LECTURE - 31

Lecture Outline

- Some queuing theory
- UNIX Fast File System
-
- *Scribe for today?*

I/O Performance

- Producer-Server Model

- Throughput vs. Response Time

- Response time and think time

- Queuing theory

  – Arrival rate, service time, utilization

  – Little's law

  – Squared coefficient of variance

  – Average residual time

  – Response time and utilization

  – M/G/1 and M/M/1 models

UNIX's Old File System

- Superblock

- Free-list

- Directory: special file – has pointer to file's inode

- Inodes have:
  - Direct points, singly indirect pointers, doubly indirect, and triply indirect pointers

- Problem: file's blocks get distributed all over the disk, deteriorating performance

- Also, block size: 512 bytes (poor performance)

UNIX Fast File System (FFS)

- Cylinder groups are defined

- Inodes are close to data blocks

- Block size: 4096 bytes

  – But, poor disk usage (close to 50% wasted)

- Idea: fragment blocks

  – But only last block of file is allowed to be fragmented

- All files of a directory are preferably in same cylinder group

- Other enhancements: long file names, file locking, symbolic links, rename, quota