


Unit-05

Disk Scheduling & File Concepts

Unit-05/Lecture-01

Disk Structure

- Disks provide the bulk of secondary storage.
- Disk drives are addressed as large 1-dimensional arrays of *logical blocks*, where the logical block is the smallest unit of transfer. The size of logical block is generally 512 bytes.
- The 1-dimensional array of logical blocks is *mapped* into the sectors of the disk sequentially.
 - Sector 0 is the first sector of the first track on the outermost cylinder.
 - Mapping proceeds in order through that track, then the rest of the tracks in that cylinder, and then through the rest of the cylinders from outermost to innermost.
- The number of sectors per track is not a constant. Therefore, modern disks are organized in zones of cylinders. The number of sectors per track is constant within a zone.




Disk I/O

- Whenever a process needs I/O to or from a disk, it issues a system call to the operating system.
 - If the desired disk drive and controller is available, the request can be serviced immediately other wise the request is placed in a queue.
 - Once an I/O completes, the OS can choose a pending request to serve next.

Disk performance Parameters

- The operating system is responsible for using hardware efficiently - for the disk drives, this means having a fast *access time* and *disk bandwidth*.
- *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.

- *Access time* has two major components
 - *Seek time* is the time for the disk are to move the heads to the cylinder containing the desired sector.
 - *Rotational latency* is the additional time waiting for the disk to rotate the desired sector to the disk head.



- Seek time is the reason for differences in performance
 - Minimize seek time
 - Seek time \approx seek distance

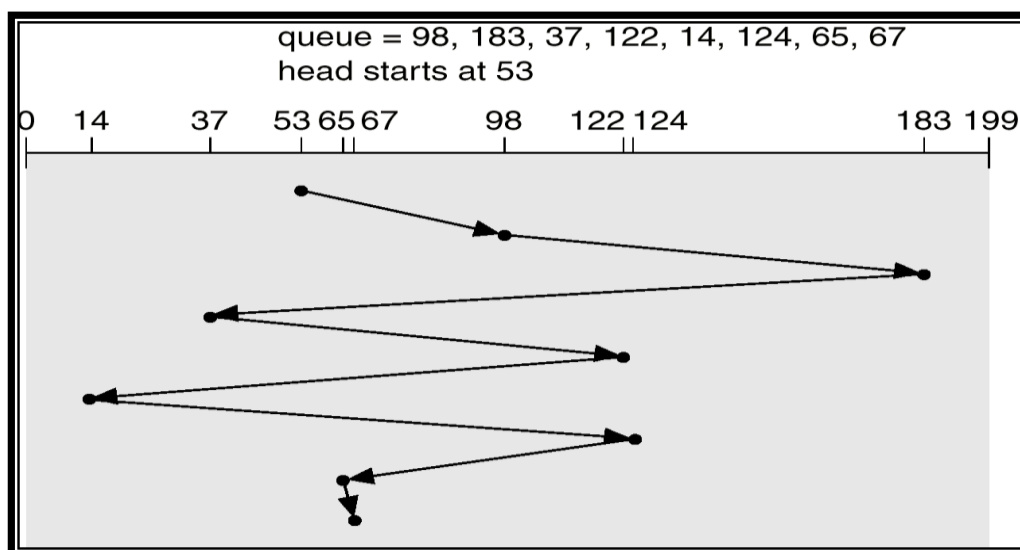
Disk Scheduling

- For a single disk there will be a number of I/O requests
- If requests are selected randomly, we will get the worst possible performance
- Several algorithms exist to schedule the servicing of disk I/O requests.
- We illustrate them with a request queue (0-199).

98, 183, 37, 122, 14, 124, 65, 67
Head pointer 53

First Come First Serve (FCFS)


- The I/O requests are served in the order in which they reach. See below (total head movement=640 cylinders)
- FCFS is a fair scheduling algorithm but not an optimal one.



Unit-05/Lecture-02


SCAN Scheduling

- 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.
- Sometimes called the *elevator algorithm*.
- Illustration shows total head movement of 208 cylinders (head is moving towards cylinder 0).




C-SCAN Scheduling

- 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.



C-Look Scheduling

- **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.



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.
- Performance depends on the number and types of requests.
- Requests for disk service can be influenced by the file-allocation method.
- 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 C-LOOK is a reasonable choice for the default algorithm

Disk Management

- *Low-level formatting, or physical formatting* - Dividing a disk into sectors that the disk controller can read and write.
- To use a disk to hold files, the operating system still needs to record its own data structures on the disk.
 - *Partition* the disk into one or more groups of cylinders.
 - *Logical formatting* or “making a file system”.
- Boot block initializes system.
 - The bootstrap is stored in ROM.

– *Bootstrap loader* program.

- Bad sectors may be managed manually. For example MS-DOS format command does a logical format and if it finds any bad sector, it writes a special value into FAT.
- *Sector sparing* method may also be used to handle bad blocks (as used in SCSI disks). The controller maintains a list of bad sectors which is updated regularly. Low level formatting also sets aside some spare sectors. The controller can be asked to replace each bad sector logically with one of the spare sectors.

S.NO	RGPV QUESTION	YEAR	MARKS
Q.1	Discuss following disk scheduling algorithm- a)SCAN b) CLOOK c)FCFS	Dec 2011	12
Q.2	Compare Look & SCAN scheduling algorithm	Dec 2012	7
Q.3	Compare the performance of C-SCAN and SCAN scheduling assuming a uniform distribution of requests. Consider the average response time and the effective bandwidth. How does performance depend on relative time of seek time and rotational latency.	Dec 2013	7

Unit-05/Lecture-03**File-System Interface**

1. File Concept
2. Access Methods
3. Directory Structure
4. File System Mounting
5. File Sharing
6. Protection

File Concept

1. Contiguous logical address space
2. Types:
 - a. Data
 - i. numeric
 - ii. character
 - iii. binary
 - b. Program

File Structure

1. None - sequence of words, bytes
2. Simple record structure
 - a. Lines
 - b. Fixed length
 - c. Variable length
3. Complex Structures
 - a. Formatted document
 - b. Relocatable load file
4. Can simulate last two with first method by inserting appropriate control characters.
5. Who decides:

a. Operating system

b. Program

File Attributes

1. **Name** – only information kept in human-readable form.
2. **Type** – needed for systems that support different types.
3. **Location** – pointer to file location on device.
4. **Size** – current file size.
5. **Protection** – controls who can do reading, writing, executing.
6. **Time, date, and user identification** – data for protection, security, and usage monitoring.
7. Information about files are kept in the directory structure, which is maintained on the disk.

File Operations

1. Create
2. Write
3. Read
4. Reposition within file – file seek
5. Delete
6. Truncate
7. $\text{Open}(Fi)$ – search the directory structure on disk for entry Fi , and move the content of entry to memory.
8. $\text{Close}(Fi)$ – move the content of entry Fi in memory to directory structure on disk.

File Types – Name, Extension

file type	usual extension	function
executable	exe, com, bin or none	read to run machine- language program
object	obj, o	compiled, machine language, not linked
source code	c, cc, java, pas, asm, a	source code in various languages
batch	bat, sh	commands to the command interpreter
text	txt, doc	textual data, documents
word processor	wp, tex, rrf, doc	various word-processor formats
library	lib, a, so, dll, mpeg, mov, rm	libraries of routines for programmers
print or view	arc, zip, tar	ASCII or binary file in a format for printing or viewing
archive	arc, zip, tar	related files grouped into one file, sometimes com- pressed, for archiving or storage
multimedia	mpeg, mov, rm	binary file containing audio or A/V information

Unit-05/Lecture-04

File Access Methods


1. Sequential Access

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

2. Direct Access

read n
write n
position to n
read next
write next
rewrite n
n = relative block number

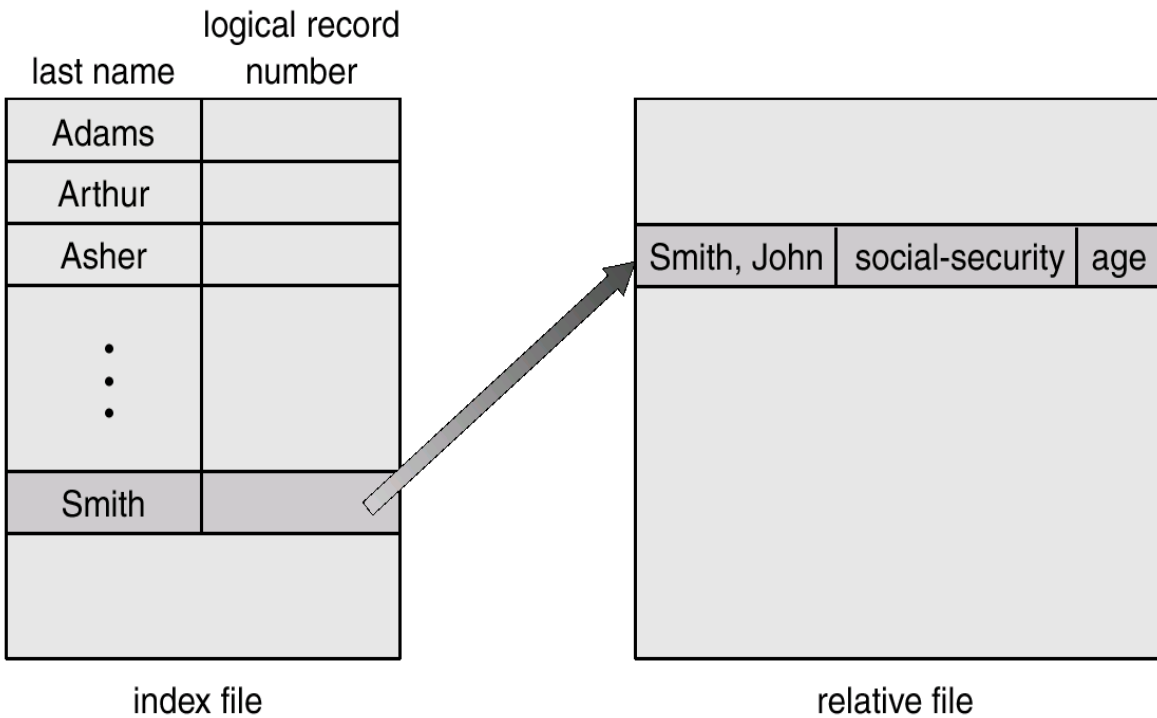
Sequential-access File



Simulation of Sequential Access on a Direct-access File

sequential access	implementation for direct access
<i>reset</i>	<i>cp = 0;</i>
<i>read next</i>	<i>read cp;</i> <i>cp = cp+1;</i>
<i>write next</i>	<i>write cp;</i> <i>cp = cp+1;</i>

Example of Index and Relative Files




Directory Structure

A collection of nodes containing information about all files.

Both the directory structure and the files reside on disk.

Backups of these two structures are kept on tapes.

A Typical File-system Organization



Information in a Device Directory

1. Name
2. Type
3. Address
4. Current length
5. Maximum length
6. Date last accessed (for archival)
7. Date last updated (for dump)
8. Owner ID (who pays)
9. Protection information (discuss later)

Operations Performed on Directory

1. Search for a file
2. Create a file
3. Delete a file
4. List a directory
5. Rename a file
6. Traverse the file system

S.NO	RGPV QUESTION	YEAR	MARKS
Q.1	Explain file access mechanism briefly	June 2011	8


Unit-05/Lecture-05

Organize the Directory (Logically) to Obtain

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


Single-Level Directory

A single directory for all users.



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

1. Efficient searching
2. Grouping Capability
3. Current directory (working directory)

a. **cd** /spell/mail/prog

b. **type** list

1. **Absolute** or **relative** path name
2. Creating a new file is done in current directory.
3. Delete a file

rm <file-name>

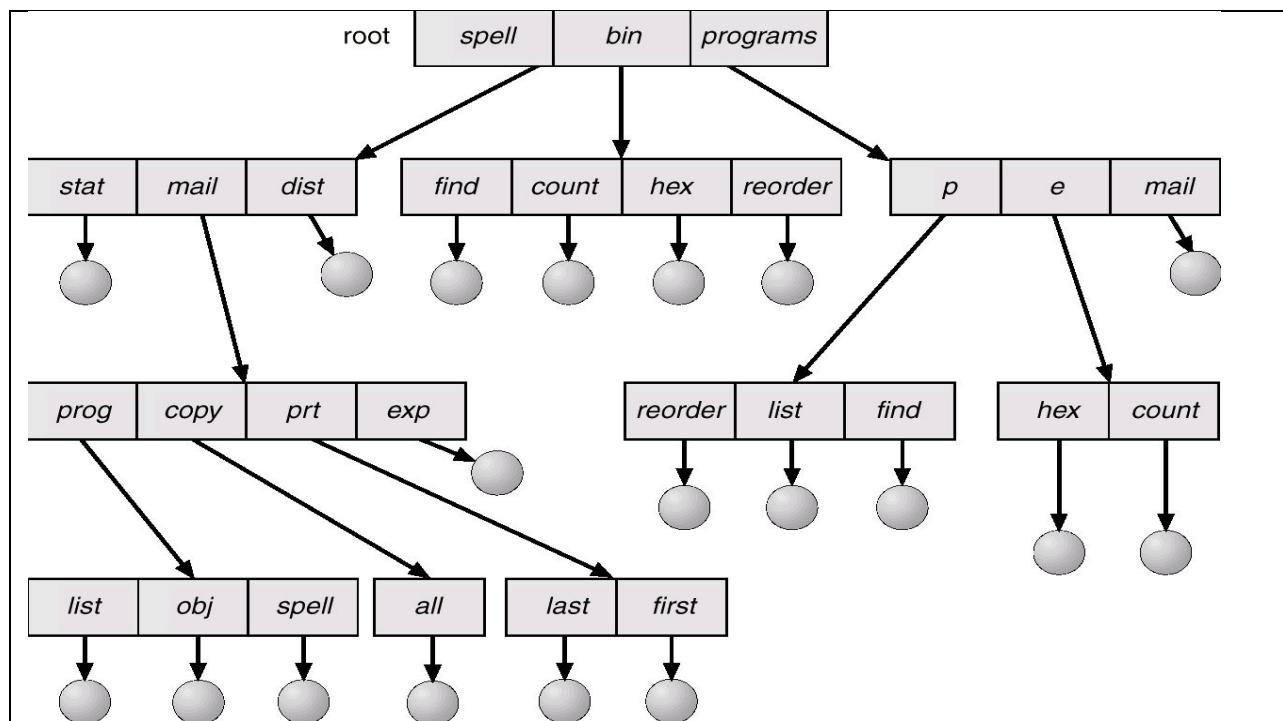
4. Creating a new subdirectory is done in current directory.

mkdir <dir-name>

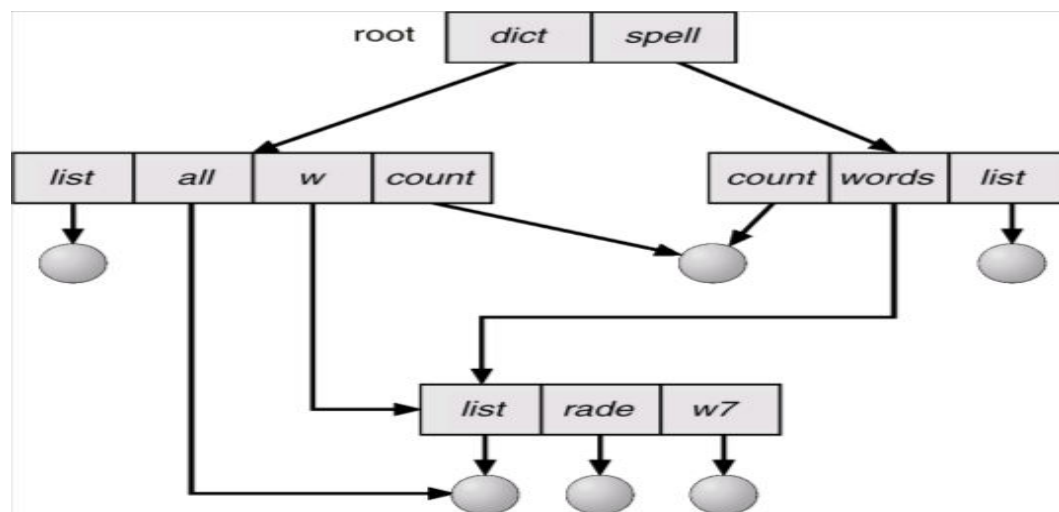
Example: if in current directory **/mail**

mkdir count

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



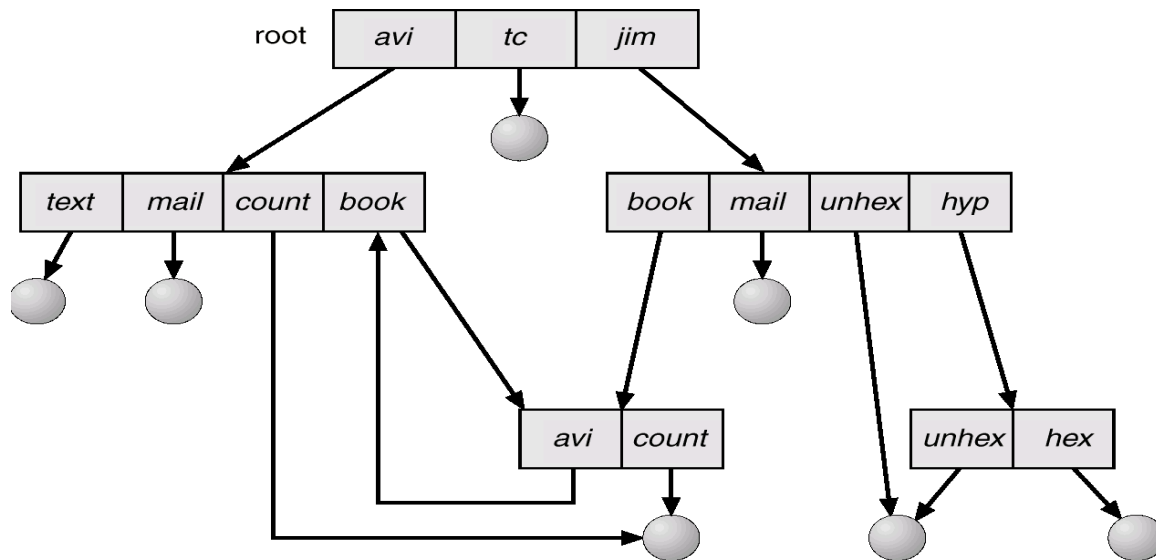
Acyclic-Graph Directories



S.NO	RGPV QUESTION	YEAR	MARKS
Q.1	Explain different schemes of defining directory structures	Dec 2011	10
Q.2	Explain directory system	Dec 2013	5

Unit-05/Lecture-06

General Graph Directory



1. How do we guarantee no cycles?

a. Allow only links to file not subdirectories.

b. Garbage collection.


c. Every time a new link is added use a cycle detection algorithm to determine whether it is OK.

File System Mounting


1. A file system must be **mounted** before it can be accessed.

2. A unmounted file system is mounted at a **mount point**.


(a) Existing. (b) Unmounted Partition



(a)



(b)



File Sharing

Sharing of files on multi-user systems is desirable.

Sharing may be done through a *protection* scheme.

On distributed systems, files may be shared across a network.

Network File System (NFS) is a common distributed file-sharing method.

Protection

1. File owner/creator should be able to control:

- a. what can be done
- b. by whom

2. Types of access

- a. Read
- b. Write
- c. Execute
- d. Append
- e. Delete
- f. List

Access Lists and Groups

Mode of access: read, write, execute

Three classes of users

RWX

a) **owner access** 7 → 1 1 1

RWX

b) **group access** 6 → 1 1 0

RWX

c) **public access** 1 → 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.

File System Implementation

1. File System Structure
2. File System Implementation
3. Directory Implementation
4. Allocation Methods
5. Free-Space Management
6. Efficiency and Performance
7. Recovery
8. Log-Structured File Systems
9. NFS


S.NO	RGPV QUESTION	YEAR	MARKS
Q.1	Explain file access mechanism briefly	Dec 2012	8

Unit-05/Lecture-07

File-System Structure

1. File structure
 - a. Logical storage unit
 - b. Collection of related information
2. File system resides on secondary storage (disks).
3. File system organized into layers.
4. *File control block* – storage structure consisting of information about a file.

Layered File System

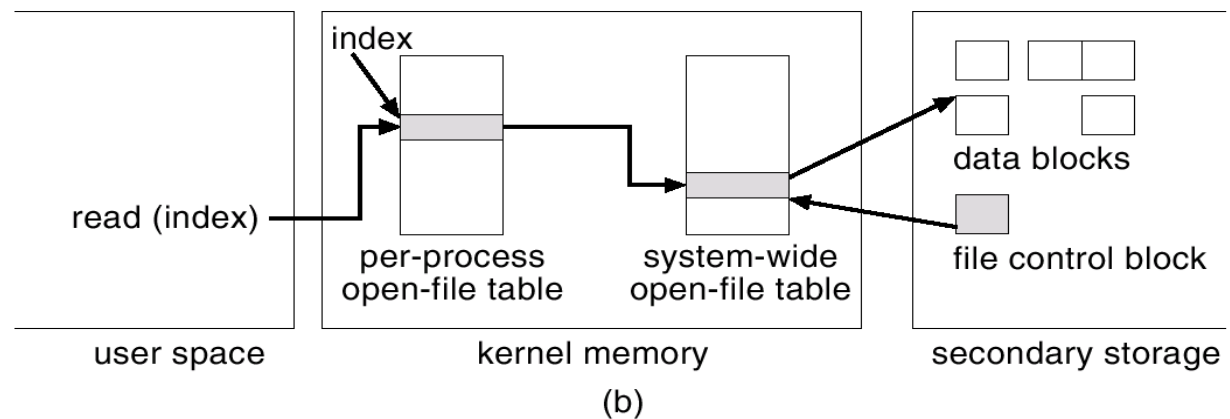



A Typical File Control Block

file permissions
file dates (create, access, write)
file owner, group, ACL
file size
file data blocks

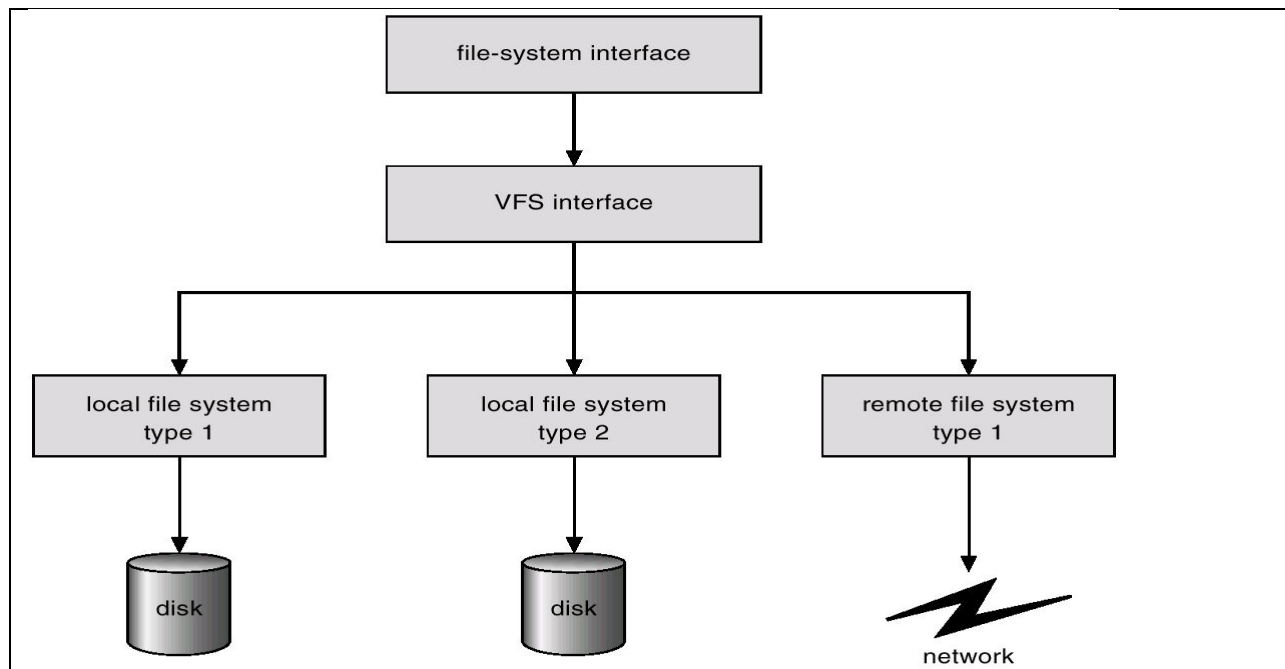
In-Memory File System Structures

1. The following figure illustrates the necessary file system structures provided by the operating systems.
2. Figure (a) refers to opening a file.
3. Figure (b) refers to reading a file.



Virtual File Systems

1. Virtual File Systems (VFS) provide an object-oriented way of implementing file systems.
2. VFS allows the same system call interface (the API) to be used for different types of file systems.
3. The API is to the VFS interface, rather than any specific type of file system.



Directory Implementation

1. Linear list of file names with pointer to the data blocks.
 - a. simple to program
 - b. time-consuming to execute
2. Hash Table – linear list with hash data structure.
 - a. decreases directory search time
 - b. *collisions* – situations where two file names hash to the same location
 - c. fixed size

S.NO	RGPV QUESTION	YEAR	MARKS
Q.1	Explain file organization & access mechanism	Dec 2013	5

Unit-05/Lecture-08

File Allocation Methods


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

- Contiguous allocation
- Linked allocation
- Indexed allocation

Contiguous Allocation

1. Each file occupies a set of contiguous blocks on the disk.
2. Simple – only starting location (block #) and length (number of blocks) are required.
3. Random access.
4. Wasteful of space (dynamic storage-allocation problem).
5. Files cannot grow.

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

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

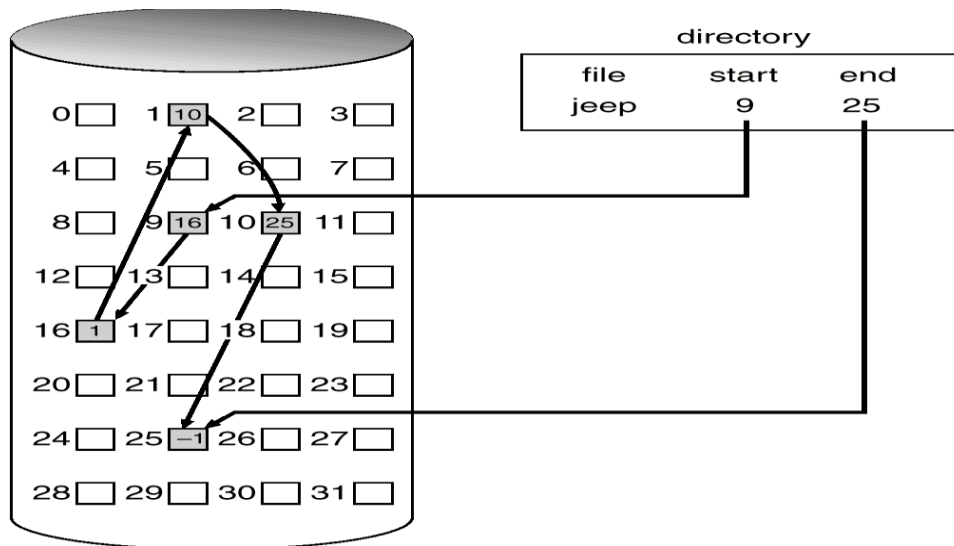
Linked Allocation

1. Each file is a linked list of disk blocks: blocks may be scattered anywhere on the disk.
2. Simple – need only starting address
3. Free-space management system – no waste of space
4. No random access
5. Mapping

Block to be accessed is the Qth block in the linked chain of blocks representing the file.


Displacement into block = $R + 1$

File-allocation table (FAT) – disk-space allocation used by MS-DOS and OS/2.




File-Allocation Table

directory entry



Indexed Allocation

1. Brings all pointers together into the *index block*.
2. Logical view.



Need index table

Random access

Dynamic access without external fragmentation, but have overhead of index block.

Mapping from logical to physical in a file of maximum size of 256K words and block size of 512 words. We need only 1 block for index table.

Q = displacement into index table

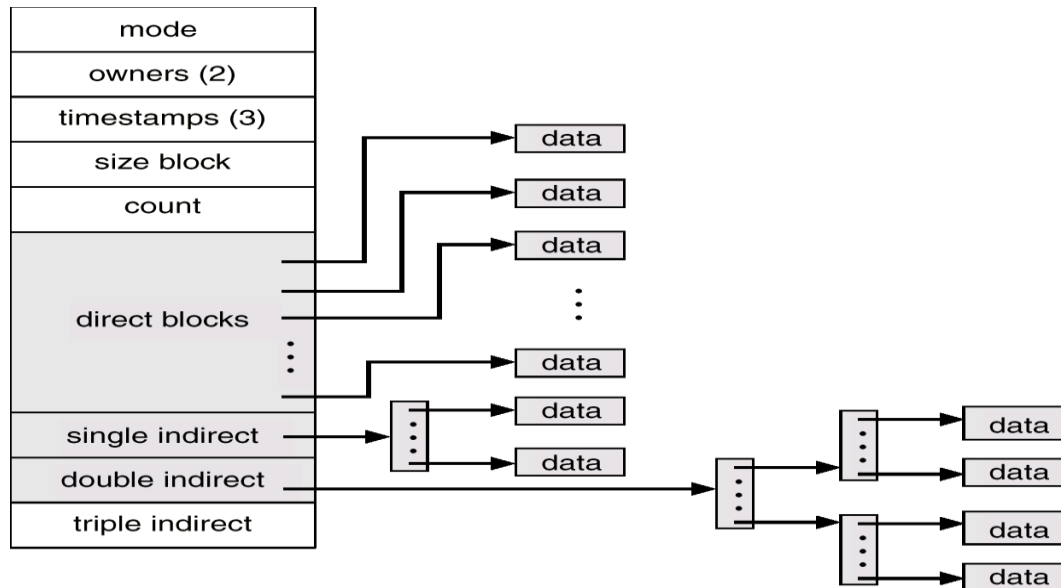
R = displacement into block

1. Mapping from logical to physical in a file of unbounded length (block size of 512 words).
2. Linked scheme – Link blocks of index table (no limit on size).
3. Mapping from logical to physical in a file of unbounded length (block size of 512 words).
4. Linked scheme – Link blocks of index table (no limit on size).
5. -level index (maximum file size is 5123)

S.NO	RGPV QUESTION	YEAR	MARKS
Q.1	Explain file allocation methods Or file organization	Dec 2013 Dec 2009	5 4

Unit-05/Lecture-09

Combined Scheme: UNIX (4K bytes per block)



Free-Space Management

Block number calculation

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

Bit map requires extra space. Example:

block size = 212 bytes

disk size = 230 bytes (1 gigabyte)

$n = 230/212 = 218$ bits (or 32K bytes)

1. Easy to get contiguous files
2. Linked list (free list)
 - a. Cannot get contiguous space easily
 - b. No waste of space

3. Grouping

4. Counting

5. Need to protect:

- o Pointer to free list
- o Bit map

** Must be kept on disk


** Copy in memory and disk may differ.

** Cannot allow for block[i] to have a situation where bit[i] = 1 in memory and bit[i] = 0 on disk.

o Solution:

** Set bit[i] = 1 in disk.
 ** Allocate block[i]
 ** Set bit[i] = 1 in memory


Linked Free Space List on Disk



Efficiency and Performance

1. Efficiency dependent on:
 - a. disk allocation and directory algorithms
 - b. types of data kept in file's directory entry
2. Performance
 - a. disk cache – separate section of main memory for frequently used blocks
 - b. free-behind and read-ahead – techniques to optimize sequential access
 - c. improve PC performance by dedicating section of memory as virtual disk, or RAM disk.

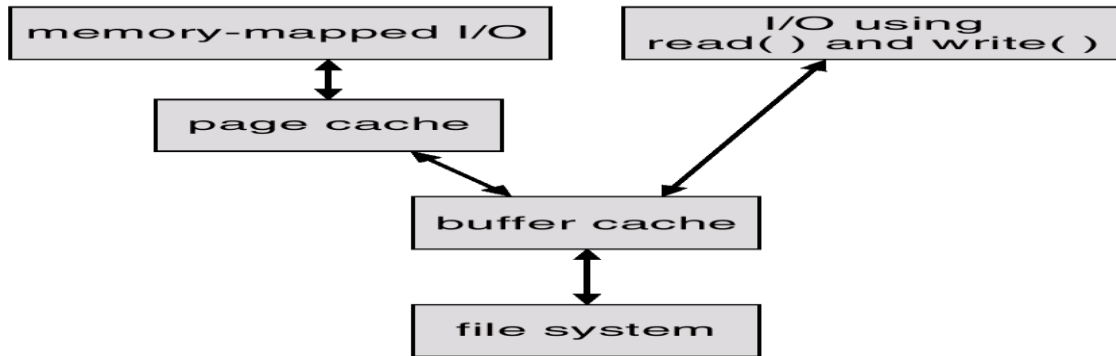
Various Disk-Caching Locations



Page Cache

1. A **page cache** caches pages rather than disk blocks using virtual memory techniques.
2. Memory-mapped I/O uses a page cache.
3. Routine I/O through the file system uses the buffer (disk) cache.
4. This leads to the following figure.


I/O Without a Unified Buffer Cache



Unified Buffer Cache

1. A unified buffer cache uses the same page cache to cache both memory-mapped pages and ordinary file system I/O.

I/O Using a Unified Buffer Cache



S.NO	RGPV QUESTION	YEAR	MARKS
Q.1	Explain file management scheme of UNIX/ LINUX operating system	Dec 2012, June 2011	7 7

REFERENCES

BOOK	AUTHOR	PRIORITY
Operating system	Silberschatz, Galvin	1
System Programming and Operating System	Dhamdhere	2