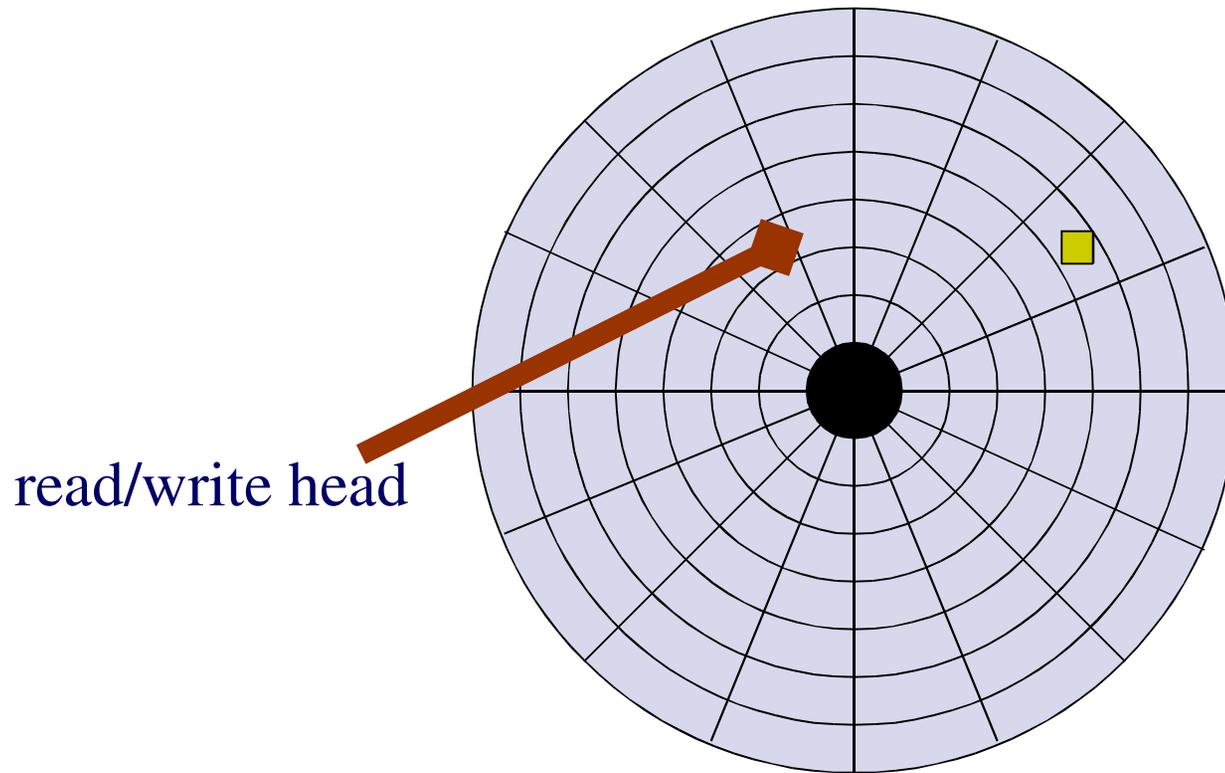
Anatomy of a Hard Drive

Let's read in a sector from the disk



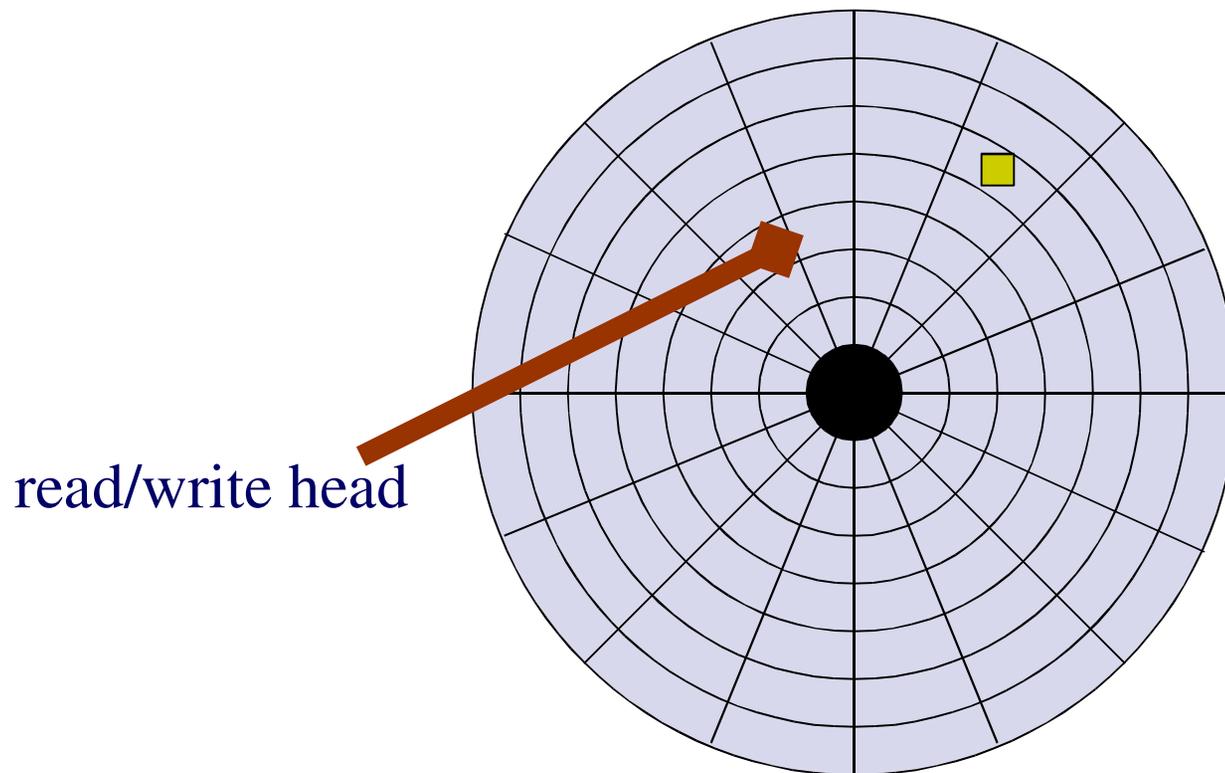
Anatomy of a Hard Drive

Let's read in a sector from the disk



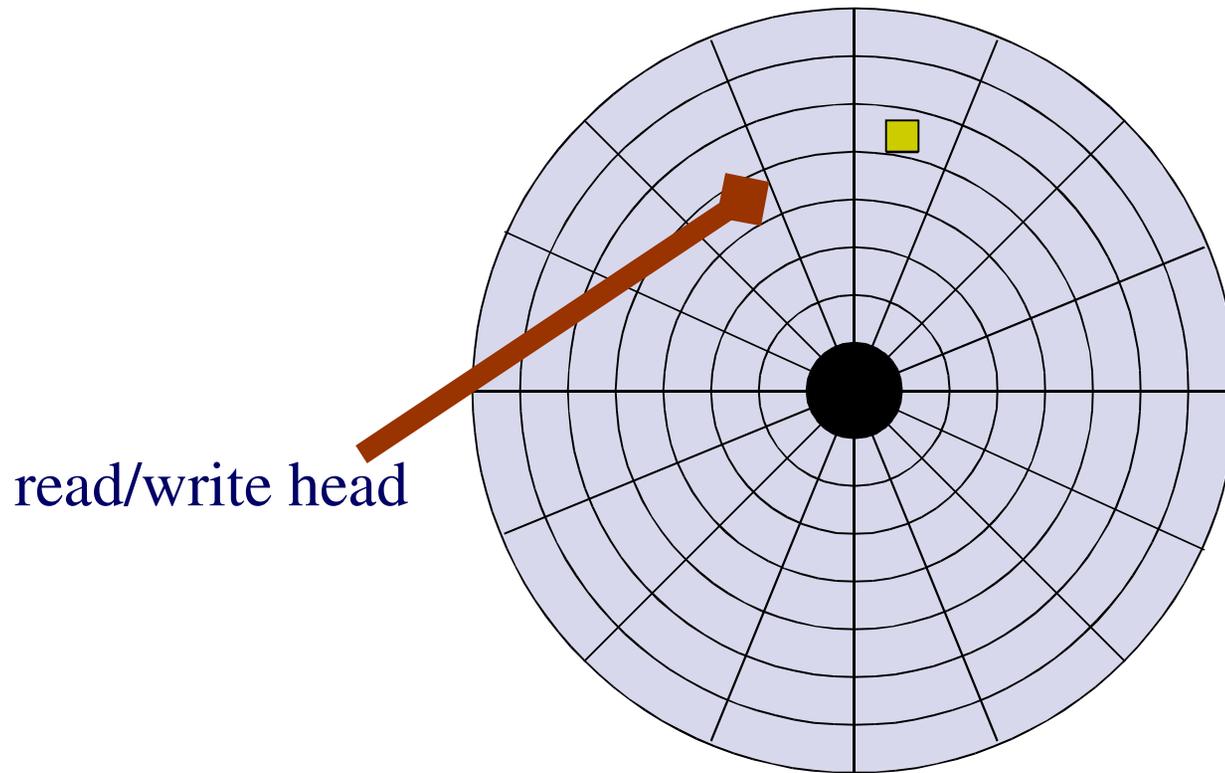
Anatomy of a Hard Drive

Let's read in a sector from the disk



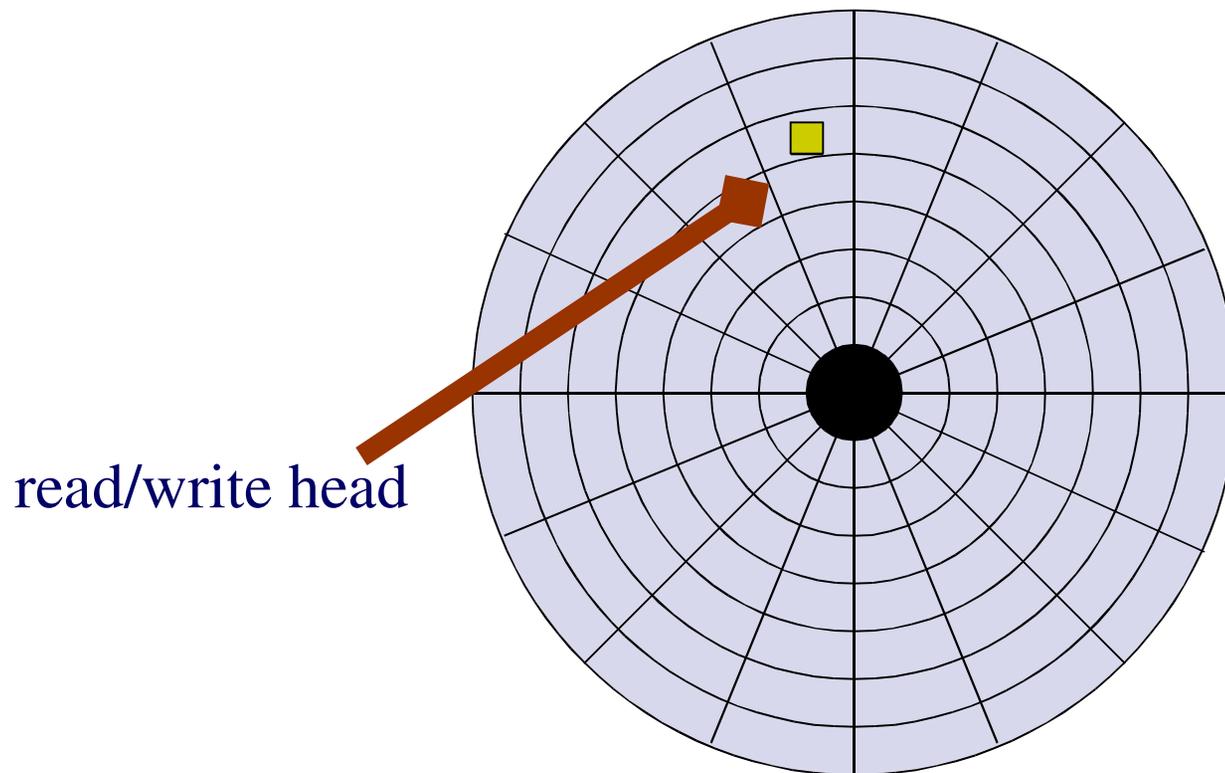
Anatomy of a Hard Drive

Let's read in a sector from the disk



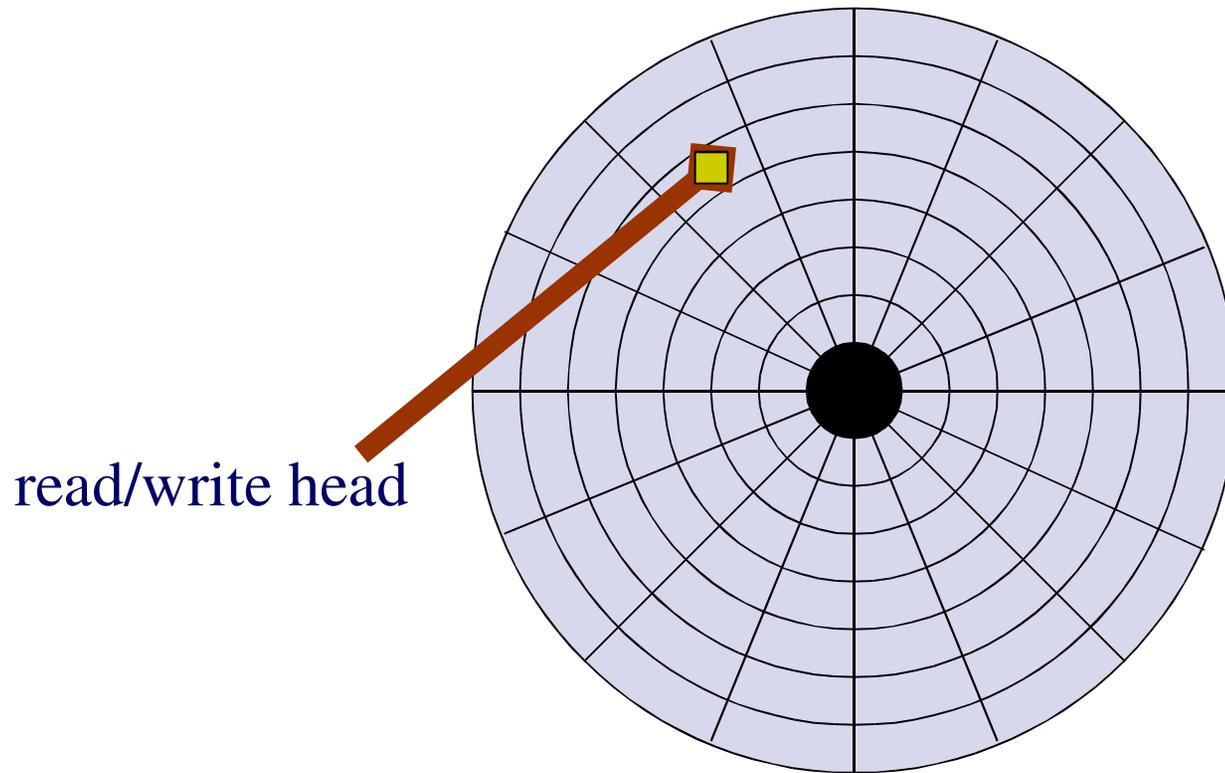
Anatomy of a Hard Drive

Let's read in a sector from the disk



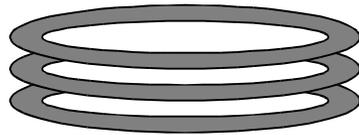
Anatomy of a Hard Drive

Let's read in a sector from the disk



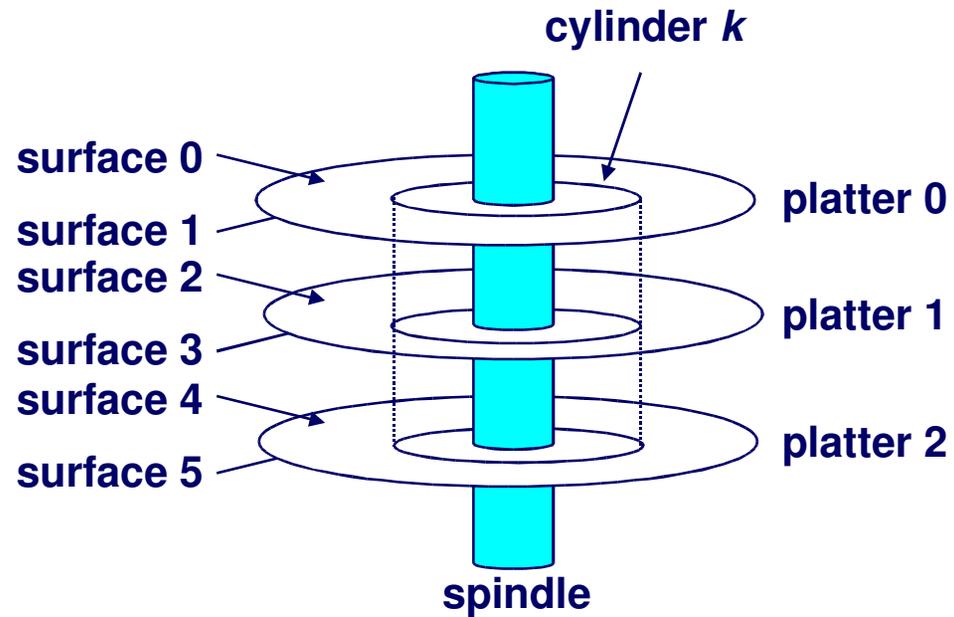
Disk Cylinder

**Matching tracks across surfaces are collectively called
a *cylinder***



Disk Cylinder

Matching tracks form a cylinder.



Cheap Access Within A Cylinder

Heads on single arm

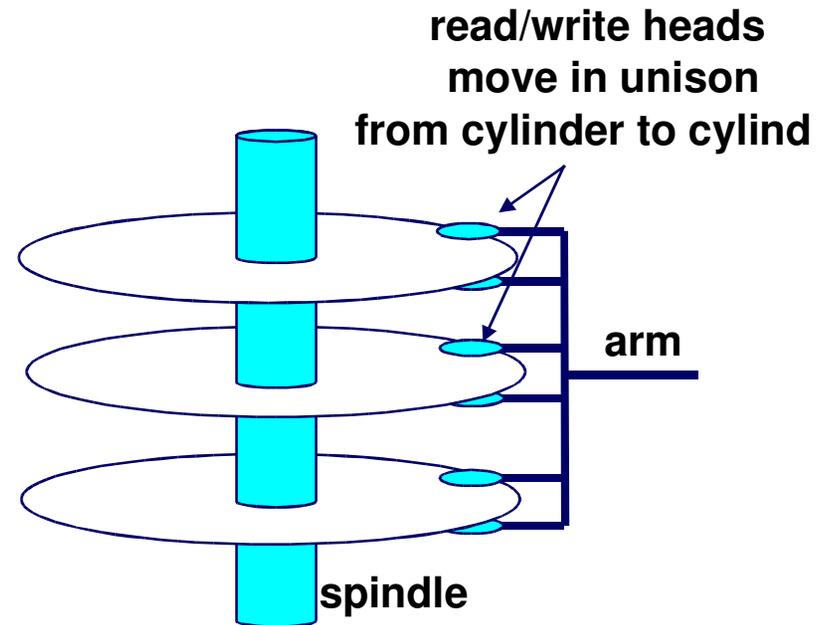
- All heads always on same cylinder

Switching heads is “cheap”

- Deactive head 3
- Activate head 4
- Wait for 1st sector header

Optimal transfer rate

- Transfer all sectors on a track
- Transfer all tracks on a cylinder
- Then move elsewhere



Anatomy of a Hard Drive

On average, we will have to move the read/write head over half the tracks

The time to do this is the “average seek time”, and is ~10ms for a 5400 rpm disk

We will also must wait half a rotation

The time to do this is rotational latency, and on a 5400 rpm drive is ~5.5ms

Anatomy of a Hard Drive

Other factors influence overall disk access time including

- **Settle time, the time to stabilize the read/write head after a seek**
- **Command overhead, the time for the disk to process a command and start doing something**

Minor compared to seek time and rotational latency

Anatomy of a Hard Drive

Total drive random access time is on the order of 15 to 20 milliseconds

- 50 ½-kilobyte transfers per second = 25 Kbyte/sec

Oh man, disks are slow

- But wait! Disk transfer rates are tens of MBytes/sec

What can we, as operating system programmers, do about this?

- Read more per seek (multi-sector transfers)
- Don't seek so randomly (“disk scheduling”)

Disk Scheduling Algorithms

The goal of a disk scheduling algorithm is to be nice to the disk

We can help the disk by giving it requests that are located close to each other on the disk

This minimizes seek time, and possibly rotational latency

There exist a variety of ways to do this

Disk Scheduling Algorithms

What the OS knows about the disk?

- Logical Block Numbers
- Interface: IDE or SCSI

What happened to sectors, tracks, etc?

- They are hidden behind the logical block numbers

How are they used?

- File systems assign logical blocks to files

First Come First Served (FCFS)

Send requests to disk as they are generated by the OS

Trivial to implement – FIFO queue in device driver

Fair

“Unacceptably high mean response time”

- File “abc” in sectors 1, 2, 3, ...
- File “def” in sectors 16384, 16385, 16386, ...
- Sequential reads: 1, 16384, 2, 16385, 3, 16386

“Fair, but cruel”

Shortest Seek Time First (SSTF)

Maintain “queue” of disk requests

Serve the request nearest to the disk arm

- Easy to estimate by difference in block numbers

Generates very fast response time

Intolerable response time *variance*, however

Why?

What Went Wrong?

FCFS - “fair, but cruel”

- Ignores position of disk arm, very slow

SSTF – good throughput, very unfair

- Ignores need to eventually scan entire disk

“Scan entire disk” - now that's an idea!

- Start disk arm moving in one direction
- Serve requests as the arm moves past them
 - No matter what order they arrived in
- When arm bangs into stop, reverse direction

SCAN

Doubly-linked ordered list of requests

- Insert according to order

Bi-directional scanning

- Direction = +1 or -1
- Tell disk: “seek to cylinder $X = \text{current} + \text{direction}$ ”
- Examine list for requests in cylinder X
- If $X == 0$ or $X == \text{max}$, direction = -direction
- Else current = X

Evaluating SCAN

Mean response time

- Worse than SSTF, better than FCFS

Response time variance

- Better than SSTF

Unfair – why?

LOOK

Just like SCAN – sweep back and forth through cylinders

Don't wait for the “thud” to reverse the scan

- Reverse when there are no requests “ahead” of the arm

Improves mean response time, variance

Still unfair though

CSCAN - “Circular SCAN”

Send requests in ascending cylinder order

When the last cylinder is reached, seek all the way back to the first cylinder

Long seek is amortized across all accesses

Variance is improved

Fair

Still missing something though...

C-LOOK

CSCAN + LOOK

Scan in one direction, as in CSCAN

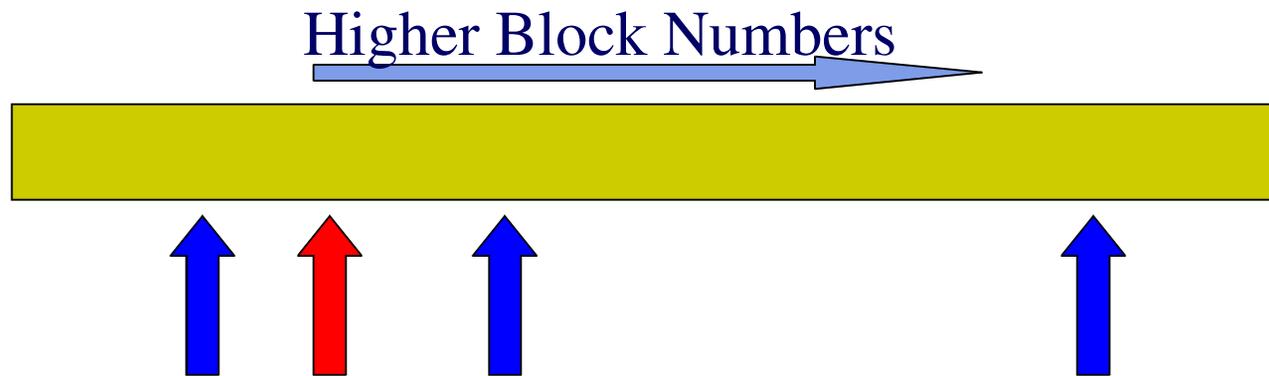
**If there are no more requests in current direction go
back to furthest request**

Very popular

C-LOOK

Blue are requests

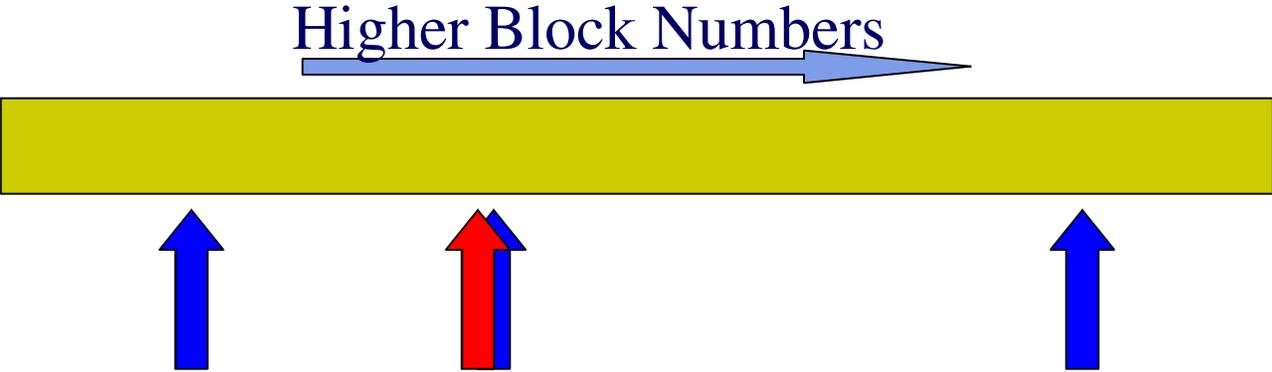
Yellow is disk



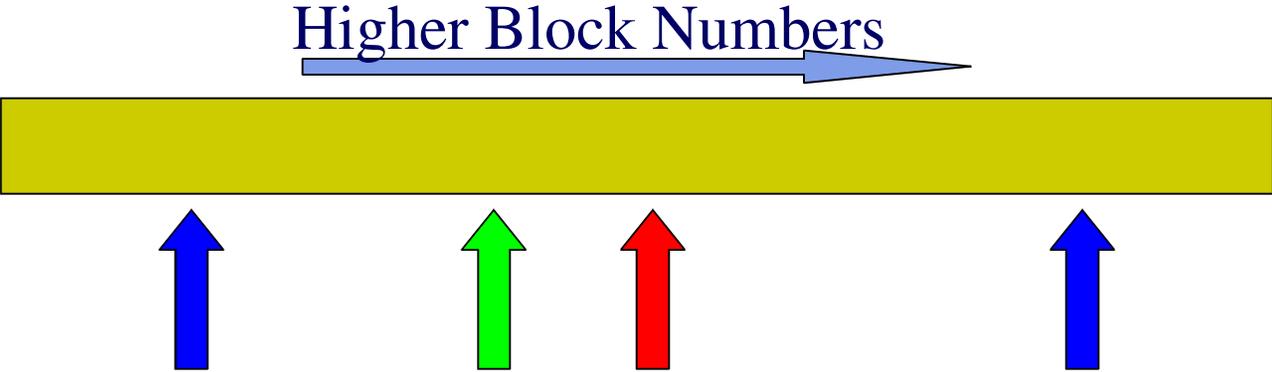
Red is disk head

Green is completed requests

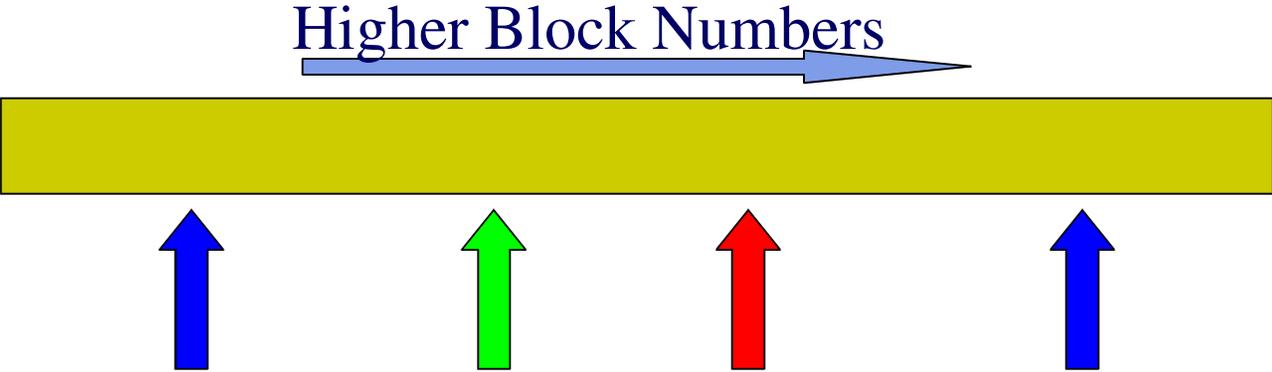
C-LOOK



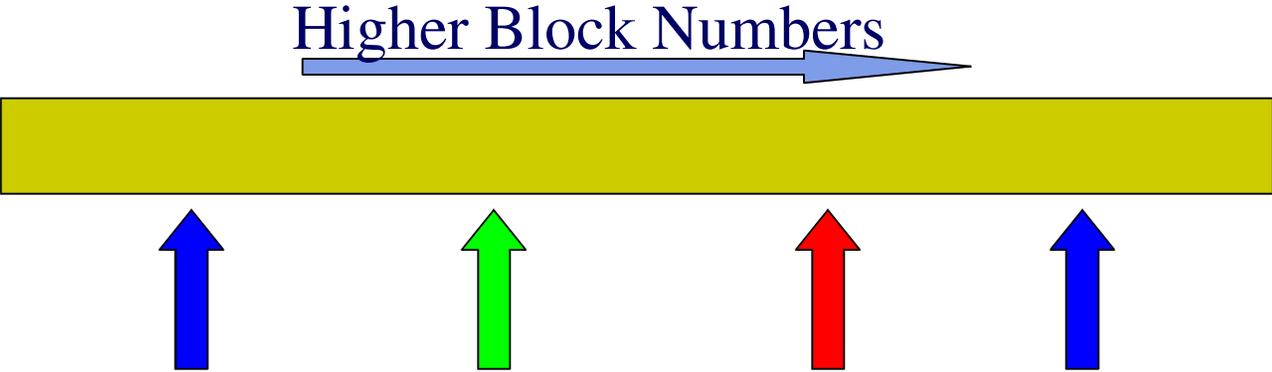
C-LOOK



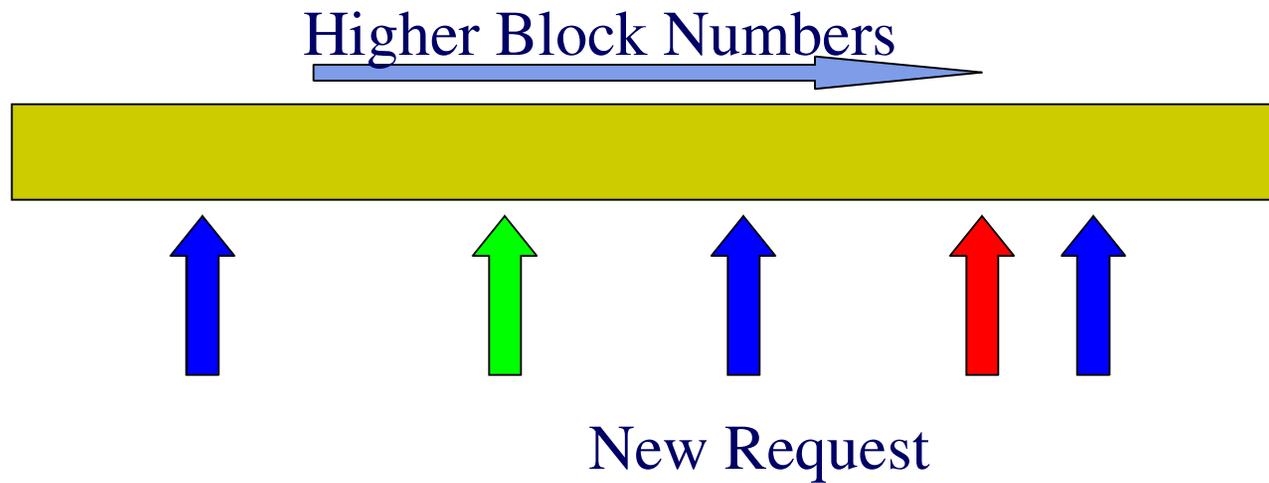
C-LOOK



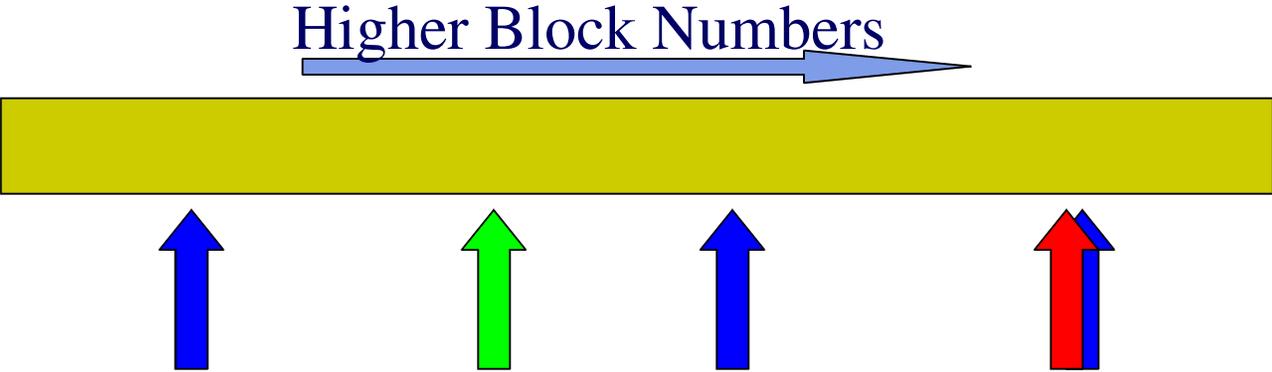
C-LOOK



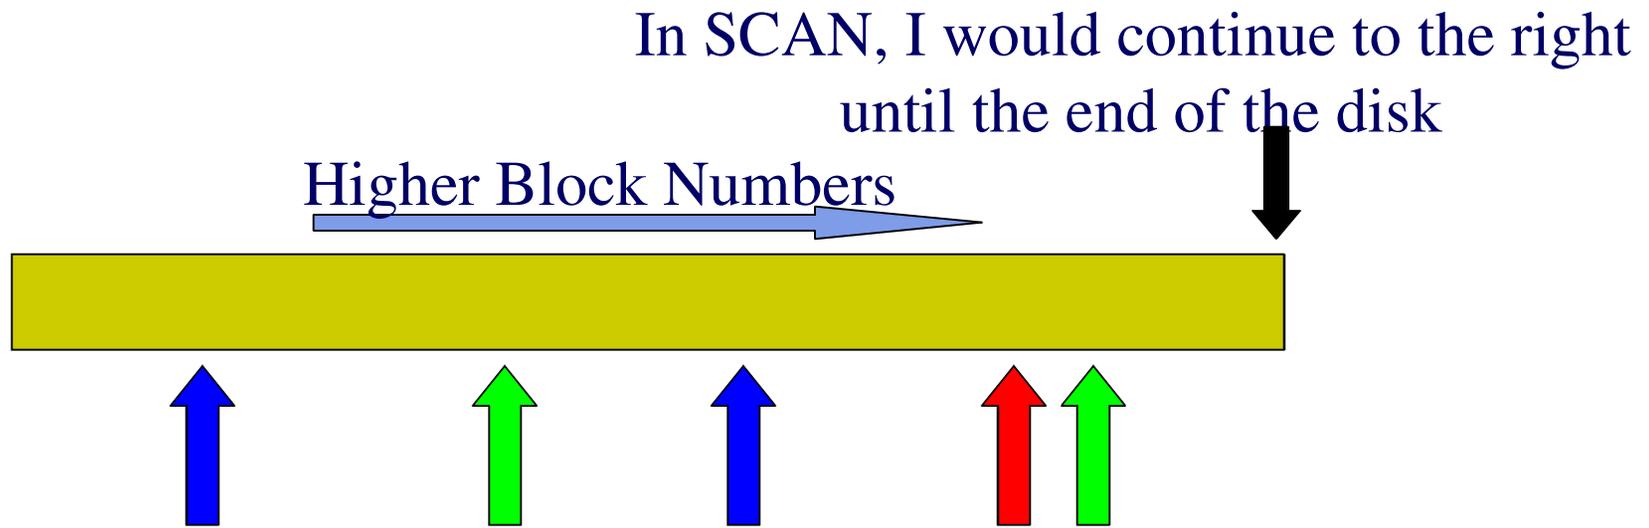
C-LOOK



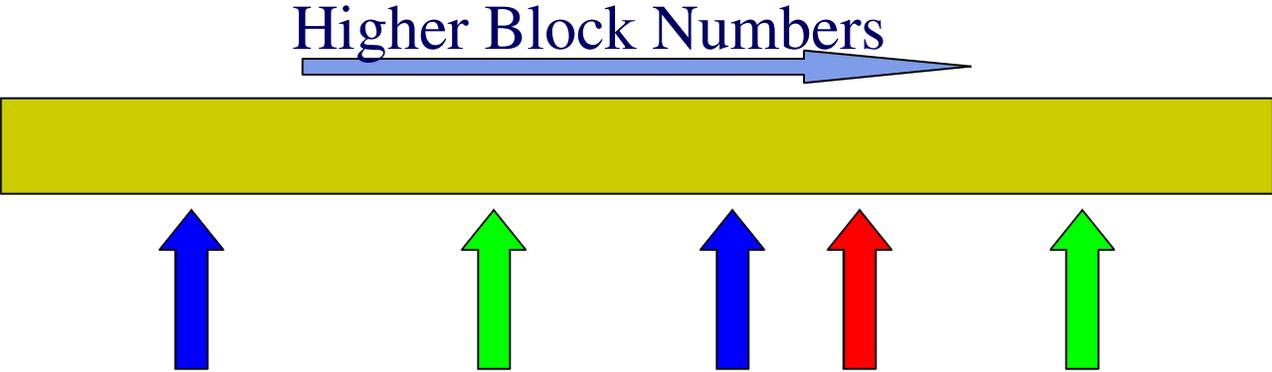
C-LOOK



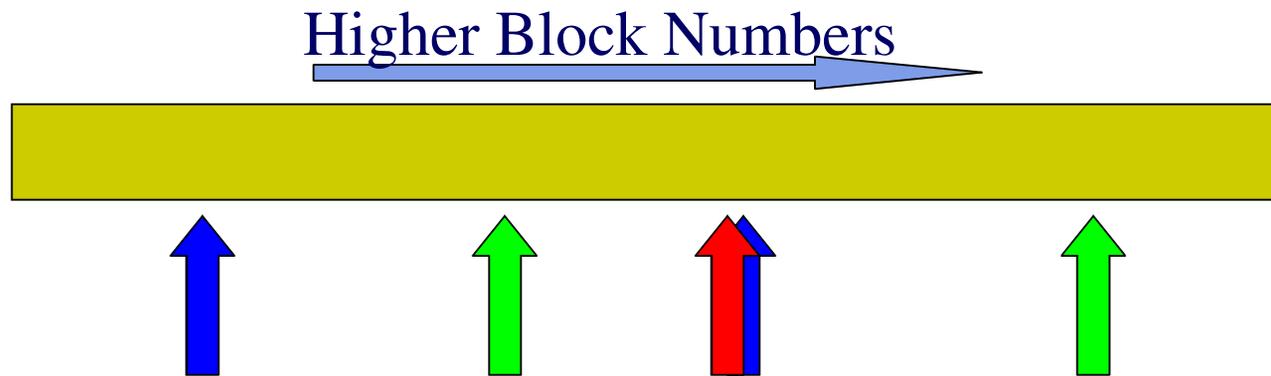
C-LOOK



C-LOOK

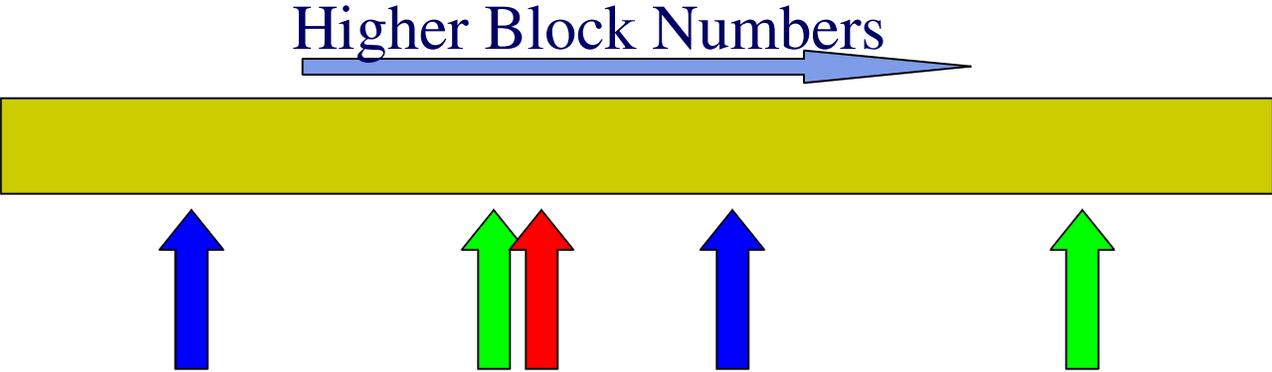


C-LOOK

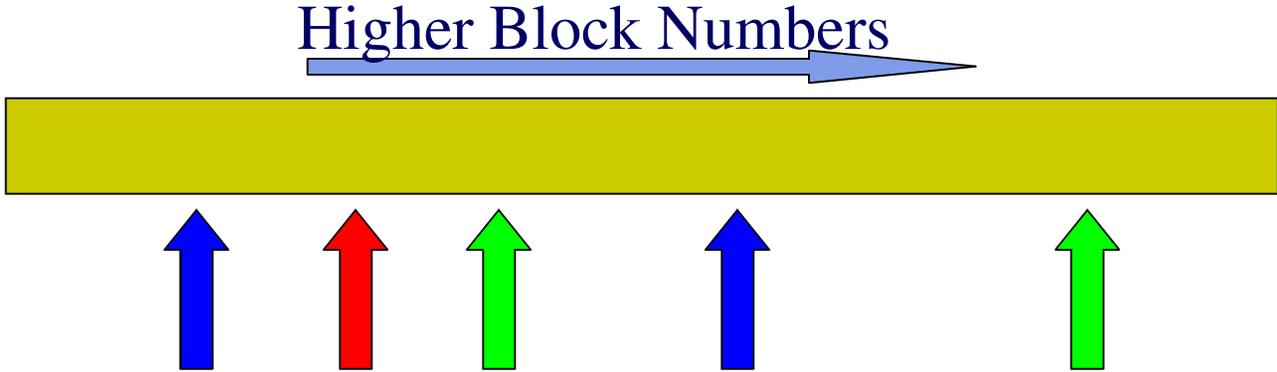


In LOOK, we would have read this request

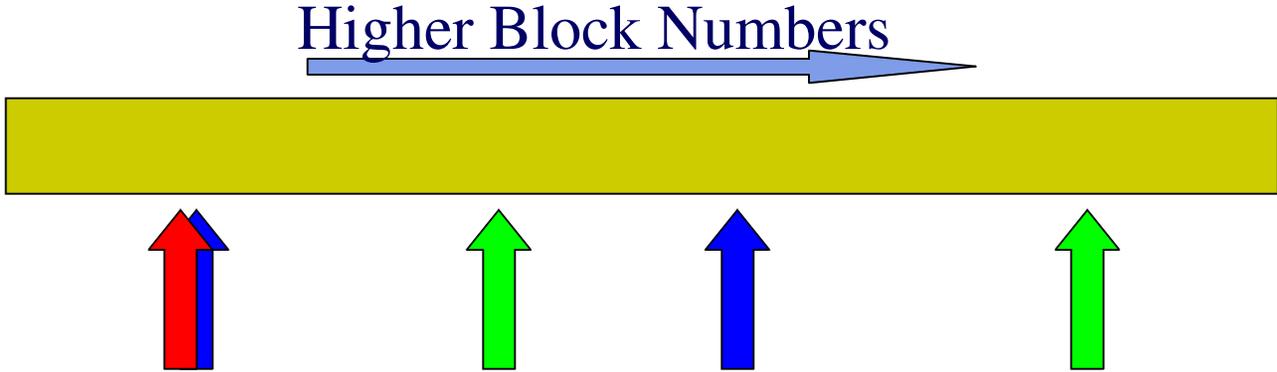
C-LOOK



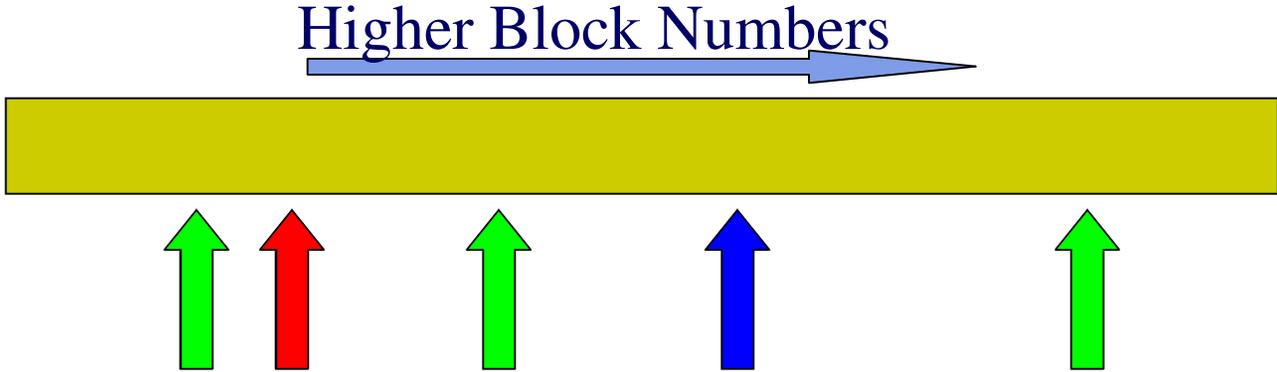
C-LOOK



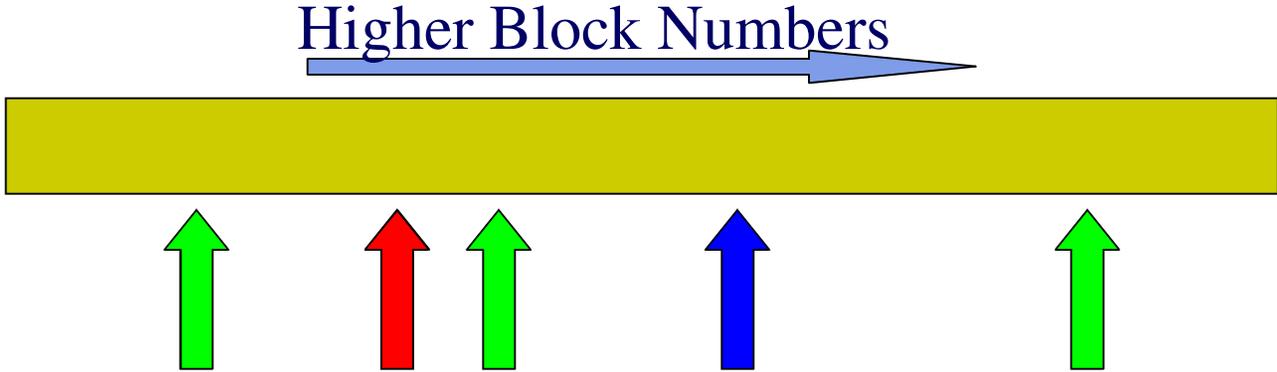
C-LOOK



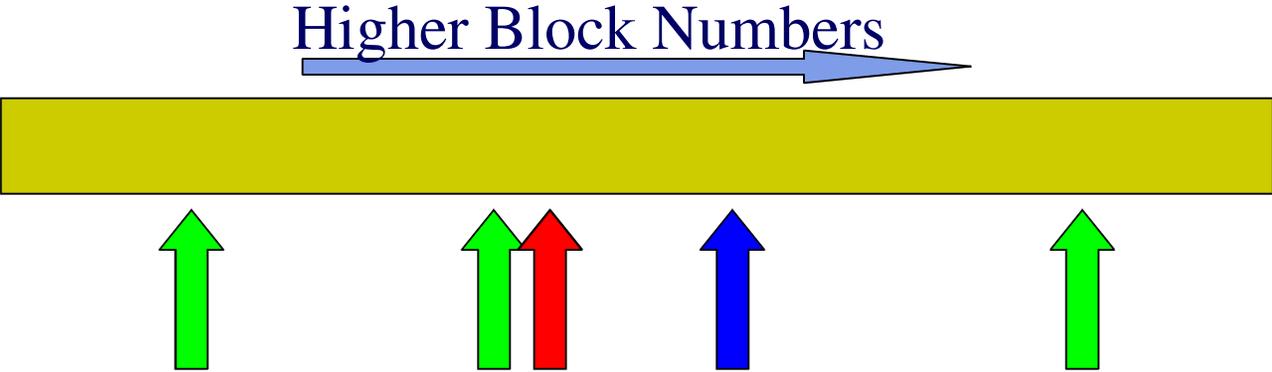
C-LOOK



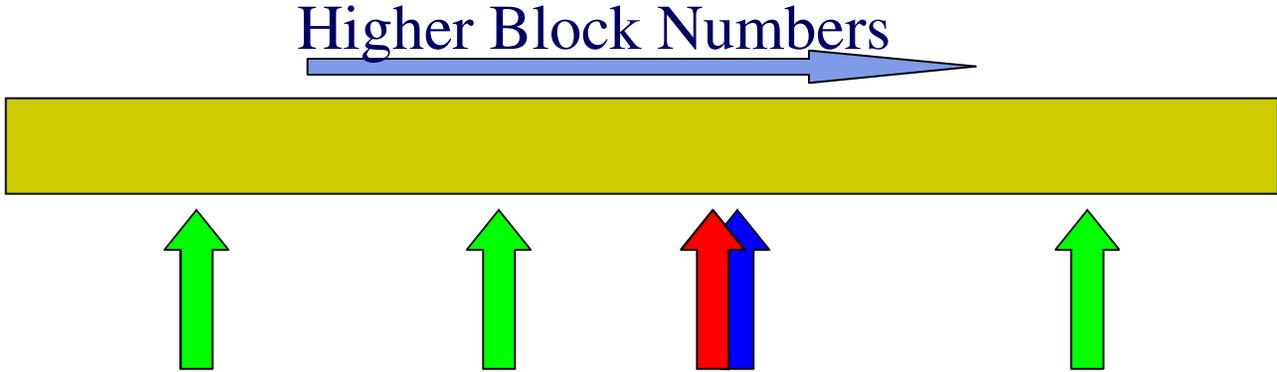
C-LOOK



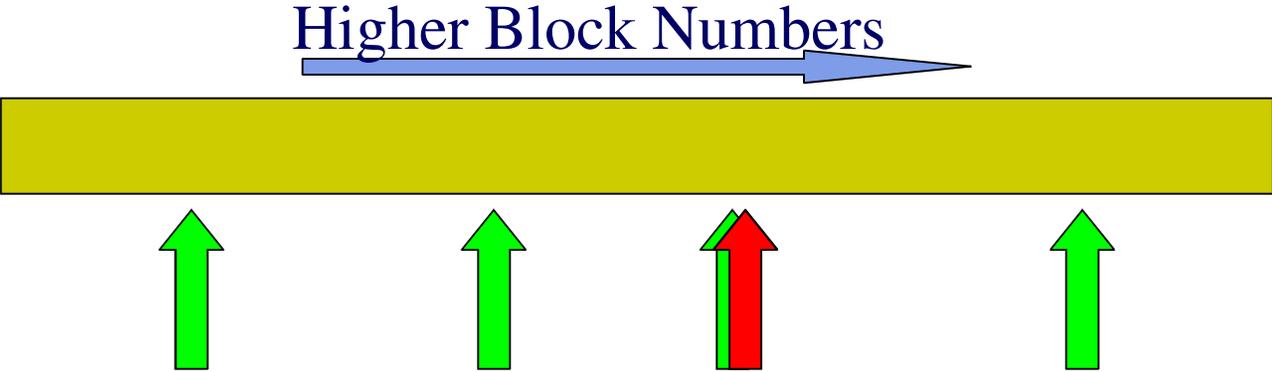
C-LOOK



C-LOOK



C-LOOK



Shortest Positioning Time First (SPTF)

Key observation

- Seek time takes a while
- But rotation time is comparable!
 - Short seeks are faster than whole-disk rotations
- What matters is *positioning* time, not seek time

SPTF is like SSTF

- Serve “temporally nearest” sector next

Challenge

- Can't estimate by subtracting sector numbers
- Must know rotation position of disk in real time!

Performs better than SSTF, but still starves requests

Weighted Shortest Positioning Time First (WSPTF)

SPTF plus fairness

Requests are “aged” to prevent starvation

- Compute “temporal distance” to each pending request
- Subtract off “age factor”
- Result: sometimes serve old request, not closest request

Various aging policies possible, many work fine

Excellent performance

Like SPTF, hard for OS to know disk status in real time

- On-disk schedulers can manage this, though...

Conclusions

Disks are complicated

Disks are very slow

FCFS is a very bad idea

- C-LOOK is ok in practice
- Disks probably do something like WSPTF internally