國立政治大學　104　學年度碩士班招生考試試題

考試科目	作業系統 8141	所別	資訊科學系 8141	考試時間	03 月 01 日(星期日) 第二 節

本測驗一共有 14 題，每一題（或子題）所佔分數的比重如標示。請用中文或英文作答。

1. **[Introduction]** Please answer the followings:
   a) (2%) What is required for two machines belonging to a cluster to cooperate to provide a highly available service?
   b) (4%) Consider a computing cluster consisting of two nodes running a database. Describe two ways in which the cluster software can manage access to the data on the disk. Discuss the benefits and disadvantages of each.

2. **[System Structures]** (5%) Protecting the operating system is crucial to ensuring that the computer system operates correctly. Provision of this protection is the reason for dual-mode operation, memory protection, and the timer. To allow maximum flexibility, however, you should also place minimal constraints on the user.
   The following is a list of instructions that are normally protected. What is the *minimal* set of instructions that must be protected?
   a) Change to user mode.
   b) Change to monitor mode.
   c) Read from monitor memory.
   d) Write into monitor memory.
   e) Fetch an instruction from monitor memory.
   f) Turn on timer interrupt.
   g) Turn off timer interrupt.

3. **[System Structures]** Please answer the followings:
   a) (3%) What system calls have to be executed by a command interpreter or shell in order to start a new process?
   b) (4%) What are the two modes of interprocess communication? What are the strengths and weaknesses of the two approaches?

4. **[Multithreaded Programming]** Consider a multiprocessor system and a multithreaded program written using the many-to-many threading model. Let the number of user-level threads in the program be more than the number of processors in the system. Discuss the performance implications of the following scenarios.
   a) (3%) The number of kernel threads allocated to the program is less than the number of processors.
   b) (3%) The number of kernel threads allocated to the program is greater than the number of processors but less than the number of user-level threads.

考試科目	作業系統 8142	所別	資訊科學系 8141	考試時間	03 月 01 日(星期日) 第二 節

5. **[Multithreaded Programming]** A system with two dual-core processors has four processors available for scheduling. A CPU-intensive application is running on this system. All input is performed at program start-up, when a single file must be opened. Similarly, all output is performed just before the program terminates, when the program results must be written to a single file. Between startup and termination, the program is entirely CPU-bound. Your task is to improve the performance of this application by multithreading it. The application runs on a system that uses the one-to-one threading model (each user thread maps to a kernel thread).
   a) (3%) How many threads will you create to perform the input and output? Explain.
   b) (3%) How many threads will you create for the CPU-intensive portion of the application? Explain.

6. **[Process Scheduling]** Consider a preemptive priority scheduling algorithm based on dynamically changing priorities. Larger priority numbers imply higher priority. When a process is waiting for the CPU (in the ready queue, but not running), its priority changes at a rate α ; when it is running, its priority changes at a rate β. All processes are given a priority of 0 when they enter the ready queue. The parameters α and β can be set to give many different scheduling algorithms.
   a) (3%) What is the algorithm that results from $\beta > \alpha > 0$?
   b) (3%) What is the algorithm that results from $\alpha < \beta < 0$?

7. **[Synchronization]** (5%) Race conditions are possible in many computer systems. Consider a banking system that maintains an account balance with two functions: `deposit(amount)` and `withdraw(amount)`. These two functions are passed the `amount` that is to be deposited or withdrawn from the bank account balance. Assume that a husband and wife share a bank account. Concurrently, the husband calls the `withdraw()` function and the wife calls `deposit()`. Describe how a race condition is possible and what might be done to prevent the race condition from occurring.

8. **[Synchronization]** Please answer the followings:
   a) (3%) Explain why implementing synchronization primitives by disabling interrupts is not appropriate in a single-processor system if the synchronization primitives are to be used in user-level programs.
   b) (3%) Show that, if the `wait()` and `signal()` semaphore operations are not executed atomically, then mutual exclusion may be violated.
   c) (4%) The implementation of mutex locks might suffer from busy waiting. Describe how to design mutex locks so that a processs waiting to acquire a mutex lock would be blocked and placed into a waiting queue until the lock became available.

考試科目	作業系統 8141 2	所別	資訊科學系 8141	考試時間	03 月 01 日(星期日) 第二 節

9. **[Deadlocks]** Please answer the followings:
   a) (3%) Consider a system consisting of four resources of the same type that are shared by three processes, each of which needs at most two resources. Show that the system is deadlock free.
   b) (3%) Assume a multithreaded application uses one reader-writer locks for synchronization. Applying the four necessary conditions for deadlock, is deadlock still possible if multiple reader-writer locks are used?
   c) (4%) Consider the version of dining-philosophers problem in which the chopsticks are placed at the center of the table and any two of them can be used by a philosopher. Assume that requests for chopsticks are made one at a time. Describe a simple rule for determining whether a particular request can be satisfied without causing deadlock given the current allocation of chopsticks to philosophers.

10. **[Memory-Management Strategies]** Please answer the followings:
   a) (3%) On a system with paging, a process cannot access memory that it does not own, why?
   b) (4%) How could the operating system allow access to other memory? Why should it or should it not?
   c) (3%) Explain why sharing a reentrant module is easier when segmentation is used than when pure paging is used

11. **[Virtual-Memory Management]** (5%) A certain computer provides its users with a virtual memory space of 2^{32} bytes. The computer has 2^{22} bytes of physical memory. The virtual memory is implemented by paging, and the page size is 4,906 bytes. A user process generates the virtual address 11123456. Explain how the system establishes the corresponding physical location. Distinguish between software and hardware operations.

12. **[Virtual-Memory Management]** Consider a demand-paging computer system where the degree of multiprogramming is currently fixed at four. The system was recently measured to determine utilization of CPU and the paging disk. The results are one of the following alternatives. For each case, what is happening? Can you increase the degree of multiprogramming to increase the CPU utilization? Is the paging helping in improving performance?

   a) (3%) CPU utilization, 13%; disk utilization, 97%
   b) (3%) CPU utilization, 87%; disk utilization, 3%
   c) (3%) CPU utilization, 13%; disk utilization, 3%

考試科目	作業系統 8141 2	所別	資訊科學系 8141	考試時間	03 月 01 日(星期日) 第二 節

13. **[Implementing File Systems]** Consider a file currently consisting of 10 blocks. Assume that the file control block (and the index block, in the case of indexed allocation) is already in memory. Calculate how many disk I/O operations are required for contiguous, linked, and indexed (single-level) allocation strategies, if, for one block, the following conditions hold. In the contiguous-allocation case, assume that there is no room to grow at the beginning but there is room to grow at the end. Also assume that the block information to be added is stored in memory. Please explicitly demonstrate for how your answers are derived.
    a) (3%) The block is added at the end.
    b) (3%) The block is removed from the end.

14. **[Mass-Storage Structure]** Requests are not usually uniformly distributed. For example, a cylinder containing the file system FAT or inodes can be expected to be accessed more frequently than a cylinder that contains only files. Suppose that you know that 50% of the requests are for a small, fixed number of cylinders.
    a) (3%) Would any of the disk-scheduling algorithms be particularly good for this case? Explain your answer.
    b) (3%) Propose a disk-scheduling algorithm that gives even better performance by taking advantage of this "hot spot" on the disk.
    c) (3%) File systems typical find data blocks via an indirect table, such as FAT in DOS or inodes in UNIX. Describe one of more ways to take advantage of this indirection to improve the disk performance.

備註	一、作答於試題上者，不予計分。 二、試題請隨卷繳交。