

※ 注意：請用 2B 鉛筆作答於答案卡，並先詳閱答案卡上之「畫記說明」。

(一) 單選題 (每題 3 分; 答錯 0 分; 答錯不倒扣)


1. Let's compile the following C code on a 64-bit machine ---

```
#include <stdio.h>
int main() {
    int *p = (int*)malloc(sizeof(int));
    int *q = (int*)malloc(sizeof(int));
    printf("%x\n", &p);
    printf("%x\n", &q);
}
```

Suppose the output of the line “printf(“%x\n”, &p)” is “bffff4f8”. What will be the output of the following line “printf(“%x\n”, &q)”?

- (A) bffff4f0
- (B) bffff4f4
- (C) bffff4f7
- (D) bffff4fc
- (E) bffff500

2. The IEEE 754 defines the following floating number format:



What is the floating number for the binary representation “1100 0001 0101 0100 0000 0000 0000 0000”

- (A) -0.38095
- (B) -0.7619
- (C) -6.625
- (D) -13.25
- (E) -26.5

(Note: the exponent bias for the single-precision floating number is 127).

3. Compiling a same C program on two different computers, A and B, as follows, which of the following is correct?

見背面

	A	B
Instruction count	10 billion	unknown
Clock rate	4 GHz	5 GHz
CPI	1.0	1.2

- (A) The instruction count on B must be greater than 9.6 billion.
- (B) The program will spend less CPU time on B.
- (C) Computer B must have higher MIPs rating.
- (D) If the program runs faster (i.e. takes shorter CPU time) on A, the instruction count on B must be greater than 10 billion.
- (E) Computer A must consume less power than B.

(The following tables are for problems 4-7)

Assume the operation times for the major functional units in a CPU are as follows:

	Memory units	ALU and adders	Register read/write	others
time	200 ps	100 ps	50 ps	0 ps

Assume the instructions are divided into 5 classes and their inclusions of the functional units during executions are:

Instruction class	Functional units used by the instruction class during an instruction execution				
I	Instruction memory	Register access	ALU	Register access	
II	Instruction memory	Register access	ALU	Memory access	Register access
III	Instruction memory	Register access	ALU	Memory access	
IV	Instruction memory	Register access	ALU		
V	Instruction memory				

4. If the above CPU is implemented in a single-cycle per instruction mechanism, what is the maximum clock rate for this CPU?
- (A) 1.67 MHz
 - (B) 1.67 GHz
 - (C) 200 MHz
 - (D) 200 GHz
 - (E) 5.0 GHz

5. If the above CPU is implemented in a multi-cycle per instruction mechanism, as shown in the second table, that is, instructions I, II, III, IV, and V take 4, 5, 4, 3, and 1 cycles, respectively, what is the maximum clock rate for this CPU?
- (A) 1.67 MHz
(B) 1.67 GHz
(C) 200 MHz
(D) 200 GHz
(E) 5.0 GHz
6. Suppose a program has the following instruction mix: I (25%), II (10%), III (45%), IV (15%), and V (5%), and suppose all the instructions are executed sequentially (i.e. non-pipeline), what is the ratio of the runtimes by the CPUs in Problems 4 over 5, assuming both are running on the maximum clock rates?
- (A) 0.33
(B) 0.79
(C) 1.27
(D) 1.32
(E) 3.0
7. Assume the average CPIs for the above instructions under a 5-stage pipeline design are: 1.5, 1, 1, 1.25, and 2, respectively, what is the speed-up of this CPU over the single-cycle implementation in Problem 4, assuming both are running on the maximum clock rates?
- (A) 1.21
(B) 2.0
(C) 2.47
(D) 3.37
(E) 5.0

(The following information about the cache of a CPU is for Problems 8-9)

- * The hit rates for instruction and data caches are 95% and 90%.
- * The CPI for this processor is 2 when cache hits.
- * The penalty for all cache misses is 50 cycles.
- * The memory load and store instructions account for the 40% of the total instructions.

8. What is the CPI when considering the cache misses?
- (A) 2.5
(B) 4.5
(C) 6.5
(D) 8.5

見背面

- (E) 10.5
9. If the CPU clock rate is quadrupled (i.e. 4X), while the memory speed remains unchanged, what will be relative performance ratio to the CPU in Problem 8 when considering cache misses? (Please note that since the CPU architecture and instruction set is unchanged, the CPI for cache hit is unchanged).
- (A) 1.3
(B) 1.75
(C) 2.6
(D) 4.0
(E) 6.5
10. Which of the following is NOT correct for pipeline hazards?
- (A) Generally speaking, a pipeline hazard happens when the next instruction cannot execute in the following clock cycle.
(B) An example of structural hazards can be the condition when two instructions are accessing the same memory in the same cycle.
(C) Data hazards arise from the dependence of one instruction on an earlier one that is still in the pipeline.
(D) Control hazards arise from the need to make a decision based on the results of one instruction (e.g. branch instruction) while others are executing.
(E) "Forwarding" is a common technique to relieve the control hazards.
11. To make better use of the computer's resources, an operating system can give
- (A) higher (B) lower (C) the same
priority to I/O-bound programs than or between the CPU-bound programs.
12. Spinlocks are most appropriate for
- (A) single-processor systems
(B) multi-processor systems
(C) systems with direct memory access controllers.
13. What is the most required hardware support for a memory management system using demand-paging?
- (A) data cache
(B) instruction cache
(C) translation look-ahead buffers
(D) Harvard architecture

14. Which property is the advantage of a stateful file service as compared with a stateless file service?
(A) increased performance
(B) robust service
(C) being suitable for world wide web service
15. We can use semaphores to deal with the n -process critical-section problem. The n processes share a semaphore, called mutex, whose value is initialized to
(A) n (B) 1 (C) 0

(The following information about the CPU scheduling for problems 16-20)

Consider a single CPU system and the following set of processes, with the length of the CPU-job time given in milliseconds:

Process	Job Time (ms)	Priority
P1	5	3
P2	1	1
P3	10	2
P4	3	3

Assume that the CPU scheduler uses First Come First Serve, Shortest-Job First, Nonpreemptive Priority (a smaller priority number implies a higher priority), and Round-Robin (quantum = 1) scheduling algorithms. We assume that the scheduling overhead is ignored. The processes are assumed to have arrived in the order P1, P2 at time 0 and P3, P4 at time 4.

16. What is the turnaround time of process P3 for scheduling algorithm Nonpreemptive Priority?
(A) 10 (B) 12 (C) 14 (D) 16 (E) 17
17. What is the turnaround time of process P3 for scheduling algorithm Round-Robin?
(A) 10 (B) 12 (C) 14 (D) 15 (E) 19
18. What is the waiting time of process P3 for scheduling algorithms Nonpreemptive Priority?
(A) 1 (B) 2 (C) 4 (D) 5 (E) 6
19. What is the waiting time of process P3 for scheduling algorithms First Come First Serve?
(A) 1 (B) 2 (C) 4 (D) 5 (E) 6

見背面

20. Which of the schedules results in the minimal average waiting time (over all processes)?
(A) First Come First Serve (B) Shortest-Job First (C) Nonpreemptive Priority (D) Round-Robin

(二) 複選題 (每題答對 4 分, 答錯 0 分; 答錯不倒扣)

21. Which of the following are correct?
(A) Programmers can compose C/C++ programs in a regular text editor.
(B) To build a program with multiple C/C++ source codes, we use compiler to compile the C/C++ codes into Byte codes.
(C) The purpose of “linker” is to link the object codes or library files into an executable.
(D) The “loader” loads the program to the memory before the execution.
(E) The executable programs compiled from C/C++ source codes are usually machine-independent.
22. Which of the following are correct?
(A) Given a same task running on two different computers, the computer that spends less CPU time on this task is said to have better CPU performance on this task.
(B) Given two different computers, the one with higher CPU clock rate must have better CPU performance on any given task.
(C) Given two CPUs that implement the same instruction set, the CPU that has smaller CPI must have better CPU performance on any given task.
(D) Given two CPUs, the one with larger internal cache must have better CPU performance on any given task.
(E) Given two different tasks running on the same computer at the same time, the task that runs longer (i.e. larger elapsed time) must account for larger CPU time.
23. Which of the following statements are correct for multi-issue processor designs?
(A) Several internal components of the processor are replicated so that multiple instructions can be launched in every pipeline stage.
(B) Static multi-issue processors all use the compiler to assist with packaging instructions and handling hazards.
(C) The Intel IA-64 architecture is an example of static multi-issue processors.
(D) In a simple dynamic multi-issue processors, instructions issue in order and the processor decides where zero, one, or more instructions can issue in a given clock cycle.
(E) The dynamic multi-issue processors are also known as VLIW (Very Long Instruction Word) processors.
24. Which of the following statements are correct for virtual memory?
(A) “Page fault” refers to the situation when a data contention occurs.
(B) The variable addresses in a C program are virtual memory addresses.
(C) The page table maps each page in virtual memory to either a page in the main memory or a page

- stored on disk.
- (D) The page table is usually implemented by SRAM.
- (E) The translation-lookaside buffer (TLB) acts as a special cache that keeps track of the recently used physical page address translations so that the page access can be made faster by not calling the address translation of the page table.
25. Which of the following are correct?
- (A) The rotation speed for a typical hard drive in modern personal computers is around 500 to 1440 RPM.
- (B) A typical serial ATA interface can transfer data at the speed of 150MB/sec.
- (C) "RAID 3" means 2 duplications of the original disk and is a way to prevent data loss out of disk crash.
- (D) The capacity for a double-layered DVD diskette is about 8~9GB.
- (E) The purpose of DMA is to transferring data between an I/O device and memory without bothering CPU.
26. In operating systems, (A) segmentation (B) paging (C) threading (D) thrashing are common ways to achieve memory protection or memory management.
27. Assume that we have a personal computer running an operating system with demand-paging. In some moments, the CPU utilization is quite low, say under 10%, while the paging disk is quite busy, say the disk utilization over 90%. If we want to increase the CPU utilization of the personal computer, we can
- (A) install a faster CPU.
- (B) install a bigger paging disk.
- (C) increase the degree of multiprogramming.
- (D) decrease the degree of multiprogramming.
- (E) install more main memory.
28. In a system with shared resources, which are necessary conditions for a deadlock to occur?
- (A) Resources are required to be mutual exclusive.
- (B) Processes are holding some resources and waiting for some resources.
- (C) The use of resources can be preempted.
- (D) Circular wait conditions occurred between processes.
- (E) The uses of resources cannot be preempted.
29. The FAT file system format
- (A) uses a kind of linked allocation techniques.
- (B) uses a kind of indexed allocation techniques.

- (C) is a kind of log-structured file systems.
- (D) is efficient for sequential access.
- (E) is efficient for direct random access.

30. Which statements are correct.

- (A) The slab memory allocation algorithm uses a separate cache for each different object type.
- (B) The buddy memory allocation technique is a memory allocation technique that aims to solve the external fragmentation problem.
- (C) The slab memory allocation algorithm can solve the internal fragmentation problem.
- (D) Segmentation fault occurs when a program attempts to access a memory location that it is not allowed to access.

