MODULE 3 - File and Device Management

PART A

1. What is a file?

A file is a named collection of related information that is recorded on secondary storage. A file contains either programs or data. A file has certain "structure" based on its type.

2. List the various file attributes.

   A file has certain other attributes, which vary from one operating system to another, but typically consist of these: Identifier, Name, Type, Location, Size, Protection, Time, Date and user identification.

3. What are the various file operations?

   The six basic file operations are: Creating a file, writing a file, reading a file, repositioning within a file, deleting a file, Truncating a file.

4. Define UFD and MFD.

   In the two-level directory structure, each user has her own user file directory (UFD). Each UFD has a similar structure, but lists only the files of a single user. When a job starts the system's master file directory (MFD) is searched. The MFD is indexed by the user name or account number, and each entry points to the UFD for that user.

5. What is a path name?

   A pathname is the path from the root through all subdirectories to a specified file. In a two-level directory structure, a user name and a file name define a path name.

6. Define seek time and latency time.

   The time taken by the head to move to the appropriate cylinder or track is called seek time. Once the head is at right track, it must wait until the desired block rotates under the readwrite head. This delay is latency time.

7. Mention the objectives of File Management System.

   The system that an operating system or program uses to organize and keep track of files. For example, a hierarchical file system is one that uses directories to organize files into a tree structure.

8. What are the contents of a typical file control block?

   File Control Block (FCB) is a file system structure in which the state of an open file is maintained.

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

9. What is meant by polling?

   Polling is the process where the computer waits for an external device to check for its readiness. The computer does not do anything else than checking the status of the device . Polling is often used with low-level hardware. Example: when a printer connected via a parallel port the computer waits until the next character has been received by the printer. These processes can be as minute as only reading 1 Byte. Polling is the continuous (or frequent) checking by a controlling device or process of other devices, processes, queues, etc.

10. State any advantages and disadvantages of placing functionality in a device controller, rather than in the kernel.

    Advantages: -

    a. Bugs are less likely to cause an operating system crash.

    b. Performance can be improved by utilizing dedicated hardware and hard-coded algorithms.

    c. The kernel is simplified by moving algorithms out of it.

    Disadvantages:

    a. Bugs are harder to fix - a new firmware version or new hardware is needed

    b. Improving algorithms likewise require a hardware update rather than just kernel or device driver update

    c. Embedded algorithms could conflict with application's use of the device, causing decreased performance.

11. How free-space is managed using bit vector implementation?

    The free-space list is implemented as a bit map or bit vector. Each block is represented by 1 bit. If the block is free, the bit is 1; if the block is allocated, the bit is 0.

12. What are the information contained in a boot control block and partition control block?

    Boot control block: Contain information needed by the system to boot an operating from that partition. If the disk does not contain an operating system, this block can be empty. It is typically the first block of a partition. In UFS, this is called the boot block.

    Partition Control block: Contains partition details, such as number of blocks in the partition, size of the blocks, free block count and free block pointers, and free FCB count and FCB pointers.

13. Define buffering.

    A buffer is a memory area that stores data while they are transferred between two devices or between a device and an application. Buffering is done for three reasons:

    a. To cope with a speed mismatch between the producer and consumer of a data stream.

    b. To adapt between devices that have different data transfer sizes.

    c. To support copy semantics for application I/O.

14. Define caching.

    A cache is a region of fast memory that holds copies of data. Access to the cached copy is more efficient than access to the original. Caching and buffering are distinct functions, but sometimes a region of memory can be used for both purposes.

15. Define spooling.

    A spool is a buffer that holds output for a device, such as printer, that cannot accept interleaved data streams. When an application finishes printing, the spooling system queues the corresponding spool file for output to the printer. The spooling system copies the queued spool files to the printer one at a time.

16. Define rotational latency and disk bandwidth.

    Rotational latency is the additional time waiting for the disk to rotate the desired sector to the disk head. Disk bandwidth is the total number of bytes transferred, divided by the time between the first request for service and the completion of the last transfer.

17. Mention the need for disk scheduling?

    In operating systems, seek time is very important. Since all device requests are linked in queues, the seek time is increased causing the system to slow down. Disk Scheduling Algorithms are used to reduce the total seek time of any request.

18. Write about low-level formatting.

    Before a disk can store data, it must be divided into sectors that the disk controller can read and write. This process is called low-level formatting or physical formatting. Low-level formatting

fills the disk with a special data structure for each sector. The data structure for a sector consists of a header, a data area, and a trailer.

19. Specify the use of boot block?

    For a computer to start running when powered up or rebooted it needs to have an initial program to run. This bootstrap program tends to be simple. It finds the operating system on the disk loads that kernel into memory and jumps to an initial address to begin the operating system execution. The full bootstrap program is stored in a partition called the boot blocks, at fixed location on the disk. A disk that has boot partition is called boot disk or system disk.

20. Write down the role of sector sparing?

    Low-level formatting also sets aside spare sectors not visible to the operating system. The controller can be told to replace each bad sector logically with one of the spare sectors. This scheme is known as sector sparing or forwarding.

21. Give details on FAT.

    FAT is a much older file-system format that is understood by many systems besides Windows, such as the software running on cameras. A disadvantage is that the FAT file system does not restrict file access to authorized users. The only solution for securing data with FAT is to run an application to encrypt the data before storing it on the file system.

22. Suppose that the disk rotates at 7200 rpm. What is the average rotational latency of the disk drive?

    7200 rpm gives 120 rotations per second. Thus, a full rotation takes 8.33 ms and the average rotational latency (a half rotation) takes 4.167 ms.

23. Write about access matrix.

    The rows of the access matrix represent domains, and the columns represent objects. Each entry in the matrix consists of a set of access rights.

24. Compare single and multi-core CPU.

    A processor that has more than one core is called Multicore Processor while one with single core is called Uni-core Processor or Uniprocessor. Multicore systems support Multithreading and Parallel Computing. Software that can run parallelly is preferred because we want to achieve parallel execution with the help of multiple cores.

25. Mention about confinement problem.

    The problem of guaranteeing that no information initially held in an object can migrate outside of its execution environment is called the confinement problem. This problem is in general unsolvable.

26. Specify the role of capability list in protection.

    A capability list for a domain is a list of objects together with the operations allowed on those objects. An object is often represented by its physical name or address, called a capability.

27. Differentiate threat and attack.

    A threat is the potential for a security violation, such as the discovery of a vulnerability, whereas an attack is the attempt to break security.

28. Give an example scenario for Logic bomb.

    Consider a program that initiates a security incident only under certain circumstances. It would be hard to detect because under normal operations, there would be no security hole. However, when a predefined set of parameters was met, the security hole would be created. This scenario is known as a **logic**

bomb. A programmer, for example, might write code to detect whether he was still employed; if that check failed, a daemon could be spawned to allow remote access, or code could be launched to cause damage to the site.

29. Write about viruses.

    A virus is a fragment of code embedded in a legitimate program. Viruses are self replicating and are designed to "infect" other programs. They can wreak havoc in a system by modifying or destroying files and causing system crashes and program malfunctions.

30. Differentiate viruses vs worms.

The primary difference between a virus and a worm is that viruses must be triggered by the activation of their host; whereas worms are stand-alone malicious programs that can self-replicate and propagate independently as soon as they have breached the system.

31. Specify the types of multiprocessor memory.

    Shared memory and Distributed memory.

32. What are the types of Path Names?

    Path names can be of two types. **Absolute path name:** Begins at the root and follows a path down to the specified file, giving the directory names on the path. **Relative path name:** Defines a path from the current directory.

33. What is Access Control List?

    The most general scheme to implement identity-dependent access is to associate with each file and directory an access control unit.

PART B

1. Compare symmetric and asymmetric encryption schemes, and discuss the circumstances under which a distributed system would use one or the other.

2. Is disk scheduling, other than FCFS scheduling, useful in a single-user environment? Explain your answer.

3. Explain the working principle of DMA with a neat sketch and justify how DMA increases the system performance.

4. Reason out why is rotational latency usually not considered in disk scheduling? How would you modify SSTF, SCAN and C-SCAN to include latency optimization?

5. Consider a system that supports the strategies of contiguous, linked, and indexed allocation. What criteria should be used in deciding which strategy is best utilized for a particular file?

6. Make a list of six security concerns for a bank's computer system. For each item on your list, state whether this concern relates to physical, human, or operating-system security.

7. In some systems, a subdirectory can be read and written by an authorized user, just as ordinary files can be.

    a. Describe the protection problems that could arise.

    b. Suggest a scheme for dealing with each of these protection problems.

8. Consider a system that supports 5,000 users. Suppose that you want to allow 4,990 of these users to be able to access one file.

a. How would you specify this protection scheme in UNIX?

b. Can you suggest another protection scheme that can be used more effectively for this purpose than the scheme provided by UNIX?

9. None of the disk-scheduling disciplines, except FCFS, is truly fair (starvation may occur).

a. Explain why this assertion is true.

b. Describe a way to modify algorithms such as SCAN to ensure fairness.

c. Explain why fairness is an important goal in a time-sharing system.

d. Give three or more examples of circumstances in which it is important that the operating system be unfair in serving I/O requests.

10. Explain why SSDs often use an FCFS disk-scheduling algorithm.

11. Describe the basic interrupt handling mechanism which enables the CPU to respond to an asynchronous event and explain the more sophisticated interrupt handling features used in modern operating system in detail.

12. Suppose that a disk drive has 5,000 cylinders, numbered 0 to 4,999. The drive is currently serving a request at cylinder 2,150, and the previous request was at cylinder 1805. The queue of pending requests, in FIFO order, is:

2069, 1212, 2296, 2800, 544, 1618, 356, 1523, 4965, 3681

Starting from the current head position, what is the total distance (in cylinders) that the disk arm moves to satisfy all the pending requests for each of the following disk-scheduling algorithms?

a. FCFS

b. SSTF

c. SCAN

d. LOOK

e. C-SCAN

f. C-LOOK

13. Consider a RAID level 5 organization comprising five disks, with the parity for sets of four blocks on four disks stored on the fifth disk. How many blocks are accessed in order to perform the following?

a. A write of one block of data

b. A write of seven continuous blocks of data

14. Discuss the relative advantages and disadvantages of sector sparing and sector slipping.

15. Discuss the various methods for allocating disk spaces to files.

16. Could you simulate a multilevel directory structure with a single-level directory structure in which arbitrarily long names can be used? If your answer is yes, explain how you can do so, and contrast this scheme with the multilevel directory scheme. If your answer is no, explain what prevents your simulation's success. How would your answer change if file names were limited to seven characters?

17. Contrast the performance of the three techniques for allocating disk blocks (contiguous, linked, and indexed) for both sequential and random file access.

18. Discuss in detail about the multi-core architecture with neat diagrams.

19. The queue of requests in FIFO is 86,147,91,177,94,150,102,175,130 What is the total head movement needed to satisfy the requests for the following Scheduling algorithms FCFS, SJF, SCAN, LOOK, C-SCAN.

20. Consider the following page reference string 1,2,3,4,2,1,5,6,2,1,2,3,7,6,3,2,1,2,3,6 Find out the number of page faults if there are 4 page frames, using the following page replacement algorithm i) LRU ii) FIFO iii) Optimal.

21. Describe the LRU page replacement algorithm, assuming there are 3 frames and the page reference string is 7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1 8 Find the number of page faults.

22. Describe the basic interrupt handling mechanism which enables the CPU to respond to an asynchronous event and explain the more sophisticated interrupt handling features used in modern operating system in detail.

23. Describe the SSTF disk scheduling algorithm using the following data. The disk head is initially at position-cylinder 53.the cylinder sequence of requests is 98, 183, 37, 122, 14, 124, 65. 67. find the total head movement.

24. Consider a logical address space of 8 pages of 1024 words each, mapped on to a physical memory of 32 frames. how many bits are there in the logical address? How many bits are there in the physical address?

25. What is virtual memory? Explain. Suppose we have a demand paged memory. The page table is held in registers. it takes 8ms to service a page fault if an empty page is available or the replaced page is not modified, and 20ms if the replaced page is modified. memory access time is 100ns. Assume that the page to be replaced is modified 70% of the time. what is the maximum acceptable page fault rate for an effective access time of no more than 200ns?

26. Memory partitions of 100kb,500 kb,200 kb,300kb,600 kb is available how would best, worst, first fit algorithm to place processes 212,417,112,426 in order. Which is the best algorithm?

27. A virtual memory system has the following specification: Size of the virtual address space=64k Size of the physical address space=4k Page size=512

Virtual Page Number	Physical Frame Number
0	0

3	1
7	2
4	3
10	4
12	5
30	6
31	7

i. Find all the virtual addresses that will generate a page fault.
ii. Compute the main memory addresses for the following virtual addresses. 24,3784,10250,30780

28. A process references 5 pages A, B, C, D, E in the following order A, B, C, D, A, E, B, C, E, D Assuming that the replacement algorithm is LRU and FIFO, find out the number of page faults during the sequence of references, starting with an empty main memory with 3 frames.

29. Consider the following directory structure:

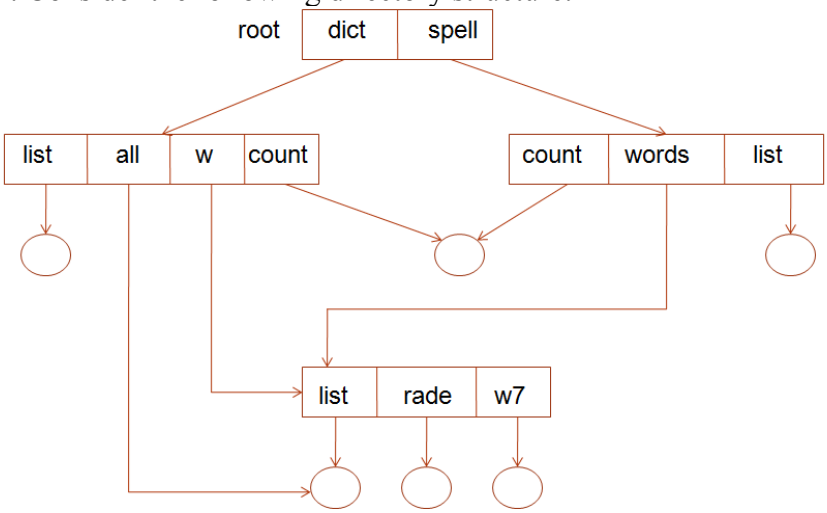

a) Does the given directory structure have shared sub directories and files? Is so, list the shared entities?
b) What is the absolute path for file ‘w7’?

Is there a cycle in this directory structure? What problems could arise if cycles are allowed to exist in the directory and how can they be resolved?

30. Describe briefly on the performance issues of multi-core processors.