

## Chapter 5

---

# Operating Systems

1. Virtual memory is
  - (a) an extremely large main memory
  - (b) an extremely large secondary memory
  - (c) an illusion of an extremely large memory
  - (d) a type of memory used in super computers
2. Spatial locality refers to the problem that once a location is referenced
  - (a) it will not be referenced again
  - (b) it will be referenced again
  - (c) a nearby location will be referenced soon
  - (d) none of the above
3. Which of the following is an example of a SPOOLED device?
  - (a) The terminal used to enter the input data for a program being executed.
  - (b) The secondary memory device in a virtual memory system.
  - (c) A line printer used to print the output of a number of jobs.
  - (d) None of the above.
4. Page fault occurs when
  - (a) the page is corrupted by application software
  - (b) the page is in main memory
  - (c) the page is not in main memory
  - (d) one tries to divide a number by 0
5. Overlay is
  - (a) a part of an operating system
  - (b) a specific memory location
  - (c) a single contiguous memory that was used in the olden days for running large programs by swapping.
  - (d) overloading the system with many user files



16. Critical region is
- (a) a part of the operating system which is not allowed to be accessed by any process
  - (b) a set of instructions that access common shared resource which exclude one another in time
  - (c) the portion of the main memory which can be accessed only by one process at a time
  - (d) none of the above
17. Kernel is
- (a) considered as the critical part of the operating system
  - (b) the software which monitors the operating system
  - (c) the set of primitive functions upon which the rest of operating system functions are built up.
  - (d) none of the above
18. With a single resource, deadlock occurs
- (a) if there are more than two processes competing for that resource
  - (b) if there are only two processes competing for that resource
  - (c) if there is a single process competing for that resource
  - (d) none of the above
19. Necessary conditions for deadlock are
- (a) non-preemption and circular wait
  - (b) Mutual exclusion and partial allocation
  - (c) both (a) and (b)
  - (d) none of the above
20. In a time-sharing operating system, when the time slot given to a process is completed, the process goes from the RUNNING state to the
- (a) BLOCKED state
  - (b) READY state
  - (c) SUSPENDED state
  - (d) TERMINATED state
21. At a particular time, the value of a counting semaphore is 10. It will become 7 after
- (a) 3 V operations
  - (b) 3 P operations
  - (c) 5 V operations and 2 P operations
  - (d) 13 P operations and 10 V operations
22. Supervisor call
- (a) is a call made by the supervisor of the system
  - (b) is a call with control functions
  - (c) are privileged calls that are used to perform resource management functions, which are controlled by the operating system.
  - (d) is a call made by someone working in root directory
23. Semaphores are used to solve the problem of
- (a) race condition
  - (b) process synchronization
  - (c) mutual exclusion
  - (d) none of the above
24. If the property of locality of reference is well pronounced in a program
- (a) the number of page faults will be more
  - (b) the number of page faults will be less

- (c) the number of page faults will remain the same  
(d) execution will be faster
- \*25. At a particular time of computation, the value of a counting semaphore is 7. Then 20 P operations and 'x' V operations were completed on this semaphore. If the final value of the semaphore is 5, x will be  
(a) 15                      (b) 22                      (c) 18                      (d) 13
26. Pre-emptive scheduling, is the strategy of temporarily suspending a running process  
(a) before the CPU time slice expires                      (b) to allow starving processes to run  
(c) when it requests I/O                      (d) none of the above
27. Mutual exclusion problem occurs  
(a) between two disjoint processes that do not interact  
(b) among processes that share resources  
(c) among processes that do not use the same resource  
(d) none of the above
28. Sector interleaving in disks is done by  
(a) the disk manufacturer                      (b) the disk controller cord  
(c) the operating system                      (d) none of the above
- \*29. Memory protection is of no use in a  
(a) single user system                      (b) non-multiprogramming system  
(c) non-multitasking system                      (d) none of the above
30. Some computer systems support dual mode operation—the user mode and the supervisor or monitor mode. These refer to the modes  
(a) by which user programs handle their data  
(b) by which the operating system executes user programs  
(c) in which the processor and the associated hardware operate.  
(d) of memory access
31. Disk scheduling involves deciding  
(a) which disk should be accessed next  
(b) the order in which disk access requests must be serviced  
(c) the physical location where files should be accessed in the disk  
(d) none of the above
- \*32. A computer system has 6 tape drives, with 'n' processes competing for them. Each process may need 3 tape drives. The maximum value of 'n' for which the system is guaranteed to be deadlock free is  
(a) 2                      (b) 3                      (c) 4                      (d) 1
33. Dirty bit is used to show the  
(a) page with corrupted data  
(b) wrong page in the memory  
(c) page that is modified after being loaded into cache memory  
(d) page that is less frequently accessed

34. Fence register is used for  
(a) CPU protection (b) memory protection  
(c) file protection (d) all of the above
35. Which of the following is a service not supported by the operating system?  
(a) Protection (b) Accounting (c) Compilation (d) I/O operation
36. The first-fit, best-fit and the worst-fit algorithm can be used for  
(a) contiguous allocation of memory (b) linked allocation of memory  
(c) indexed allocation of memory (d) all of the above
37. Which of the following are single-user operating systems?  
(a) MS-DOS (b) UNIX (c) XENIX (d) OS/2
38. In Round Robin CPU scheduling, as the time quantum is increased, the average turn around time  
(a) increases (b) decreases (c) remains constant (d) varies irregularly
39. In a multiprogramming environment  
(a) the processor executes more than one process at a time  
(b) the programs are developed by more than one person  
(c) more than one process resides in the memory  
(d) a single user can execute many programs at the same time.
40. Which of the following are true?  
(a) A re-entrant procedure can be called any number of times.  
(b) A re-entrant procedure can be called even before the procedure has not returned from its previous call.  
(c) Re-entrant procedures cannot be called recursively.  
(d) Re-entrant procedures can be called recursively.
- \*41. In a paged memory, the page hit ratio is 0.35. The time required to access a page in secondary memory is equal to 100 ns. The time required to access a page in primary memory is 10 ns. The average time required to access a page is  
(a) 3.0 ns (b) 68.0 ns (c) 68.5 ns (d) 78.5 ns
42. A state is safe if the system can allocate resources to each process (up to its maximum) in some order and still avoid deadlock.  
Which of the following are true?  
(a) Deadlocked state is unsafe.  
(b) Unsafe state may lead to a deadlock situation.  
(c) Unsafe state must lead to a deadlock situation.  
(d) Deadlocked state is a subset of unsafe state.
43. The size of the virtual memory depends on the size of the  
(a) data bus (b) main memory  
(c) address bus (d) none of the above

- \*44.** In a multi-user operating system, 20 requests are made to use a particular resource per hour, on an average. The probability that no requests are made in 45 minutes is  
(a)  $e^{-15}$                       (b)  $e^{-5}$                       (c)  $1 - e^{-5}$                       (d)  $1 - e^{-10}$
- 45.** In which of the following scheduling policies does context switching never take place?  
(a) Round-robin                      (b) Shortest job first  
(c) Pre-emptive                      (d) First-cum-first-served
- 46.** In which of the following directory systems, is it possible to have multiple complete paths for a file, starting from the root directory?  
(a) Single level directory                      (b) Two level directory  
(c) Tree structured directory                      (d) Acyclic graph directory
- 47.** Suppose that a process is in 'BLOCKED' state waiting for some I/O service. When the service is completed, it goes to the  
(a) RUNNING state                      (b) READY state  
(c) SUSPENDED state                      (d) TERMINATED state
- 48.** In a system that does not support swapping  
(a) the compiler normally binds symbolic addresses (variables) to relocatable addresses.  
(b) the compiler normally binds symbolic addresses to physical addresses.  
(c) the loader binds relocatable addresses to physical addresses.  
(d) binding of symbolic addresses to physical addresses normally takes place during execution.
- 49.** To obtain better memory utilization, dynamic loading is used. With dynamic loading, a routine is not loaded until it is called for. For implementing dynamic loading,  
(a) special support from hardware is essential  
(b) special support from operating system is essential  
(c) special support from both hardware and operating system are essential  
(d) user programs can implement dynamic loading without any special support from the operating system or the hardware.
- 50.** Which of the following is true?  
(a) The linkage editor is used to edit programs which have to be later linked together.  
(b) The linkage editor links object modules during compiling or assembling.  
(c) The linkage editor links object modules and resolves external references between them before loading.  
(d) The linkage editor resolves external references between the object modules during execution time.
- 51.** Which of the following is true?  
(a) Overlays are used to increase the size of physical memory.  
(b) Overlays are used to increase the logical address space.  
(c) When overlays are used, the size of a process is not limited to the size of physical memory.  
(d) Overlays are used whenever the physical address space is smaller than the logical address space.

**The next 5 questions are based on the following information.**

Consider a set of 5 processes whose arrival time, CPU time needed and the priority are given below:

<i>Process Priority</i>	<i>Arrival Time (in ms)</i>	<i>CPU Time Needed (in ms)</i>	<i>Priority</i>
P1	0	10	5
P2	0	5	2
P3	2	3	1
P4	5	20	4
P5	10	2	3

Note: Smaller the number, higher the priority.

- \*52. If the CPU scheduling policy is FCFS, the average waiting time will be  
 (a) 12.8 ms            (b) 8 ms            (c) 16 ms            (d) none of the above
- \*53. If the CPU scheduling policy is SJF, the average waiting time (without pre-emption) will be  
 (a) 12.8 ms            (b) 6.8 ms            (c) 17 ms            (d) none of the above
- \*54. If the CPU scheduling policy is SJF with pre-emption, the average waiting time will be  
 (a) 8 ms            (b) 14 ms            (c) 5.6 ms            (d) none of the above
- \*55. If the CPU scheduling policy is priority scheduling without pre-emption, the average waiting time will be  
 (a) 12.8 ms            (b) 11.8 ms            (c) 10.8 ms            (d) none of the above
- \*56. If the CPU scheduling policy is priority scheduling with pre-emption, the average waiting time will be  
 (a) 19 ms            (b) 7.6 ms            (c) 8 ms            (d) none of the above
57. In partitioned memory allocation scheme, the  
 (a) best fit algorithm is always better than the first fit algorithm.  
 (b) first fit algorithm is always better than the best fit algorithm.  
 (c) superiority of the first fit and best-fit algorithms depend on the sequence of memory requests.  
 (d) none of the above
58. Cascading termination refers to termination of all child processes before the parent terminates  
 (a) normally  
 (b) abnormally  
 (c) normally or abnormally  
 (d) none of the above
59. For implementing a multiprogramming operating system  
 (a) special support from processor is essential  
 (b) special support from processor is not essential  
 (c) cache memory must be available  
 (d) more than one processor must be available

- \*60. Consider a system having ' $m$ ' resources of the same type. These resources are shared by 3 processes  $A, B, C$ , which have peak time demands of 3, 4, 6 respectively. The minimum value of ' $m$ ' that ensures that deadlock will never occur is
- (a) 11                      (b) 12                      (c) 13                      (d) 14
- \*61. A system has 3 processes sharing 4 resources. If each process needs a maximum of 2 units then, deadlock
- (a) can never occur                      (b) may occur  
(c) has to occur                      (d) none of the above
- \*62. ' $m$ ' processes share ' $n$ ' resources of the same type. The maximum need of each process doesn't exceed ' $n$ ' and the sum all their maximum needs is always less than  $m + n$ . In this set up deadlock
- (a) can never occur                      (b) may occur  
(c) has to occur                      (d) none of the above
- \*63. A process refers to 5 pages,  $A, B, C, D$  and  $E$  in the order -  $A; B; C; D; A; B; E; A; B; C; D; E$ . If the page replacement algorithm is FIFO, the number of pages which transfer with an empty internal store of 3 frames is
- (a) 8                      (b) 10                      (c) 9                      (d) 7
64. Distributed systems should
- (a) meet prescribed time constraints                      (b) aim better resource sharing  
(c) aim better system utilization                      (d) aim low system overhead
65. The main function of shared memory is to
- (a) use primary memory efficiently                      (b) do intra process communication  
(c) do inter process communication                      (d) none of the above
66. Which of the following is the most suitable scheduling scheme in a real-time operating system?
- (a) round-robin                      (b) first-come-first-served  
(c) pre-emptive scheduling                      (d) random scheduling
- \*67. In Question number 63, if the number of available page frames is increased to 4 then the number of page transfers
- (a) decreases                      (b) increases  
(c) remains the same                      (d) none of the above
68. 'Aging' is
- (a) keeping track of cache contents  
(b) keeping track of what pages are currently residing in the memory.  
(c) keeping track of how many times a given page is referenced.  
(d) increasing the priority of jobs to ensure termination in a finite time.
- \*69. If there are 32 segments, each of size 1 kbytes, then the logical address should have
- (a) 13 bits                      (b) 14 bits                      (c) 15 bits                      (d) 16 bits



- \*70. Disk requests come to a disk driver for cylinders in the order 10, 22, 20, 2, 40, 6 and 38, at a time when the disk drive is reading from cylinder 20. The seek time is 6 ms per cylinder. The total seek time, if the disk arm scheduling algorithm is first-come-first-served is  
(a) 360 ms                      (b) 850 ms                      (c) 900 ms                      (d) none of the above
71. In question 70, if the scheduling algorithm is the closest cylinder next, then the total seek time will be  
(a) 360 ms                      (b) 876 ms                      (c) 850 ms                      (d) 900 ms
- \*72. A certain moving arm disk storage with one head has following specifications:  
Number of tracks / recording surface = 200  
Disk rotation speed = 2400 rpm  
Track storage capacity = 62500 bits  
The average latency time (assume that the head can move from one track to another only by traversing the entire track) is  
(a) 2.5 s                      (b) 2.9 s                      (c) 3.1 s                      (d) 3.6 s
73. Memory protection is normally done by the  
(a) processor and the associated hardware                      (b) operating system  
(c) compiler                      (d) user program
74. Which of the following scheduling algorithms gives minimum average waiting time?  
(a) FCFS                      (b) SJF                      (c) Round-robin                      (d) Priority
75. In question number 72, the transfer rate will be  
(a) 2.5 Mbits/s                      (b) 4.25 Mbits/s                      (c) 1.5 Mbits/s                      (d) 3.75 Mbits/s
76. In a paged segmented scheme of memory management, the segment table itself must have a page table because  
(a) the segment table is often too large to fit in one page  
(b) each segment is spread over a number of pages  
(c) segment tables point to page tables and not to the physical location of the segment  
(d) the processor's description base register points to a page table
77. Which of the following page replacement algorithms suffers from Belady's anomaly?  
(a) Optimal replacement                      (b) LRU  
(c) FIFO                      (d) Both optimal replacement and FIFO
- \*78. Which of the following scheduling policy is well suited for a time-shared operating system?  
(a) Shortest job first                      (b) Round robin  
(c) First-come-first-serve                      (d) Elevator
- \*79. The address sequence generated by tracing a particular program executing in a pure demand paging system with 100 records per page, with 1 free main memory frame is recorded as follows. What is the number of page faults?  
0100, 0200, 0430, 0499, 0510, 0530, 0560, 0120, 0220, 0240, 0260, 0320, 0370.  
(a) 13                      (b) 8                      (c) 7                      (d) 10

\*80. A computer system has 4 K word cache organized in a block-set-associative manner, with 4 blocks per set, 64 words per block. The number of bits in the SET and WORD fields of the main memory address format is

- (a) 15, 4                      (b) 6, 4                      (c) 7, 2                      (d) 4, 6

\*81. A computer installation has 1000 K of main memory. The jobs arrive and finish in the following sequence.

Job 1 requiring 200 K arrives  
 Job 2 requiring 350 K arrives  
 Job 3 requiring 300 K arrives  
 Job 1 finishes  
 Job 4 requiring 120 K arrives  
 Job 5 requiring 150 K arrives  
 Job 6 requiring 80 K arrives

Among best fit and first fit, which performs better for this sequence?

- (a) First fit                      (b) Best fit  
 (c) Both perform the same      (d) None of the above

82. A memory page containing a heavily used variable that was initialized very early and is in constant use is removed, when the page replacement algorithm used is

- (a) LRU                      (b) FIFO                      (c) LFU                      (d) none of the above

83. Consider the following heap.

JOB1-50	Free-150	JOB2-300	FREE-350	JOB3-600
---------	----------	----------	----------	----------

The sequence of requests for blocks of sizes 300, 25, 125, 50 can be satisfied if we use

- (a) either first fit or best fit policy      (b) first fit but not best fit  
 (c) best fit but not first fit              (d) none of the above

84. Consider the resource allocation graph in Fig. 5.1.




Fig. 5.1

This system is in a deadlock state. This remark is

- (a) true
- (b) false
- (c) impossible to determine
- (d) unpredictable

85. Which of the following is a safe sequence?

- (a) P0, P1, P2, P3
- (b) P1, P0, P3, P2
- (c) P2, P0, P1, P3
- (d) None of the above

86. A demand paging system, with page table held in registers, takes 5 ms to service a page fault if an empty page is available, or if the page to be replaced is not dirty. It takes 15 ms if the replaced page is dirty. Memory access time is 1  $\mu$ s. Assume we want an effective access time of 2  $\mu$ s and that the page to be replaced is dirty 60% of the time. What is the approximate maximum acceptable page fault rate to meet this access time requirement?

- (a) 0.1%
- (b) 1.0%
- (c) 2.5%
- (d) 0.01%

87. Consider a computer with 8 Mbytes of main memory and a 128 K cache. The cache block size is 4 K. It uses a direct mapping scheme for cache management. How many different main memory blocks can map onto a given physical cache block?

- (a) 2048
- (b) 256
- (c) 64
- (d) None of the above

88. Which of the following applications are well suited for batch processing?

- (a) Process control
- (b) Video game control
- (c) Preparing pay bills of employees
- (d) Preparing mailing addresses

\*89. Locality of reference implies that the page reference being made by a process

- (a) will always be to the page used in the previous page reference.
- (b) is likely to be one of the pages used in the last few page references.
- (c) will always be one of the pages existing in memory.
- (d) will always leads to a page fault.

90. The correct matching for the following pairs

- |                          |                 |
|--------------------------|-----------------|
| (A) Disk scheduling      | (1) Round robin |
| (B) Batch Processing     | (2) SCAN        |
| (C) Time sharing         | (3) LIFO        |
| (D) Interrupt processing | (4) FIFO        |

is:

- (a) A-3, B-4, C-2 and D-1
- (b) A-4, B-3, C-2 and D-1
- (c) A-2, B-4, C-1 and D-3
- (d) A-2, B-1, C-4 and D-3

91. Thrashing

- (a) reduces page I/O
- (b) decreases the degree of multiprogramming
- (c) implies excessive page I/O
- (d) improves the system performance

92. Dirty bit for a page in a page table

- (a) helps avoid unnecessary writes on a paging device
- (b) helps maintain LRU information
- (c) allows only read on a page
- (d) none of the above

\*93. Each process  $P_i$ ,  $i = 1, 2, 3, \dots, 9$  is coded as follows.

```
repeat
    P(mutex)
    { critical section }
    V(mutex)
forever
```

The code for  $P_{10}$  is identical except that it uses  $V(mutex)$  instead of  $P(mutex)$ . What is the largest number of processes that can be inside the critical section at any moment?

- (a) 1                      (b) 2                      (c) 3                      (d) none of the above

94. When an interrupt occurs, an operating system

- (a) ignores the interrupt  
 (b) always changes the state of the interrupted process after processing the interrupt.  
 (c) always resumes execution of the interrupted process after processing the interrupt  
 (d) may change the state of the interrupted process to "blocked" and schedule another process.

### Answers

- |          |             |          |          |          |
|----------|-------------|----------|----------|----------|
| 1. c     | 2. c        | 3. c     | 4. c     | 5. c     |
| 6. c     | 7. b        | 8. a     | 9. a     | 10. d    |
| 11. d    | 12. b, c    | 13. a    | 14. b    | 15. b    |
| 16. b    | 17. c       | 18. d    | 19. c    | 20. b    |
| 21. b, d | 22. c       | 23. b, c | 24. b, d | 25. c    |
| 26. a    | 27. b       | 28. c    | 29. d    | 30. c    |
| 31. b    | 32. a       | 33. c    | 34. b    | 35. c    |
| 36. a    | 37. a, c    | 38. d    | 39. c    | 40. b, d |
| 41. c    | 42. a, b, d | 43. c    | 44. d    | 45. b, d |
| 46. d    | 47. b       | 48. a, c | 49. d    | 50. c    |
| 51. c    | 52. a       | 53. b    | 54. c    | 55. c    |
| 56. b    | 57. c       | 58. c    | 59. b    | 60. a    |
| 61. a    | 62. a       | 63. c    | 64. b    | 65. c    |
| 66. c    | 67. b       | 68. d    | 69. c    | 70. d    |
| 71. a    | 72. a       | 73. a    | 74. b    | 75. a    |
| 76. b    | 77. c       | 78. b    | 79. c    | 80. d    |
| 81. a    | 82. b       | 83. b    | 84. b    | 85. c    |
| 86. d    | 87. c       | 88. c, d | 89. b    | 90. c    |
| 91. c    | 92. a       | 93. c    | 94. d    |          |

## Explanations

15. Deadlock occurs when each of the 3 user processes hold one resource and make simultaneous demand for another. If there are 4 resources one of the 3 user processes will get the fourth instance of the resource and relinquish one or both of the resource(s) it is currently holding after using.
25. Each P operation will decrease the semaphore value by 1 and V operation increases it by 1. If  $x$  is 18, then 7 P operations will make semaphore value 0. If this is followed by 7 V operations the value comes back to 7. So, after 18 P and 18 V operations, the value of the semaphore will be 7. The remaining 2 P operations result in the semaphore value 5.
29. Even in a non-multiprogramming system, memory protection may be used, when, for example, spooling is being used.
32. 2 process can never lead to deadlock as the peak time demand of 6 (3 + 3) tape drives can be satisfied. But 3 processes can lead to a deadlock if each hold 2 drives and then demand one more.
41.  $0.35 \times 10 + (1 - 0.35) \times 100 = 68.5$  ns
44. The arrival pattern is a Poisson distribution.  
 $P(k \text{ requests}) = e^{-\mu T} (\mu T)^k / k!$   
 Here  $k = 0$ ,  $\mu = 20$ ,  $T = 3/4$ .  
 So required probability is  $e^{-15}$
52. It is  $0 + 10 + (15 - 2) + (18 - 5) + (38 - 10)$  divided by 5, i.e., 12.8 ms.
53. It is  $8 + 0 + 3 + 15 + 8$  divided by 5, i.e., 6.8 ms.
54. It is  $10 + 3 + 0 + 15 + 0$  divided by 5, i.e., 5.6 ms.
55. It is  $30 + 0 + 3 + 3 + 18$  divided by 5, i.e., 10.8 ms.
56. It is  $30 + 3 + 0 + 5 + 0$  divided by 5, i.e., 7.6 ms.
60. Having 11 resources ensures that at least 1 process will have no pending request. This process after using will release the resources and so deadlock can never occur.
61. At least one process will be holding 2 resources in case of a simultaneous demand from all the processes. That process will release the 2 resources, thereby avoiding any possible deadlock.
62. Using Banker's algorithm, one can show that one process has to acquire all its needed resources. This process, after completing its task, will release all its resources, thereby avoiding any possible deadlock.
63. The first 3 references A, B, C fills the internal storage with A, B, C in 3 page transfers. Now the next reference D results in a page fault. So, page A is downloaded and D takes its place after a page transfer. So, the internal store has D, B and C. The next reference is A—results in a page fault. So, a page transfer takes place and swaps B and A. Continuing this way, we find totally 9 page transfers are necessary.
67. Refer Qn 63. Applying the same logic, we find the required number of page transfer is 10. So, increasing the number of pages need not necessarily reduce the number of page faults. It is the actual sequences of references that decides.

69. To specify a particular segment, 5 bits are required (since  $2^5 = 32$ ). Having selected a page, to select a particular byte one needs 10 bits (since  $2^{10} = 1$  K byte). So, totally  $5 + 10 = 15$  bits are needed.
70. The disk drive has to traverse totally 146 cylinders (verify). So, seek time is  $6 \times 146 = 876$  ms.
72. To cover  $2400 \times 62500$  bits, 60 s are needed. Average latency time is the time needed to traverse 100 tracks i.e.,  $100 \times 62500$  bits, which is 2.5 s.
78. As this scheduling policy has a better average response time for the interactive users.
79. When it tries to access 0100, it results in a page fault as the memory is empty right now. So, it loads the second page (which has the addresses 100-199). Trying to access 200 will result in a page fault, as it is not in memory right now. So the third page with the addresses from 200 to 299 will replace the second page in memory. Trying to access 430 will result in another page fault. Proceeding this way, we find trying to access the addresses 0510, 0120, 0220, and 0320 will all result in page faults. So, altogether 7 page faults.
80. There are 64 words in a block. So the 4 K cache has  $(4 \times 1024)/64 = 64$  blocks. Since 1 set has 4 blocks, there are 16 sets. 16 sets needs 4 bits for representation. In a set there are 4 blocks, which needs 2 bits. Each block has 64 words. So, the word field has 6 bits.
81. The memory configuration after the arrival of the jobs 1, 2 and 3 and the termination of job 1 can be depicted as:

FREE-200	JOB2-350	JOB3-300	FREE-150
----------	----------	----------	----------

First fit algorithm will allocate the FREE-200 slot for job 4. But best fit algorithm will allocate the FREE-150 slot for job 4. The memory configuration for the first fit and best fit will be

JOB4-120	FREE-80	JOB2-350	JOB3-300	FREE-150
----------	---------	----------	----------	----------

and

FREE-200	JOB2-350	JOB3-300	JOB4-120	FREE-30
----------	----------	----------	----------	---------

respectively. When job 5 arrives, it will be allotted the FREE-150 slot by the first fit algorithm and the FREE-200 slot by the best fit algorithm. The memory allocation table for the first fit and best fit will be

JOB4-120	FREE-80	JOB2-350	JOB3-300	JOB5-150
----------	---------	----------	----------	----------

and

JOB5-150	FREE-50	JOB2-350	JOB3-300	FREE-30
----------	---------	----------	----------	---------

When Job 6 arrives, it will be allotted the FREE-80 slot by the first fit algorithm. The best-fit algorithm will find no room to store Job 5 as the needed 80 K, is not available contiguously. So, it has to wait till a job terminates. So, the first-fit algorithm performs better in this case.

89. Locality of reference is based on the fact that a page that is referenced is likely to be referenced again in the near future.

93. Let the mutex be initialized to 1. Any one of the 9 processes  $P_i$ ,  $i = 1, 2, 3, \dots, 9$  can get into the critical section after executing  $P(\text{mutex})$  which decrements the mutex value to 0. At this time  $P_{10}$  can enter into the critical section as it uses  $V(\text{mutex})$  instead of  $P(\text{mutex})$  to get into the critical section. As a result of this, mutex will be incremented by 1. Now any one of the 9 processes  $P_i$ ,  $i = 1, 2, 3, \dots, 9$  (excepting the one that is already inside the critical section) can get into the critical section after decrementing the mutex to 0. None of the remaining processes can get into the critical section.

If the mutex is initialized to 0, only 2 processes can get into the critical section. So the largest number of processes is 3.