CS311 - COMPUTER SYSTEMS, I

Homework #8 - Due: Friday, December 14, at 4 pm

Emphasis: The IO System; Parallel Architectures; Performance

1. A certain disk has the following parameters:

   • 4 surfaces
   • 4096 tracks per surface
   • Each track consists of 1024 sectors.
   • The size of one sector is 512 bytes.
   • The disk rotates at 7500 RPM (revolutions per minute).
   • The average seek time is 12 ms.

   a. What is the total storage capacity of this disk in bytes?

   b. In gigabytes? (Remember 1 gig = 1024 meg; 1 meg = 1024 K; and 1 K = 1024)

   c. What is the average access time? (Give your answer in ms)

   d. What is the data transfer rate (once the head is correctly positioned)? (Give your answer in bytes / second)

   e. Express your answer to the previous problem in MB / second. (Remember (b) above)

   f. How long (best case) would it take to transmit the entire contents of this disk over Gordon's original Internet connection (57,600 <u>bits</u> per second)? (Give an approximate answer in days - consider only the speed of the network connection, which is so much slower than the disk as to make the time for the latter negigible.)

2. Recall that on homework #3 you calculated the following storage requirements for various types of data:

   a. A 500 page book, if each page contains an average of 2000 characters = 1 million bytes

   b. A 100 x 100 pixel black and white image = 1250 bytes

   c. A 100 x 100 pixel color image, using 24-bit color = 30,000 bytes

   d. One minute of sound, sampled at 22 kHz using 8-bit samples = 1,320,000 bytes

   Now determine, for each case, how many items the disk in the preceeding problem could store if it were entirely dedicated to this one task. (Of course, it cannot storing a fraction of an item is meaningless, so give the greatest integer ≤ the computed value)

3. Suppose the disk of the previous two problems is used for a file system that is clustered using 8KB blocks, each consisting of 16 contiguous sectors. What is the average time needed to access and then read (or write) an 8KB block of data?

4. A program repeatedly performs a three-step process: It reads in a block of data from the disk, does some processing on that data, and then writes out the result as another block elsewhere on the same disk. The average total time needed to access and transfer a block of data is 20 ms. No other program is using the disk or the processor. The processing step takes 20 million clock cycles, and the clock rate is 400 MHz. What is the overall speed of the system in blocks processed per second?

   a. If there is no overlapping of IO and computation

   b. If disk IO can be performed in parallel with computation.

   c. If disk IO can be performed in parallel with computation, and the CPU is replaced with a 1.0 GHz CPU.

   d. For the last case (overlap of IO and computation, and a much faster CPU), how does the overall speedup compare with what you might expect given the much faster CPU? Why?

   e. Suppose you now replace the single disk with two otherwise identical disks - one used only for reading input data, and one used only for writing output data (and still using the faster CPU). Now what is the overall system speed? Have we now achieved the speedup one would expect from using a faster CPU?

5. Consider the task of doubling all the elements of a 8192 element vector, which might be done on a SISD machine by the following code. Suppose each iteration of the loop requires 20 clock cycles - so the whole process needs 163,840 clocks.

```
for (int i = 0; i < 8192; i ++)
  x[i] *= 2.0;
```

   a. Now suppose, instead, that the process is done on a pipelined processor having a single pipeline that supports memory to memory vector operations. Assume that each computation (fetch `x[i]` , double it, store result) still takes 20 clock cycles, but that the pipeline, once full, can produce one result each cycle. How many clock cycles total will be needed for this case?

   b. Now suppose, instead, that the process is done on a SIMD processor having 64 ALU's, each with its own private memory. Suppose each ALU uses 20 clock cycles to process one array element. Suppose, further, that the array elements are evenly divided across the memories of the various ALU's. How many clock cycles total will be needed for this case?

6. Suppose a collection of 5 benchmark programs was run on each of three different systems, yielding the results shown below. (Results are reported as time in milliseconds).

Program	System A	System B	System C
1	45	125	75
2	300	275	350
3	250	100	200
4	400	300	500
5	800	1200	700

   a. Compare the performance of the three systems using the arithmetic mean.

   b. Repeat the comparison, but using the geometric mean, with times normalized to System A.

   c. Are there any surprises? Explain