國立交通大學 107 學年度碩士班考試入學試題

科目：計算機系統(1103)　　考試日期:107 年 2 月 2 日 第 3 節

系所班別：資訊聯招　　第 1 頁,共 8 頁

【不可使用計算機】＊作答前請先核對試題、答案卷(試卷)與准考證之所組別與考科是否相符！！

一、複選題（80%），共二十題，每題全對得 4 分。答對一個選項得 1 分（答對兩個選項得 2 分，以此類推），答錯一個選項倒扣 2 分（答錯兩個選項倒扣 4 分，最多扣至該題 0 分為止），整題未作答不給分。（舉例：每題有 a、b、c、d 四個選項，某題答案為 a、b，而作答時選 a、b、c，則三個選項正確一個錯誤，該題得 3-2=1 分）請使用答案卡作答

1. Consider high-level, assembly, machine languages, compiler, and assembler:
   (a) There are many different high-level languages; programs written in any of them can theoretically be executed directly on the hardware of all computers manufactured by different venders.
   (b) There are many different assembly languages; programs written in any of them can only be executed on specific hardware or a series/family of hardware of a computer manufactured by a specific manufacturer.
   (c) Assembler reads in a program written in high-level language, and produces a corresponding program in assembly language.
   (d) A machine language directly relates to a specific hardware design, hence even for different CPUs such as Pentium, Pentium II up to Pentium IV in the same series, different models/versions may very possibly need to use different machine languages.

2. This problem concerns the five classic components of a computer in computer organization.
   (a) Processor, also called CPU (central processing unit), is one of the five classic components.
   (b) The component that can really "compute" to generate a new result (such as adding x and y to produce z) is the datapath.
   (c) Operating system and compiler are two of the five classic components.
   (d) A touch screen can serve as two on the five classic components.

3. Do you know how we present the performance of a computer system? Try to answer the following questions:
   (a) Two popular parameters used to represent performance are *response time* and *throughput*; both of them use the same measuring unit.
   (b) The parameter an individual user cares more about is the throughput.
   (c) The parameter a data center manager cares more about is the throughput.
   (d) Using more CPUs in a computer system may not help in reducing the response time.

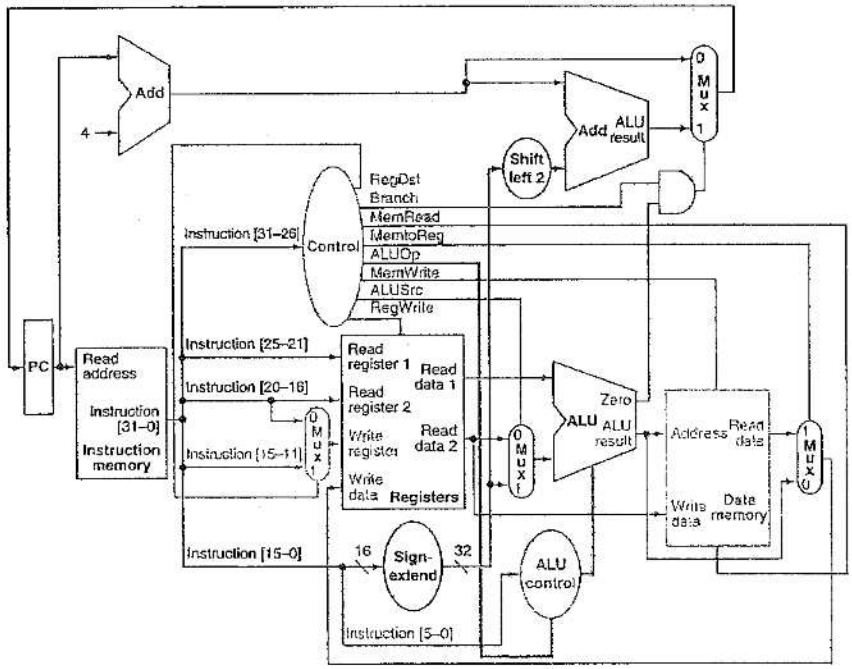
4. We study the memory hierarchy. Assume that a hierarchical memory contains these hierarchies: magnetic disk, DRAM main memory, and cache memories (1 to 3 levels).
   (a) Using many hierarchies enlarges the effective memory capacity, while this will slow down the average memory access time.
   (b) Cache memory uses the so-called cache block or cache line, which is commonly many bytes in size, to take advantage of program's spatial locality.
   (c) Since the cache miss rate is assumed to be the lower the better, if the electronic technology allows, any cache should be made as large as possible.
   (d) Cache memories are usually hardware controlled, and the OS may not even need to know their existence.

5. We continue to look into the memory system design. Now let the hierarchical memory system contains a magnetic disk drive, a DRAM main memory, a cache memory, and a register set in the CPU.
   (a) All above mentioned memories are random access memories.
   (b) If the computer system uses virtual memory system, usually it is the operating system that manages the use of physical memory using the virtual memory manager.
   (c) How the registers should be used is totally determined by the C program that is now executing.
   (d) For access performance, a cache memory usually decreases its set size when its capacity becomes larger. While the number of pages in a virtual memory space is huge, but the adopted set size is

typically the largest possible. The purposes of these two seemingly contradicting schemes are both for the same access performance purpose.

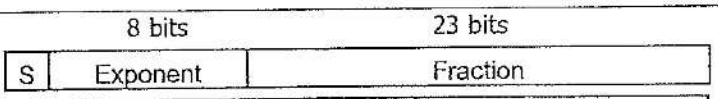
6. Given the single-cycle data path in the figure below and consider MIPS instructions. Assume the latencies for logic blocks in the figure are listed as table below. Please ignore the latency for control decoder, control signal delay, or other unspecified blocks. Which of following statements about the cycle time are correct?
   (a) The time necessary for executing a *beq* instruction is 230ps.
   (b) The time necessary for executing a *R-type* instruction is 210ps.
   (c) The time necessary for executing a *lw* instruction is 330ps.
   (d) The minimal clock cycle time to support all above three instructions, including *beq*, *lw*, and *R-type* is 360ps.

I-mem	Add	Regs	ALU	Sign-extend	D-mem
100ps	50ps	30ps	50ps	10ps	120ps



7. Given a MIPS CPU for running a program. Assume the PC (program counter) currently contains address 0x00000040. Which of following statements are correct?
   (a) It is possible to use a single ***beq*** to get to address `0x00080040`.
   (b) It is possible to use a single ***beq*** to get to address `0xFFFFFF00`.
   (c) It is possible to use a single ***j*** to get to address `0x08000040`.
   (d) It is possible to use a single ***j*** to get to address `0xFFFFFF00`.

8. Given an IEEE754 single-precision representation format in the table (on the next page) and the bias value as 127. Which of the following statements are correct? (Let "*" be the multiplication operator)
   (a) The number -1.1_2*2^{-1} has the single precision representation as:
   Sign = -1, Exponent = $0111\ 1110_2$, Fraction = $100\ 0000\ 0000\ 0000\ 0000\ 0000_2$
   (b) The number 101.101_2*2^{-130} has the single precision representation as:
   Sign = 0, Exponent = 00000000_2, Fraction = $010\ 1101\ 0000\ 0000\ 0000\ 0000_2$
   (c) The smallest positive single-precision normalized number is: 1.0_2*2^{-126}.
   (d) The smallest positive single-precision de-normalized number is 1.0_2*2^{-127}.



Values (single precision)		Meaning
Exp.	Fraction	
0	0	0
0	Non-zero	+/- de-normalized number
1-254	Anything	+/- floating-point number
255	0	+/- infinity
255	Non-zero	NaN (Not a number)

9. Consider a pipeline CPU with five stages: IF, ID, EX, MEM, WB. Stall cycles due to mis-predicted branches increase the CPI of pipeline. Given a code sequence with instruction mix as: R-type: 50%, Beq: 25%, lw & sw: 25%. Assume branch prediction accuracy is 20% and there are no data hazards and cache misses. It is clear that the system suffers from extra CPIs due to mis-predicted branches. Which of the following statements about the extra CPI are correct?
   (a) The extra CPI is 0.4 if the mis-predicted branch outcomes are corrected in MEM stage.
   (b) The extra CPI is 0.6 if the mis-predicted branch outcomes are corrected in MEM stage.
   (c) The extra CPI is 0.4 if the mis-predicted branch outcomes are corrected in ID stage.
   (d) The extra CPI is 0.2 if the mis-predicted branch outcomes are corrected in ID stage.

10. Given a static 2-issue pipelined processor, where a single 2-issue packet can contain one ALU/Branch and one load/store instructions. Assume the original code sequence below will be executed for 9 iterations. By unrolling 3 copies of the code with proper register renaming, the scheduled code is listed as following. To obtain an accurate running result with the unrolled code sequence, which of following statements are correct.
   (a) The (A),(B),(C),(D),(E) are 4($s1), 8($s1), 0($s1), 4($s1), 8($s1), respectively.
   (b) The (A),(B),(C),(D),(E) are -4($s1), -8($s1), 0($s1), -4($s1), -8($s1), respectively.
   (c) The (A),(B),(C),(D),(E) are 8($s1), 4($s1), 12($s1), 8($s1), 4($s1), respectively.
   (d) Compared to original code sequence, the unrolled code has a speedup equal to 6/5.

Original code sequence

```
L: lw   $t0, 0($s1)
   addu $t0, $t0, $s2
   sw   $t0, 0($s1)
   addi $s1, $s1, -4
   bne  $s1, $zero, L
```

Unrolled code sequence

	ALU/Branch	Load/store
L:	addi $s1, $s1, -12	lw $t0, 0($s1)
	nop	lw $t1, **(A)**
	addu $t0, $t0, $s2	lw $t2, **(B)**
	addu $t1, $t1, $s2	sw $t0, **(C)**
	addu $t2, $t2, $s2	sw $t1, **(D)**
	bne $s1, $zero, loop	sw $t2, **(E)**

11. A multithread program runs on a multiprocessor system. Consider the following statements regarding the thread model of the program. Which one(s) are generally true?
   (a) If the program execution time decreases as the number of CPUs increases, then the thread model is *not* many-to-one.
   (b) If a synchronous I/O blocks all threads of the program, then the thread model is one-to-one.
   (c) If increasing the number of CPUs reduces the program execution time, then the program is *not* I/O-bound.
   (d) Each thread of the program receives a larger share of CPU time with the many-to-one thread model

than it does with the one-to-one thread model.

12. Consider the three process states: *ready*, *running*, and *waiting*. Which one(s) of the following statements are generally true?
    (a) The CPU scheduler selects a process among waiting processes to run.
    (b) If a state transition from running to ready is possible, then the CPU scheduler is preemptive.
    (c) Context switch on a process is required when the state of the process changes from waiting to ready.
    (d) A process enters the waiting state after it makes a system call.

13. Operating system implementation separates policy from mechanism. An operating system provides multiple policies to solve a problem so that the system administrator can select the most appropriate one according to the system workload. Which one(s) of the following actions requires a policy?
    (a) Selecting the next process to run
    (b) Looking up the interrupt vector table to determine the address of an interrupt service routine
    (c) Deciding the service order of pending I/O requests
    (d) Finding a victim page for replacement

14. Which one(s) of the following phenomena are major problems of First-In First-Out CPU scheduling algorithm?
    (a) Long waiting time
    (b) Low I/O utilization
    (c) Process starvation
    (d) Low CPU utilization

15. Which one(s) of the following are possible consequences of deadlock avoidance?
    (a) Involuntary process termination
    (b) More memory leak
    (c) Reduced thread-level parallelism
    (d) Priority inversion (a low-priority thread keeps running when a high-priority thread is ready)

16. Which of the following statements are correct?
    (a) For a system with 64-bit logical address space, a two-level paging scheme is appropriate.
    (b) Although the inverted page table decreases the amount of time needed to search the table, it increases the amount of memory needed to store each page table when a page reference occurs.
    (c) Both the first-fit and best-fit strategies for memory allocation suffer from internal fragmentation.
    (d) The TLB is associative, high-speed memory in which each entry consists of a key and a value.

17. Select the following statements that are correct.
    (a) Port scanning is not an attack but rather a means for a cracker to detect a system's vulnerabilities to attack.
    (b) A Trojan Horse is a code segment that misuses its environment.
    (c) A digital certificate contains a public key digitally signed by the user.
    (d) In a message-authentication code (MAC), a cryptographic checksum is generated from the message using a one-way hash function where no secret key is involved.

18. Which of the following statements are correct?
    (a) DES, AES, RC4 are all symmetric encryption algorithms.
    (b) AES, RC4 are block ciphers.
    (c) In an asymmetric cipher, both keys for a user much kept secret.
    (d) In UNIX systems, a user password in the cleartext form is stored in a password file.

19. Given a C program:

```
#include <stdio.h>
#define BUFFER_SIZE 256
int main(int argc, char *argv[])
{
    char buffer[BUFFER_SIZE];
    if (argc < 2)
        return -1;
    else {
        strcpy(buffer, argv[1]);
        return 0
    }
}
```

(a) The program is lack of bounds checking.
(b) Buffer-overflow vulnerability exists
(c) The return address can be replaced
(d) Using strncpy() in the program gives the same result as using strcpy().

20. Which of the following statements are correct?
(a) For the page replacement algorithms, the page-fault rate always decreases as the number of allocated frames increase.
(b) RAID structure has been used to improve reliability. Rebuild performance varies with the RAID level used. Rebuilding is easiest for RAID 0.
(c) While a process is calling a system call, a software interrupt is invoked to switch from the user mode to the kernel mode.
(d) CPU utilization increases as degree of multiprogramming increases before thrashing occurs.

二、題組（20%），共四個題組，每一題組內所有小題全對得 5 分，答錯任一小題或未作答得 0 分。題組 A 包含題號 21~23，題組 B 包含題號 24~27，題組 C 包含題號 28~30，題組 D 包含題號 31~33。

A. For a 32-bit byte-addressable virtual memory space, the physical main memory is 2GB with page frame size of 4KB; the L2 cache is 512KB 4-way set-associative with block size of 64B; and the L1 cache is 32KB direct-mapped with block size of 32B.

21. We first look at the memory spaces:
(a) The number of pages is 4G and the number of page frames is 2G.
(b) The number of pages is 2G and the number of page frames is 4G.
(c) The number of pages is 1M and the number of page frames is 512K.
(d) The number of pages is undetermined and the number of page frames is 512K.

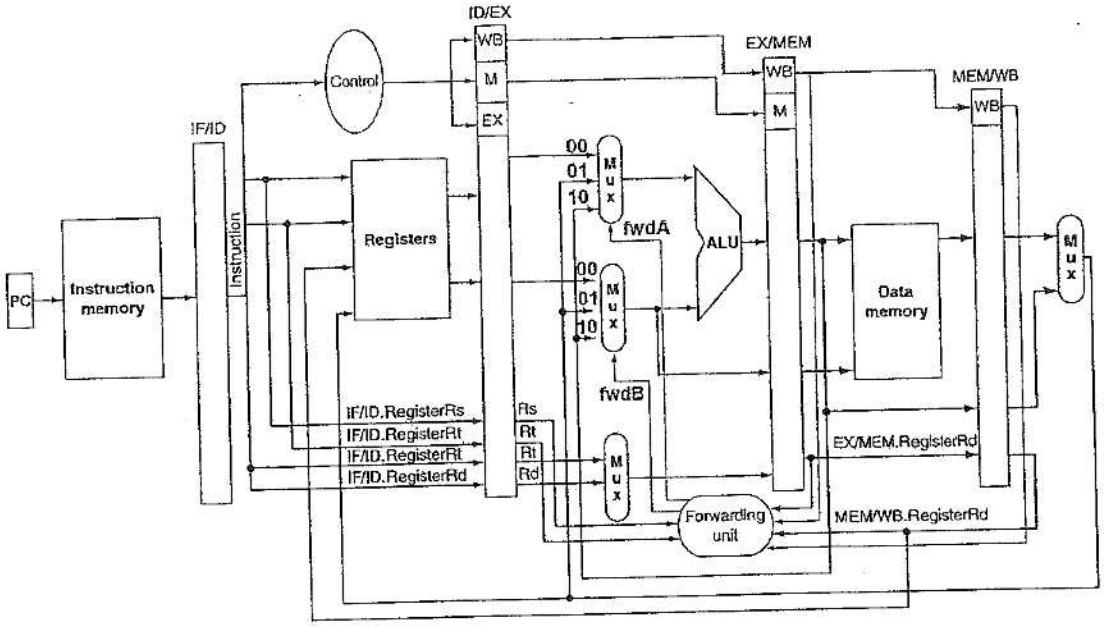
22. Then let the L2 cache be physically addressed, and describe how L2 cache interprets the memory address sent to it:
(a) The tag contains 14 bits, the index 11 bits, and the byte offset 6 bits.
(b) The tag contains 15 bits, the index 11 bits, and the byte offset 6 bits.
(c) The tag contains 14 bits, the index 12 bits, and the byte offset 5 bits.
(d) The tag contains 17 bits, the index 10 bits, and the byte offset 5 bits.

23. Let the individual access times of the L1 cache, L2 cache and main memory be 1ns, 5ns, and 300ns; and the local miss rates of them be 20%, 5%, and 0%, respectively. The average memory access time is hence:
(a) 306.0ns

(b) 65.0ns
(c) 5.0ns
(d) 2.6ns

B. Given a five-stage pipeline CPU in the figure below, where data forwarding mechanisms have been designed. For running the following code sequence on this pipeline CPU, please answer the questions below.

```
1. add $2, $10, $11
2. add $1, $12, $13
3. add $1, $1, $2
4. sub $3, $1, $2
```



24. When the 3rd instruction, **add**, is running at **EX** stage, what value should be set for the control signal **fwdA**?
(a) 00
(b) 01
(c) 10
(d) don't care

25. When the 3rd instruction, **add**, is running at **EX** stage, what value should be set for the control signal **fwdB**?
(a) 00
(b) 01
(c) 10
(d) don't care

26. When the 4th instruction, **sub**, is running at **EX** stage, what value should be set for the control signal **fwdA**?
(a) 00
(b) 01
(c) 10
(d) don't care

27. When the 4th instruction, **sub**, is running at **EX** stage, what value should be set for the control signal **fwdB**?
(a) 00
(b) 01
(c) 10
(d) don't care

C. **Ten** students want to study in a quiet room. However, there are only **four** seats in the room. A student must wait outside if there is not any available seat in the room. When one or more seats are available, a student enters the room, takes a free seat, and begins to study. When he finishes studying, he releases the seat and leaves the room. No two students can share the same seat. The following code fragment shows a solution to this problem. Answer the following questions:

```
seat[4]={F,F,F,F};// seat[0..3], F (free) T (occupied)
semaphore S1= __X__ ; // see questions below
semaphore S2=1;
student()
{
    wait(S1);       // try to enter the room
    wait(S2);       // try to find a free seat
    for(int i=0;i<4;i++)
       if(seat[i]==F)
           break; // found a free seat
    assert(i!=4); // a definite error
    seat[i]=T;      // take the seat
     __Y__ ;        // see questions below
    study();        // begin to study
    seat[i]=F;      // release the seat
     __Z__ ;        // see questions below
}
```

28. What is the value *X* (the initial value of semaphore S1)?
(a) 0
(b) 1
(c) 10
(d) 4

29. What is the statement *Y*?
(a) signal(S1);
(b) signal(S2);
(c) wait(S1);
(d) wait(S2);

30. What is the statement *Z*?
(a) signal(S1);
(b) signal(S2);
(c) wait(S1);
(d) wait(S2);

D. Given one-way hash function h(), suppose Allen's RSA public key, secret key, one-time AES key are Kpa, Ksa, and Kda, respectively. Bob's RSA public key, and secret key are Kpb and Ksb, respectively. Let $\{M\}_K$ represent the encryption of M using key K. Also let M || N represent the concatenation of two

data items M and N.

31. To achieve authentication only, which will be the best way for Allen to send a big file M to Bob?
(a) $\{M\}_{Ksa}$
(b) $M \| \{M\}_{Ksa}$
(c) $\{h(M)\}_{Ksa}$
(d) $M \| \{h(M)\}_{Ksa}$

32. To achieve confidentiality only, which will be the best way for Allen to send the file M to Bob?
(a) $\{M\}_{Kpb}$
(b) $\{M\}_{Kda}$
(c) $\{M\}_{Kda} \| \{K_{da}\}_{Kpb}$
(d) $\{M\}_{Kpb} \| \{K_{da}\}_{Kpb}$

33. To achieve both confidentiality and authentication, which will be the best way for Allen to send the file M to Bob?
(a) $\{\{M\}_{Ksa}\}_{Kpb}$
(b) $\{\{M\}_{Kpb}\}_{Ksa}$
(c) $\{M\}_{Kda} \| \{\{K_{da}\}_{Ksa}\}_{Kpb}$
(d) $\{M \| \{h(M)\}_{Ksa}\}_{Kda} \| \{K_{da}\}_{Kpb}$