國立陽明交通大學 112 學年度碩士班考試入學招生試題

科目：計算機系統(1103)　　考試日期：112 年 2 月 6 日 第 3 節

系所班別：資訊聯招　　第 1 頁，共 9 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 1 分(答錯兩個選項倒扣 2 分，最多扣至本科目計算機系統 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 3−1=2 分)，請用答案卡作答

1. Which of the following statements on page replacement is (are) true?
   (a). FIFO page replacement algorithm may suffer from Bélády's anomaly.
   (b). With LRU page replacement, adding page frames may increase the number of page faults.
   (c). The LRU page replacement algorithm achieves the minimum number of page faults.
   (d). Increasing the size of TLB (translation lookaside buffer) reduces the number of page replacements.

2. Assume that we create a 64 MB shared memory for processes X and Y on an Intel x86 machine. Which of the following statements is (are) true?
   (a). The shared memory must be mapped to the same virtual address range in both X and Y.
   (b). If we set the shared memory pages in process X as read-only, process Y will not be able to write to the shared memory.
   (c). The shared memory can be swapped to a backing store (e.g., the disk).
   (d). It is possible that less than 64 MB of physical memory space is actually allocated for the shared memory.

3. Which of the following statements is (are) true about direct memory access (DMA)?
   (a). It allows the application programs to access the main memory directly.
   (b). A DMA transfer will slow down the execution of a CPU-intensive program.
   (c). We can use device interrupts to synchronize DMA transfers.
   (d). We can use polling to synchronize DMA transfers.

4. To move file f from /x/f to /y/f, which of the following step(s) are required (cannot be skipped)?
   (a). Make an entry for f in directory y.
   (b). Allocate data blocks on the disk for /y/f.
   (c). Copy the file content of /x/f to the data blocks of /y/f.
   (d). Erase the entry for f in directory x.

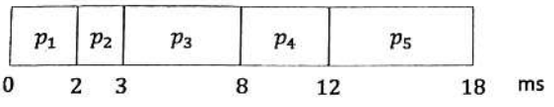
5. Assume that a file system has no fragmentation to begin with (the disk blocks of each file are contiguous). If we want to have some files with non-contiguous data blocks on the file system, which of the following operation(s) are essential?
   (a). Changing a file name.
   (b). Deleting a file.

(c). Creating a new file and writing some data to the file.
(d). Overwriting the data of an existing file.

6. Which statement(s) about process are correct?

(a). When a process creates a child process, it can either continue its execution or wait for the child process to complete before continuing.
(b) A process may run out its stack memory by spawning too many processes as it runs.
(c) In a multithreaded multicore system, it is possible for two threads of the same process to run simultaneously in two different cores.
(d) A deep recursive function call can consume a lot of memory from the heap section of the process.

7. Consider the following CPU scheduling result in a single-processor system. If all the five processes complete within 18 ms, which statement(s) are correct?



(a) This result is not possible with round-robin scheduling.
(b) This result is possible if non-preemptive SJF is used and processes do not arrive at the same time.
(c) This result is possible if preemptive SJF is used and processes do not arrive at the same time.
(d) If both the preemptive SJF and the non-preemptive SJF yield this result, the waiting time of p_4 must be zero.

8. Consider the following C code for creating two threads. Which statement(s) are correct?

(a) The output is always 20000.
(b) It is possible that the output is 19999.
(c) The output is fixed to 20000 if we change the order of the two calls to `pthread_join`.
(d) The first thread has a higher priority than the second one in thread scheduling.

```
// header files here
int count = 0;

void *mythread(void *arg) {
  for (int i=0; i<10000; i++) {
    // wait for some time
    count += 1;
  }
  return(NULL);
}

int main() {
  pthread_t p1, p2;
  pthread_create(&p1, NULL, mythread, NULL);
  pthread_create(&p2, NULL, mythread, NULL);
  pthread_join(p2, NULL); pthread_join(p1, NULL);
  printf("%d", count);
```

```
    return 0;
}
```

9. Which statement(s) about condition variable are correct?

(a) A call to `pthread_cond_wait()` causes the calling thread to busy wait on a condition variable.
(b) `pthread_cond_wait()` includes a mutex lock as a parameter.
(c) `pthread_cond_signal()` function signals one thread that is waiting on the condition variable.
(d) The thread that calls `pthread_cond_signal()` aborts its execution right after the return of the call, after which CPU scheduler gives the execution right to the awakened thread.

10. Consider a system with m resources and 4 threads: T_0, T_1, T_2, and T_3. Let x_i and y_i denote the maximum number of resources requested and the current number of resources held by thread T_i, respectively, for all i, $0 \le i \le 3$. Which of the following statement(s) are correct?

	Maximum Needs	Current Needs
T_0	x_0	y_0
T_1	x_1	y_1
T_2	x_2	y_2
T_3	x_3	y_3

(a) The system state is deadlocked if $\sum_{i=0..3}(x_i - y_i) > m - (y_0 + y_1 + y_2 + y_3)$.
(b) The system state is deadlocked if $(x_2 - y_2) > m - (y_0 + y_1 + y_2 + y_3)$.
(c) The system state is deadlocked if $(x_i - y_i) > m - (y_0 + y_1 + y_2 + y_3)$ for all i, $0 \le i \le 3$.
(d) The system state is safe if $(x_i - y_i) \le m - (y_0 + y_1 + y_2 + y_3)$ for some i, $0 \le i \le 3$.

11. Two CPUs are implemented to execute exactly the same ISA. Assume that CPU A is a single-cycle implementation while CPU B is a five-stage pipelined implementation. The instruction and data memory devices are zero wait state memory for both CPU. Which of the following statements are true when both CPUs runs the same program C with finite execution time?

(a). The execution time of CPU B is shorter than that of CPU A.
(b). The critical path of CPU B can always be made shorter than that of CPU A.
(c). CPU B will need more circuit resources than CPU A does.
(d). The CPI of CPU A is higher than the CPI of CPU B.

12. To remove hazards for a single-issue in-order execution pipelined processor, which of the following techniques may require help from the compilers (assuming that you only write programs using high-level languages)?

(a). Delayed branch.
(b). Dynamic branch prediction with branch history tables.
(c). Load-use hazard removal.
(d). Forwarding.

13. Modern ISAs often contain a pair of instructions for multi-thread and multi-core synchronization. In

MIPS, this pair of instructions includes a special load called a load linked, `ll`; and a special store called a store conditional, `sc`. Which of the following statements are true?

(a). The instruction `ll` reads and locks a memory cell such that no other threads can change its value until an `sc` instruction executed by the same thread changes its value.

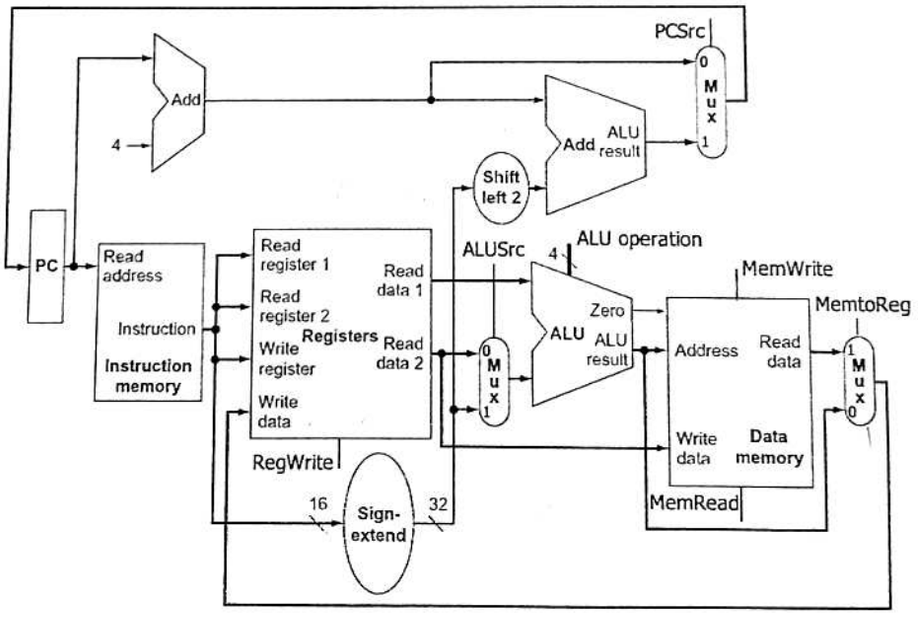
(b). The execution of the `sc` instruction will always be successful if no context switch occurs after the `ll` instruction.

(c). The following code segment has a bug in protecting the memory pointed to by `s1`:

```
try:
    ll    $t0, 0($s1)
    addi  $t1, $t0, 1
    bnez  $t1, skip
    sc    $t1, 0($a1)
    beqz  $t0, try
skip:
```

(d). The `ll` and `sc` instruction pairs perform shared resource protection more efficiently than the atomic swap instruction used in, for example, the ARM ISA. Especially when there are a lot of concurrent processes.

14. The following figure shows a simple implementation of a MIPS processor. Which of the following statements are true?



(a). The following code segment can be used to detect a stuck-at-0 fault of MemtoReg:

```
addi  $t1, zero, 1
sw    $t1, 0(zero)
lw    $t0, 0(zero)
slti  $t0, $t0, 1
```

(b). The CPU will not suffer the load-use hazards.

(c). The instruction memory does not have to be a synchronous memory (i.e. driven by a clock signal).

(d). The register file does not have to be a synchronous component.

15. Assume that we have a 5-stage processor that has neither data forwarding nor branch prediction. It

has separate instruction and data memory. Bubbles will be inserted when there are hazards. The processor always resolves the branch target at stage 2 (the decode stage). Which of the following statements are true when the following code segment is executed?

```
      addi   $t1, zero, zero
loop: addi   $t1, $t1, 1
      lw     $s0, 2000($t1)
      subi   $t2, $s0, 2
      sw     $t2, 4000($t1)
      bne    $t1, $t3, loop
```

(a). The bubble(s) inserted for the `subi` instruction can be removed by forwarding.
(b). If `$t3` is 100, the CPI will be larger than 2.5 for the execution of the code segment.
(c). Excluding the adders for the PC, the processor needs at least two more adders. Note that a subtractor is counted as an adder.
(d). Adding a static branch predictor increases the CPI for this code segment when `$t3` is 100.

16. Which of following statements about cache are true?

(a) A virtually-addressed cache access is typically faster than that of a physically-addressed cache.
(b) High associativity in a cache can reduce compulsory misses.
(c) Cache performance is of less importance in faster processors because the processor speed compensates for the high memory access time.
(d) Larger blocks reduce the miss rate by taking advantage of spatial locality

17. The following table shows the instruction type breakdown of a given application executed in a CPU. Assume the CPU has a 2GHz clock rate. Which of following statements are correct?

(a) The execution time of this program is 1024ms.
(b) We can make the program run two times faster by only improving the CPI of L/S instructions.
(c) We can make the program run two times faster by only improving the CPI of INT instructions.
(d) If we improve the CPU such that CPI of INT and FP instructions is reduced by 40% and the CPI of L/S and branch is reduced by 30%, the execution time of running this program will become 342ms.

	FP instr.	INT instr.	L/S instr.	Branch instr.
Number of instr.	80×10^6	240×10^6	160×10^6	32×10^6
CPI	1	1	4	2

18. Given a processor with two caches: an I-cache and a D-cache, and the miss penalty is 200 cycles for all misses. Assume the I-cache miss rate is 2%, D-cache miss rate is 4%, and the percentage of all *lw* and *sw* instructions is 50%. If the processor has a CPI of 2 without any memory stalls. Which of following statements are true when memory stall is considered? (Note, each statement is independent.)

(a) The CPI will become 10 if memory stall is considered.
(b) If we double the clock rate (but the memory access time to handle a cache miss does not change), the CPI will become 18.
(c) If we double the clock rate (but the memory access time to handle a cache miss does not change), the machine will have a speedup of 9/10.

(d) If a L2 cache is added only for I-cache and its miss penalty is 20 clock cycles. Assume global instruction miss rate to main memory became 0.5% after L2 cache is added. The CPI will become 3.4.

19. The system has a direct-mapped cache with total data size of 8192 bytes and block size of 16 bytes. Assume the cache starts out empty. For the following piece of code, which of following statements are correct? (Note: Assume that **int** is of 32-bits (4 bytes); variables *x* and *i* are both in registers and the array A is in memory.)

(a) There will be 256 cache misses in total.
(b) All the misses are compulsory misses.
(c) The cache hit-rate is 0.25.
(d) If we double the block size, the miss rate will be halved.

```
int A[256];
int x = 0;
for (int i = 0; i < 256; i++ ) {
    x += A[i];
}
```

20. The following table shows two processors and their read/write operations on two different variables *a* and *b* (initially *a* = *b* = 0). After the operations, which of following statements are correct?

(a) The value of *a* can be 1, 2 or 3 if the system ensures cache coherence.
(b) The value of *a* can be 1, 2, or 3 if the system does not ensure cache coherence.
(c) The value of *b* can be 1, 2, or 3 if the system ensures cache coherence.
(d) The value of *b* can be 1, 2, or 3 if the system does not ensure cache coherence.

Processor 1	Processor 2
a++;　*b*++;	*a* =2;　*b* = *b*+2;

二、題組(20%)，共四個題組，各題組下的小題為單選，每一題組內所有小題完全答對得 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 31~34，請用答案卡作答

題組 A

Assume that we have four disks (disk0~3), and each disk has 1TB of space. We use the notation diskX[i] to denote the i^{th} sector on diskX. Now, we form a RAID5 array with all four disks.

21. Which of the following number is closest to the maximum capacity (the maximum amount of data you

can store on the array) of the RAID5 array?

(a). 1TB

(b). 2TB

(c). 3TB

(d). 4TB

22. What is the maximum number of disk failures that the RAID5 array can tolerate?

(a). 1

(b). 2

(c). 3

(d). 4

23. In the worst case, how many disks will be involved in writing a single byte to the RAID5 array?

(a). 1

(b). 2

(c). 3

(d). 4

題組 B

Consider three processes P, Q, and R using message-passing for communications. Their behaviors are illustrated below. Let BS = blocking send, NS = nonblocking send, BR = blocking receive, and NR = nonblocking receive.

Process P	**Process** Q	**Process** R
...	...	...
send(Q, msg1)	send(R, msg3)	send(P, msg5)
receive(R, msg2)	receive(P, msg4)	receive(Q, msg6)

24. Which type of naming is used? (a) direct (b) indirect (c) insufficient evidence to judge (d) none above.
25. Which data structure is needed to support this type of communications? (a) mailbox (b) port (c) semaphore (d) none above.
26. Which condition will cause deadlock in this example? (a) BS (b) BR (c) both BS and BR (d) none above.
27. If deadlock is not possible, which condition will ensure each process to proceed only after receiving its message? (a) BS (b) BR (c) NS (d) NR.

題組 C

We want to translate a C function `my_call()` into the MIPS assembly language. The function my_call() calls another function `func()` which is declared as "`int func(int s, int t);`". The translated assembly code for the function `my_call()` is as follows:

```
my_call:
   addi $sp,$sp,-12
   sw   $ra,8($sp)
   sw   $s1,4($sp)
   sw   $s0,0($sp)
   move  $s1,$a2
   move  $s0,$a3
   jal   func
   move $a0,$v0
   add   $a1,$s0,$s1
   jal   func
   lw   $ra,8($sp)
   lw   $s1,4($sp)
   lw   $s0,0($sp)
   addi  $sp,$sp,12
   jr   $ra
```

28. How many parameters does the function `my_call()` have?

(a). 1

(b). 2

(c). 3

(d). 4

29. Which of the following C code could be the body of `my_call()`?

(a). `{return func(func(a,b), c+d);}`

(b). `{return func(a,b) + func(c,d);}`

(c). `{return func(a+b, func(a+b,c));}`

(d). `{return func(a, func(c+d, b));}`

30. The tail-call optimization is a technique that when a caller simply returns the value that it gets from the callee, we can avoid allocating a procedure frame for the last function call. How many instructions can we save if the tail-code optimization is used to optimize the assembly code?

(a). 1

(b). 2

(c). 3

(d). 4

題組 D

Consider a two-way set associative cache with four-word bocks and LRU replacement. Assume the cache has a total data size of 32 words and 24-bit byte address is used. Please answer each question below.

31. What's the size (in bits) of tag field in one block used in this cache?

(a) 23 (b) 22 (c) 19 (d) 18

32. What's the total size (in bits) of this cache, including data, tag, and valid fields?

(a) 1024 (b) 1100 (c) 1176 (d) 1184

33. Assume the cache is initially empty. Given a series of memory references: 24, 8, 25, 10, and 41 (which are word *addresses*), how many cache miss will be for the references?

(a) 1 (b) 2 (c) 3 (d) 4

34. For your answer of above question, how many cache misses are *compulsory miss*?

(a) 1 (b) 2 (c) 3 (d) 4