國立臺灣科技大學
九十三學年度碩士班考試試題

系所組別：資訊工程系
科　　目：作業系統

總分 100 分

1. Consider reading a file using a single-threaded file server and a multithreaded server. It takes 12 msec to get a request for work, dispatch it, and do the rest of the necessary processing, assuming that the data needed are in the block cache. If a disk operation is needed, as is the case one-third of the time, an additional 24 msec is required, during which time the thread sleeps. How many requests/sec can the server handle if it is single-threaded? How many requests/sec can the server handle if it is multithreaded? (10 %)

2. A soft real-time system has four periodic events with periods of 100, 150, 200, and 250 msec each. Suppose that the four events require 60, 30, 10, and x msec of CPU time, respectively. What is the largest value of x for which the system is schedulable? (10 %)

3. A system has p processes and r identical resources. Each process needs a maximum of m resources. What condition must hold to make the system deadlock free? (10 %)

4. A local area network is used as follows. The user issues a system call to write data packets to the network. The operating system 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 20 megabits/sec. The receiving network controller stores each bit a microsecond after it is sent. When the last bit arrives, the destination CPU is interrupted, and the kernel copies the newly arrived packet to a kernel buffer to inspect it. Once it has figured out which user the packet is for, the kernel copies the data to the user space. If we assume that each interrupt and its associated processing takes 0.5 msec, that packets are 1024 bytes (ignore the headers), and that copying a byte takes 1 microsec, what is the maximum rate, in bytes per second, 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 acknowledgement comes back. Ignore the acknowledgement time. (10 %)

5. How much cylinder skew is needed for a 6000-rpm disk with a track-to-track seek time of 1 msec? The disk has 240 sectors of 2 KB each on each track. (10 %)

國立臺灣科技大學
九十三學年度碩士班考試試題

系所組別：資訊工程系
科　　目：作業系統

6. (10%)Johnny is running the following experiment to study the effectiveness of multiprogramming in her operating system: He has chosen a fixed, fairly large, program **P**. The experiment involves concurrently running **N** processes, each executing **P** on the identical inputs (so the behavior should be identical) for values of **N** ranging from 1 to 100, and making a variety of measurements. One of the quantities that Johnny is measuring is the number of blocks read from disk. When N=1, the process reads M blocks from disk. Johnny expects, therefore, that in general N processes will read **N*M** blocks. What he finds is quite different. For small values of **N**, the number of blocks read is much smaller than **N*M**; in fact, it's barely larger then **M**, regardless of **N**. For large value of **N**, the number of blocks read is much larger than **N*M**.
   A. Explain why for small values of **N**, the number of blocks read is much smaller than **N*M**.
   B. Explain why for large value of **N**, the number of blocks read is much larger than **N*M**
7. (10%)Assume you have a fast-wide SCSI-II disk that operates at 20MB/sec. The particulars are 7200RPM, 512 B/sector, 160 sectors/track, 7000 cylinders, 20 tracks/cylinder and an average seek time of 8msec. We want to estimate the effective transfer rate for a random-access workload that reads individual sectors that are scattered across the disk.
   A. What is the average rotational latency?
   B. What is the sector read speed?
   C. What is the I/O rate per second for random sector access?
   D. What is the effective transfer rate for random sector access?
8. (10%)A machine that is using virtual memory has 32 bit virtual addresses and 256 MB physical memory. The page size is 4 KB. There are two programs, **A** and **B**, which will be run at different times. Program **A** has a size of 512MB and Program **B** has a size of 16 MB. Assume each page table entry requires 8 bytes of storage in memory.
   A. Compute the required page table size that program **A** requires, if the system is using inverted page tables. Assume each program requires a different inverted page table and assume the whole physical memory is dedicated to the program (although this may not be realistic).
   B. Do the above question for program **B**.
   C. Compute the required page table size for program **A** if the system is using *twolevel page tables*. Assume top-level page table index is 11 bits, and second-level page table index is 9 bits.
   D. Do the above for program **B**.
9. (10%)Consider the two-dimensional array A:

   int A[][] = new int[100][100];

   where A[0][0] is stored at location 200, in a paged memory system with pages of size 200 bytes. A small process (whose size is 100 bytes) resides in page 0 (locations 0 through 99) for manipulating the matrix A; thus every instruction fetch will be from page 0. Assume 3 page frames are allocated for the program (page frame 1, page frame 2, page frame 3). How many

國立臺灣科技大學
九十三學年度碩士班考試試題

系所組別：資訊工程系
科　　目：作業系統

page faults will be generated by the following array-initialization loops, using LRU replacement, and assuming page frame 1 has the process in it, and the other two are initially empty? Briefly explain your answers.

A.

```
for (int j = 0; j < 100; j++)
    for (int i = 0; i < 100; i++)
        A[i][j] = 0;
```

B.

```
for (int i = 0; i < 100; i++)
    for (int j = 0; j < 100; j++)
        A[i][j] = 0;
```

10. (10%)Assume (i) an empty physical memory of size *100 frames* at the beginning, (ii) the page frames are 4Kbytes, (iii) the page table entries are initialized after page faults. Consider the case where the page faults occur in pages *0, 50, 1, 51, 2, ...* What physical address does virtual address 8626 map to in physical memory?