

## Unit I: OVERVIEW AND PROCESS MANAGEMENT

### Important 2-Mark Questions:

- 

#### Define an Operating System.1...

- 

An operating system is software that manages computer hardware and provides an environment for application programs to run, acting as an intermediary between the user and hardware<sup>78</sup>. It also provides a basis for application programs<sup>8</sup>.

- 

#### What is a process?1...

- 

A **process** is a program in execution<sup>111</sup>. It is the unit of work in a system, consisting of both operating-system processes (executing system code) and user processes (executing user code)<sup>9</sup>...

- 

#### Differentiate between multiprogramming and time-sharing.1...

- 

**Multiprogramming** increases CPU utilization by organizing jobs (code and data) so that the CPU always has something to execute, keeping several jobs in memory simultaneously<sup>1416</sup>. **Time-sharing** is an extension of multiprogramming where CPU scheduling algorithms rapidly switch between jobs, creating the illusion that each job is running concurrently, thus providing quick response time to interactive users<sup>13</sup>...

- 

#### List two components shared by threads belonging to the same process.118

- 

Threads within the same process share their **code section**, **data section**, and other operating-system resources like **open files** and **signals**<sup>18</sup>.

- 

#### What is Inter-Process Communication (IPC)?1...

-

**Inter-Process Communication (IPC)** refers to mechanisms that allow cooperating processes to exchange information with each other<sup>1</sup>.... This can occur between processes on the same computer or across a network<sup>2</sup>.

**Important 16-Mark Questions:**

- 

**Elaborate on the components of a computer system and explain the fundamental role of an operating system as a resource manager and control program.**<sup>14</sup>...

- 

A computer system consists of hardware, the operating system, application programs, and users<sup>22</sup>. The operating system manages computer hardware and provides a basis for application programs<sup>8</sup>. From the computer's perspective, the operating system acts as a **resource allocator**, managing CPU time, memory space, file-storage space, and I/O devices, allocating them to specific programs and users efficiently and fairly<sup>23</sup><sup>26</sup>. It also functions as a **control program**, managing the execution of user programs to prevent errors and improper use of the computer, with particular concern for I/O devices<sup>24</sup>. Key aspects include multiprogramming for CPU utilization and time-sharing for user responsiveness<sup>14</sup>....

- 

**Explain the concept of a process, its various states, and the information maintained in a Process Control Block (PCB). Discuss the essential operations that an operating system performs on processes.**<sup>1</sup>...

- 

A **process** is an instance of a program in execution<sup>11</sup>. As a process executes, it transitions through various states: **New, Running, Waiting, Ready, and Terminated**<sup>3</sup><sup>13</sup><sup>4</sup>. Each process is represented in the operating system by a **Process Control Block (PCB)**, which contains information such as the process state, program counter, CPU registers, CPU-scheduling information, memory-management information, accounting information, and I/O status information<sup>28</sup>.... The operating system is responsible for critical process management activities, including **creating and deleting both user and system processes, suspending and resuming processes, and providing mechanisms for process synchronization and communication**<sup>1</sup>....

- 

**Discuss the various multi-threading models, such as Many-to-One, One-to-One, and Many-to-Many. Explain the common issues related to multithreaded programming, including fork() and exec() system calls, and thread cancellation.**<sup>1</sup>...

-

A **thread** is a basic unit of CPU utilization, comprising a thread ID, program counter, register set, and stack, sharing code and data with other threads of the same process<sup>18</sup>. Multi-threading models describe the relationship between user threads and kernel threads:

- 

**Many-to-One Model:** Maps many user-level threads to a single kernel thread. Efficient for thread management in user space, but a blocking system call by one thread blocks the entire process, and threads cannot run in parallel on multiprocessors<sup>37</sup>.

- 

**One-to-One Model:** Maps each user-level thread to a kernel thread. Provides greater concurrency and allows multiple threads to run in parallel on multiprocessors. However, creating many user threads can be resource-intensive due to kernel overhead<sup>37</sup>. Windows XP, Solaris (since Solaris 9), and Linux use this model<sup>39,44</sup>.

- 

**Many-to-Many Model:** Multiplexes many user-level threads to a smaller or equal number of kernel threads. Allows developers to create many user threads and enables parallel execution on multiprocessors. When a thread blocks, the kernel can schedule another thread<sup>38</sup>.

- 

Key **threading issues** include:

- 

**fork() and exec() System Calls:** When fork() is called, some UNIX systems offer two versions: one that duplicates all threads and one that duplicates only the invoking thread<sup>41</sup>. exec() typically replaces the entire process, including all threads<sup>41</sup>.

- 

**Thread Cancellation:** The process of terminating a thread before it has completed its task, which can be asynchronous (terminates immediately) or deferred (checks periodically if it should terminate)<sup>40</sup>.

---

## Unit II: PROCESS SCHEDULING AND SYNCHRONIZATION

**Important 2-Mark Questions:**

- 

List any two CPU scheduling criteria.<sup>245</sup>

- 

CPU scheduling criteria include **CPU utilization, throughput, turnaround time, waiting time, and response time**<sup>4546</sup>.

- 

**What is the critical-section problem?**<sup>247</sup>

- 

The **critical-section problem** is to design a protocol that cooperating processes can use to ensure that when one process is executing in its critical section (where shared data is accessed), no other process is allowed to execute in its critical section<sup>247</sup>.

- 

**Define a semaphore.**<sup>248</sup>

- 

A **semaphore** is a synchronization tool that provides a simple but effective solution to the critical-section problem. It is an integer variable accessed only through two atomic operations: `wait()` and `signal()`<sup>4849</sup>.

- 

**List the four necessary conditions for deadlock.**<sup>250</sup>

- 

The four necessary conditions for deadlock are **mutual exclusion, hold and wait, no preemption, and circular wait**<sup>250</sup>.

- 

**What is a safe state in deadlock avoidance?**<sup>5152</sup>

- 

A system is in a **safe state** if there exists a **safe sequence** of processes for the current allocation of resources. A safe sequence is a sequence of all the processes in the system such that for each process  $P_i$ , the resources that  $P_i$  can still request can be satisfied by the currently available resources plus the resources held by all processes  $P_j$ , where  $j < i$ <sup>5152</sup>.

**Important 16-Mark Questions:**

-

**Explain various CPU scheduling algorithms, including First-Come, First-Served (FCFS), Shortest-Job-First (SJF), Priority Scheduling, Round Robin (RR), and Multilevel Feedback Queues. Provide examples and discuss their trade-offs.2...**

◦

**CPU scheduling** determines which process in the ready queue is allocated the CPU62.

▪

**FCFS:** Processes are executed in the order they arrive. Simple to implement but can lead to long waiting times for short processes (convoy effect)46....

▪

**SJF:** Associates with each process the length of its next CPU burst and schedules the process with the shortest next CPU burst. Optimal for minimizing average waiting time but difficult to implement as burst times are hard to predict34....

▪

**Priority Scheduling:** Each process is assigned a priority, and the CPU is allocated to the process with the highest priority. Can suffer from starvation if low-priority processes never get to run5563.

▪

**Round Robin (RR):** Similar to FCFS but adds preemption using a time quantum. Each process gets a small unit of CPU time (time quantum) and if it doesn't complete within the quantum, it's preempted and added to the end of the ready queue56....

▪

**Multilevel Feedback Queues:** Allows processes to move between queues based on their CPU burst characteristics, favoring I/O-bound processes over CPU-bound processes. Highly configurable but complex to tune60....

◦

Trade-offs involve criteria like CPU utilization, throughput, turnaround time, waiting time, and response time45.

•

**Discuss the process synchronization problem, its requirements, and explain solutions using synchronization hardware (TestAndSet, Swap) and semaphores. Illustrate with classical problems of synchronization like the Bounded-Buffer Problem or Readers-Writers Problem.2...**

◦

**Process Synchronization** is essential when processes cooperate and share data, to ensure data consistency and avoid race conditions<sup>75</sup>.... The **critical-section problem** ensures that only one process executes its critical section at a time<sup>47</sup>. Solutions must satisfy **mutual exclusion, progress, and bounded waiting**<sup>81</sup>.

◦

**Synchronization Hardware:** Modern systems provide special atomic hardware instructions like **TestAndSet()** or **Swap()** that can test and modify a word or swap contents of two words as an uninterruptible unit, enabling simple mutual exclusion<sup>76</sup>....

◦

**Semaphores:** A synchronization tool using wait() (decrement) and signal() (increment) operations to control access to shared resources<sup>48</sup>. Binary semaphores (mutex locks) are used for mutual exclusion<sup>49</sup>.

◦

**Classical Problems of Synchronization:**

▪

**Bounded-Buffer Problem:** Producers produce items and consumers consume them using a shared buffer of fixed size, requiring synchronization to prevent overflow/underflow and concurrent access<sup>72</sup>.

▪

**Readers-Writers Problem:** Multiple readers can access shared data concurrently, but only one writer can access the data at a time, and no reader can access while a writer is writing<sup>73</sup>.

•

**Define deadlock and elaborate on the various methods for handling deadlocks, including prevention, avoidance (Banker's Algorithm), detection, and recovery.**<sup>2</sup>...

◦

A **deadlock** is a situation where two or more processes are blocked indefinitely, waiting for each other to release resources<sup>2</sup>. It occurs when four conditions hold simultaneously: **mutual exclusion, hold and wait, no preemption, and circular wait**<sup>50</sup>.

◦

**Methods for Handling Deadlocks:**<sup>90</sup>

▪

**Deadlock Prevention:** Ensures that at least one of the four necessary conditions for deadlock cannot hold:

- 

**Mutual Exclusion:** Not required for sharable resources (e.g., read-only files)<sup>91</sup>.

- 

**Hold and Wait:** Processes must request all resources before starting execution, or release all held resources before requesting new ones<sup>91</sup>.

- 

**No Preemption:** If a process holding resources requests another that cannot be immediately allocated, it must release all held resources. Preempted resources are added to the list of resources for which the process is waiting<sup>92</sup>.

- 

**Circular Wait:** Impose a total ordering of all resource types and require each process to request resources in an increasing order of enumeration<sup>93</sup>.

- 

**Deadlock Avoidance:** Requires the operating system to be given in advance information about which resources a process will request<sup>107</sup>. The **Banker's Algorithm** is a common avoidance scheme that determines if granting a request leaves the system in a safe state<sup>5152</sup>.

- 

**Deadlock Detection:** Allows the system to enter a deadlock state and then detects it, typically using a **wait-for graph**. It involves an algorithm that periodically checks for cycles in the graph<sup>101....</sup>

- 

**Recovery from Deadlock:** Once a deadlock is detected, the system must recover, often by **process termination** (aborting one or more deadlocked processes) or **resource preemption** (taking resources from one process and giving them to another)<sup>104....</sup>

-----

### Unit III: MEMORY MANAGEMENT

#### Important 2-Mark Questions:

- 

**What is swapping?**<sup>3...</sup>

- 

**Swapping** is a memory management technique where a process can be temporarily removed from main memory (swapped out) to a backing store (disk) and then brought back into memory (swapped in) to continue execution<sup>109</sup>.... This is often managed by a medium-term scheduler to improve process mix or free up memory<sup>110</sup>.

- 

**Distinguish between logical and physical address space.**<sup>90</sup>

- 

A **logical address** is an address generated by the CPU<sup>90</sup>. A **physical address** is an address seen by the memory unit (i.e., the one loaded into the memory-address register)<sup>90</sup>. In systems with virtual memory, the logical address space can be much larger than the physical address space.

- 

**Define internal and external fragmentation.**<sup>105</sup>

- 

**Internal fragmentation** occurs when allocated memory is slightly larger than requested memory, and the unused part within the allocated block is still assigned to the process and cannot be used by other processes<sup>105</sup>. **External fragmentation** occurs when total memory space exists to satisfy a request, but it is not contiguous; memory is fragmented into small, non-contiguous holes<sup>105</sup><sup>112</sup>.

- 

**What is demand paging?**<sup>3</sup>...

- 

**Demand paging** is a virtual memory technique that loads pages into physical memory only when they are needed or "demanded" during program execution, rather than loading an entire process at startup<sup>114</sup><sup>115</sup>. This reduces I/O and memory requirements for processes<sup>113</sup>.

- 

**Explain thrashing in virtual memory.**<sup>3</sup>...

- 

**Thrashing** occurs when a process spends more time paging (swapping pages between memory and disk) than executing instructions<sup>117</sup>. This happens when a process does not

have enough frames allocated to it to hold all the pages it is actively using (its working set), leading to a very high page-fault rate and significant performance degradation<sup>116117</sup>.

### Important 16-Mark Questions:

- 

**Describe various memory management techniques such as contiguous memory allocation, pure paging, and pure segmentation. Compare and contrast their advantages and disadvantages, including issues like fragmentation and sharing.**<sup>3...</sup>

- 

**Memory Management** is crucial for sharing memory among multiple processes to improve CPU utilization and response time<sup>84</sup>.

- 

**Contiguous Memory Allocation:** Each process is loaded into a single, contiguous section of memory<sup>93</sup>. This can be implemented with fixed-sized or variable-sized partitions<sup>9596</sup>.

- 

*Advantages:* Simple to implement.

- 

*Disadvantages:* Suffers from **external fragmentation**, where available memory is scattered in small non-contiguous blocks, making it difficult to allocate large processes<sup>105112</sup>. Sharing code is difficult.

- 

**Paging:** Divides physical memory into fixed-size blocks called **frames** and logical memory into blocks of the same size called **pages**<sup>3119</sup>. A **page table** maps logical pages to physical frames<sup>119</sup>.

- 

*Advantages:* Eliminates external fragmentation, allows non-contiguous allocation, and facilitates sharing of common code (e.g., pure code)<sup>105....</sup>

- 

*Disadvantages:* Can suffer from **internal fragmentation** (last page may not be full)<sup>105</sup>. Increased context-switch time due to page table reloads<sup>99</sup>.

-

**Segmentation:** Supports the user's view of memory as a collection of variable-sized logical units called **segments**, each with a name and length<sup>3103</sup>. Logical addresses consist of a segment name and an offset<sup>103</sup>.

- 

*Advantages:* Reflects the user's logical view of program, facilitates sharing of code and data between processes, and supports dynamic memory allocation.

- 

*Disadvantages:* Suffers from **external fragmentation**<sup>105</sup>. Memory allocation is more complex due to variable segment sizes.

- 

**Segmentation with Paging:** Combines segmentation and paging, where each segment is further divided into pages<sup>3105</sup>. This scheme typically maps segments to a paged memory space<sup>105121</sup>.

- 

*Advantages:* Combines the benefits of both (logical view of segmentation, no external fragmentation of paging)<sup>105</sup>.

- 

*Disadvantages:* Increased complexity and overhead for address translation.

- 

**Explain the concept of virtual memory and demand paging. Describe and compare different page replacement algorithms (e.g., FIFO, LRU, Optimal, Clock algorithm) and discuss how their performance is evaluated.**<sup>3...</sup>

- 

**Virtual Memory** is a technique that allows the execution of processes that are not entirely in main memory, separating the user's logical memory from physical memory<sup>3122</sup>. It enables processes to run larger than physical memory and supports multiprogramming<sup>122</sup>.

- 

**Demand Paging** is the primary implementation of virtual memory, loading pages into memory only when they are actively referenced<sup>114</sup>. A **page fault** occurs when a referenced page is not in memory, triggering the operating system to load it from disk<sup>126</sup>.

-

**Page Replacement Algorithms** are needed when a page fault occurs and there are no free frames in memory<sup>3123</sup>. The goal is to select a page to replace to minimize page faults:

- 

**First-In, First-Out (FIFO):** Replaces the page that has been in memory the longest. Simple but may remove actively used pages<sup>135</sup>....

- 

**Optimal Page Replacement (OPT):** Replaces the page that will not be used for the longest period of time. Guarantees the lowest page-fault rate but is impossible to implement in practice as it requires future knowledge<sup>136</sup>....

- 

**Least Recently Used (LRU):** Replaces the page that has not been used for the longest period of time. A good approximation of optimal but difficult to implement precisely due to tracking usage history<sup>138</sup>....

- 

**LRU-Approximation Algorithms (e.g., Clock / Second-Chance):** Use a reference bit to approximate LRU, giving pages a "second chance" before replacement<sup>140</sup>.... Solaris uses a modified two-hand clock algorithm<sup>120</sup>.

- 

**Performance Evaluation:** Algorithms are typically evaluated by running them on a given **page reference string** and counting the number of **page faults**<sup>134145</sup>. **Effective access time** considers the memory access time and the page-fault rate<sup>146</sup>. Simulations with trace tapes or queueing models can also be used<sup>67147</sup>.

- 

**Discuss the allocation of frames in virtual memory. Explain the phenomenon of thrashing, its causes, and various techniques an operating system can employ to detect and eliminate this problem.**<sup>3</sup>...

- 

**Allocation of Frames** concerns how to distribute the fixed number of available memory frames among competing processes in a demand-paged system<sup>3</sup>....

- 

**Fixed Allocation:** Each process is given a fixed number of frames, either equally or proportionally based on size/priority<sup>149</sup>.

-

**Dynamic Allocation:** Allows processes to gain or lose frames dynamically based on their current needs. This includes **global replacement** (any page in memory is a candidate for replacement) and **local replacement** (only pages belonging to the faulting process are candidates)<sup>156</sup>.

◦

**Thrashing:** A situation where a system spends most of its time swapping pages in and out, rather than doing useful work<sup>3117</sup>. This occurs when the total demand for frames by all processes exceeds the total number of available physical frames ( $D > m$ )<sup>116151</sup>. Each process lacks enough memory to hold its **working set** (the set of pages actively used), leading to continuous page faults<sup>157</sup>.

◦

#### **Detecting and Eliminating Thrashing:**

▪

**Page-Fault Frequency (PFF):** Monitor the page-fault rate. If it's too high, the process needs more frames; if too low, it has too many frames<sup>154</sup>.

▪

**Working-Set Model:** Dynamically allocates enough frames to a process to accommodate its current locality of reference<sup>112....</sup> By estimating the working-set size (WSS) for each process, the system can ensure  $D \leq m$  (total demand less than or equal to available frames)<sup>112158</sup>. If  $D$  exceeds  $m$ , some processes may need to be suspended (swapped out) to reduce the degree of multiprogramming<sup>158</sup>.

-----

## **Unit IV: FILE SYSTEMS**

### **Important 2-Mark Questions:**

•

**Define the file concept.**<sup>3...</sup>

◦

The **file concept** is a uniform, logical view of information storage provided by the operating system, abstracting from the physical properties of storage devices to define a logical storage unit<sup>140159</sup>.

•

**List two common file access methods.**<sup>3128</sup>

- 

Two common file access methods are **sequential access** (data is accessed in order, one record after another) and **direct access** (records are accessed directly without reading preceding records)<sup>128</sup>.

- 

**What is file system mounting?**<sup>3...</sup>

- 

**File system mounting** is the process by which an operating system makes a file system on a storage device (volume) available for use, typically by attaching it to a specific directory (mount point) in the existing file system hierarchy<sup>3....</sup>

- 

**Distinguish between mandatory and advisory file locking.**<sup>161</sup>

- 

With **mandatory file locking**, the operating system enforces the lock: if a process acquires an exclusive lock, no other process can access the locked file until it's released<sup>161</sup>. With **advisory file locking**, the operating system does not enforce the lock; it's up to software developers to explicitly acquire and release locks before accessing the file<sup>161</sup>. Windows typically uses mandatory locking, while UNIX systems use advisory locks<sup>161</sup>.

- 

**What is free-space management in file systems?**<sup>3...</sup>

- 

**Free-space management** is the process by which the operating system keeps track of all the free disk blocks available for allocation to files. This is crucial for efficient disk utilization<sup>3....</sup>

**Important 16-Mark Questions:**

- 

**Explain the concept of a file, its attributes, and different access methods. Describe various directory structures used in operating systems, illustrating with diagrams.**<sup>3...</sup>

- 

A **file** is an abstract data type defined and implemented by the operating system, representing a sequence of logical records<sup>140159</sup>. **File attributes** include name, identifier, type, location, size, protection, time, and user identification<sup>167</sup>.

-

## Access Methods:

- 

**Sequential Access:** Data is processed in order, useful for tapes and many file processing tasks<sup>128</sup>.

- 

**Direct Access:** Records can be read or written rapidly in any order, useful for databases<sup>128</sup>.

- 

Other methods include **indexed access** (built upon direct access, using an index to quickly locate records)<sup>143</sup>.

- 

**Directory Structure:** Organizes and provides information about all files in the system<sup>164</sup>. The directory acts as a symbol table mapping file names to directory entries<sup>132</sup>. Common structures include:

- 

**Single-Level Directory:** All files in one directory. Simple but suffers from naming conflicts and difficulty in grouping files for multiple users<sup>133</sup>.

- 

**Two-Level Directory:** Each user has their own User File Directory (UFD) for their files, resolving naming conflicts between users but not for files from different projects by the same user<sup>179</sup>.

- 

**Tree-Structured Directory:** A hierarchical structure where directories can contain files or subdirectories, allowing users to create their own subdirectories. Offers efficient searching and grouping but links may cause issues<sup>180</sup>.

- 

**Acyclic-Graph Directory:** Allows directories to share files and subdirectories, providing flexibility but introducing challenges like dangling pointers (when shared files are deleted)<sup>182184</sup>.

- 

**General Graph Directory:** Allows cycles, making directory traversal and garbage collection more complex<sup>185</sup>.

-

**Discuss how file systems are implemented, focusing on their layered structure, directory implementation, different allocation methods (contiguous, linked, indexed), and free-space management techniques.**

◦

File systems have a **layered design**. Key layers include **I/O control** (device drivers), **basic file system** (reads/writes physical blocks), **file-organization module** (maps logical blocks to physical blocks), and **logical file system** (manages metadata like directories and FCBs).

◦

**On-disk structures** include the boot control block, volume control block, directory structure, and File-Control Blocks (FCBs) for each file. **In-memory structures** include a system-wide open-file table and a per-process open-file table.

◦

**Allocation Methods** (how disk blocks are allocated to files):

▪

**Contiguous Allocation:** Each file occupies a contiguous set of blocks on disk. Simple for sequential and direct access, but suffers from external fragmentation and difficulty in growing files.

▪

**Linked Allocation:** Each block contains a pointer to the next block. No external fragmentation, but only supports sequential access and susceptible to pointer loss. The **File-Allocation Table (FAT)** improves linked allocation by storing pointers in a table on disk.

▪

**Indexed Allocation:** Each file has an index block (inode in UNIX) containing pointers to all its data blocks. Supports direct access, no external fragmentation, but overhead of index block and potential for large index blocks for very large files. This can be extended with multi-level indexing (single, double, triple indirect blocks).

◦

**Free-Space Management:** Methods to keep track of available disk blocks:

▪

**Bit Map:** A bit vector where each bit represents a disk block (1=allocated, 0=free). Efficient for finding contiguous blocks but requires significant memory for large disks.

- 

**Linked List:** Free disk blocks are linked together, with the first free block pointing to the next, and so on<sup>206</sup>. Simple but inefficient for finding contiguous blocks.

- 

**Grouping/Counting:** Variations to improve efficiency by storing pointers to multiple free blocks or counting consecutive free blocks<sup>206</sup>.

- 

**Elaborate on file sharing and protection mechanisms in multi-user operating systems. Discuss access control lists and capabilities.**<sup>3...</sup>

- 

**File Sharing:** In multi-user systems, sharing files is highly desirable for collaboration and resource efficiency<sup>136</sup>. Issues include consistency semantics (e.g., UNIX consistency semantics for shared access)<sup>215227</sup>.

- 

**File Protection:** Mechanisms to control who can access files and how they can be accessed<sup>164....</sup>

- 

**Access Control:** The system mediates file sharing, either by default allowing access or requiring explicit grants<sup>138</sup>. Most systems use a **protection domain** (a set of objects and the access rights to those objects)<sup>222</sup>.

- 

**Access Control List (ACL):** For each object (file), a list specifies the domains and the access rights for each domain<sup>223228</sup>. It explicitly lists which users or groups have what permissions (e.g., read, write, execute)<sup>218219</sup>. UNIX uses a simplified ACL with owner, group, and other permissions<sup>220</sup>.

- 

**Capability List:** For each domain (process or user), a list of objects that can be accessed along with the operations allowed on those objects<sup>223....</sup> Capabilities are unforgeable tickets that permit certain operations<sup>226</sup>.

- 

**Lock-Key Mechanism:** A compromise where each object has a list of "locks" (bit patterns) and each domain has a list of "keys." A process can access an object if its domain has a key matching one of the object's locks<sup>224</sup>.

---

## Unit V: I/O SYSTEMS

### Important 2-Mark Questions:

- 

**List two categories of I/O devices.**215

- 

I/O devices broadly fit into categories such as **storage devices** (e.g., disks, tapes), **transmission devices** (e.g., network cards, modems), and **human-interface devices** (e.g., screen, keyboard, mouse)215.

- 

**What is the purpose of a device driver?**117214

- 

A **device driver** is a kernel module that encapsulates the details and oddities of specific hardware devices, presenting a uniform device-access interface to the I/O subsystem214. It translates high-level commands into low-level, hardware-specific instructions117.

- 

**Define buffering in I/O systems.**229

- 

**Buffering** is a kernel service that uses memory to temporarily store data during I/O transfers, accommodating data transfer rate mismatches, providing copy semantics for application data, and facilitating scatter-gather operations229230.

- 

**What is disk scheduling?**4...

- 

**Disk scheduling** refers to the operating system's management of the order in which disk I/O requests are serviced, with the goal of improving performance (e.g., reducing access time and increasing bandwidth)4....

- 

**What is RAID?**4...

-

**RAID** (Redundant Arrays of Inexpensive Disks) is a mass-storage structure that uses multiple disks to provide data redundancy (for reliability) and/or improve performance (through data striping)4....

### **Important 16-Mark Questions:**

- 

**Describe the architecture of I/O hardware, including I/O ports, buses, and device controllers. Explain the role of interrupts and Direct Memory Access (DMA) in I/O operations.**4...

- 

**I/O Hardware** consists of various components:

- 

**I/O Devices:** Vary widely in function and speed (e.g., mouse, disk, network card)17....

- 

**I/O Port:** A collection of registers used for communication between the CPU and device controllers216237.

- 

**Bus:** A common set of wires connecting the CPU, main memory, and I/O devices189236.

- 

**Device Controller:** An electronic component that operates a port, bus, or device, containing control registers, a data buffer, and special-purpose processors189.... It performs data transfer between the device and its local buffer238.

- 

**I/O Operations:**

- 

**Polling:** The host CPU repeatedly checks the busy bit of the device controller to see if it's ready for the next command239. Simple but can waste CPU cycles246.

- 

**Interrupts:** A hardware mechanism where a device controller asserts a signal on the interrupt request line, causing the CPU to stop its current task, save its state, and jump to an interrupt handler routine to service the device217.... This allows the CPU to perform other work while waiting for I/O248. Modern systems use interrupt priority levels218....

- 

**Direct Memory Access (DMA):** For large data transfers, a DMA controller is used to transfer blocks of data directly between device buffers and main memory without CPU intervention, reducing CPU overhead and improving performance<sup>189</sup>.... The CPU is only interrupted once per block or burst to acknowledge completion<sup>243249</sup>.

- 

**Explain the application I/O interface and the kernel I/O subsystem, including services like buffering, caching, spooling, device reservation, and error handling. Discuss I/O performance considerations.**<sup>4</sup>...

- 

**Application I/O Interface:** Provides a standard, uniform way for applications to access I/O devices, abstracting away hardware differences<sup>191</sup>. Common categories include block I/O, character-stream I/O, memory-mapped file access, network sockets, and timers<sup>193</sup>. System calls (read(), write(), open(), close()) are the primary interface<sup>194</sup>.... Both blocking and nonblocking I/O are supported<sup>196197</sup>.

- 

**Kernel I/O Subsystem:** Coordinates an extensive collection of services available to applications and other parts of the kernel<sup>250267</sup>.

- 

**I/O Scheduling:** Determines a good order to execute I/O requests, maintaining a wait queue for each device to improve overall system efficiency and response time<sup>198199</sup>.

- 

**Buffering:** Using memory to hold data during transfers to cope with speed mismatches or differences in transfer unit sizes<sup>229</sup>.

- 

**Caching:** Storing copies of data in faster storage (e.g., main memory) for quicker access. A unified buffer cache can manage both process pages and file data using virtual memory techniques<sup>251264</sup>.

- 

**Spooling:** Buffering output for devices that cannot accept interleaved data streams (e.g., printers), with each application's output spooled to a separate disk file and then printed sequentially<sup>202</sup>.

-

**Device Reservation:** Providing explicit mechanisms for exclusive device access for devices that cannot multiplex requests (e.g., tape drives)<sup>203</sup>.

- 

**Error Handling:** Detecting and compensating for transient failures (e.g., network overload, disk read errors) and handling permanent failures, often providing error codes to applications<sup>204205</sup>.

- 

**Protection:** Ensuring that user programs cannot directly access I/O devices or memory-mapped I/O ports without proper authorization<sup>163206</sup>.

- 

**I/O Performance:** A major factor in overall system performance<sup>269</sup>. It places heavy demands on the CPU (device-driver code, context switches), memory bus (data copies), and interrupt-handling mechanisms. Optimizing performance involves balancing CPU, memory, bus, and I/O performance<sup>256257</sup>.

- 

**Discuss Mass-Storage Structure, including various disk-scheduling algorithms (e.g., FCFS, SSTF, SCAN, C-SCAN, LOOK, C-LOOK), disk management (formatting, boot blocks, bad-block recovery), and RAID levels.**<sup>4...</sup>

- 

**Mass-Storage Structure:** Disks are the primary secondary-storage devices, structured as large one-dimensional arrays of logical blocks<sup>170....</sup>

- 

**Disk Scheduling Algorithms:** Manage the order of disk I/O requests to optimize performance, primarily by minimizing seek time<sup>231....</sup>

- 

**FCFS (First-Come, First-Served):** Simplest, processes requests in the order they arrive<sup>284306</sup>.

- 

**SSTF (Shortest-Seek-Time-First):** Selects the request with the minimum seek time from the current head position. Can lead to starvation<sup>285306</sup>.

-

**SCAN (Elevator Algorithm):** The disk arm moves from one end of the disk to the other, servicing requests along the way, then reverses direction<sup>286</sup><sup>306</sup>.

- 

**C-SCAN (Circular-SCAN):** Similar to SCAN but when the arm reaches one end, it immediately returns to the other end without servicing requests on the return trip, ensuring more uniform wait times<sup>287</sup><sup>306</sup>.

- 

**LOOK/C-LOOK:** Variations of SCAN/C-SCAN where the arm only travels as far as the furthest request in each direction before reversing, instead of going all the way to the end of the disk<sup>288</sup><sup>306</sup>.

- 

### **Disk Management:**

- 

**Disk Formatting:** **Low-level formatting** divides the disk into sectors for the controller to read/write<sup>278</sup><sup>289</sup>. **Partitioning** divides the disk into one or more logical partitions<sup>307</sup><sup>308</sup>.

- 

**Boot Blocks:** Special blocks on disk storing the **bootstrap program** (bootstrap loader) that loads the operating system kernel into memory at system startup<sup>307</sup>....

- 

**Bad-Block Recovery:** Mechanisms to handle sectors that are corrupted (bad sectors), often by **sector sparing** (remapping bad sectors to spare sectors) or **sector slipping**<sup>292</sup><sup>311</sup>.

- 

**Swap-Space Management:** Management of disk space used as a backing store for virtual memory, typically using raw disk partitions for efficiency<sup>4</sup>....

- 

### **RAID (Redundant Arrays of Inexpensive Disks):**<sup>4</sup>...

- 

**RAID levels** combine disk striping with parity bits to achieve varying cost-performance trade-offs in terms of reliability and data transfer rates<sup>234</sup>.... Examples include RAID Level 0 (striping, no redundancy), RAID Level 1 (mirroring), and RAID Level 5 (striping with distributed parity)<sup>296</sup>....

-

## Case Study: Implementation of Distributed File System in Cloud OS / Mobile OS.4

◦

While the syllabus mentions "Cloud OS / Mobile OS," the provided sources don't delve into specific case studies for these. However, they discuss **Distributed File Systems (DFS)** concepts which are applicable6....

◦

**Distributed File Systems (DFS):** A file-service system whose users, servers, and storage devices are dispersed among the sites of a distributed system312320.

◦

**Goals/Benefits:** Provide access to remote resources, improve computation speed, and enhance data availability and reliability295312. They aim for **transparency**, making remote resources appear local322323.

◦

### Implementation Issues:

▪

**Naming and Transparency:** How files are named and how location transparency (users unaware of where files are stored) or location independence (file names don't indicate physical storage) is achieved317322.

▪

**Remote File Access:** Strategies like transferring the entire file (data migration) or just accessing parts on demand, often involving caching at the client or server300315.

▪

**Stateful vs. Stateless Service:** Whether the server maintains client state information .

▪

**File Replication:** Maintaining multiple copies of files for availability and reliability .

▪

**Consistency Semantics:** Defining the consistency of shared files (e.g., UNIX consistency semantics)215.

◦

**Challenges:** Dealing with network failures, data consistency, concurrency control, and ensuring robustness298312.

◦

**Example (AFS):** The Andrew File System (AFS) is a distributed file system that focuses on client-side caching and shared name spaces for scalability and performance, addressing heterogeneity issues<sup>316</sup>.