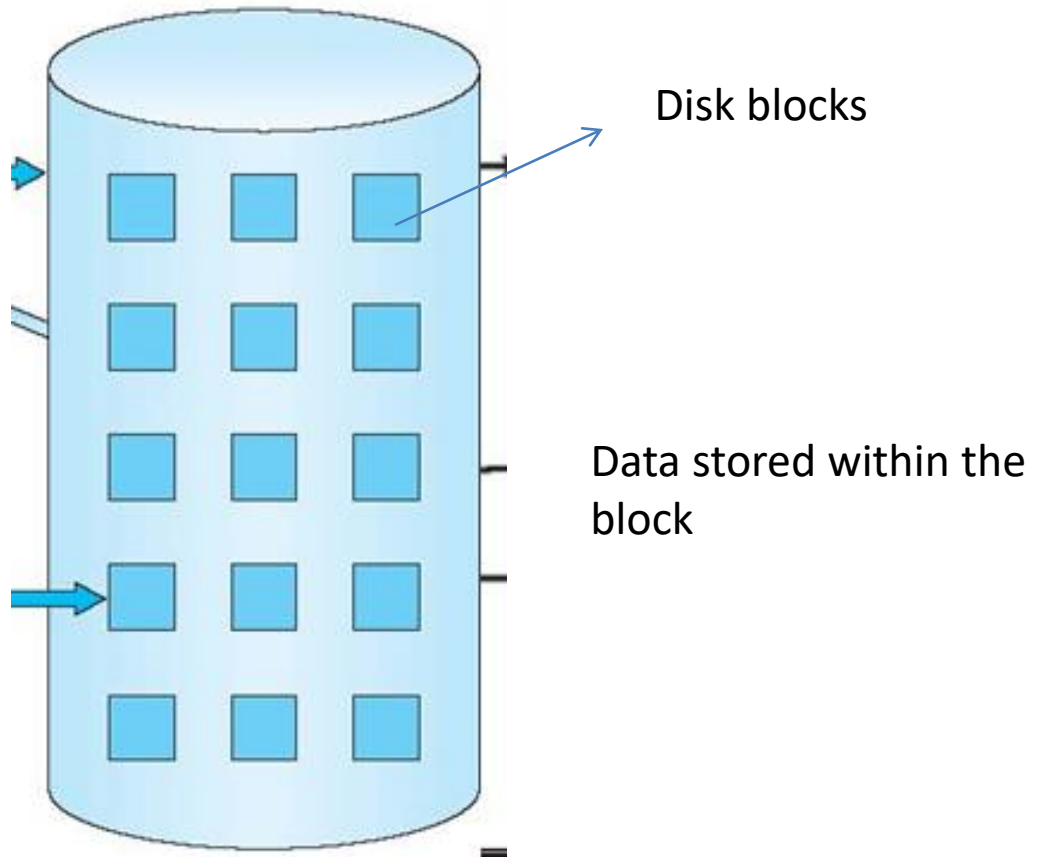


# File Management

# Objectives

- ***File Systems*** : File system structure, allocation methods (contiguous, linked, indexed), free-space management (bit vector, linked list, grouping), directory implementation (linear list, hash table)
- ***Disk Management*** : disk structure, disk scheduling (FCFS, SSTF, SCAN,C-SCAN) , disk reliability, disk formatting, boot block, bad blocks.




# File-System Structure

- File system
  - Provide efficient and convenient access to disk
  - Easy access to the data (store, locate and retrieve)
- Two aspects
  - User's view
    - Define files/attributes, operations, directory
  - Implementing file system
    - Data structures and algorithms to map logical view to physical one

# Disk Layout


- Files stored on disks.
- Disks broken up into one or more partitions, with separate file system on each partition

# A Typical File-system Organization

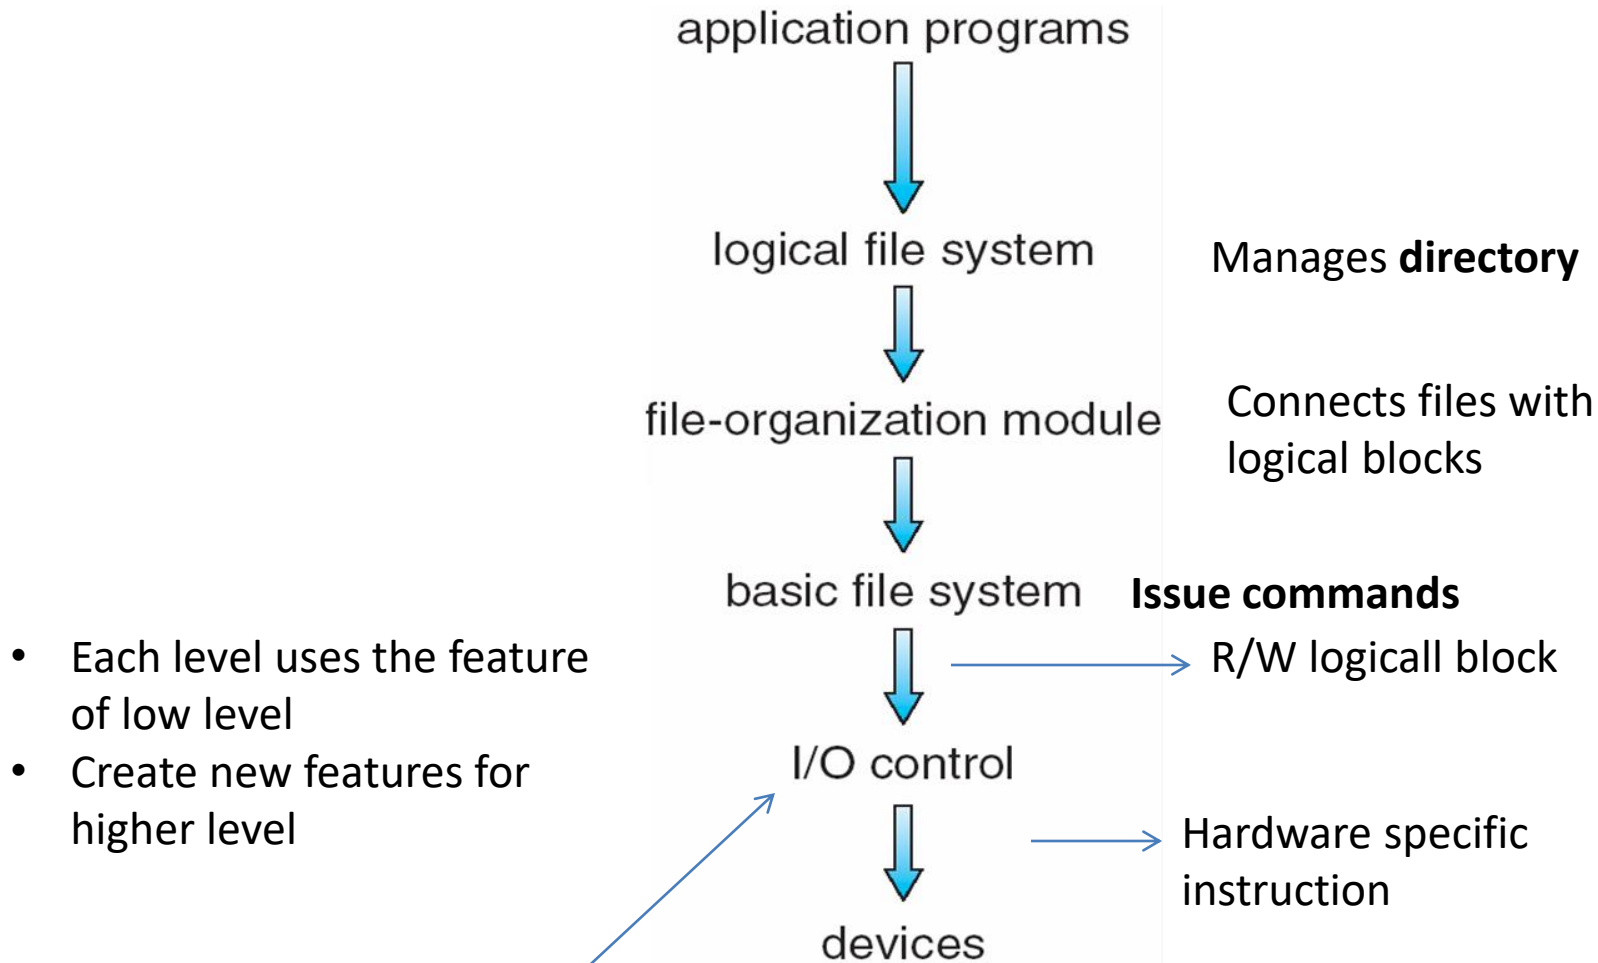


# Directory Structure

- A collection of nodes containing information about all files



# Layered File System



- Each level uses the feature of low level
- Create new features for higher level

Device driver, transfer information between memory/disk



# File System Layers

**I/O control layer** consists of device drivers manage I/O devices at the I/O control layer

- Given commands like “read block 34 into memory location 1060” outputs low-level hardware specific commands to hardware controller

**Basic file system** Issues commands with logical block address

Disk scheduling

Buffering

**File organization module** understands files, logical blocks

- Connects files with logical block #
- Manages free space, disk allocation

# File System Layers (Cont.)

- **Logical file system** manages metadata information
  - Translates file name into file number, file access, permission, location by maintaining file control blocks (**inodes** in Unix)
  - Directory management
- Layering useful for reducing complexity and redundancy, but adds overhead and can decrease performance
  - Shares the I/O control and basic FS
- Many file systems, sometimes many within an operating system
  - Each with its own format (CD-ROM is ISO 9660; Unix has **UFS**, **FFS**; Windows has FAT, FAT32, NTFS as well as floppy, CD, DVD, Linux has more than 40 types, with **extended file system** ext2 and ext3 leading; plus distributed file systems, etc)

# A Typical File Control Block

file permissions

file dates (create, access, write)


file owner, group, ACL

file size

file data blocks or pointers to file data blocks

# Directory Structure

- A collection of nodes containing information about all files




# File system data structures

- On disk
- In memory

# Disk Layout

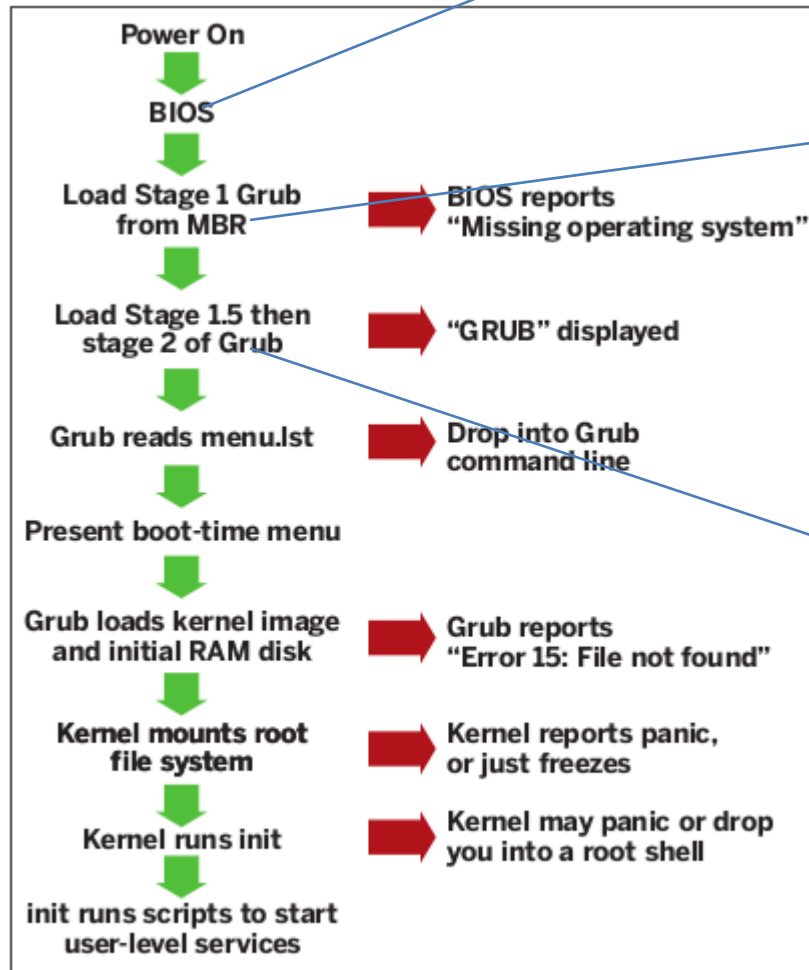
- Files stored on disks. Disks broken up into one or more partitions, with separate file system on each partition
- Sector 0 of disk is the Master Boot Record
- Used to boot the computer
- End of MBR has partition table. Has starting and ending addresses of each partition.
- One of the partitions is marked active in the partition table

# Disk Layout



# Booting sequence

- Checks system integrity
- Loads and executes the boot loader



- Master Boot Record
- Occupies the first sector of the disk
- Specifies the disk blocks numbers where second stage boot loader resides


Grand Unified Boot Loader



# Disk Layout

- **Boot control block** contains info needed by system to boot OS from that partition
  - Needed if partition contains OS, usually first block of partition
- **Partition control block (superblock, master file table)** contains partition details
  - Total # of blocks, # of free blocks, block size, free block pointers or array
- Directory structure organizes the files
  - Names and FCBs
- Per-file **File Control Block (FCB)** contains many details about the file
  - Unix (UFS) Inode number, permissions, size, dates
  - Windows (NTFS) stores into in master file table using relational DB structures

# Disk Layout



# A Typical File Control Block

file permissions

file dates (create, access, write)

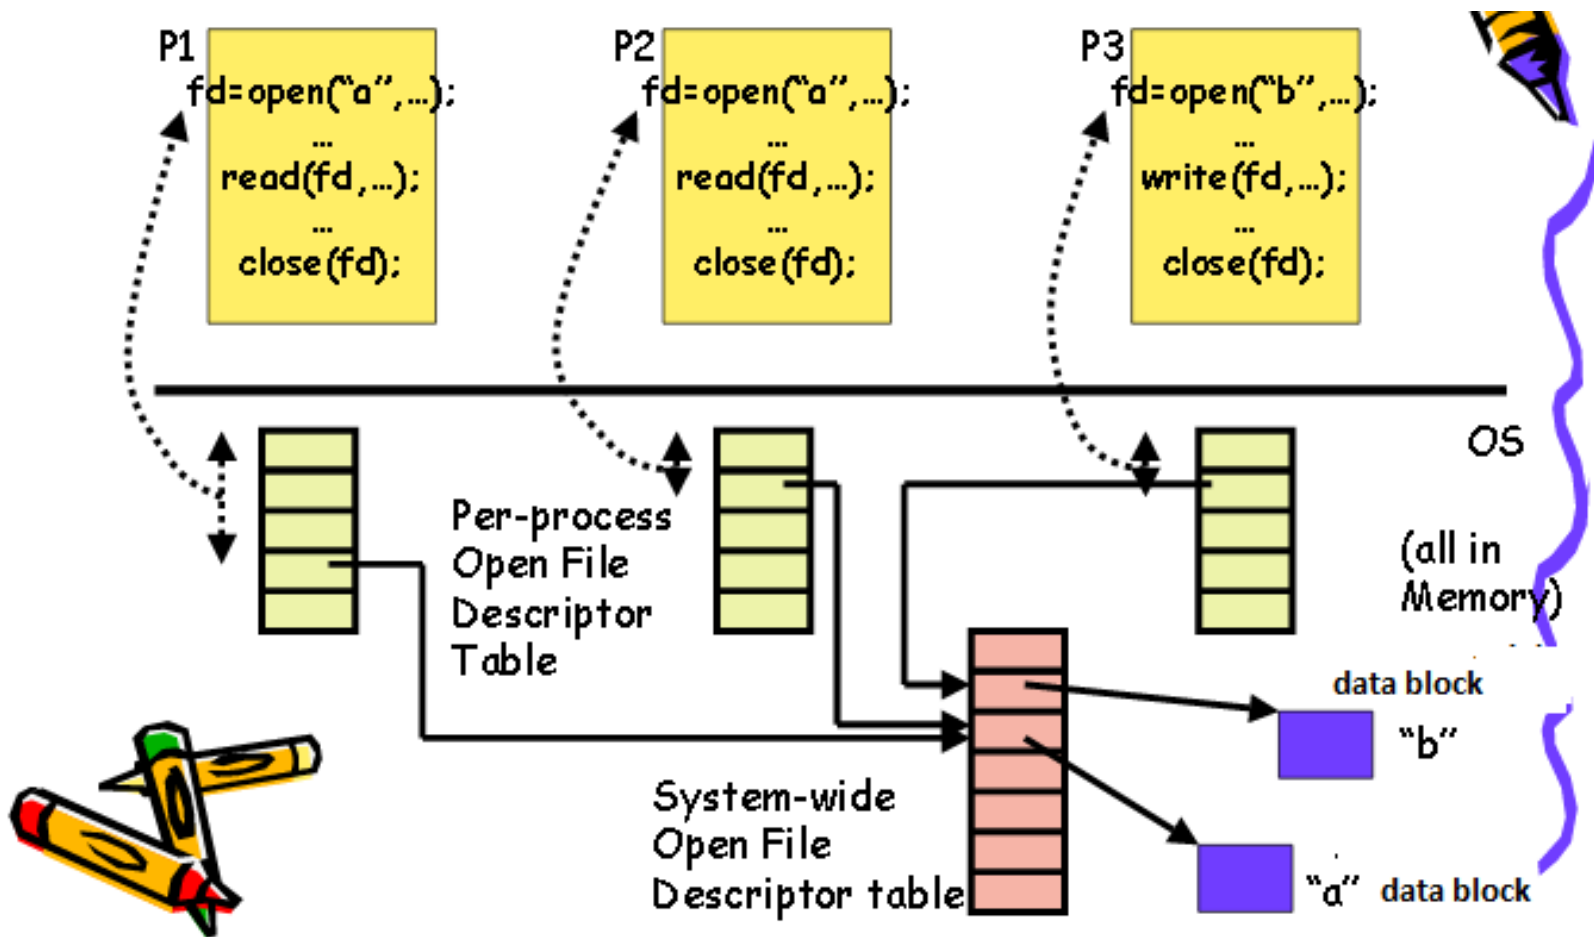
file owner, group, ACL

file size

file data blocks or pointers to file data blocks

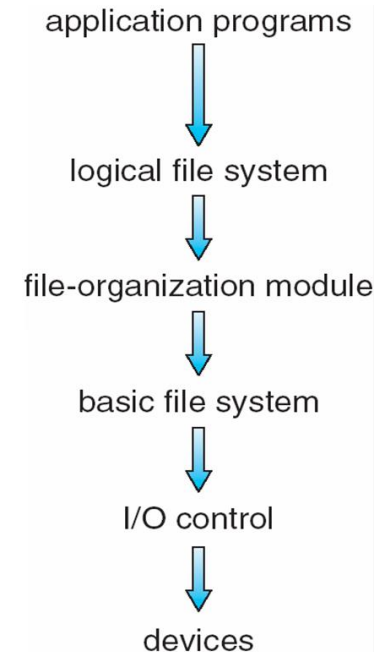
# In-Memory File System Structures


- In memory directory structure holds the directory information of recently accessed directories
- **System-wide-open file** contains a copy of FCB for each opened file
- **Per-process open file table**: contains pointer to appropriate entry in the **system wide open file table**





# File handling

- Create a new file
  - Application program calls the logical file system
- Logical file system
  - Allocates a new FCB
  - Reads the appropriate directory into memory
  - Updates directory with new filename and FCB
  - Write it back to disk
- Using the file (I/O)
  - Open() [filename]
  - Directory is searched
  - FCB is copied into **system wide open file table**
  - Entry made to **Per-process open file table**
    - Pointer to the system table entry
    - File descriptor#





# In-Memory File System Structures





- A process closes a file
  - Per process table entry removed
  - System table count decremented
- All processes closed the file
  - Updated file info is copied back to disk
  - System wide open file table entry removed

# Directory Implementation

- **Linear list** of file names with pointer to the data blocks
  - Simple to program
  - Time-consuming to execute
    - Linear search time
  - New file creation / deletion
  - Linked list
- Cache the frequently accessed entry
- Binary search to speedup directory search
  - Could keep ordered alphabetically
  - or use B+ tree

- **Hash Table** – hash data structure
  - Hash value computed from filename
  - Decreases directory search time
  - Insertion and deletion simple
  - **Collisions** – situations where two file names hash to the same location
    - Chaining
- Hash table of fixed size
- Performance depends on hash function

# Allocating Blocks to files


An allocation method refers to how disk blocks are allocated for files

- Most important implementation issue
- Methods
  - Contiguous allocation
  - Linked list allocation
  - Linked list using table
  - Indexed

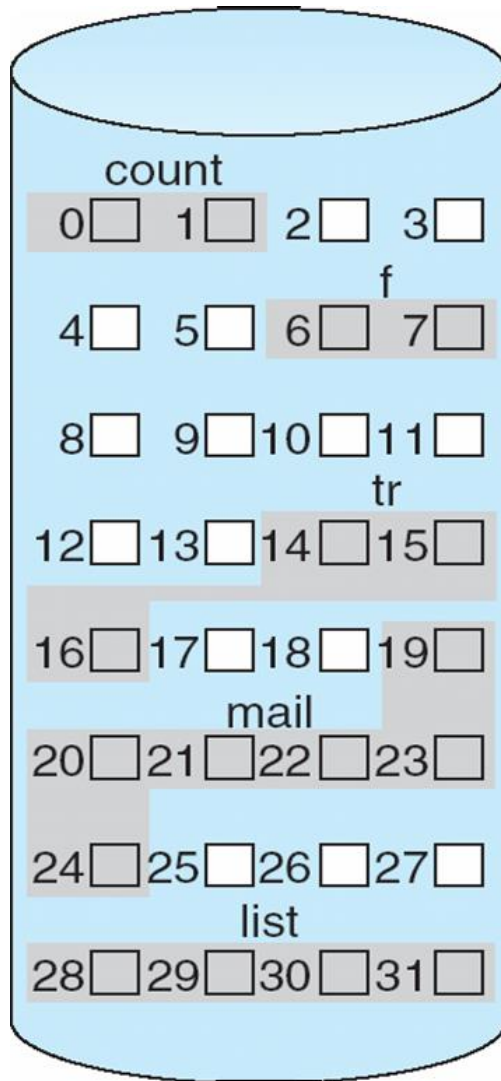
# Allocation Methods - Contiguous

- **Contiguous allocation** – each file occupies set of contiguous blocks
- Blocks are allocated  $b, b+1, b+2, \dots$ 
  - Best performance in most cases
  - Simple – only starting location (block #) and length (number of blocks) are required (directory)
- Easy to implement
- Read performance is great. Only need one seek to locate the first block in the file. The rest is easy.
- Accessing file is easy
  - Minimum disk head movement
  - Sequential and direct access

# Moving-head Disk Mechanism



# Contiguous Allocation of Disk Space



directory

file	start	length
count	0	2
tr	14	3
mail	19	6
list	28	4
f	6	2


- Problems
- Finding space for file
  - Satisfy the request of size  $n$  from the list of holes
  - External fragmentation
    - Need for **compaction routine**
    - **off-line (downtime)** or **on-line**
- Do not know the file size a priori
  - Terminate and restart
  - Overestimate
  - Copy it in a larger hole
  - Allocate new contiguous space (Extent)




# Extent-Based Systems

- Here, a contiguous chunk of space is allocated initially.
- Then, if that amount proves not to be large enough
  - another chunk of contiguous space, known as an **extent**, is added.
- The location of a file's blocks is then recorded as a
  - First block and a block count,
  - plus a link to the first block of the next extent.
- Many newer file systems (i.e., Veritas File System) use a modified contiguous allocation scheme
- A file consists of one or more extents

# Contiguous Allocation



(a)




(b)

(a) Contiguous allocation of disk space for 7 files.

(b) The state of the disk after files D and F have been removed.

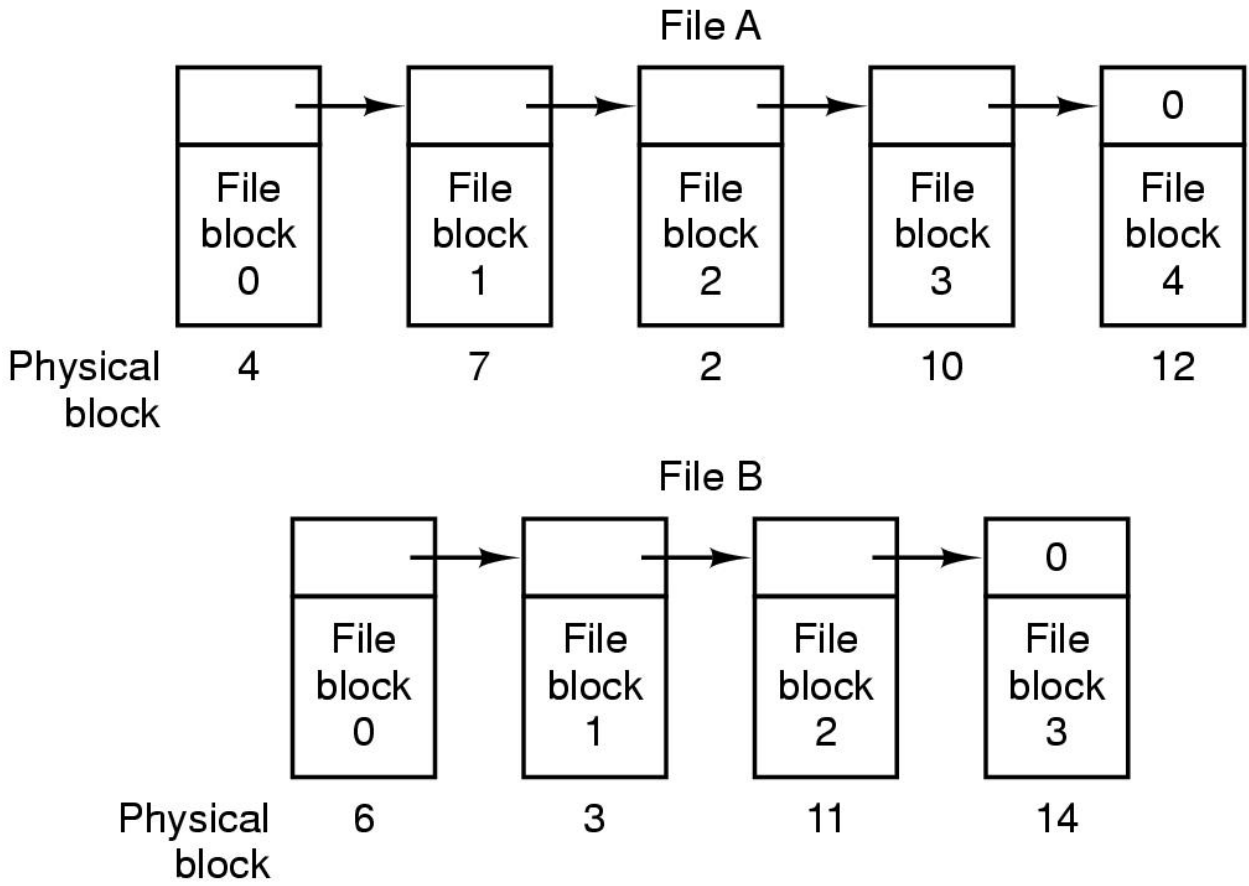
# Linked Allocation

- Each file is a linked list of disk blocks: blocks may be scattered anywhere on the disk




- Each block contains pointer to next block
- File ends at nil pointer

# Linked List Allocation



Storing a file as a linked list of disk blocks.

# Linked Allocation



# Linked Allocation

- Free blocks are arranged from the free space management
- No external fragmentation
- Files can continue to grow

## Disadvantage

1. Effective only for sequential access  
Random/direct access (i-th block) is difficult

2. Space wastage

If block size 512 B

Disk address 4B

Effective size 508B

3. Reliability

Lost/damaged pointer

Bug in the OS software and disk hardware failure

Double linked


## Solution: Clusters

- Improves disk access time (head movement)
- Decreases the link space needed for block
- Internal fragmentation

# Allocation Methods - Linked

- **Linked allocation** – each file a linked list of blocks
  - No compaction, external fragmentation
  - Free space management system called when new block needed
  - Reliability can be a problem
  - Locating a block can take many I/Os and disk seeks

# Moving-head Disk Mechanism





# Linked Allocation

## FAT (File Allocation Table) variation

Beginning of partition has table, indexed by block number

Much like a linked list, but faster on disk and cacheable


New block allocation simple

Section of the disk at the beginning of the partition contains FAT table

Unused blocks => 0



**Many head movements**

**Better random access**



# File-Allocation Table

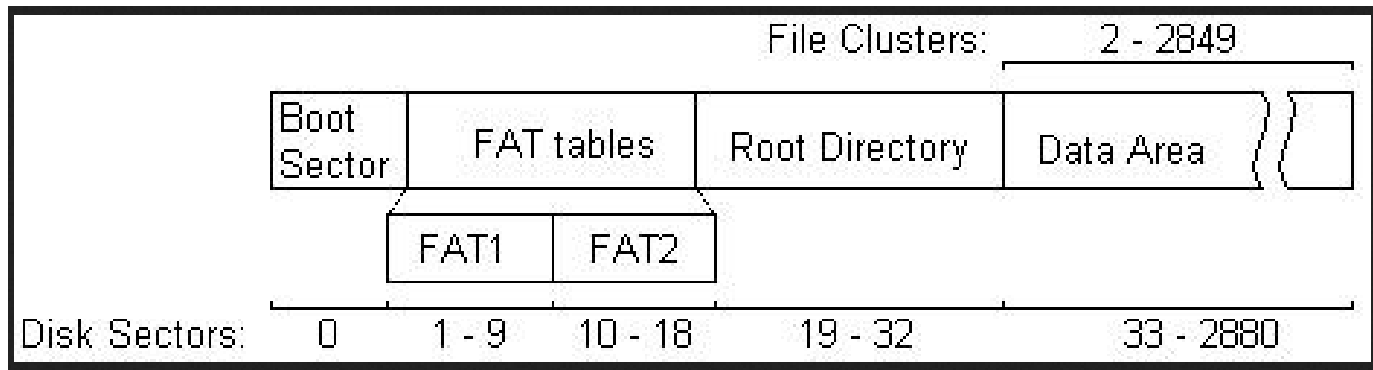
directory entry



MS DOS

Caching of FAT16

# DOS




# Allocation Methods - Indexed

- **Indexed allocation**

- Each file has its own **index block**(s) of pointers to its data blocks

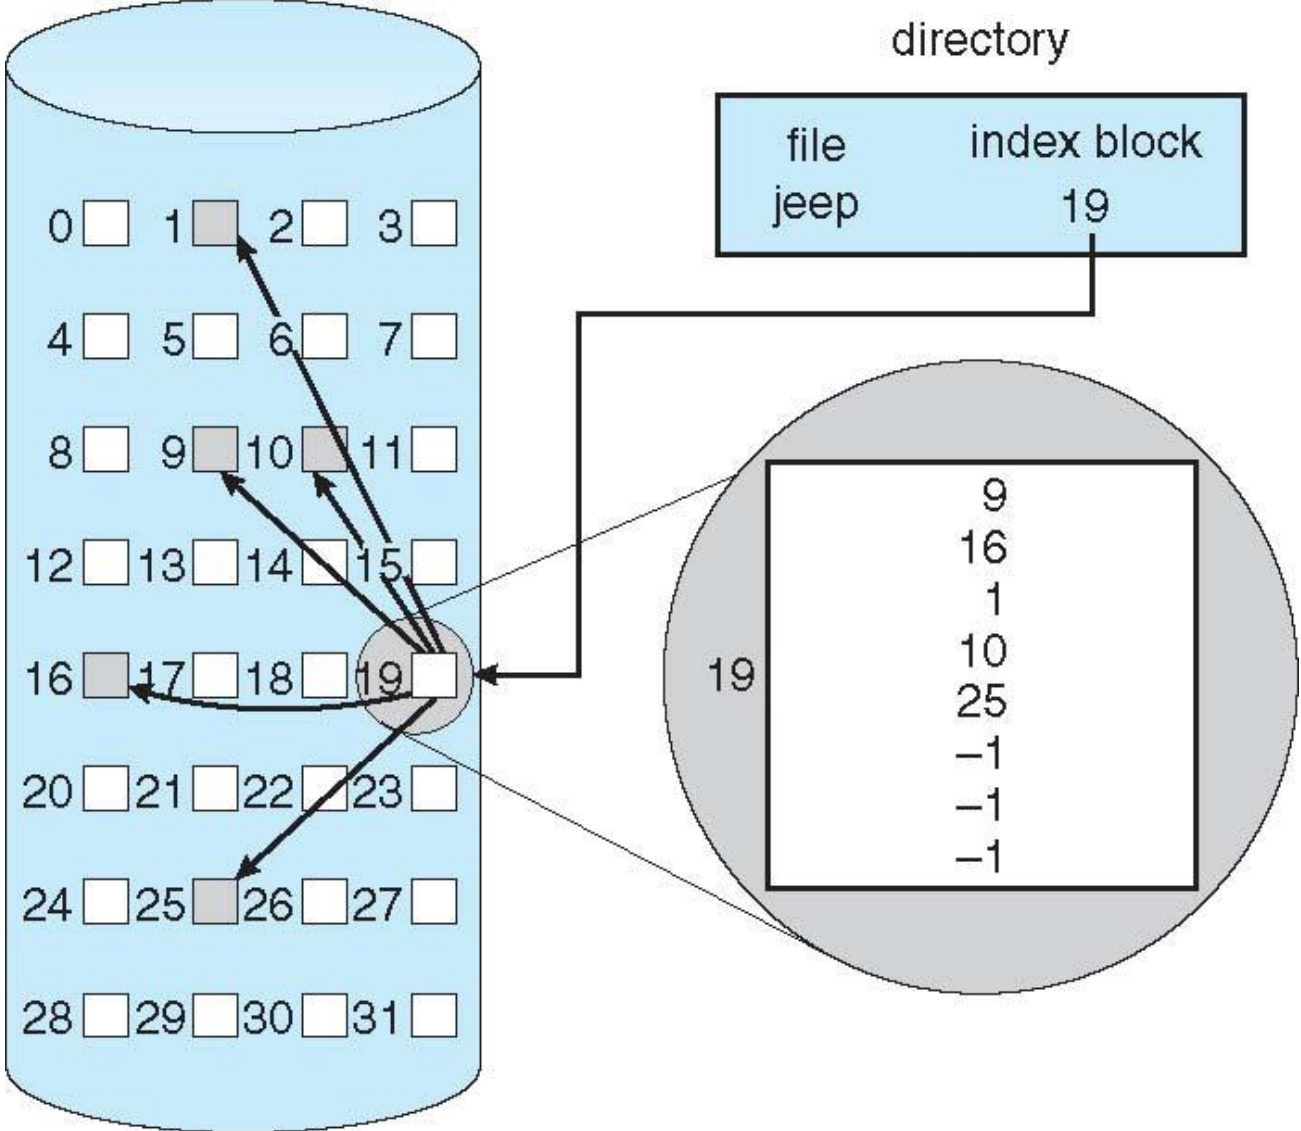
- Directory contains address of the index block



index table

- Logical view

# Example of Indexed Allocation



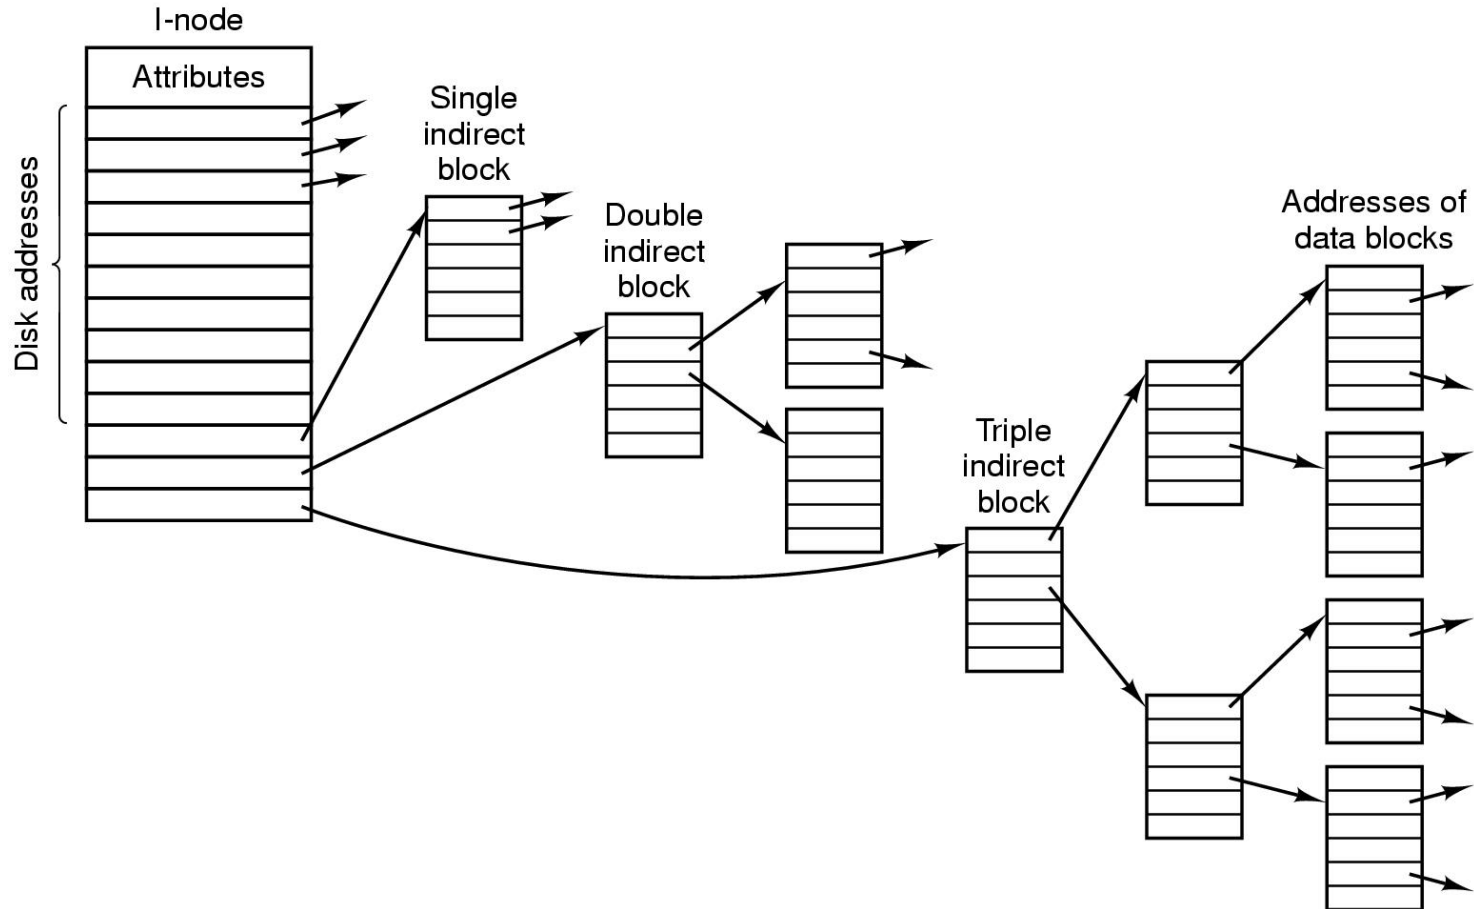
# Indexed Allocation

- Efficient random access without external fragmentation,
- Size of index block
  - One data block
- Overhead of index block
  - Wastage of space
  - Small sized files

# Linked scheme

- Linked scheme – Link blocks of index table (no limit on size)
- Multilevel index

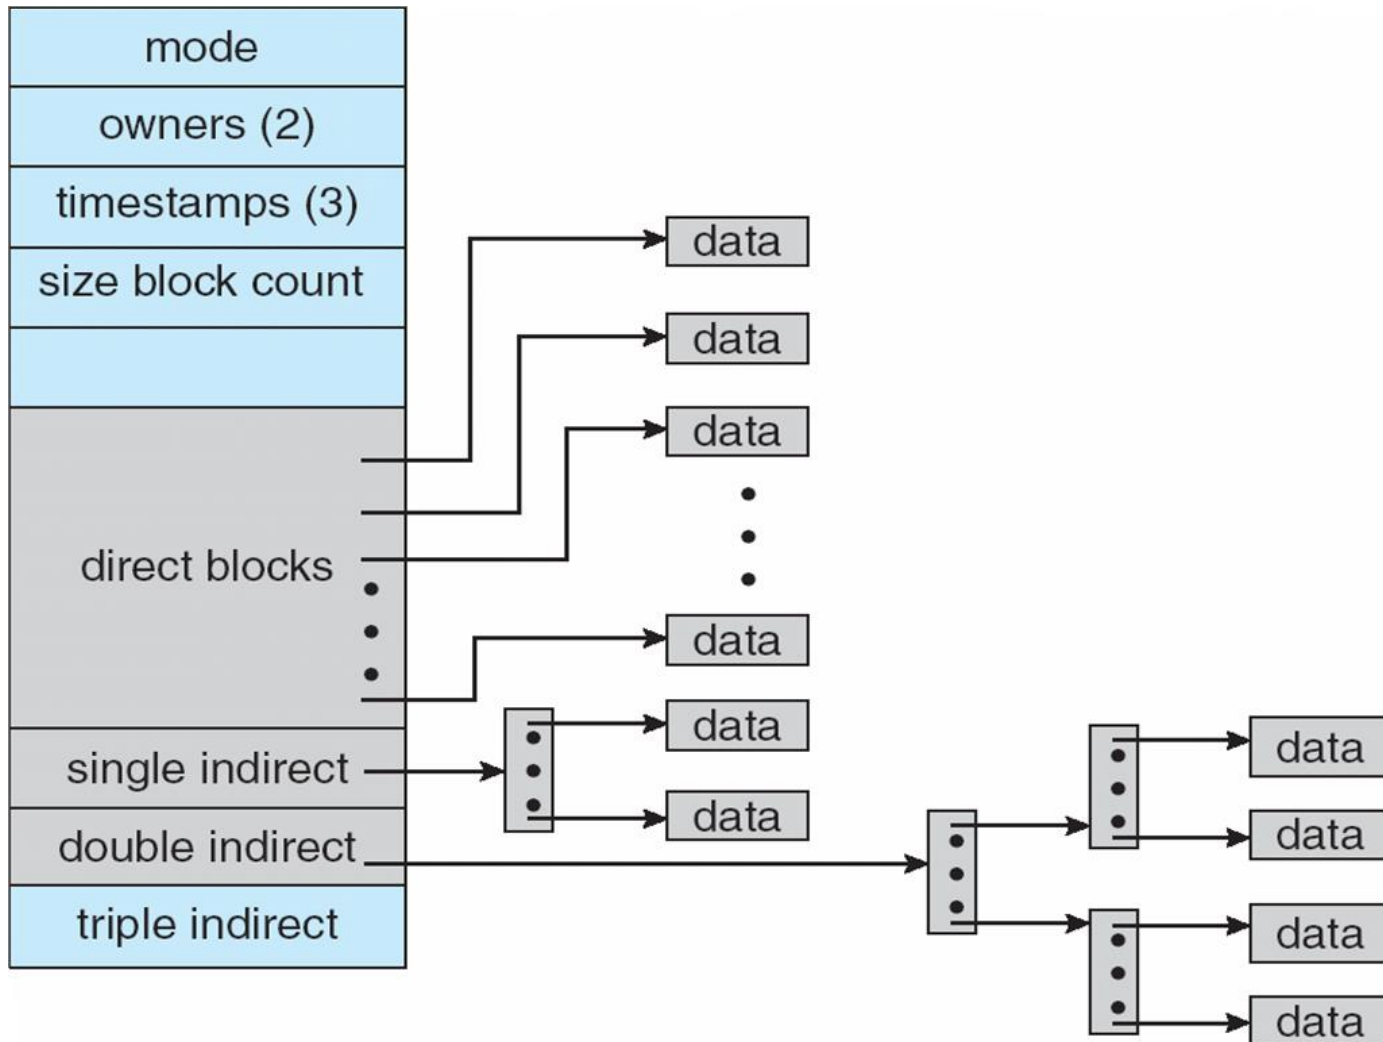
# The UNIX File System



A UNIX i-node.




# Combined Scheme: UNIX UFS (4K bytes per block, 32-bit addresses)




# File-Allocation Table

directory entry



FAT16

16 bits



MS DOS

Maximum size of the disk supported by FAT26?

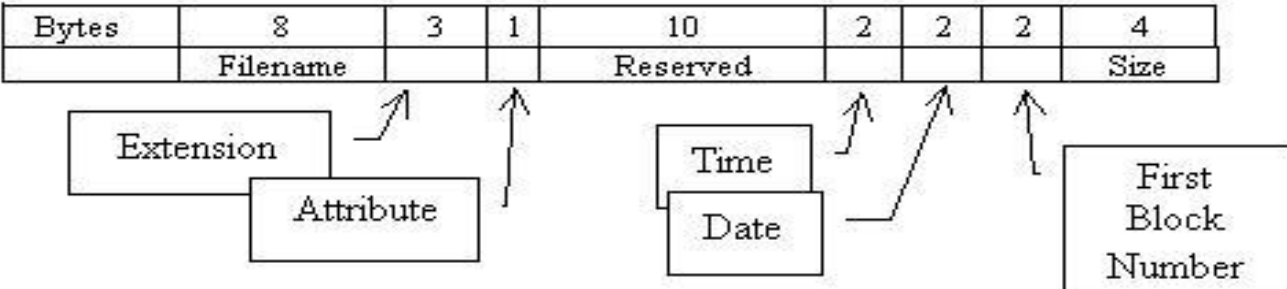
no. of disk blocks

FAT

# Implementing Directories


- OS uses path name supplied by the user to locate the directory entry
- Stores attributes
- Directory entry specifies block addresses by providing
  - Number of first block (contiguous)
  - Number of first block (linked)
  - Number of i-node

# Implementing MS DOS Directories



games	attributes
mail	attributes
news	attributes
work	attributes

(a)




(b)

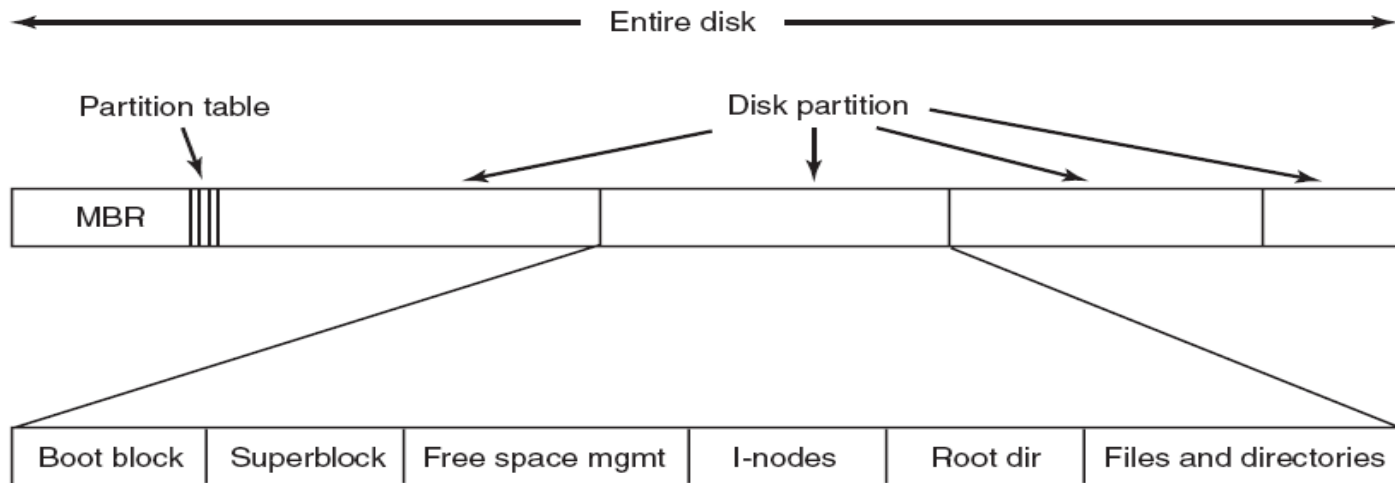
Data structure containing the attributes

Each entry 32 bytes long

# The UNIX File System



# Disk Layout



# The UNIX File System

Root directory

1	.
1	..
4	bin
7	dev
14	lib
9	etc
6	usr
8	tmp

Looking up  
usr yields  
i-node 6

I-node 6  
is for /usr

Mode size times
132

I-node 6  
says that  
/usr is in  
block 132

Block 132  
is /usr  
directory

6	.
1	..
19	dick
30	erik
51	jim
26	ast
45	bal

/usr/ast  
is i-node  
26

I-node 26  
is for  
/usr/ast

Mode size times
406

I-node 26  
says that  
/usr/ast is in  
block 406

Block 406  
is /usr/ast  
directory


26	.
6	..
64	grants
92	books
60	mbox
81	minix
17	src

/usr/ast/mbox  
is i-node  
60

The steps in looking up */usr/ast/mbox*.

# Free-Space Management

- File system maintains **free-space list** to track available blocks
- **Bit vector** or **bit map** ( $n$  blocks)
- Each block is represent by 1 bit



$$\text{bit}[i] = \begin{cases} 1 \Rightarrow \text{block}[i] \text{ free} \\ 0 \Rightarrow \text{block}[i] \text{ occupied} \end{cases}$$

Simple and Efficient to find first free blocks or  $n$  consecutive free blocks

**Mac**

CPUs have instructions to return offset within word of first “1” bit



Bit map

Block number calculation=

(number of bits per word) \*  
(number of 0-value words) +  
offset of first 1 bit



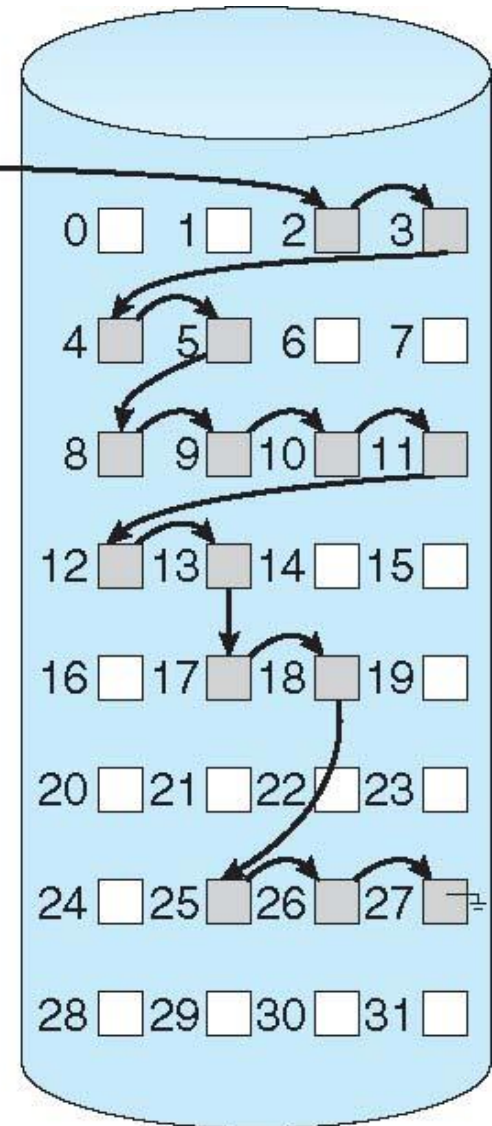
# Free-Space Management

- Bit map requires extra space
  - Example:
    - block size = 4KB =  $2^{12}$  bytes
    - disk size =  $2^{40}$  bytes (1 terabyte)
    - Number of blocks  $n = 2^{40}/2^{12} = 2^{28}$  bits (or 256 MB)
    - if clusters of 4 blocks -> 64MB of memory
- Keep the vector in main memory

# Linked Free Space List on Disk

- Link together all the free disk blocks
- Keep a pointer to the first free block
- Cannot get contiguous space easily
  - Traverse the list
- Generally require first free block

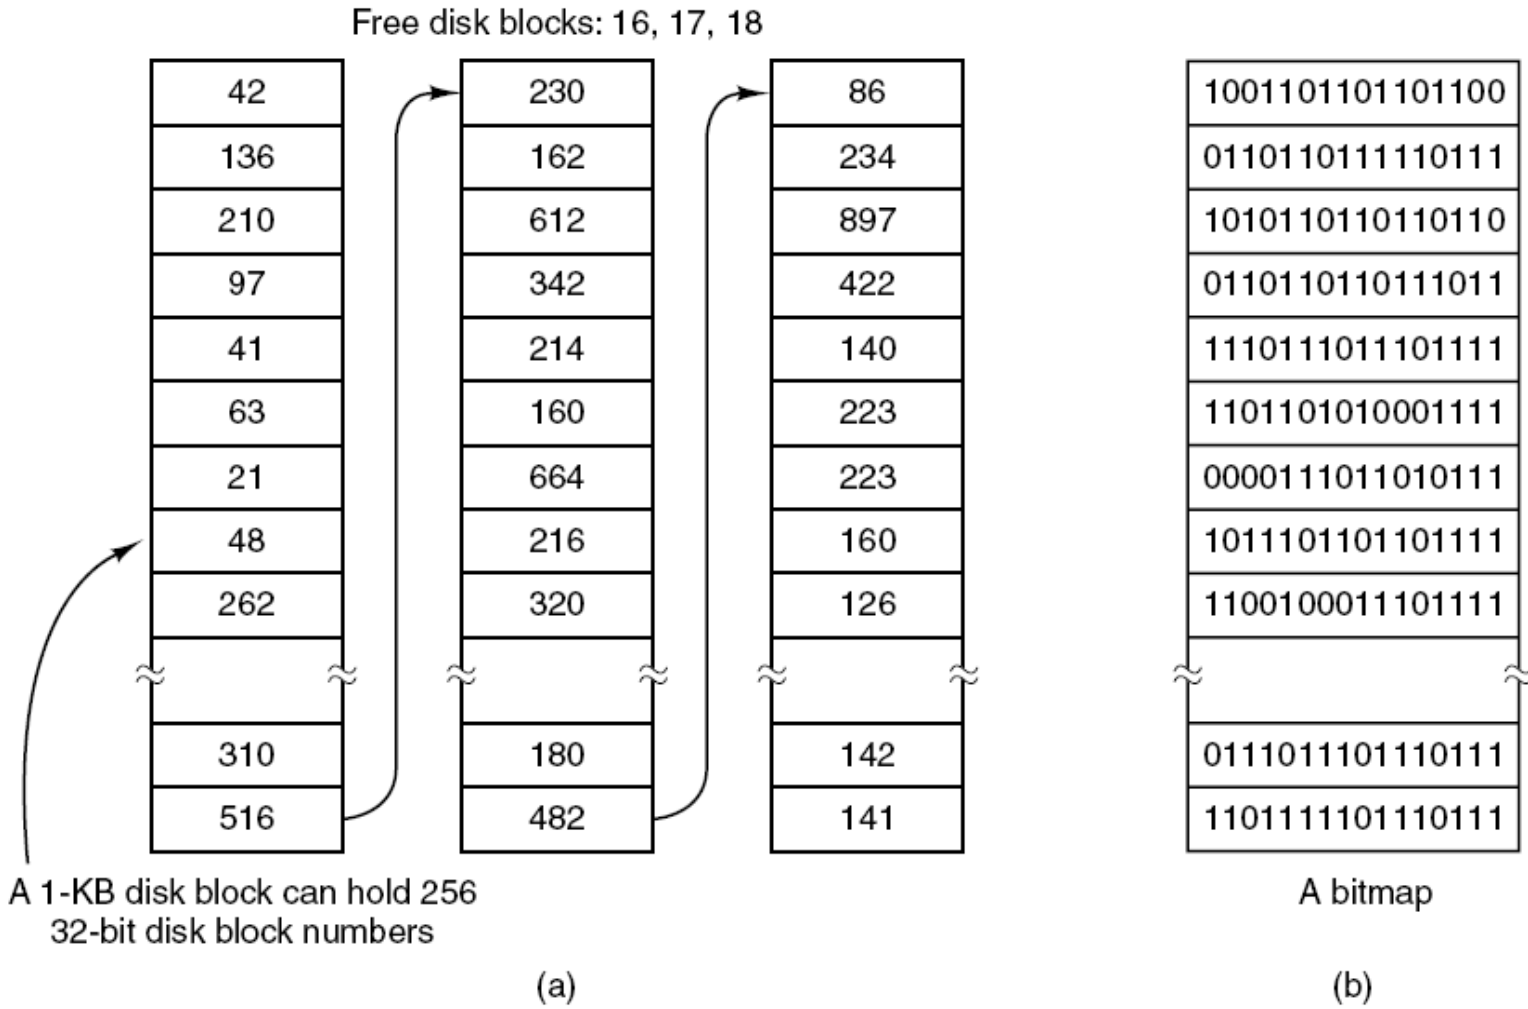
free-space list head



# Free-Space Management

- Grouping
  - Reserve few disk blocks for management
  - Modify linked list to store address of next  $n-1$  free blocks in first free block, plus a pointer to next block that contains free-block-pointers
- Counting

# Keeping Track of Free Blocks (1)



Storing the free list using (a) Grouping (b) Bitmap.

1KB block

16 bits block number


Each block holds 511 free blocks

20M disk needs 40 blocks for free list


How many for bit map?

# Free Space for FAT

directory entry




MS DOS



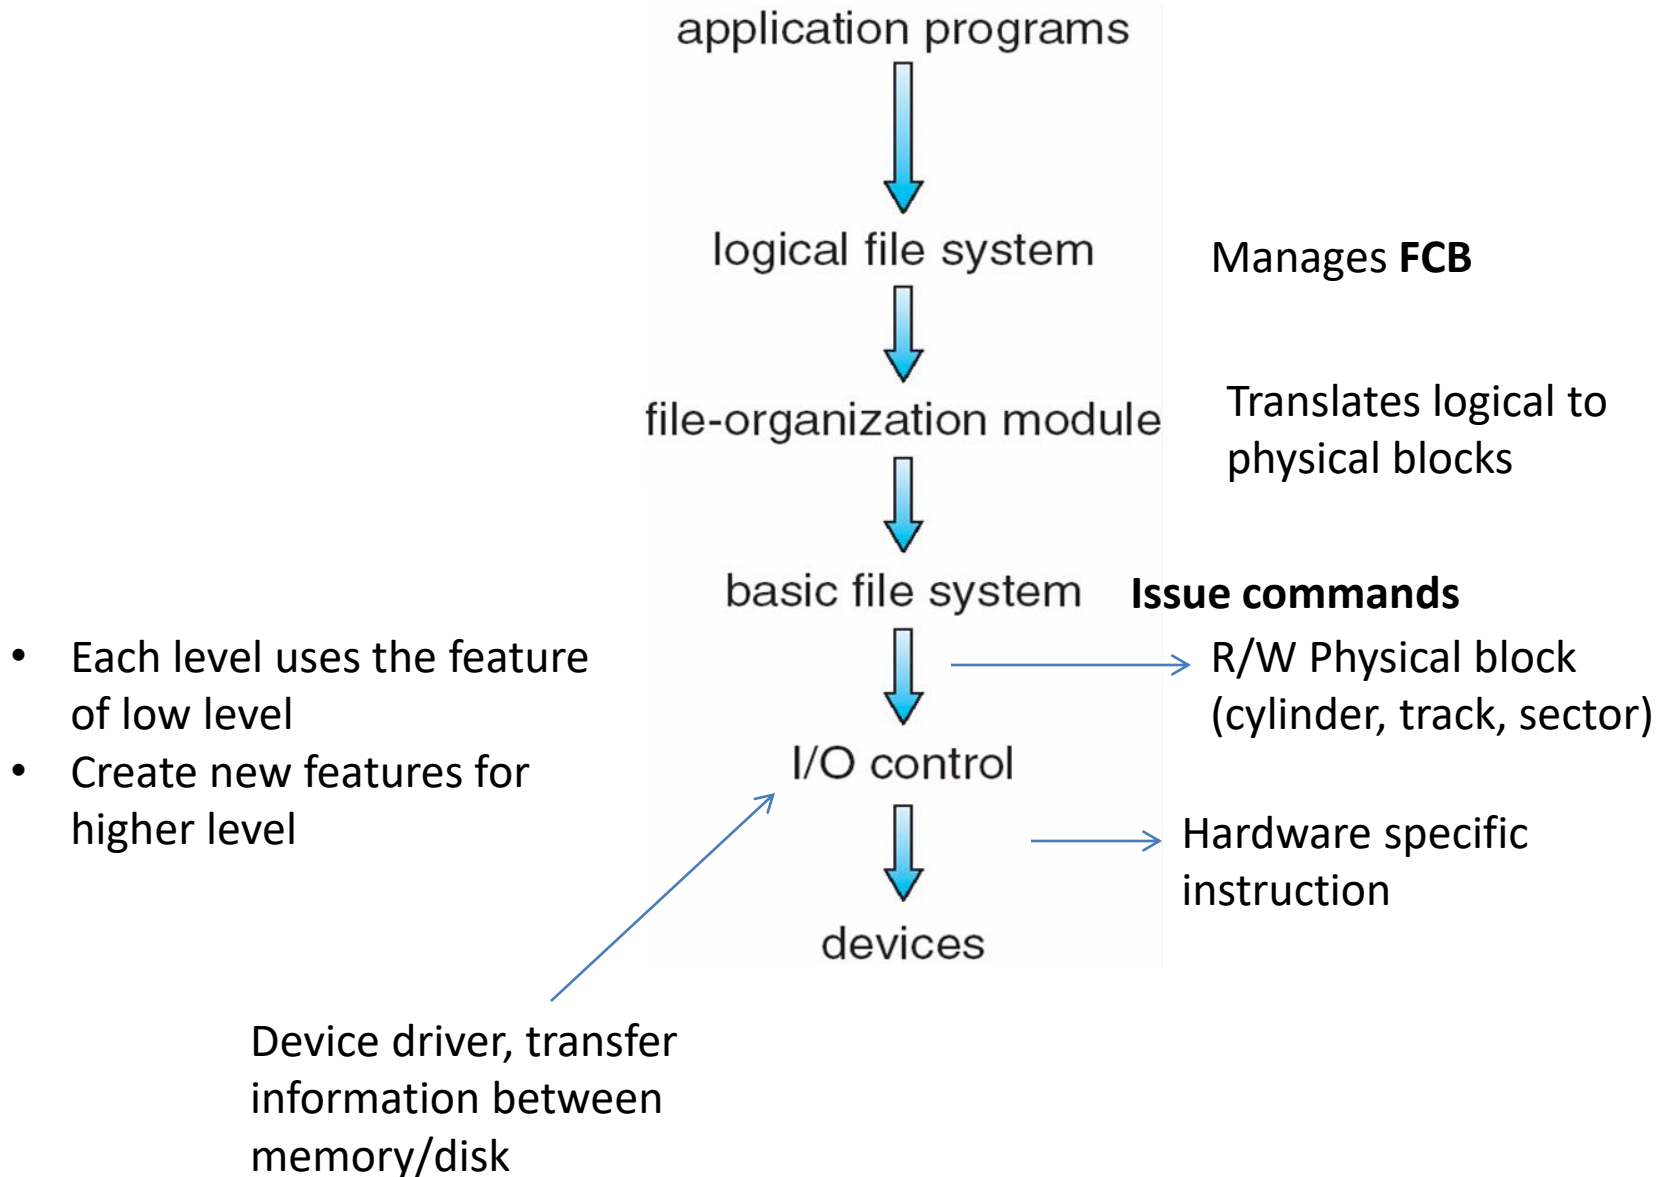
# Demand Paging Example

- Memory access time = 200 nanoseconds
- Average page-fault service time = 8 milliseconds
- $EAT = (1 - p) \times 200 + p (8 \text{ milliseconds})$   
 $= (1 - p) \times 200 + p \times 8,000,000$   
 $= 200 + p \times 7,999,800$
- If one access out of 1,000 causes a page fault, then  
EAT = 8.2 microseconds.  
This is a slowdown by a factor of 40!!
- If want performance degradation < 10 percent
  - $220 > 200 + 7,999,800 \times p$   
 $20 > 7,999,800 \times p$
  - $p < .0000025$
  - < one page fault in every 400,000 memory accesses

**Better utilization of swap space**



# Layered File System

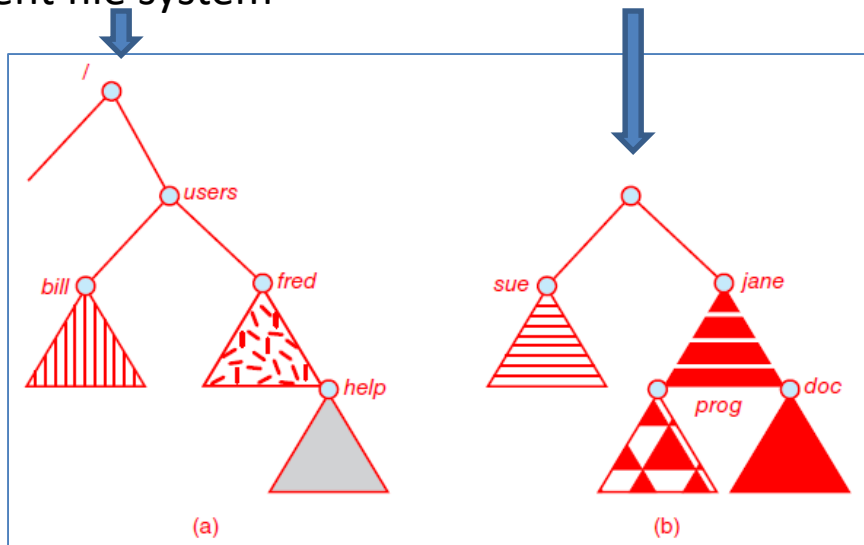




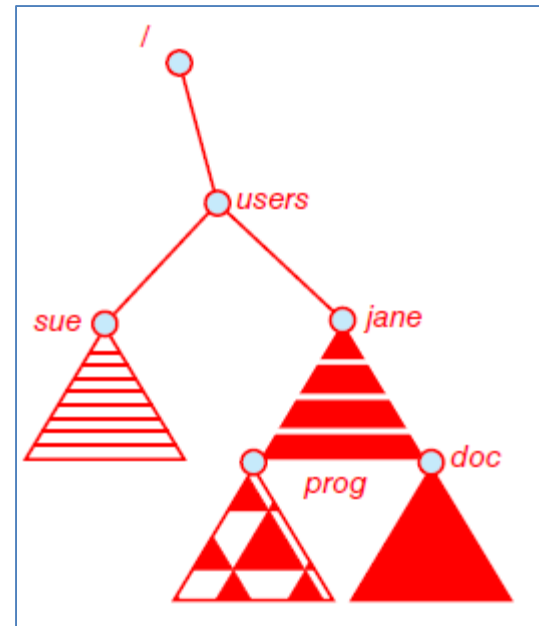
# Mounting

- **Partition** must be mounted with **file system** before it can be available to processes on the system
- The operating system is given the **name of the disk** and the **mount point**—the location within the file structure where the partition is to be attached

Current file system



Effects of mounting the partition over /users



# Various possibilities

- Systems impose restrictions for mounting.
- For example, a system may disallow a mount over a directory that contains files
- It may make the mounted file system available at that directory and hide the directory's existing files until the file system is unmounted
- System may allow the same file system to be mounted repeatedly, at different mount points
- it may only allow one mount per file system.

## **Example Mac**

- Whenever the system encounters a disk for the first time (either at boot time or while the system is running),
- the macOS operating system searches for a file system on the device.
- If it finds one, it automatically mounts the file system under the /Volumes directory,

# Problem

Consider the organization of a Unix file represented by i-node. Assume that there are 10 direct block pointers, and one singly, doubly and triply indirect pointers in each i-node. Assume that the disk block size is 4KB. Disk block pointer is 4 bytes.

- (i). What is the maximum file size supported by the system?
- (ii) Assuming i-node of a file is available in the main memory, how many disk accesses are required to access the byte in position 54, 423,956?