請按題目順序作答 !!

[PART I] Operating System

NOTE: all the following questions require you to discuss your answers to receive full credits. You will NOT receive any credit if you only provide single word answer such as Yes or NO. To provide sound and complete answers, you may need to provide your own assumptions. Hence, you are advised to think thoroughly before writing your answers.

Bob is starting up a new computation service company to provide internet services, called *OnTime-or-Free*. The company claims that the service is charged only when the submitted computation is completed on time, i.e., 200ms or 100ms. During the Chinese New Year holiday in 2011, Bob worked in his own studio to build BCC, *Bob's Cloud Computing environment*, from scratch so as to start his own business in Spring 2011. Bob has to select the right combination for computing hardware and operating systems to provide high performance computing services with low cost. He plans to build a computer cluster; each of the computers is a commercial off-the-shelf computer and is connected to each other by computer networks.

1. (10 pts) While choosing the processors to build the computing node, Bob has two choices: symmetric multiple processors (SMP) systems and multicore processor systems.
    a. (4 pts) Both simultaneous multithreading (SMT) on multiprocessor system and multi-threading on multicore processors provide concurrent execution for threads. Please define these two threading models.
    b. (6 pts) Suppose that each SMP processor has *K* processors and each multicore processor has *K* cores. Please advise Bob to select one of the above two threading models for better performance when
        i. (3 pts) more than 90% of the threads in the applications conduct floating-point computation.
        ii. (3 pts) 50% of the threads in the applications conduct floating-point computation and the others conduct integer computation.

    Please discuss your answers.

2. (8 pts) On each computing node with paging, there are *M* processes to compete for *K* resources such as disks, pipes and sockets. Suppose each process needs *N*

見背面

resources at the same time to complete its task. Please answer the following questions.

a) (4 pts) Suppose the resources are all non-sharable, K=15, and N=3. For which values of M is the system deadlock free? Show your calculation.

b) (4 pts) Suppose the resources are all sharable, K=15, and N=3. Please complete the following figure to show the relationship between M and the system throughput. Explain your answer.

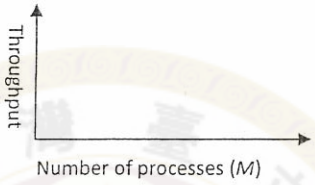

3. (12 pts) To prevent the processes from being blocked forever, Bob provides the following example code segment to his customers, where the *write()* function should wait only until the timer previously set by *alarm()*, i.e., 3 seconds, expires.

```
int fd; char buffer[100];
...                    // code ignored
alarm( 3 );            // set a timer for 3 seconds
write( fd, buffer, 100 );   // write 100 bytes from buffer to
the opened
                       // file descriptor fd
...                    // code ignored
```

a) (3 pts) Please define race condition.

b) (3 pts) Please describe the cause of the race condition in the above code segment. Such race condition could block the processes forever.

c) (6 pts) Please give two possible ways to avoid the race condition mentioned above. Your solutions must (1) prevent the processes from being blocked forever due to the *write()* function and (2) avoid polling.

4. (12 pts) After thoroughly thinking, Bob decides to use multicore processors to build BCC. To configure the operating system for best performance, Bob studies the potential workload and finds that each thread takes 200ms in average to complete its computation.

a. (3 pts) Suppose the thread is computation intensive. Please discuss if the scheduler on each multicore processor should migrate this thread to

接次頁

another idle core when the allocated core is now overloaded and the thread already finishes 50ms of computation.

b. (3 pts) Suppose the thread is IO intensive. Please discuss if the scheduler should migrate this thread to another idle core when the allocated core is now overloaded and the thread already finishes 50ms of computation.

c. (6 pts) Bob's customers ask him to put a limit on response time. Please help Bob give two CPU scheduling algorithms for each core to minimize average response time for the threads on each core and support real-time applications, respectively.

5. (8 pts) Both distributed operating systems and network operating systems can be used to build the could computing environment. Which of the above operating systems is more suitable when

a. (4 pts) the cloud servers are built with homogeneous computing hardware;

b. (4 pts) each of the servers prefers to manage their local resources.

Please discuss your answers to receive full credits.

[PART II] Computer Architecture

Note: [Part II]題目，請務必算出答案，勿只列式子,若無法整除，算至小數後第一位 (採四捨五入)

6. (20 pts) You are given a bus-based multiprocessor system which implements a snooping-based cache coherency protocol. There are two processors P1, P2 in the system. Each processor has its own L1 cache. The memory system specification is given below:

- L1 Cache: physical addressed, 64KB, direct-mapped, 32B block size, write-back cache
- Virtual Memory: 8K page, 1 GB virtual address space
- TLB: fully associative, 128 entries
- 64 MB of addressable physical memory.

x1, x2, and x3 represent one-word variables. The physical addresses (represented in hex) of x1, x2 and x3 are 0x00AF0804, 0x00AF0808, and 0x00AE0804, respectively. Answer the following questions considering the possible events: TLB-hit/miss, page-table hit/miss, L1-hit/miss, write-back, invalidate.

a. (3 pts) What are the sizes of the tag and data arrays of the TLB?

見背面

b. (2 pts) Please write down the correct order of the events listed above for detecting a page fault. Note that some events may not occur.

c. (15 pts) Assume TLB hits for all the memory requests. X1 and x2 are in the L1 caches of P1 and P2. Show the bus activities for each memory request.

Time	P1	P2	Bus activity
1	Write x1		
2		Read x2	
3	Write x1		
4		Read x3	
5	Read x2		

7. (5 pts) Multiple choice (全對才有分)

Which of the following statements about pipelining are correct?

(a) Two instructions with data dependency will cause data hazards.

(b) Assume a 5-stage pipeline of MIPS architecture. All the data hazards in the following code sequence can be avoided by data forwarding.

```
add $4, $2, $3   // $4 is the destination register
add $5, $4, $2
lw   $6, 8($5)
add $8, $6,$4
```

(c) The best speedup from pipelining is equal to the number of pipeline stages.

(d) Pipelining helps both instruction throughput and latency.

(e) Once an exception occurs, all instructions in pipeline need to be flushed.

(f) Control hazard can be resolved with dynamic branch prediction.

8. (15 pts) The following figure shows a typical I/O system. The storage contains HDDs with a read/write bandwidth of 128 MB/sec and the average seek and rotational latency of 2ms. The bandwidth of the front side bus, memory channels, and PCIe are denoted in the figure. Additional system configurations are listed below:

- Each core sustains 3 billion instructions per second
- The user program uses 200,000 instructions per I/O operation
- The operating system uses 100,000 instructions per I/O operation.

Note that all the numbers in the question are represented in base 10 (i.e., 1K = 1000). Please answer the following question assuming the workload of 64 KB

接次頁

random reads.

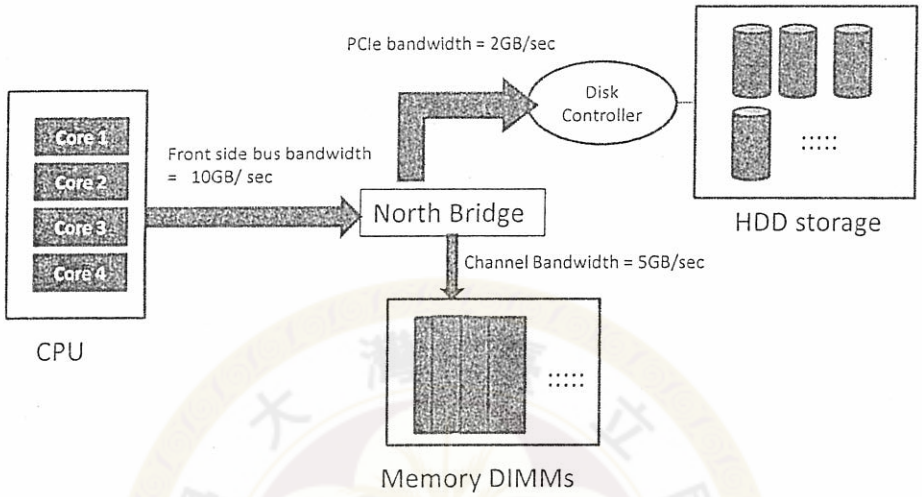

(a) (2 pts) What is the maximum sustained I/O rate that 4 cores can deliver?

(b) (3 pts) What is the maximum sustained I/O rate of 8 HDDs can deliver ignoring disk conflicts and assuming the disk controller is not the bottleneck?

(c) (5 pts) What is the maximum sustained I/O rate that the described I/O system can deliver with 8 HDDs? Please explain your answer.

(d) (5 pts) How many HDDs can the described I/O system accommodate?

9. (10 pts) You are considering adding a register-memory addressing mode to the MIPS instruction set. With this new addressing mode, there would be 5% increase in cycle time and the CPI of the ALU instruction type (including both the register-register and register-memory types) becomes 2.

(a) (2 pts) What is the average CPI of the old implementation?

(b) (3 pts) What type of instructions could be reduced?

(c) (5 pts) If half of the instructions that you answer in question (b) can be reduced with the new addressing mode, what is the average CPI of the new implementation? Will you decide to add the new addressing mode? Please explain your answer.

OP	Freq	CPI
ALU	40%	1
Load	20%	3
Store	20%	3
Branch	20%	2

試題隨卷繳回