Reg. No.

Question Paper Code : 57247

8/6/16 FN

B.E/B.Tech. DEGREE EXAMINATION, MAY/JUNE 2016

Fifth Semester

Electronics and Instrumentation Engineering

CS 6401 – OPERATING SYSTEMS

(Common to Fourth Semester Computer Science Engineering and Information Technology, Instrumentation and Control Engineering and Sixth Semester Electronics and Communication Engineering and Medical Electronics)

(Regulations 2013)

Time : Three Hours | Maximum : 100 Marks

Answer ALL questions.

PART – A (10 × 2 = 20 Marks)

1. What are the advantages of peer-to-peer systems over client-server systems ?
2. What is the purpose of system programs ?
3. Under what circumstances is user level threads is better than the kernel level threads ?
4. What is the meaning of the term busy waiting ?
5. Name two differences between logical and physical addresses.
6. How does the system detect thrashing ?
7. Why is rotational latency usually not considered in disk scheduling ?
8. How does DMA increase system concurrency ?
9. State the components of the Linux system.
10. Define the function of Caching-only serves.

PART – B (5 × 16 = 80 Marks)

11. (a) (i) Distinguish between the client-server and peer-to-peer models of distributed systems. (8)

(ii) Describe three general methods for passing parameters to the operating system with example. (8)

OR

(b) (i) How could a system be designed to allow a choice of operating systems from which to boot ? What would the bootstrap program need to do ? (8)

(ii) Describe the differences between symmetric and asymmetric multiprocessing. What are three advantages and one disadvantage of multiprocessor systems ? (8)

12. (a) (i) It is possible to have concurrency but not parallelism ? Explain. (8)

(ii) Consider a system consisting of four resources of the same type that are shared by three processes, each of which needs at most two resources. Show that the system is deadlock free. (8)

OR

(b) (i) Describe the actions taken by a kernel to context-switch between processes. (8)

(ii) Provide two programming examples in which multithreading does not provide better performance than a single-threaded solution. (8)

13. (a) (i) Describe a mechanism by which one segment could belong to the address space of two different processes. (8)

(ii) Why are segmentation and paging sometimes combined into one scheme ? Explain them in detail with example. (8)

OR

(b) (i) Under what circumstances do page faults occur ? Describe the actions taken by the operating system when a page fault occurs. (8)

(ii) Discuss situations in which the least frequently used (LFU) page replacement algorithm generates fewer page faults than the least recently used (LRU) page-replacement algorithm. Also discuss under what circumstances the opposite holds good. (8)

14. (a) (i) Why is it important to balance file-system I/O among the disks and controllers on a system in a multitasking environment ? (8)

(ii) Discuss the advantages and disadvantages of supporting links to files that cross mount points. (8)

OR

(b) (i) Explain why logging metadata updates ensures recovery of a file system after a file-system crash. (8)

(ii) Could a RAID level 1 organization achieve better performance for read requests than a RAID level 0 organization ? If so, how ? (8)

15. (a) (i) Why is live migration possible in virtual environments but much less possible for a native operating system ? (8)

(ii) What are the primary goals of the conflict-resolution mechanism used by the Linux kernel for loading kernel modules ? (8)

OR

(b) Explain the step-by-step procedure for setting up a Linux multifunction server. (16)

(b) What is deadlock? What are the necessary conditions for deadlock to occur? Explain the deadlock prevention method of handling deadlock.

Consider the following information about resources in a system.

(i) There are two classes of allocatable resource labeled R1 and R2

(ii) There are two instances of each resource

(iii) There are four processes labeled p1 through p4

(iv) There are some resource instances already allocated to processes as follows:

- One instance of R1 held by p2, another held by p3
- One instance of R2 held by p1, another held by p4

(v) Some processes have requested additional resources, as follows:

- p1 wants one instance of R1
- p3 wants one instance of R2

(1) Draw the resource allocation graph for this system

(2) What is the state (runnable, waiting) of each process? For each process that is waiting indicate what it is waiting for.

(3) Is this system deadlocked? If so, state which processes are involved. If not, give an execution sequence that eventually ends, showing resource acquisition and release at each step.

Reg. No. : | | | | | | | | | | | | |

Question Paper Code : 71678

B.E./B.Tech. DEGREE EXAMINATION, APRIL/MAY 2017

Fourth/Fifth/Sixth Semester

Computer Science and Engineering

CS 6401 — OPERATING SYSTEMS

(Common Electronics and Communication Engineering, Electronics and Instrumentation Engineering, Instrumentation and Control Engineering, Medical Electronics Engineering, Information Technology)

(Regulations 2013)

Time : Three hours Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. Consider a memory system with a cache access time of 10 ns and a memory access time of 110 ns — assume the memory access time includes the time to check the cache. If the effective access time is 10% greater than the cache access time, what is the hit ratio H?

2. What are the objectives of operating systems?

3. "Priority inversion is a condition that occurs in real time systems where a low priority process is starved because higher priority processes have gained hold of the CPU" — Comment on this statement.

4. Differentiate single threaded and multi-threaded processes.

5. What is the difference between a user-level instruction and a privileged instruction? Which of the following instructions should be privileged and only allowed to execute in kernel mode?

(a) Load a value from a memory address to a general-purpose register.

(b) Set a new value in the program counter (PC) register.

(c) Turn off interrupts.

6. Will optimal page replacement algorithm suffer from Belady's anomaly? Justify your answer.

7. Suppose that the disk rotates at 7200 rpm.

   What is the average rotational latency of the disk drive?

8. Differentiate between file and directory.

9. Mention any two features of Linux file systems.

10. Enlist the advantages of using kernel modules in Linux.

PART B — (5 × 13 = 65 marks)

11. (a) (i) Explain the concept of multiprocessor and Multicore organization. (7)

    (ii) Discuss about direct memory access. (6)

Or

(b) (i) Explain the various structures of an operating system. (8)

(ii) Describe system calls and system programs in detail with neat sketch. (5)

12. (a) Consider the following set of processes, with the length of the CPU — burst time in given ms:

Process	Burst Time	Arrival time
P1	8	0.00
P2	4	1.001
P3	9	2.001
P4	5	3.001
P5	3	4.001

Draw four Gantt charts illustrating the execution of these processes using FCFS, SJF, priority and RR (quantum=2) scheduling. Also calculate waiting time and turnaround time for each scheduling algorithms. (13)

Or

(b) What is a race condition? Explain how a critical section avoids this condition. What are the properties which a data item should possess to implement a critical section? Describe a solution to the Dining philosopher problem so that no races arise. (13)

13. (a) Discuss the given Memory Management techniques with diagrams

(i) Partition Allocation Methods. (7)

(ii) Paging and Translation Look-aside Buffer. (6)

Or

(b) (i) Describe about free space management on I/O buffering and blocking. (7)

(ii) Discuss the concept of buddy system allocation with neat sketch. (6)

14. (a) (i) Discuss about the various file access methods. (7)

(ii) With neat sketch explain about the

(1) Directory structure

(2) File sharing (6)

Or

(b) (i) Explain about Kernel I/O subsystem and transforming I/O to hardware operations. (7)

(ii) On a disk with 1000 cylinders, numbers 0 to 999, compute the number of tracks the disk arm must move to satisfy the entire request in the disk queue. Assume the last received was at track 345 and the head is moving towards track 0. The queue in FIFO order contains requests for the following tracks. 123, 874, 692, 475, 105 and 376. Find the seek length for the following scheduling algorithm.

(1) SSTF (2) LOOK (3) CSCAN (6)

15. (a) Explain the concepts of domain name system and multifunction server. (13)

Or

(b) Write short notes on LINUX kernel and virtualization with neat sketch. (13)

PART C — (1 × 15 = 15 marks)

16. (a) What do you mean by term synchronization? What is Semaphore? Explain how semaphore can used as synchronization tool. Consider a coke machine that has 10 slots. The producer is the delivery person and the consumer is the student using the machine. It uses the following three semaphores:

semaphore mutex

semaphore fullBuffer /* Number of filled slots */

semaphore emptyBuffer /* Number of empty slots */

(i) Write pseudo code for delivery_person( ) and student( )

(ii) What will be the initial values of the semaphores?

(iii) Write a solution that guarantees the mutual exclusion and has no deadlocks.

Or

Reg. No. :

07/05/18

Question Paper Code :.40905

B.E./B.Tech. DEGREE EXAMINATION, APRIL/MAY 2018
Fourth/Fifth/Sixth Semester
Computer Science and Engineering
CS 6401 – OPERATING SYSTEMS
(Common to : Electronics and Communication Engineering/Electronics and Instrumentation Engineering/Instrumentation and Control Engineering/Medical Electronics/Information Technology)
(Regulations 2013)

Time : Three Hours | Maximum : 100 Marks

Answer ALL questions

PART – A **(10×2=20 Marks)**

1. What is the difference between trap and interrupt ?
2. Mention the purpose of system calls.
3. What are the benefits of synchronous and asynchronous communication ?
4. Give an programming example in which multithreading does not provide better performance than a single-threaded solutions.
5. Define external fragmentation.
6. What are the counting based page replacement algorithm ?
7. State the typical bad-sector transactions.
8. What is the advantage of bit vector approach in free space management ?
9. List the advantages and disadvantage of writing an operating system in high-level language such as C.
10. What is handle ? How does a process obtain a handle ?

PART – B (5×13=65 Marks)

11. a) State the operating system structure. Describe the operating-system operations in detail. Justify the reason why the lack of a hardware-supported dual mode can cause serious shortcoming in an operating system ?

(OR)

b) i) Give reason why caches are useful. What problems do they solve ? What problems do they cause ? If a cache can be made as large as the device for which it is caching why not make it that large and eliminate the device ? **(8)**

ii) Describe the major activities of operating system with regards to file management. **(5)**

12. a) Describe the difference among short-term, medium-term and long-term scheduling with suitable example.

(OR)

b) Explain the differences in the degree to which the following scheduling algorithms discriminate in favor of short processes :

i) RR

ii) Multilevel feedback queues.

13. a) Explain why sharing a reentrant module is easier when segmentation is used than when pure paging is used with example.

(OR)

b) Discuss situation under which the most frequently used page replacement algorithm generates fewer page faults than the least recently used page-replacement algorithm. Also discuss under which circumstances the opposite holds.

14. a) What are the various disk space allocation methods. Explain any two in detail.

(OR)

b) State and explain the FCFS, SSTF and SCAN disk scheduling with examples.

15. a) i) Under what circumstance would an user process request an operation that results in the allocation of a demand-zero memory region. **(8)**

ii) Describe an useful application of the no-access page facility provided in Window XP. **(5)**

(OR)

b) i) What optimization were used to minimize the discrepancy between CPU and I/O speeds on early computer systems. **(8)**

ii) What manages cache in Windows XP ? How is cache managed ? **(5)**

PART – C (1×15=15 Marks)

16. a) Consider a system consisting of 'm' resources of the same type being shared by 'n' processes. Resource can be requested and released by processes only one at a time. Show that the system is deadlock free if the following two conditions hold :

i) The maximum need of each process is between 1 and m resources.

ii) The sum of all maximum needs is less than m + n.

(OR)

b) Consider the following set of processes, with the length of the CPU burst given in milliseconds :

Process	Burst Time	Priority
P_1	10	3
P_2	1	1
P_3	2	3
P_4	1	4
P_5	5	2

The process are assumed to have arrived in the order P_1, P_2, P_3, P_4, P_5 all at time 0.

i) Draw Gantt charts that illustrate the execution of these processes using the scheduling algorithms FCFS (smaller priority number implies higher priority) and RR (quantum = 1). **(10)**

ii) What is the waiting time of each process for each of the scheduling algorithms ? **(5)**

———————

PART C — (1 × 15 = 15 marks)

16. (a) (i) Consider the following system snapshot using data structures in the Banker's algorithm, with resources A, B, C, and D, and process P0 to P4:

	Max	Allocation	Need	Available
	ABCD	ABCD	ABCD	ABCD
P0	6012	4001	????	3211
P1	1750	1100		
P2	2356	1254		
P3	1653	0633		
P4	1656	0212		

Using Banker's algorithm, answer the following questions.

(1) What are the contents of the Need matrix? (2)

(2) Is the system in a safe state? Why? (3)

(3) If a request from process P4 arrives for additional resources of (1,2,0,0,), Can the Banker's algorithm grant the request immediately? Show the new system state and other criteria.(4)

(ii) Consider a demand-paged computer system where the degree of multiprogramming is currently fixed at four. The system was recently measured to determine utilization of CPU and the paging disk. The results are one of the following alternatives. For each case, what is happening? Can the degree of multiprogramming be increased to increase the CPU utilization? Is the paging helping?

(1) CPU utilization 13 percent; disk utilization 97 percent

(2) CPU utilization 87 percent; disk utilization 3 percent

(3) CPU utilization 13 percent; disk utilization 3 percent. (6)

Or

(b) Explain in detail about

(i) File allocation methods. (6)

(ii) Process synchronization in OS. (9)

Reg. No. : | | | | | | | | | | | | |

Question Paper Code : 52862

B.E./B.Tech. DEGREE EXAMINATIONS, APRIL/MAY 2019.

Fourth/Fifth/Sixth Semester

Computer Science and Engineering

CS 6401 – OPERATING SYSTEMS

(Common to Electronics and Communication Engineering/Electronics and Instrumentation Engineering/Instrumentation and Control Engineering/Medical Electronics/Information Technology)

(Regulation 2013)

(Also common to PTCS 6401 – Operating Systems for B.E. (Part-Time) for Third Semester – Computer Science and Engineering – Regulations 2014)

Time : Three hours Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. Define operating system. Name the objectives of operating systems.

2. Elucidate system boot.

3. When a process creates a new process using the fork ( ) operation, which of the following state is shared between the parent process and the child process?

(a) Stack

(b) Heap

(c) Shared memory segments.

4. What resources are used when a thread is created? How do they differ from those used when a process is created?

5. Under what circumstances do page faults occur? Describe the actions taken by the operating system when a page fault occurs.

6. Consider a logical address space of eight pages of 1024 words each, mapped onto a physical memory of 32 frames.

(a) How many bits are there in the logical address?

(b) How many bits are there in the physical address?

7. Why do some systems keep track of the type of a file, while others leave it to the user or simply do not implement multiple file types? Which system is "better?"

8. Why is it advantageous for the user for an operating system to dynamically allocate its internal tables? What are the penalties to the operating system for doing so?

9. The Linux kernel does not allow paging out of kernel memory. What effect does this restriction have on the kernel's design?

10. List down the benefits of virtualization.

PART B — (5 × 13 = 65 marks)

11. (a) (i) List five services provided by an operating system. Explain how each provides convenience to the users. (8)

(ii) Explain also in which cases it would be impossible for user-level programs to provide these services. (5)

Or

(b) Describe evolution of operating system. (13)

12. (a) (i) Write the difference between user thread and kernel thread. (8)

(ii) Palm OS provides no means of concurrent processing. Discuss three major complications that concurrent processing adds to an operating system. (5)

Or

(b) (i) What is the average turnaround time for the following processes using

(1) FCFS

(2) SJF non-preemptive

(3) Preemptive SJF. (9)

Process	Arrival Time	Burst Time
P1	0.0	8
P2	0.4	4
P3	1.0	1

(ii) With example elucidate livelock. (4)

13. (a) Draw the diagram of segmentation memory management scheme and explain its principle. (13)

Or

(b) When do page faults occur? Consider the reference string:

1, 2, 3, 4, 2, 1, 5, 6, 2, 1, 2, 3, 7, 6, 3, 2, 1, 2, 3, 6.

How many page faults and page fault rate occur for the FIFO, LRU and Optimal replacement algorithms, assuming one, two, three, four page frames? (13)

14. (a) (i) In a variable partition scheme, the operating system has to keep track of allocated and free space. Suggest a means of achieving this. Describe the effects of new allocations and process terminations in your suggested scheme. (5)

(ii) Consider a file currently consisting of 100 blocks. Assume that the file control block (and the index block, in the case of indexed allocation) is already in memory. Calculate how many disk I/O operations are required for contiguous, linked, and indexed (single-level) allocation strategies, if, for one block, the following conditions hold. In the contiguous-allocation case, assume that there is no room to grow in the beginning, but there is room to grow in the end. Assume that the block information to be added is stored in memory.

(1) The block is added at the beginning.

(2) The block is added in the middle.

(3) The block is added at the end.

(4) The block is removed from the beginning.

(5) The block is removed from the middle.

(6) The block is removed from the end. (8)

Or

(b) Consider a disk queue with requests for I/O to blocks on cylinders

98, 183, 37, 122, 14, 124, 65, 67

If the disk head is start at 53, then find out the total head movement with respect to FCFS, SSTF, SCAN, C-SCAN and LOOK scheduling. (13)

15. (a) (i) Explain the Components of Linux System with neat sketch. (6)

(ii) Write the Various System Administrator Roles in LINUX OS. (7)

Or

(b) How to Install and Configuring Network Services in LINUX?

b) i) Consider the atomic fetch-and-set x, y instruction unconditionally sets the memory location x to 1 and fetches the old value of x in y without allowing any intervening access to the memory location x. Consider the following implementation of P and V functions on a binary semaphore.

```
void P (binary_semaphore *s) {
unsigned y;
unsigned *x = & (s- > value);
do {
   fetch-and-set x, y;
} while (y);
}
void V (binary_semaphore *s) {
 S- >value = 0;
}
```

Write whether the implementation may or may not work if context switching is disabled in P. (4)

ii) Consider a situation where we have a file shared between many people.

If one of the people tries editing the file, no other person should be reading or writing at the same time, otherwise changes will not be visible to him/her.

However if some person is reading the file, then others may read it at the same time.

a) What kind of situation is this ? (3)

b) Consider the following problem parameters to solve this situation. (8)

Problem parameters :

1) One set of data is shared among a number of processes.
2) Once a writer is ready, it performs its write. Only one writer may write at a time.
3) If a process is writing, no other process can read it.
4) If at least one reader is reading, no other process can write.
5) Readers may not write and only read.

Reg. No. : | | | | | | | | | | | | |

Question Paper Code : 50387 24/11/2017 AN

B.E./B.Tech. DEGREE EXAMINATION, NOVEMBER/DECEMBER 2017
Fourth/Fifth/Sixth Semester
Computer Science and Engineering
CS 6401 – OPERATING SYSTEMS
(Common to : Electronics and Communication Engineering/Electronics and Instrumentation Engineering/Instrumentation and Control Engineering/Medical Electronics/Information Technology)
(Regulations 2013)

Time : Three Hours Maximum : 100 Marks

Answer ALL questions

PART – A (10×2=20 Marks)

1. Mention the objectives of an operating systems.
2. What is SYSGEN and system boot ?
3. Name and draw five different process states with proper definition.
4. Elucidate mutex locks with its procedure.
5. Write about swapping. Let us assume the user process is of size 1MB and the backing store is a standard hard disk with a transfer rate of 5MBPS. Calculate the transfer rate.
6. Consider the following page-reference string :

   1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12.

   How many page faults and page fault ratio would occur for the FIFO page replacement algorithm ? Assuming there is four frames.
7. Suppose that the disk rotates at 7200 rpm.
   a) What is the average rotational latency of the disk drive ?
   b) Identify seek distance can be covered in the time ?

8. Enlist different types of directory structure.

9. List the advantages of Linux OS.

10. Write the purpose of using virtualization.

PART – B (5×13=65 Marks)

11. a) Explain Cache memory and its mapping. (13)

(OR)

b) Describe evolution of operating system. (13)

12. a) i) What is a process ? Discuss components of process and various states of a process with the help of a process state transition diagram. (8)

ii) Write the difference between user thread and kernel thread. (5)

(OR)

b) i) What is the average turnaround time for the following processes using

a) FCFS (3)

b) SJF non-preemptive. (3)

c) Preemptive SJF. (3)

Process	Arrival Time	Burst Time
P1	0.0	8
P2	0.4	4
P3	1.0	1

ii) With example elucidate livelock. (4)

13. a) Draw the diagram of segmentation memory management scheme and explain its principle. (13)

(OR)

b) When do page faults occur ? Consider the reference string :

1, 2, 3, 4, 2, 1, 5, 6, 2, 1, 2, 3, 7, 6, 3, 2, 1, 2, 3, 6.

How many page faults and page fault rate occur for the FIFO, LRU and optimal replacement algorithms, assuming one, two, three, four page frames ? (13)

14. a) i) In a variable partition scheme, the operating system has to keep track of allocated and free space. Suggest a means of achieving this. Describe the effects of new allocations and process terminations in your suggested scheme. (5)

ii) What are different allocation methods in disk storage ? Explain with neat sketch. (8)

(OR)

b) Consider a disk queue with requests for I/O to blocks on cylinders

98, 183, 37, 122, 14, 124, 65, 67

If the disk head is start at 53, then find out the total head movement with respect to FCFS, SSTF, SCAN, C-SCAN and LOOK scheduling. (13)

15. a) i) Explain the components of Linux system with neat sketch. (6)

ii) Write the various system administrator roles in LINUX OS. (7)

(OR)

b) i) How to install and configuring network services in LINUX. (9)

ii) Describe the benefits of virtualization in LINUX OS. (4)

PART – C (1×15=15 Marks)

16. a) Consider the following system snapshot using data structures in the Banker's algorithm, with resources A, B, C and D and process P0 to P4 :

	Max	Allocation	Need	Available
	A B C D	A B C D		A B C D
P0	6 0 1 2	4 0 0 1		3 2 1 1
P1	1 7 5 0	1 1 0 0		
P2	2 3 5 6	1 2 5 4		
P3	1 6 5 3	0 6 3 3		
P4	1 6 5 6	0 2 1 2		

Using Banker's algorithm, answer the following questions :

a) How many resources of type A, B, C and D are there ? (2)

b) What are the contents of the need matrix ? (3)

c) Is the system in a safe state ? Why ? (3)

d) If a request from process P4 arrives for additional resources of (1, 2, 0, 0,), can the Banker's algorithm grant the request immediately ? Show the new system state and other criteria. (7)

(OR)

01/11/18

Reg. No. : | | | | | | | | | | | | |

Question Paper Code : 20363

B.E./B.Tech. DEGREE EXAMINATION, NOVEMBER/DECEMBER 2018.

Fourth Semester

Computer Science and Engineering

CS 6401 – OPERATING SYSTEMS

(Common to Electronics and Communication Engineering, Electronics and Instrumentation Engineering, Instrumentation and Control Engineering, Medical Electronics and Information Technology)

(Regulations 2013)

(Also common to PTCS 6401 — Operating System B.E. (Part-Time) Third Semester Computer Science and Engineering – Regulations 2014)

Time : Three hours Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. Some computer systems do not provide a privileged mode of operation in hardware. Is it possible to construct a secure operating system for these computer systems?

2. Can traps be generated intentionally by a user program? If so, for what purpose?

3. What is the meaning of the term busy waiting?

4. Can a multithreaded solution using multiple user-level threads achieve better performance on a multiprocessor system than on a single-processor system?

5. Under what circumstances would a user be better off using a timesharing system rather than a PC or single-user workstation?

6. State the effect of Thrashing in an operating system.

7. Mention the significance of LDT and GDT in segmentation.

8. List the major attributes and operations of a file system.

9. Do FAT file system advantageous? Justify your answer.

10. Mention the importance of a Kernel in LINUX Operating system.

PART B — (5 × 13 = 65 marks)

11. (a) (i) Explain the various types of system calls with an example for each. (8)

(ii) Discuss the functionality of system boot with respect to an Operating System. (5)

Or

(b) Discuss the essential properties of the following types of systems

(i) Time sharing systems. (4)

(ii) Multi-processor systems. (4)

(iii) Distributed Systems. (5)

12. (a) (i) Explain why interrupts are not appropriate for implementing synchronization primitives in multiprocessor systems. (7)

(ii) Compute the average waiting time for the processes using non-preemptive SJF scheduling algorithm. (5)

Process	Arrival time	Burst time
P_1	0.0	7
P_2	2.0	4
P_3	4.0	1
P_4	5.0	4

Or

(b) Discuss how the following pairs of scheduling criteria conflict in certain settings.

(i) CPU utilization and response time. (4)

(ii) Average turnaround time and maximum waiting time. (5)

(iii) I/O device utilization and CPU utilization. (4)

13. (a) Compare paging with segmentation in terms of the amount of memory required by the address translation structures in order to convert virtual addresses to physical addresses.

Or

(b) Most systems allow programs to allocate more memory to its address space during execution. Data allocated in the heap segments of programs is an example of such allocated memory. What is required to support dynamic memory allocation in the following schemes?

(i) Contiguous memory allocation (4)

(ii) Pure segmentation (5)

(iii) Pure paging. (4)

14. (a) Consider a file system where a file can be deleted and its disk space reclaimed while links to that file still exist. What problems may occur if a new file is created in the same storage area or with the same absolute path name? How can these problems be avoided?

Or

(b) Illustrate an application that could benefit from operating system support for random access to indexed files.

15. (a) UNIX coordinates the activities of the kernel I/O components by manipulating shared in-kernel data structures, whereas Windows NT uses object-oriented message passing between kernel I/O components. Discuss three pros and three cons of each approach.

Or

(b) Discuss virtualization techniques used in different operating systems.

PART C — (1 × 15 = 15 marks)

16. (a) Which of the following scheduling algorithms could result in starvation?

(i) First-come, first-served (5)

(ii) Shortest job first (5)

(iii) Round robin (5)

Detail with Justification.

Or

(b) Outline a solution using semaphores to solve dinning philosopher problem. (15)

PART – C (1×15=15 Marks)

16. a) Consider the following resource allocation graph :

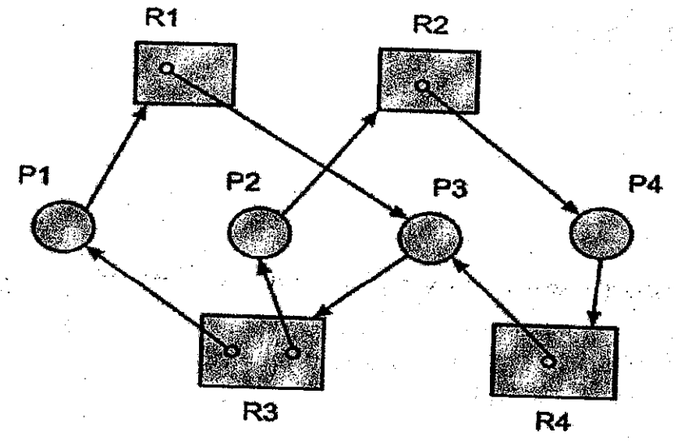

Determine if there is a deadlock. If so, indicate the processes and resources involved. Show how the deadlock can be resolved through addition of resources. If not, argue why this is the case, i.e. there is no deadlock. In either case, provide a feasible sequence of processes to show completion.

(OR)

b) Explain in detail how UNIX Virtual File System has been Implemented.

18/11/19/AN

Reg. No. :

Question Paper Code : 91397

B.E./B.Tech. DEGREE EXAMINATIONS, NOVEMBER/DECEMBER 2019
Fourth/Fifth/Sixth Semester
CS6401 – OPERATING SYSTEMS
(Common to : Computer Science and Engineering/Electronics and Communication Engineering/Electronics and Instrumentation Engineering/Instrumentation and Control Engineering/Medical Electronics/Information Technology)
(Regulations 2013)
(Also Common to PTCS 6401 – Operating Systems – for Third Semester Computer Science and Engineering – Regulations 2014)

Time : Three Hours Maximum : 100 Marks

Answer ALL questions

PART – A (10×2=20 Marks)

1. What are the pros and cons of Microkernals Operating system Structures ?
2. Why it is important for scheduler to distinguish I/O bound programs from CPU bound programs ?
3. List out the various process states available.
4. What is semaphore ? Explain the two primitive operations of a semaphore.
5. State the differences between static and dynamic memory allocation.
6. When does a page fault occur ? Explain various page replacement strategies/ algorithms.
7. Suppose that a disk drive has 5000 cylinders, numbered 0 to 4999. The drive is currently serving a request at cylinder 143 and the previous was at cylinder 125. The queue of pending requests, in FIFO order is: 86, 1470, 913, 1774, 948, 1509, 1022, 1750, 130. 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 SCAN.
8. What type of file accessing method can be used for batch and payroll applications ?
9. Define latency, transfer and seek time with respect to disk I/O.
10. List out the features of Linux OS.

PART – B (5×13=65 Marks)

11. a) Describe the features of the following types of operating systems structures

i) MS-DOS Layer Structure

ii) Layered Approach

iii) Microkernel System Structure

iv) Modules.

(OR)

b) What are system calls ? How do system calls help user programs to interact with the OS ? Explain.

12. a) Assume some OS needs to schedule four processes using different scheduling algorithms. For each process, the following information shows its burst time (processing time), priority (The lower the number, the higher the priority) and arrival time.

Processes	Burst Time	Priority	Arrival Time
P1	12	3	0
P2	6	4	2
P3	4	1	4
P4	18	2	6

What is the Average Waiting Time of those processes for each of the following schedule algorithms ? (Draw a Gantt Chart for each algorithm.)

i) First Come First Serve (FCFS)

ii) Non-preemptive Shortest Job First (NP-SJF)

iii) Preemptive Shortest Job First (P-SJF)

iv) Priority Scheduling

v) Round-Robin (scheduling time quantum is 5 time units)

(OR)

b) What is a critical section ? Explain readers and writers problem with semaphore.

13. a) A computer system has a 36-bit virtual address space with a page size of 8 K, and 4 bytes per page table entry.

i) How many pages are in the virtual address space ?

ii) What is the maximum size of addressable physical memory in this system ?

iii) If the average process size is 8 GB, would you use a one-level, two-level or three-level page table. Why ?

(OR)

b) Explain the following with example :

i) Thrashing (6)

ii) Page replacement algorithm. (7)

14. a) Consider a disk with a rotational rate of 10,000 RPM, an average seek time of 8 ms, and an average of 500 sectors per track. Estimate the average time to read a random sector from disk. Do this by summing the estimates of the seek time, rotational latency and transfer time.

(OR)

b) Elaborate on the various File allocation methods.

15. a) A multicore processor is said to be sequentially consistent if all loads and stores appear to occur in some global total order that is consistent with program order in every core. Sadly, most modern processors are not sequentially consistent: memory accesses can appear to occur in different orders from the perspective of different cores or even in circular order

i) Give a possible reason why different cores might see stores in different orders.

ii) All non-sequentially consistent machines provide special (expensive) instructions that can, when desired, be used to force a memory access to be seen everywhere at once, after all previous accesses of the same core and before all subsequent accesses of the same core. Suggest how the programmer, language, and/or compiler might use such instructions to achieve the illusion of sequential consistency.

(OR)

b) Discuss how the following pairs of scheduling criteria conflict in certain settings.

i) CPU utilization and response time

ii) Average turnaround time and maximum waiting time

iii) I/O device utilization and CPU utilization.