COS214 Tutorial 6

Roberto Togneri, 2000

1. (a) What is "device independence"?
   (b) In which of the four I/O software layers is each of the following done.
      (i) Computing the track, sector, and head for a disk read.
      (ii) Maintaining a cache of recently used blocks.
      (iii) Writing commands to the device registers.
      (iv) Checking to see if the user has permission to use the device.
      (v) Converting binary integers to ASCII for printing.

2. A user process is accessing data from a large file and performs 0.5 seconds of computation for each 128K of data read. The disk rotation speed is 5000 rpm, average seek time is 2 msec, and data transfer rate is 10 MB/sec. Memory to memory data transfer rate is 100 nsec per 4-byte word. Calculate the CPU time used by the user process for every 5 minutes wall-clock time, and the amount of memory used by kernel buffers, assuming:
   (a) no buffering
   (b) single buffer
   (c) double buffer

   Assume the process is the only active one in the system and the file is stored contiguously on the disk.

3. A local area network (LAN) is used as follows. The user issues a system call to write to the network. The OS then copies the data to a kernel buffer. Then it copies the data to the network controller board. When all the bytes are safely inside the controller, they are sent over the network at a rate of 10 Mbits/sec. When the last bit arrives, the destination CPU is interrupted, and the kernel copies the new data to a kernel buffer to inspect it. Once it has figured out which user they are for, the kernel copies the data to the user space. If we assume that each interrupt and its associated processing takes 1 msec, that packets are 1024 bytes (ignore the headers), and that copying a byte takes 1 μsec, what is the maximum rate at which one process can pump data to another? Assume that the sender is blocked until the work is finished at the receiving side and an acknowledgment comes back. For simplicity, assume that the time to get the acknowledgment back is so small it can be ignored.

4. (a) A floppy disk has 40 cylinders (an old one anyway, most now have 80). A seek takes 6 msec per cylinder moved. If no attempt is made to put the blocks of a file close to each other, two blocks that are logically consecutive (i.e. follow each other in the file) will be about 13 cylinders apart on average. If, however, the operating system makes an attempt to cluster related blocks, the mean interblock distance can be reduced to 2 cylinders (for example). How long does it take to read a 100 block file in both cases, if the rotational latency is 100 msec and the transfer time is 25 msec per block?

   (b) Why are output files for the printer normally spooled on disk before being printed, instead of being printed directly from the application program?

5. Disk requests come in to the disk driver for cylinders 10, 22, 20, 2, 40, 6 and 38, in that order (Initially, the arm is at cylinder 20). A seek takes 6 msec per cylinder moved. How much seek time is needed for

   (a) First-Come, First Served (FCFS).
   (b) Shortest Seek First (SSF).
   (c) SCAN elevator algorithm (initially moving upwards)

6. A personal computer salesman visiting a university in Tasmania remarked during a sales pitch that his company had devoted substantial effort to making their version of UNIX very fast. As an example, he noted that their disk driver used the elevator algorithm and also queued multiple requests within a cylinder in sector order. A student, Harriet Hacker, was impressed and bought one. She took it home and wrote a program to randomly read 10,000 blocks spread across the disk. To her amazement, the performance measured was identical to first come first served. Was the salesman lying?