CSC 341 Operating Systems Spring '18

Final

Time allowed for this test is 120 min.

This is a closed textbook, no-notes test.

You are allowed to access this questionnaire on a classroom workstation via the Internet browser at the address provided by the professor.

Except for that, NO ACCESS is allowed to any software, Internet, computers, calculators, iPads, iPhones, cellphones, removable memories, discs, etc., or to any information pertaining to this course.

INSTRUCTIONS - READ CAREFULLY

There are 50 questions on this exam. Questions 1 through 48 are multiple-choice questions. Questions 49 through 50 are essay-type questions. Additionally, question 51 is an optional make-up question for the midterm.

For all multiple-choice questions, use scantron form 882-E like this:

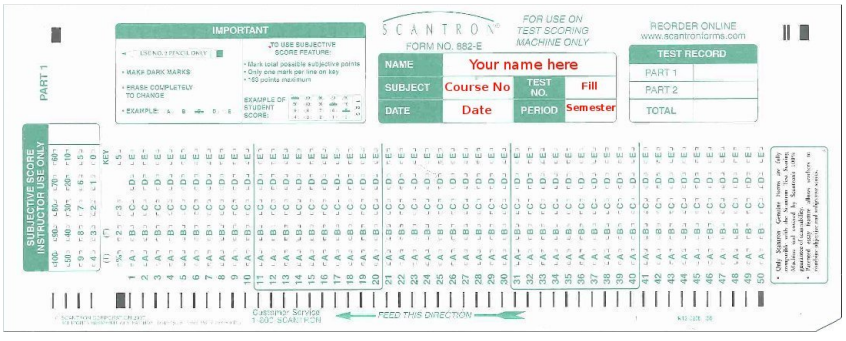

For all essay-type questions, use scratch paper provided by the professor. Do not use your own scratch paper for anything at any time.

Each multiple-choice question is worth 1 point of credit. Each essay-type question is worth 5 points of credit. The make-up question is worth 4 points of credit.

Pick one answer for each multiple-choice question and mark it clearly on your scantron. If none of the choices provided for any question seems correct, or more than one choice provided seems correct, chose the answer that, in your opinion, is the closest to the correct one.

Some questions are repeated as "Same as above" questions in order to allow partial credit for some partially correct answers. Although the correct answer is always the same for both, you may chose different answers for each of the two in the case you are not sure which one is correct. However, by doing so you will get less than the maximum score.

Show all work for full credit on each essay-type question.

THE QUESTIONS BEGIN HERE

PART I: MULTIPLE-CHOICE QUESTIONS

1. Under which one of the following circumstances *preemption* takes place?

   (A) A running process was sent to a waiting queue.

   (B) A ready process was sent to a waiting queue.

   (C) A running process was sent to the ready queue.

   (D) A waiting process was sent to the ready queue.

   (E) None or more than one of the above indicates preemption.

2. Same question as above.

3. Which one of the following optimization criteria were established in order to maximize utilization of CPU?

(A) minimizing throughput

(B) maximizing turnaround time

(C) maximizing waiting time

(D) minimizing CPU utilization

(E) None or more than one of the above.

4. Same question as above.

5. What is the main purpose of the RR scheduling algorithm?

   (A) To secure equal access to CPU for all processes.

   (B) To maximize CPU utilization.

   (C) To approximate SJF scheduling algorithm.

   (D) To prevent endless loops.

   (E) All of the above.

6. The following is a scenario for scheduling of four processes P_1, P_2, P_3, P_4 visualized on a Gantt chart:

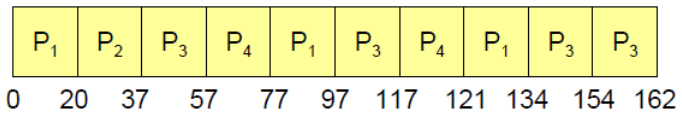

   What is the average waiting time in this scenario? Assume that all the processes arrived in the ready queue at time 0 and that at time 162 the ready queue became empty.

   (A) less than 70.

   (B) **CORRECT** 73.

   (C) 76.

   (D) 77.

   (E) More than 80.

7. Same question as above.

8. Which of the following pertains to a *logical address* in main memory?

   (A) An address that is imaginary and is not physically equal to any address that is actually stored in any register or buffer.

   (B) Every address that is received by the main memory logic.

   (C) Every address that is stored in the address register.

   (D) An address that is a logical consequence of asserted conditions.

   (E) None of the above.

9. Consider implementation of paged memory organization visualized on the figure below.

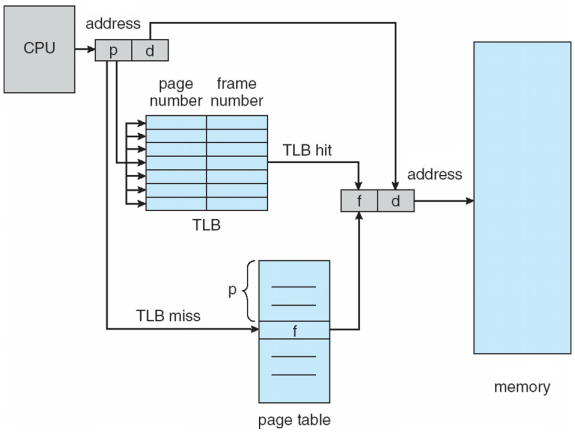

   Assuming that the page table is stored in memory that has average access time = 1 unit of time, the associative lookup is ε units of time, and the TLB hit ratio is α, what is the *effective access time* to the memory for that implementation? Indicate the most accurate formula.

   (A) $1 + \varepsilon$

   (B) $1 + \varepsilon\alpha$

   (C) $2 - (1 - \varepsilon)\alpha$

   (D) $2 + \varepsilon - \alpha$

   (E) None of the above.

10. Same question as above.

11. Assuming that the array declared as

```
20   int [][] data = new int [128][128];
```

is stored one row per page in virtual memory that was allocated 1 frame, what is the *minimum* number of page faults produced by the following memory-resident program?

```
23                  for (int i = 0; i <128; i++)
24                    for (int j = 0; j < 128; j++)
25                          data[j][i] = 0;
```

(A) 0

(B) 1

(C) 128

(D) 256

(E) 16,384

12. Given the following contents of the page table, and assuming the page size = 256 words, what is the physical address of the word at the logical address 517?

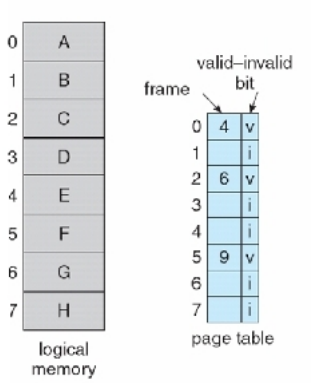

(A) 523

(B) 1,283

(C) 1,541

(D) 2,307

(E) invalid address

13. Same question as above.

In questions 14 through 19, using the diagram below, identify six steps of handling a page fault in demand paging.

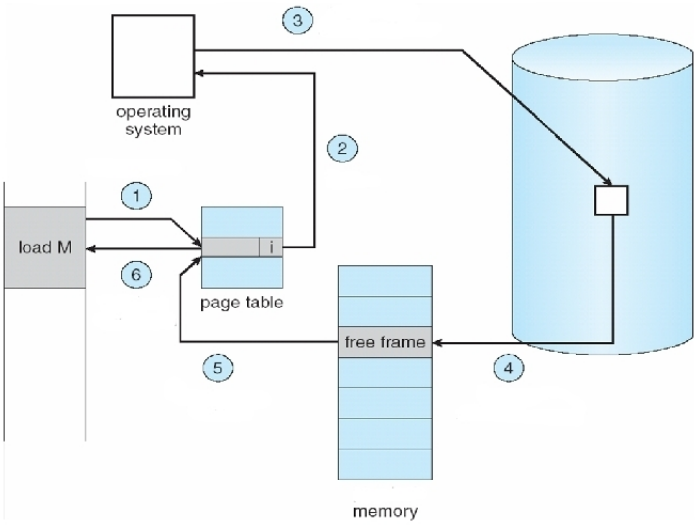

14. Step 1 is:

    (A) Write address into page table

    (B) Fetch the contents indicated by the arrow to the Instruction counter

    (C) Reference logical address

    (D) Issue a System Call

15. Step 2 is:

    (A) Copy one entry to OS memory

    (B) Issue an interrupt (a trap) to OS

    (C) This is not a step: it indicates that the page table is under control of OS

    (D) OS will execute instruction load M on behalf of the user process

16. Step 3 is:

    (A) Swap OS out to the system disk file

    (B) Write a page on the disk

(C) Find page on the disk

(D) Swap the waiting process out from the main memory

17. Step 4 is:

    (A) Link a disk block to a free frame

    (B) Swap the waiting process into a free frame

    (C) Bring in the missing page

    (D) Mark a free frame for the use of OS

18. Step 5 is:

    (A) Write the free frame into page table

    (B) Redirect calls from the disk to the page table

    (C) Free an entry in the page table

    (D) Update and reset the page table

19. Step 6 is:

    (A) Copy an entry in the page table into the main memory

    (B) Restart instruction load M

    (C) Override instruction load M with a system call

    (D) redirect page table references as indicated by the arrow

20. Given memory-access time of 100 nanoseconds and page-fault service time of 60 microseconds, what is the maximal probability of page fault that would make the effective access time for a demand-paged memory not larger that 110 nanoseconds?

    (A) 0.06

    (B) 0.00017

    (C) 0.0000001

    (D) 0.0000000001

21. Same question as above.

22. Given the following address sequence:

    0100, 0101, 0232, 0311, 0404, 0100, 0102, 0103, 0233, 0252, 0532, 0104, 0100, 0101, 0233, 0312, 0405, 0532,

    determine which of the page replacement algorithms (out of those discussed in class) reveals Belady anomaly when the number of frames is increased from 3 to 4? **Assume that the page size is 100**.

    (A) FIFO
    (B) LRU
    (C) Optimal
    (D) All of the above
    (E) None of the above

23. Same question as above.

24. Which of the following schemes is an example of application of optimal page-replacement algorithm?

    (A)
    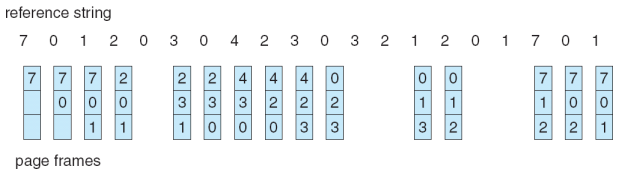

    (B)
    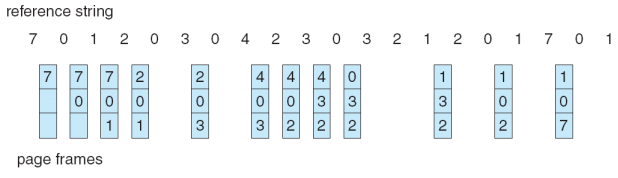

    (C)
    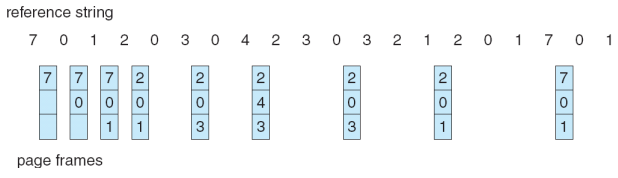

    (D) None of the above.

25. Same question as above.

26. Assume that the so-called *links* to subdirectories are allowed in tree-structured directories. Such an arrangement may result in the so-called *general-graph directories*, an example of which is visualized on the following picture:

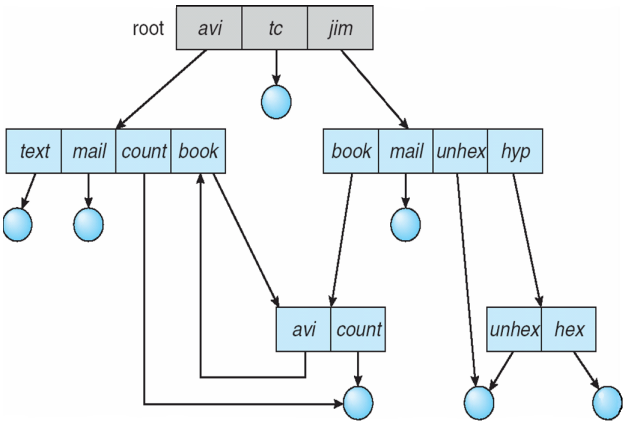

In cases like that, what is the computationally fastest way of traversing the entire directory?

(A) Traverse the directory following its tree structure *and* the links.

(B) Traverse the directory following its tree structure *but not* the links.

(C) Run the cycle detection program before traversing the directory in order to detect falling into an endless cycle during the traversal.

(D) Run the cycle detection program while traversing the directory in order to detect falling into an endless cycle during the traversal.

(E) Limit the number of directories that will be accessed during the traversal.

27. Same question as above.

28. The following picture visualizes a physical structure of magnetic disk:

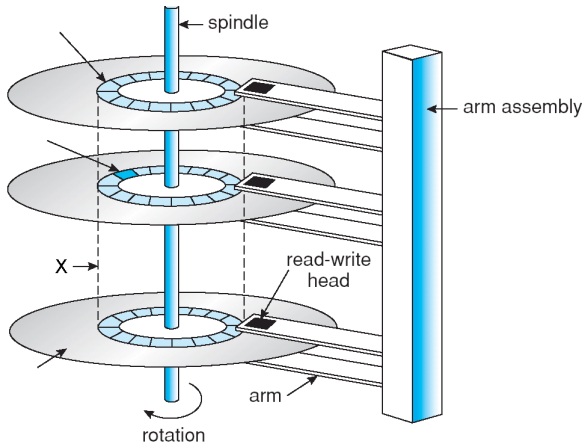

What is the name of the part of disk's physical addressing scheme indicated with a letter X on that picture?

(A) distance from the spindle

(B) sector

(C) track

(D) capacity limit

(E) none of the above

29. Same question as above.

30. What is disk *average latency*?

(A) An average time to move the read/write head from one track to another.

(B) An average time to perform the readback after write.

(C) Half time of one disk revolution.

(D) An average time to adjust the angular speed of the disk.

(E) An average delay of a late access to a disk block.

31. Same question as above.

32. Assuming that the only limitation on speed of moving the read/write had from one cylinder to another is a maximal linear acceleration, how are the seek times for

- moving the read/write head from cylinder 4 to cylinder 16, and

- moving the read/write head from cylinder 16 to cylinder 64

related?

(A) They are approximately the same

(B) The former is about half of the latter.

(C) The former is about $\frac{1}{4}$ of the latter.

(D) The former is about $\frac{1}{16}$ of the latter.

33. Same question as above.

34. Assume that file F of size of 5 GB has a multi-level index allocation of disk blocks, as discussed in class. Assuming that all file blocks and the index blocks are kept on the disk and not in the main memory, and assuming that one index block can hold up to 1,024 disk block addresses, what is the number of disk accesses necessary to perform a read from file F in the worst case?

(A) 1.

(B) 2.

(C) 3.

(D) 4.

(E) 5 or more.

35. Same question as above.

36. Assume that file F of size of 1 MB has a linked allocation of disk blocks, as discussed in class. Assuming that all file blocks are kept on the disk and not in the main memory and no FAT table is used, and assuming that one file block can hold up to 4,092 bytes of data plus one disk block address, what is the number of disk accesses necessary to perform a read from file F in the worst case?

(A) 1.

(B) 2.

(C) 257.

(D) 4,092.

(E) 1 million or more.

37. Same question as above.

38. What is the name of the disk scheduling algorithm depicted below?

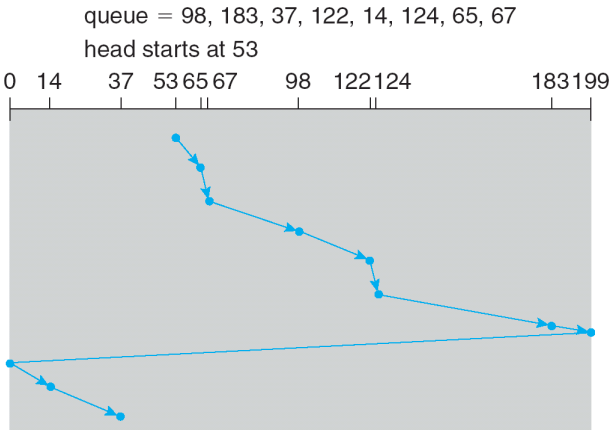

(A) FCFS

(B) SSTF

(C) SCAN

(D) C-SCAN

(E) C-LOOK

39. What is the name of the disk scheduling algorithm depicted below?

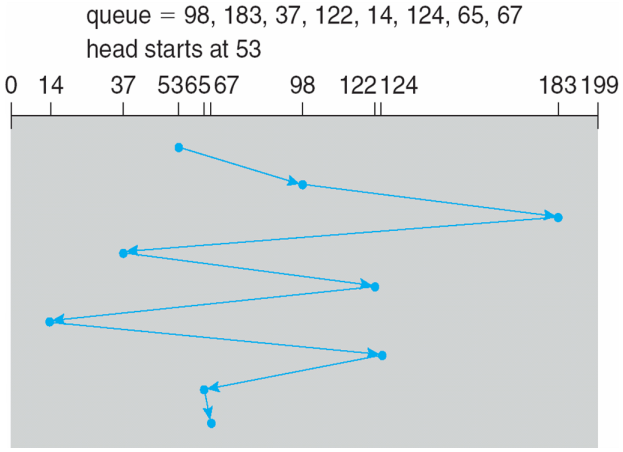

(A) FCFS

(B) SSTF
(C) SCAN
(D) C-SCAN
(E) C-LOOK

40. What is the name of the disk scheduling algorithm depicted below?

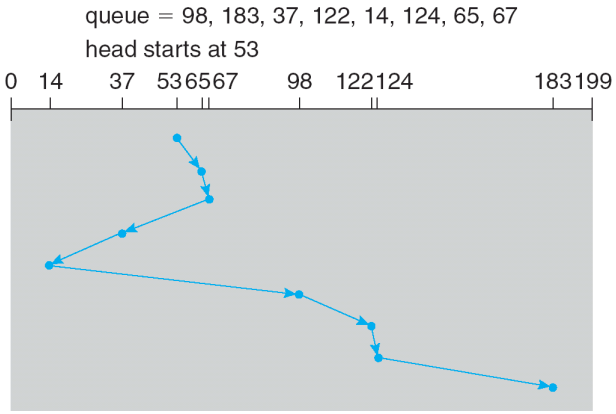

(A) FCFS
(B) SSTF
(C) SCAN
(D) C-SCAN
(E) C-LOOK

41. What is the name of the disk scheduling algorithm depicted below?

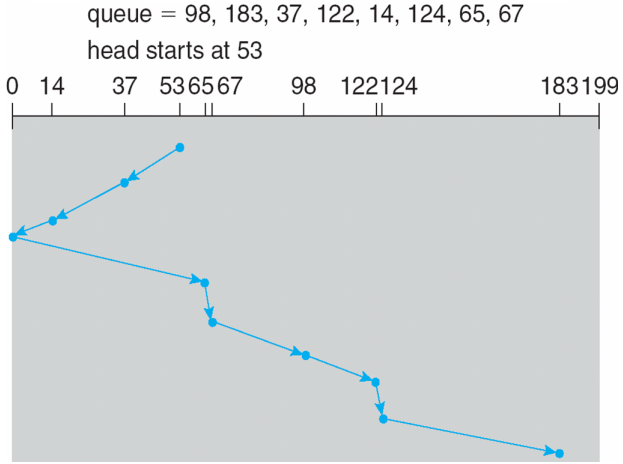

(A) FCFS

(B) SSTF

(C) SCAN

(D) C-SCAN

(E) C-LOOK

42. What is the name of the disk scheduling algorithm depicted below?

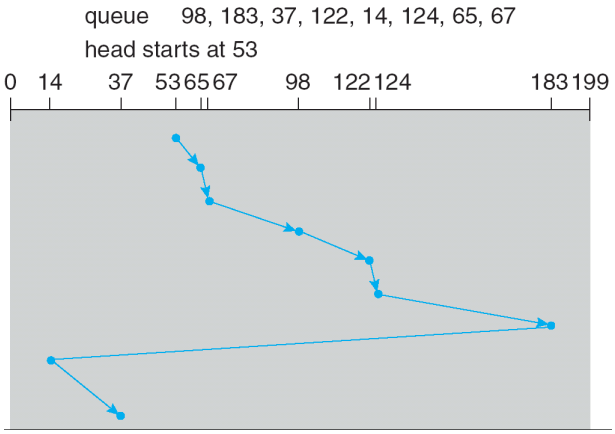

(A) FCFS

(B) SSTF

(C) SCAN

(D) C-SCAN

(E) C-LOOK

43. What is the value of the radius r of the innermost track of a magnetic disk with 32 sectors per track that yields the maximum capacity per platter if the radius R of the outermost track is 3.5 inches?

(A) as close to 0 inches as technologically practical

(B) 0.75 inches

(C) 1.75 inches

(D) 2.75 inches

(E) 3.5 inches.

44. Same question as above.

45. Which of the following disc allocation methods is well-suited for transfers of large amounts of data by a direct memory access (DMA) controller?

    (A) Contiguous allocation.

    (B) Linked allocation.

    (C) Indexed allocation.

    (D) More than one of the above

    (E) None of the above.

46. Same question as above.

47. What is the size of the bit vector needed to manage free space (assuming that the bit-vector free disk space management is being used) of a 1 gigabyte disk with 4,096-byte blocks?

    (A) 1 megabyte

    (B) 256 kilobytes

    (C) 32 kilobytes

    (D) 4,096 bytes

    (E) None of the above.

48. Same question as above.

PART II: ESSAY-TYPE QUESTIONS

49. (A) (1 point) Give the formula (used and derived in class) for the average time $\mathtt{W_{avg}}$ a process waits in the ready queue between its arrival and subsequent departure, given the number $\mathtt{n}$ (of processes) and, for each process $\mathtt{P_i}$ where $\mathtt{i} = 1, ..., \mathtt{n}$, $\mathtt{P_i}$'s arrival time $\mathtt{T_{arrive}}(\mathtt{P_i})$ and departure time $\mathtt{T_{depart}}(\mathtt{P_i})$. Assume that process $\mathtt{P_n}$ is the one that departs the last from the ready queue.

    (B) (4 points) Derive the formula you gave in question 49A.

50. Calculate the capacity of a maximally utilized 3-zone CAV disk, with one surface only, that rotates at 7,200 RPM and whose innermost track accommodates 32 sectors, 512 bytes each. Assume that the maximum (as opposed to actual) capacity of adjacent tracks differs by exactly 8 bytes. A simplistic and not-to-scale scheme of the disk is shown below.

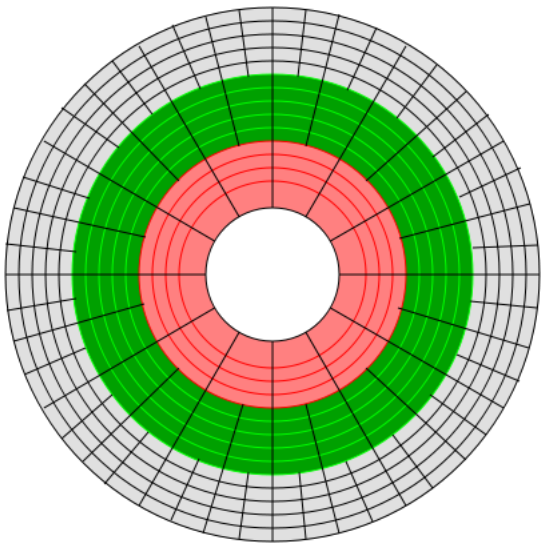

PART III: A MAKE-UP QUESTION FOR THE MIDTERM

The only effect of this make-up question is that its score will add to your previous score on the Midterm. In particular, it will not carry over to scores of quizzes and final.

51. Let each resource type be assigned a unique integer number that is distinct from the number assigned to any other resource type. Prove that the following policy prevents deadlocks resulting from processes' requests for resource allocation:

A process is not allowed to request a resource whose resource-type number is less than or equal to the resource-type number of any resource that is currently being held by the said process.

Hint: Use a resource allocation graph.

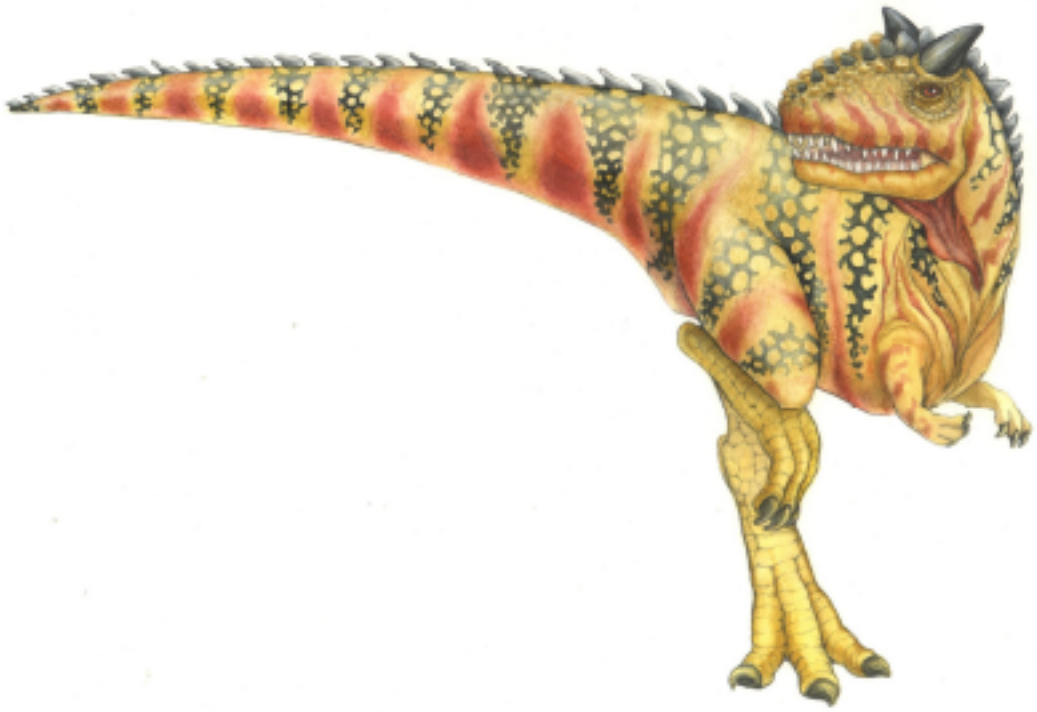

THAT'S ALL, FOLKS.