

# CMSC421: Principles of Operating Systems

**Nilanjan Banerjee**

*Assistant Professor, University of Maryland*

Baltimore County

nilanb@umbc.edu

<http://www.csee.umbc.edu/~nilanb/teaching/421/>

**Principles of Operating Systems**

**Acknowledgments: Some of the slides are adapted from Prof. Mark Corner and Prof. Emery Berger's OS course at Umass Amherst**

## Announcements


- Homework 3 is out (due nov 27<sup>th</sup>)
- Travelling next week: Dr. Joshi will teach on M/W

# Access Lists and Groups

- Mode of access: read, write, execute
- Three classes of users

a) owner access	7	⇒	RWX 1 1 1
b) group access	6	⇒	RWX 1 1 0
c) public access	1	⇒	RWX 0 0 1

- Ask manager to create a group (unique name), say G, and add some users to the group.
- For a particular file (say *game*) or subdirectory, define an appropriate access.




## Access Control - chmod

- Can read bits via ls, set bits via chmod

```
elnux14> ls -l ack.scm
-rw-r----- 1 emery fac 197 Feb 25 15:19 ack.scm
elnux14> chmod -r ack.scm
elnux14> ls -l ack.scm
--w----- 1 emery fac 197 Feb 25 15:19 ack.scm
elnux14> cat ack.scm
cat: ack.scm: Permission denied
```

# Access Control Lists (ACLs) in Windows

- ACLs are more expressive
  - Specify different rights per user or group
  - Opinion: one of the biggest UNIX problems



# Access Methods

- **Sequential Access**


read next  
write next  
reset  
no read after last write  
(rewrite)

- **Direct Access**


read  $n$   
write  $n$   
position to  $n$   
    read next  
    write next  
rewrite  $n$

$n$  = relative block number

# Sequential-access File



# Example of Index and Relative Files






# Directories

- Directory - just special file
  - Contains metadata, filenames
  - Store pointers to files
- Typically **hierarchical tree**
  - odd exposure of data structure to user

# A Typical File-system Organization



## Operations Performed on Directory


- Search for a file
- Create a file
- Delete a file
- List a directory
- Rename a file
- Traverse the file system

# Organize the Directory (Logically) to Obtain

- Efficiency - locating a file quickly
- Naming - convenient to users
  - Two users can have same name for different files
  - The same file can have several different names
- Grouping - logical grouping of files by properties, (e.g., all Java programs, all games, ...)

# Single-Level Directory

- A single directory for all users




Naming problem

Grouping problem

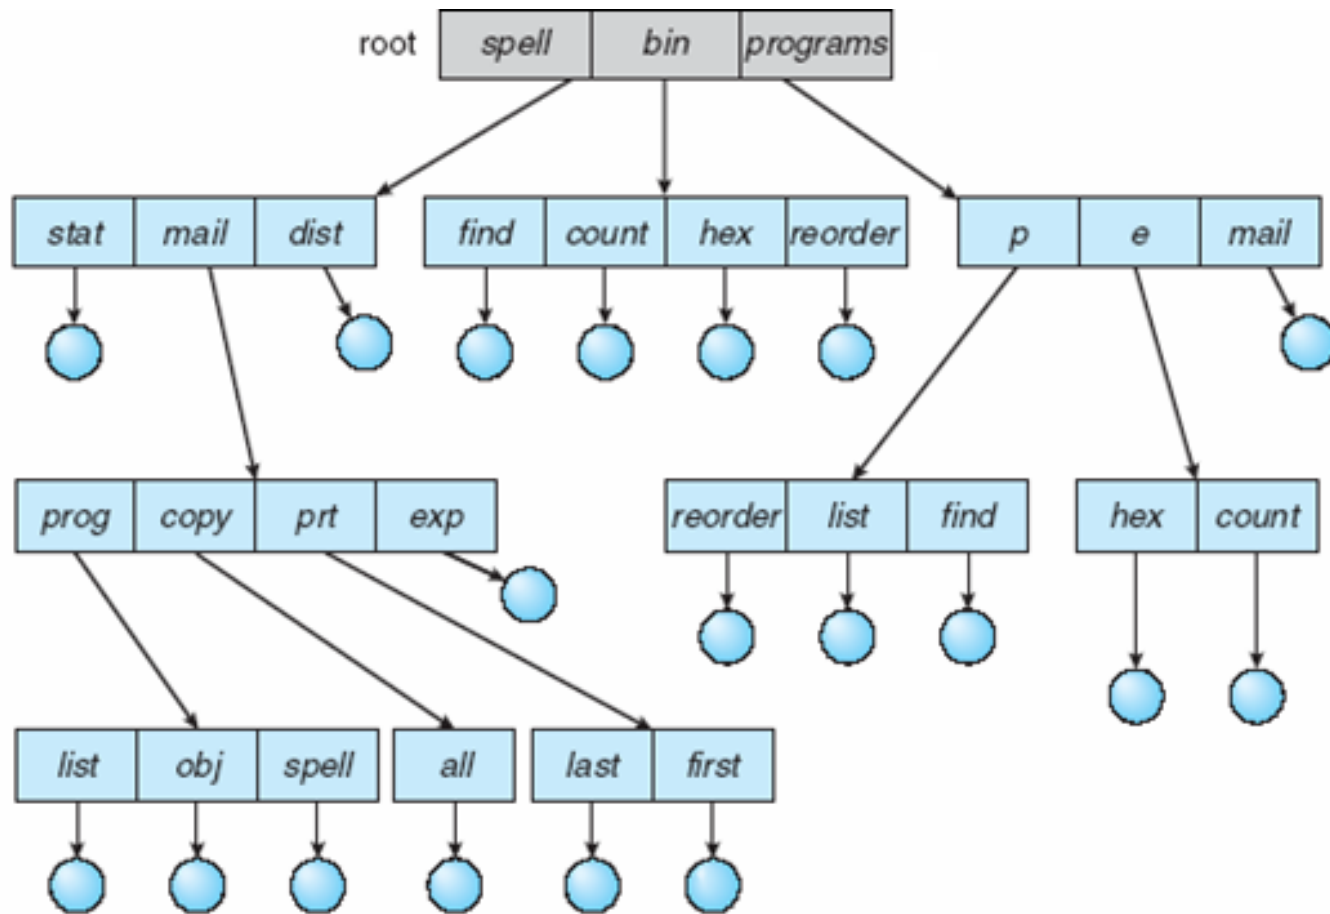
# Two-Level Directory

- Separate directory for each user



- Path name
- Can have the same file name for different user
- Efficient searching
- No grouping capability

# Tree-Structured Directories



## Tree-Structured Directories (Cont.)

- Efficient searching
- Grouping Capability
- Current directory (working directory)
  - `cd /spell/mail/prog`



## Tree-Structured Directories (Cont)

- **Absolute** or **relative** path name
- Creating a new file is done in current directory
- Delete a file


**rm <file-name>**

- Creating a new subdirectory is done in current directory

**mkdir <dir-name>**

Example: if in current directory **/mail**

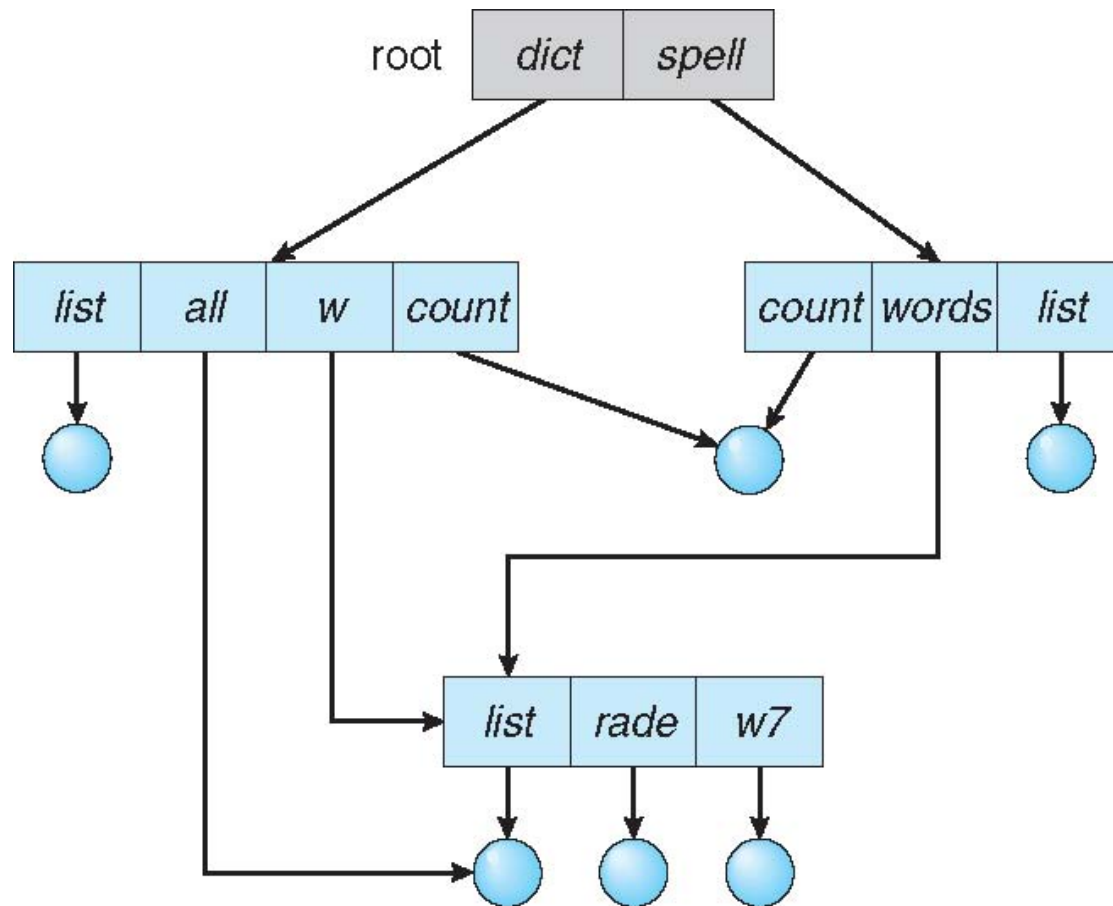
**mkdir count**



Deleting “mail” ⇒ deleting the entire subtree rooted by “mail”

# Acyclic-Graph Directories

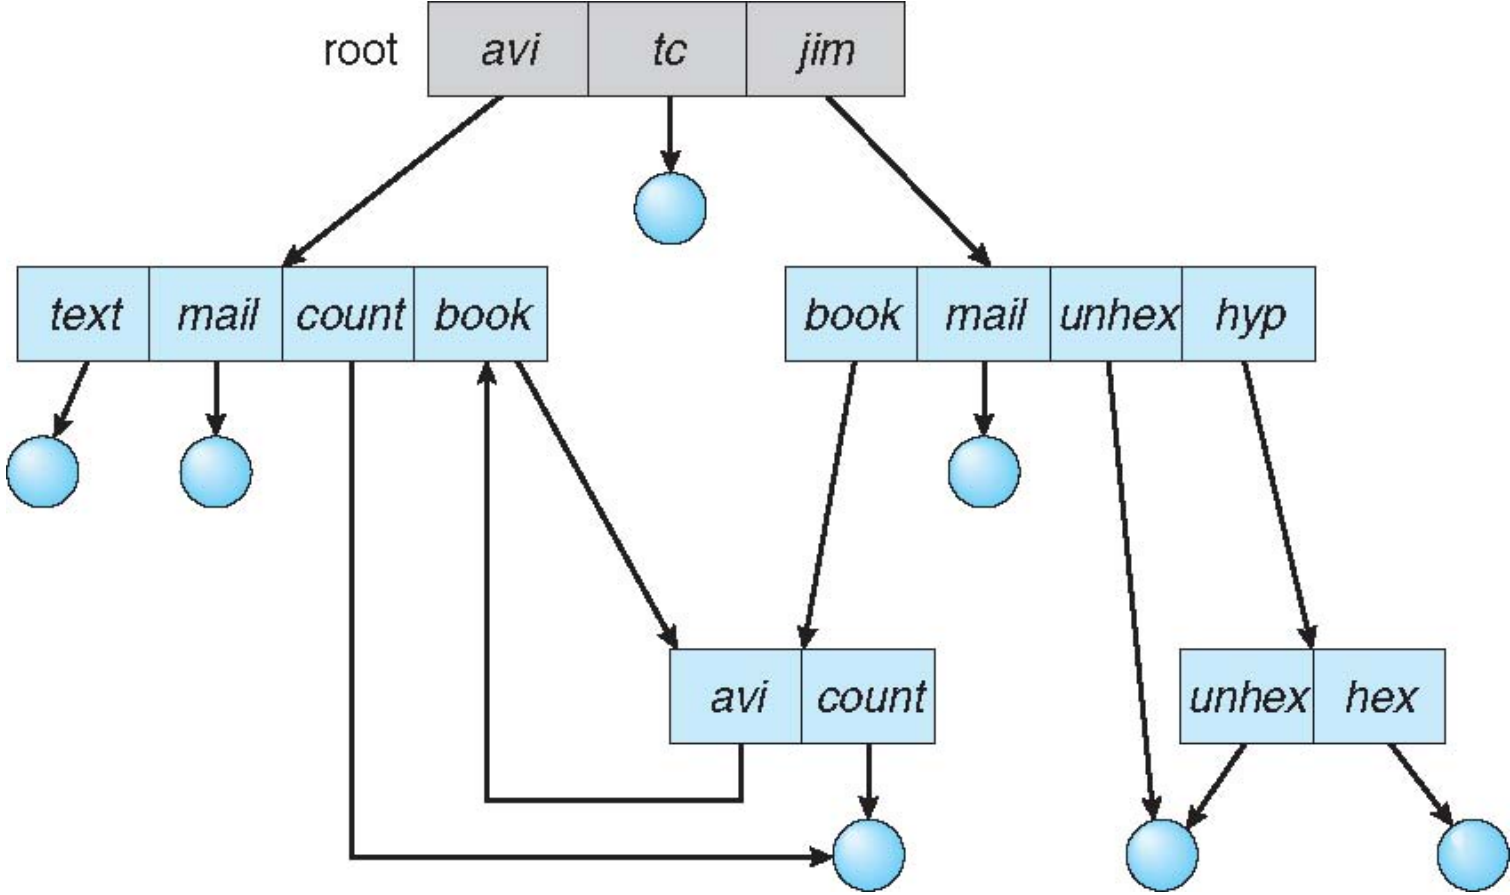
- Have shared subdirectories and files (how do you accomplish this?)



## Acyclic-Graph Directories (Cont.)

- Two different names (aliasing)
- If *dict* deletes *list*  $\Rightarrow$  dangling pointer  
Solutions:
  - Backpointers, so we can delete all pointers  
Variable size records a problem
  - Entry-hold-count solution
- New directory entry type
  - **Link** - another name (pointer) to an existing file
  - **Resolve the link** - follow pointer to locate the file

# General Graph Directory



## General Graph Directory (Cont.)

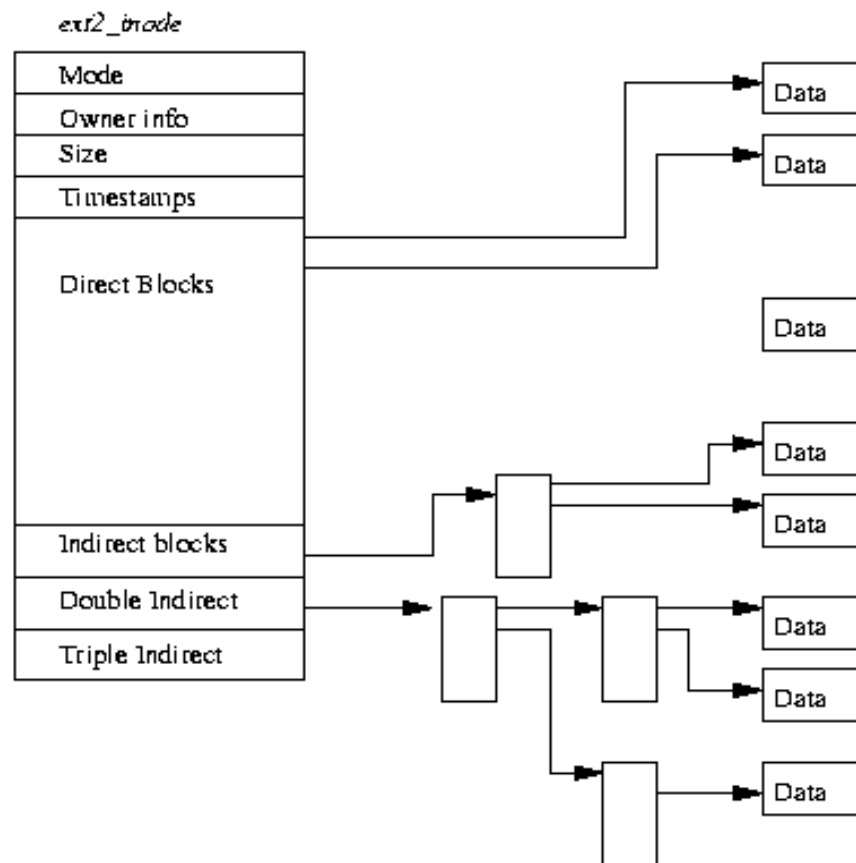
- How do we guarantee no cycles?
  - Allow only links to file not subdirectories
  - Garbage collection
  - Every time a new link is added use a cycle detection algorithm to determine whether it is OK

## 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

# Inodes (An example of a FCB)

- On disk data structure
  - Describes where all the bits of a file are

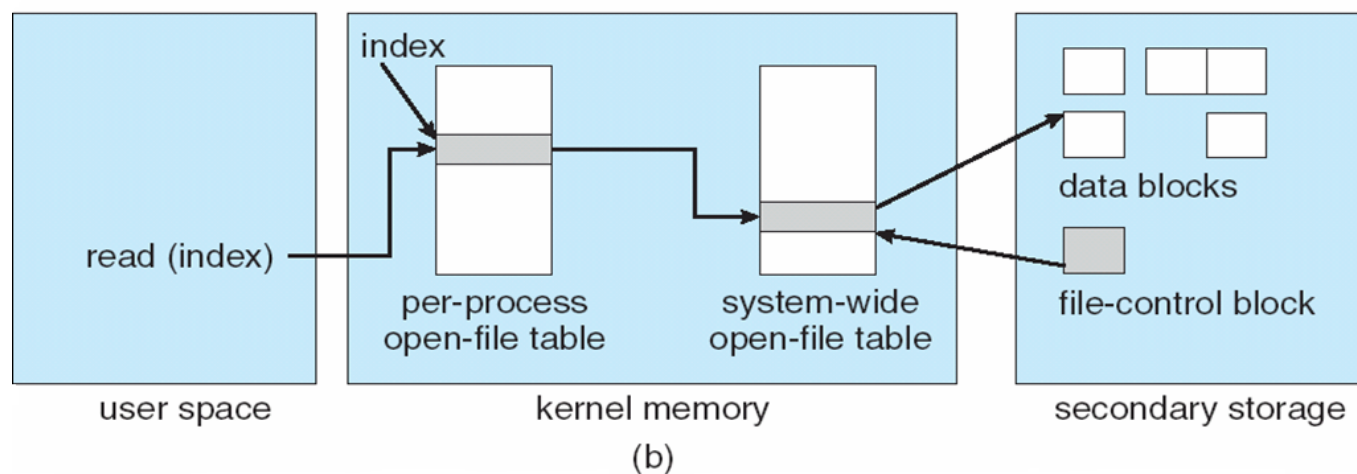



# Directories

- Directory - just special file
  - Contains metadata, filenames
    - pointers to inodes
- Typically **hierarchical tree**
  - odd exposure of data structure to user



# In-Memory File System Structures



# Directory Implementation


- **Linear list** of file names with pointer to the data blocks
  - Simple to program
  - Time-consuming to execute
    - Linear search time
    - Could keep ordered alphabetically via linked list or use B+ tree
- **Hash Table** - linear list with hash data structure
  - Decreases directory search time
  - **Collisions** - situations where two file names hash to the same location
  - Only good if entries are fixed size, or use chained-overflow method

## Allocation Methods - Contiguous

- An allocation method refers to how disk blocks are allocated for files:
- **Contiguous allocation** - each file occupies set of contiguous blocks
  - Best performance in most cases
  - Simple - only starting location (block #) and length (number of blocks) are required
  - Problems include finding space for file, knowing file size, external fragmentation, need for **compaction off-line (downtime)** or **on-line**

# Contiguous Allocation

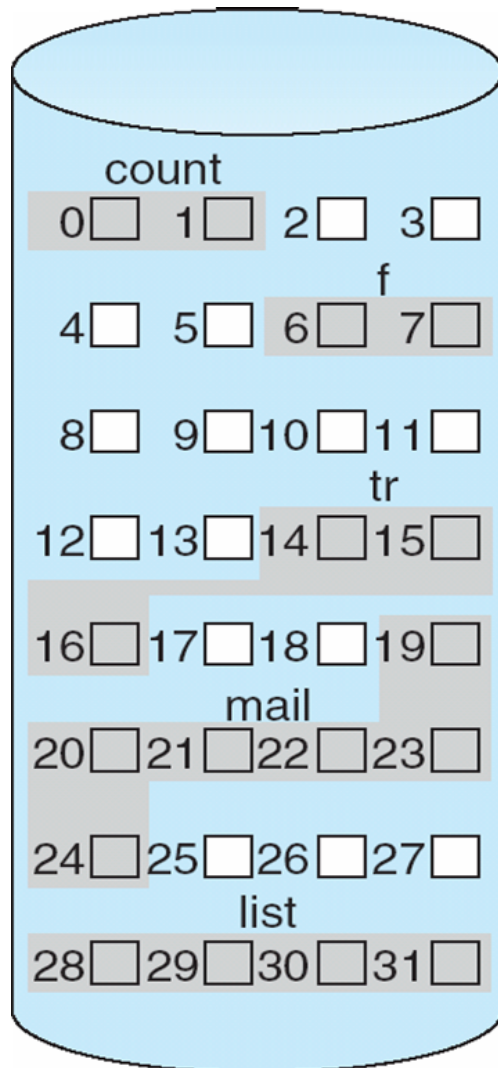
- Mapping from logical to physical



Block to be accessed =  $Q + \text{starting address}$

Displacement into block =  $R$

# Contiguous Allocation of Disk Space



directory

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

# Extent-Based Systems


- Many newer file systems (i.e., Veritas File System) use a modified contiguous allocation scheme
- Extent-based file systems allocate disk blocks in extents
- An **extent** is a contiguous block of disks
  - Extents are allocated for file allocation
  - A file consists of one or more extents

## Allocation Methods - Linked

- **Linked allocation** - each file a linked list of blocks
  - File ends at nil pointer
  - No external fragmentation
  - Each block contains pointer to next block
  - No compaction, external fragmentation
  - Free space management system called when new block needed
  - Improve efficiency by clustering blocks into groups but increases internal fragmentation
  - Reliability can be a problem
  - Locating a block can take many I/Os and disk seeks
- **FAT (File Allocation Table) variation**
  - Beginning of volume has table, indexed by block number
  - Much like a linked list, but faster on disk and cacheable
  - New block allocation simple


# Linked Allocation

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







# Linked Allocation



# File-Allocation Table

directory entry




# Allocation Methods - Indexed

- Indexed allocation

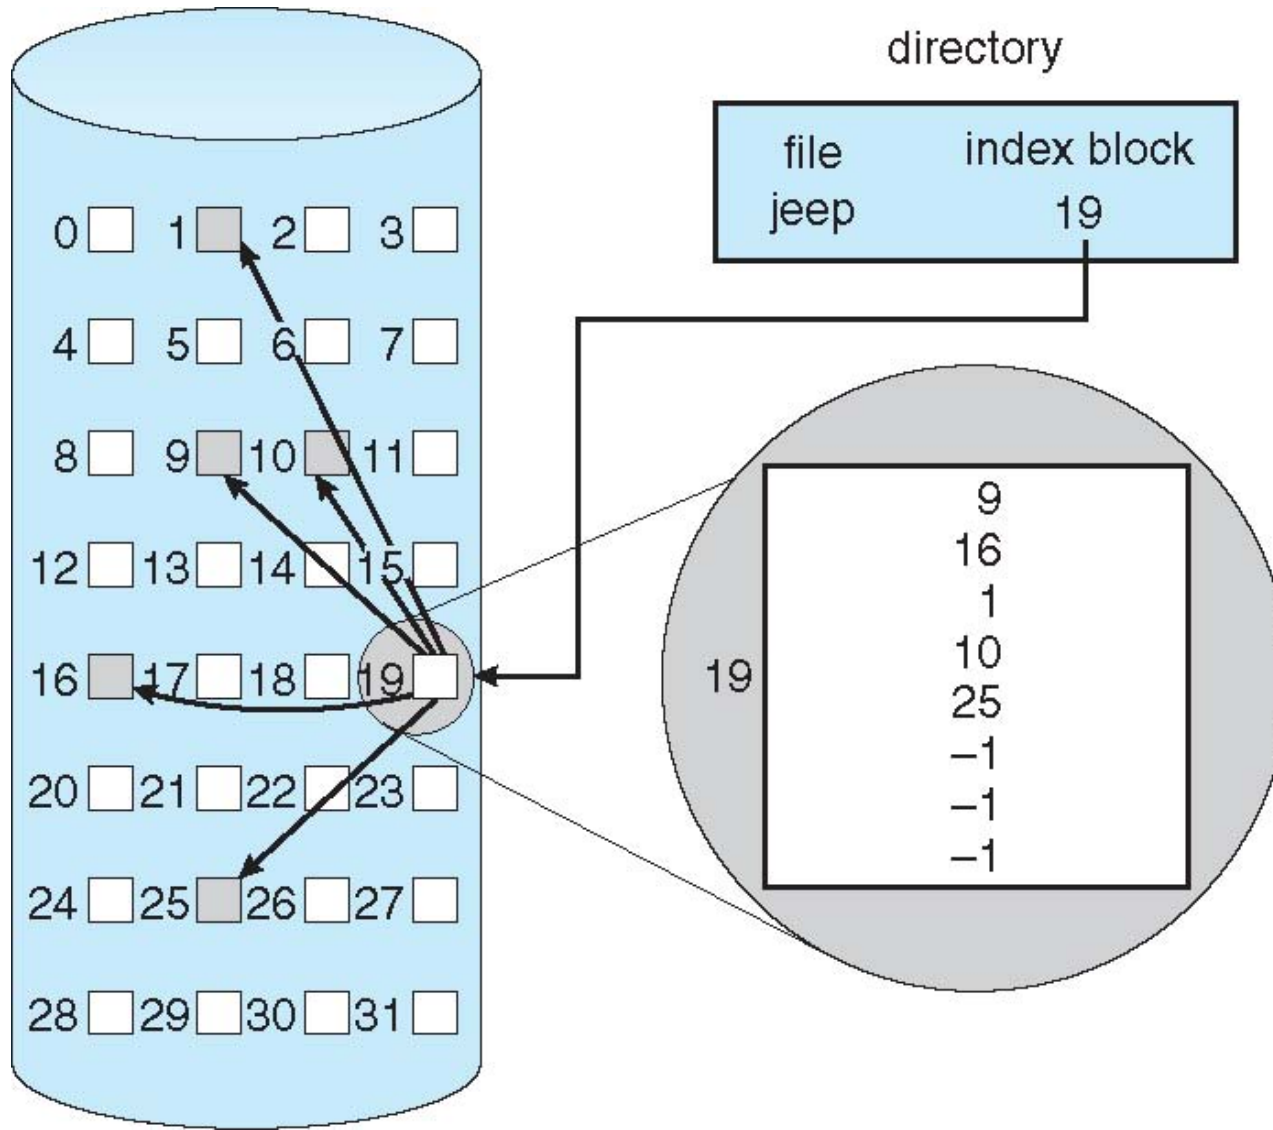
- Each file has its own **index block(s)** of pointers to its data blocks

- Logical view



index table

# Example of Indexed Allocation



**Lets chat about project 2**