

25/11/13

(12)

B.E./B.Tech. END SEMESTER EXAMINATIONS, NOV./DEC. 2013
DEPARTMENT OF INFORMATION SCIENCE AND TECHNOLOGY

Branch : Information Technology

V Semester

CS 9023-UNIX INTERNALS

Time : 3 Hrs

Max: 100 Marks

Part A (10 x 2 = 20 Marks)

Answer all questions

1. Give the layout of the UNIX file system.
2. Does the bread algorithm always initiates a disk read? Justify.
3. List any two additional fields of the in-core copy of the inode. Why are they not present in the disk inode?
4. What does the namei algorithm do? Name two system calls that use it.
5. What are kernel data structures modified by the OPEN system call? How are they modified?
6. Compare an unnamed pipe with a named pipe.
7. Define the context of a process.
8. What is the process scheduling policy in UNIX?
9. What is a swap map?
10. List the ways by which processes communicate with each other in UNIX.

Part B (5 x 16 = 80 Marks)

11. Discuss with diagrams the five scenarios the Unix kernel may follow while allocating a buffer for a disk block.
 12. (a) (i) Explain the structure of a disk inode. Explain the algorithm that allocates an in-core inode given a disk inode. (10)

(ii) Assume that there are 16 inodes per block and block 2 is the beginning of the inode list, compute the logical disk block that contains inode number 40. Assuming that the size of the disk inode is 32 bytes, compute the byte offset of the inode number 40 in that block. (6)
- OR
- (b) (i) Explain the structure of the super block. Explain the algorithm that assigns a disk inode to a newly created file. (10)

(ii) Explain how free disk block numbers are maintained in the super block. What is the main difference between the free inode list and free block list of the super block? (6)

13 (a) Explain the algorithm which implements the OPEN system call. Mention clearly the kernel data structures used and how its fields are used / updated. When does the algorithm return an error? What are the low level algorithms used by it?

OR

(b) Explain the algorithm which implements the PIPE system call. Mention clearly the kernel data structures used and how its fields are used / updated. When does the algorithm return an error? What are the low level algorithms used by it?

14 . (a) (i) Explain with a diagram, the various state transitions that a process undergoes during its lifetime in a Unix system. (8)

(ii) Explain any **ONE** algorithm that manipulates region mentioning input(s), output(s), purpose and system call(s) that uses it. (8)

OR

(b) (i) Explain how the fork system call is implemented in Unix. (10)

(ii) Explain the algorithm for Process Scheduling. (6)

15. (a) (i) Explain the data structures used in the demand paging memory management system of Unix. (8)

(ii) When does the system incur validity page fault? Explain clearly how it is handled. (8)

OR

(b) Write notes on the (i) I/O subsystem of UNIX. (8)

(ii) Interprocess Communication in UNIX. (8)