

國立臺北科技大學九十六學年度碩士班招生考試

系所組別：1810 資訊工程系碩士班甲組

第一節 作業系統 試題

第一頁 共四頁

注意事項：

1. 本試題共四題，配分共 100 分。
2. 請標明大題、子題編號作答，不必抄題。
3. 全部答案均須在答案卷之答案欄內作答，否則不予計分。

一、Multiple Choices (3% each, 42% total)

1. A major task of operating systems is to synchronize concurrent threads for accessing a shared resource. One problem in controlling the threads is called critical-section. Which of the following concepts can be used to solve the critical-section problem?
 - I. Mutual exclusion
 - II. Race condition
 - III. Semaphore
 - IV. Atomic Instruction

(A) III only
(B) I, II, III, and IV
(C) I, III, and IV
(D) I and IV
(E) I, II, and IV
2. Which of the following descriptions regarding the multithread models are correct?
 - I. In One-to-One model, only one thread can access the kernel at a time on multiprocessors.
 - II. In Many-to-One model, the entire process will block if a thread makes a blocking system call.
 - III. In Many-to-Many model, the number of kernel threads can be specific to either a particular application or a particular machine.
 - IV. In One-to-One model, developers can create as many user threads as necessary.
 - V. The two-level model is a variation of the Many-to-Many model.

(A) II and III
(B) II, III, and IV
(C) I, II, IV, and V
(D) II, III, and V
(E) I, III, and V

3. Which of the following approaches can be used to deal with the deadlock problem?
 - I. We can use a protocol to prevent or avoid deadlocks, ensuring that the system will never enter a deadlock state.
 - II. We can allow the system to enter a deadlock state, detect it, and recover.
 - III. We can ignore the problem altogether and pretend that deadlocks never occur in the system.
 - IV. We can combine the approaches in I and II and create another approach to deal with deadlock.

(A) I and II
(B) I, II, and III
(C) I, III, and IV
(D) I and III
(E) I, II, III, and IV
4. Which of the following descriptions about CPU scheduling algorithms is correct?
 - I. The FCFS scheduling algorithm is nonpreemptive.
 - II. The SJF scheduling algorithm depends on the length of the next CPU burst of a process, rather than its total length.
 - III. The SJF is a variation of the priority scheduling algorithm.
 - IV. The round-robin scheduling is similar to FCFS, but preemption is added to switch between processes.
 - V. A major problem with preemptive priority scheduling algorithm is indefinite blocking or starvation.

(A) I, III, and IV
(B) I, II, III, IV, and V
(C) II, III, and IV
(D) I, III, and V
(E) I, II, III, and IV
5. Which of the following statements regarding the use of page in memory allocation are correct?
 - I. The use of page can solve the problem of external fragmentation.
 - II. The use of page requires hardware support to be practically feasible.
 - III. A small page size can improve the internal fragmentation and reduce the overhead of page-table entry.
 - IV. The two-level paging algorithm is one way to structure pages in order support a large logical-address space.

(A) I and II
(B) I, II, and III
(C) I, II, and IV
(D) II and III
(E) I, II, III, and IV
6. Which of the following statements regarding the directory structures and file sharing are incorrect?
 - I. The two-level directory structure has a separate directory for each user.
 - II. The tree-structured directory prohibits the sharing of files through different directory entries.

注意：背面尚有試題

- III. The acyclic-graph directory structure allows directories to share subdirectories and files.
- IV. In UNIX, a shared file implemented using symbolic links cannot be deleted since it will result in dangling links.
- V. The UNIX keeps a reference count for a shared file implemented using nonsymbolic links (or hard links).

- (A) I only
- (B) II and III
- (C) IV and V
- (D) IV only
- (E) V only

7. Which of the following statements regarding the allocation of disk space are incorrect?

- I. The contiguous allocation algorithm suffers from the problems of external fragmentation and size declaration.
- II. Both the linked allocation and the indexed allocation effectively support direct file access.
- III. The pointer overhead of the index block is generally greater than the pointer overhead of linked allocation.
- IV. The inode in UNIX is a variation of indexed allocation while the FAT in MS-DOS is a variation of linked allocation.

- (A) II only
- (B) II and III
- (C) III only
- (D) II, III, and IV
- (E) III and IV

8. Which of the following statements regarding the disk scheduling are incorrect?

- I. The C-SCAN algorithm is a variant of SCAN designed to provide a more uniform wait time.
- II. The shortest-seek-time-first (SSTF) may cause a starvation problem. However, it can improve the rotational latency.
- III. The performance of a disk scheduling algorithm can be greatly influenced by the file-allocation method.
- IV. In general, the SCAN and C-SCAN perform better for systems that place a heavy load on the disk.

- (A) I only
- (B) II only
- (C) III only
- (D) II and III
- (E) II, III, and IV

9. Which of the following statements regarding the serializability are correct?

- I. A serial schedule guarantees that each transaction is executed atomically.
- II. A nonserial schedule does not necessarily imply an incorrect execution.
- III. One protocol that ensures serializability is the two-phase locking protocol.
- IV. The timestamp-ordering protocol ensures conflict serializability since conflicting operations are processed in time-stamp order.

- (A) I, III, and IV
- (B) II and IV
- (C) I and III
- (D) II and III
- (E) I, II, III, and IV

10. Which of the following statements regarding the virtual memory are correct?

- I. The effects of thrashing can be limited if local replacement algorithm is used.
- II. To restart a process as soon as possible, the desired page of the process can be read into a free frame from the free-frame pool before the victim page is written out.
- III. The LRU algorithm does not exhibit Belady's anomaly.
- IV. The design of virtual memory requires considering the size of page. To have a better virtual memory design, the page size should be large if page table size and I/O time are considered.

- (A) I, II, III, and IV
- (B) I, II, and III
- (C) I, III, and IV
- (D) I and IV
- (E) II and III

11. Which of the following statements regarding the segmentation are correct?

- I. Segmentation is a memory management scheme that supports the user view of memory.
- II. Segmentation can cause an internal fragmentation.
- III. A particular advantage of segmentation is the association of protection with the segment.
- IV. Segmentation scheme allows the sharing of code or data among different processes.

- (A) I, II, III, and IV
- (B) I, II, and III
- (C) I, III, and IV
- (D) I, II, and IV
- (E) II and III

12. In a dual-mode (system mode and user mode) system, which of the following components are necessary?

- I. A bit field in the hardware to indicate the current mode
- II. Privileged instructions
- III. Operating system running in the system mode
- IV. Some user programs in the system mode

- (A) II, III, and IV
- (B) I, II, and III
- (C) I, II, III, and IV
- (D) I, III, and IV
- (E) II and III

13. Which of the following instructions are usually implemented as privileged instructions?

- I. Software interrupt
- II. Test-and-set instructions
- III. I/O instructions
- IV. Floating point instructions
- V. MMU manipulation instructions

- (A) I, II, and V
- (B) I and III
- (C) I, II, III, IV, and V
- (D) III and V
- (E) II and IV

14. Which of the following descriptions are incorrect for reentrant function design?

- I. Static data can be declared in the function.
- II. Global variable can be used in the function. However, it must be protected in a critical section.
- III. The function can not return a pointer to static data.
- IV. The function works only on the data provided to it by the caller.
- V. Only reentrant function can be called by the function.

- (A) I and II
- (B) III and IV
- (C) II and V
- (D) III and V
- (E) I and III

二、Fill-in-the-blank (43% total)

1. One soft real-time system has five periodic tasks, with periods of 50, 100, 150, 200, and 300 msec each. Assume these tasks require 10, 25, 15, 40, and C msec of CPU time, respectively. What is the largest value of C for the system to be schedulable? _____ msec (3%)
2. Assume a process is allocated with 4 frames that are initially empty. Give a page reference sequence 7, 0, 1, 2, 0, 3, 0, 4, 2, 3, 0, 3, 2, 1, 2, 0. (3% each, 6% total)
 - (1) What is the number of page faults for the LRU page replace algorithm? _____
 - (2) What is the number of page faults for the Optimal page replace algorithm? _____
3. Consider a paging system with the page table stored in memory. (3% each, 6% total)
 - (1) If a memory access takes 200 nanoseconds, how long does a paged memory reference take? _____ ns
 - (2) If we add associative registers (a TLB), and 75 percent of all page table references are found in associative registers, what is the effective memory reference time? Assume that finding a page-table entry in the TLB takes zero time. _____ ns

4. Which of the requirements for a solution to the critical-section problem are not satisfied for the following program? _____ (2%)

```

/* lock is a shared data and initialized to 0 */
while (1) {
    key = 1;
    while (key)
        atomic_swap(lock, key);
    ... critical section...
    lock = 0;
    ...remainder section...
}
    
```

5. An incomplete program for the Bounded-Buffer problem is shown below. Two threads *producer()* and *consumer()* were implemented in this program. Assume that there are N buffers. And three semaphores *empty*, *full*, and *mutex* are used and initialized to the values N , 0 , and 1 , respectively. Please complete this program by placing the item number ($a \sim f$) of the following semaphore operations into the uncompleted slots of the program. (2% each, 16% total)

- a) *wait(empty)*
- b) *signal(empty)*
- c) *wait(full)*
- d) *signal(full)*
- e) *wait(mutex)*
- f) *signal(mutex)*

<pre> producer() { ...produce an item in p... (1) _____; (2) _____; ...add p to buffer... (3) _____; (4) _____; } </pre>	<pre> consumer() { (5) _____; (6) _____; ...remove an item from buffer to c... (7) _____; (8) _____; ...consume c... } </pre>
---	--

6. For the components listed below, answer the following questions by selecting proper item numbers, such as a , b , c , etc. Note that the answers can include single or multiple items. (2% each, 10% total)

- a) Software interrupt
- b) Hardware interrupt
- c) MMU
- d) TLB
- e) Cache
- f) Register
- g) Timer
- h) Real time clock
- i) DMA
- j) Atomic instructions

- (1) Which are the essential components for the preemptive multi-tasking to work? _____
- (2) Which components are used to implement system calls for Linux? _____
- (3) Which components are generally used to solve race condition problems on both single and multi-processor systems? _____
- (4) Which components can be used to increase the CPU utilization in accessing peripheral devices for a multi-programming system? _____
- (5) What are the essential components to implement virtual memory? _____

注意：背面尚有試題

CS2-7

三、What is the cause of thrashing? How does the system detect thrashing? Describe the working-set model and the page-fault frequency approaches to prevent thrashing. (10%)

四、The following example defines the address for an interrupt status register and then checks to see if an interrupt has been triggered. This kind of code segment is often seen in the implementations of operating systems for hardware data references.

```
...  
/* define the address of the IRQ status register */  
#define IRQ_STATUS (*(volatile unsigned *)0xff002c00)  
  
...  
/* check if IRQ#3 is triggered? */  
if (IRQ_STATUS & 0x08) {  
    ...  
}
```

Please explain why *volatile* (of C language) is used in defining the address? (5%)