國立陽明交通大學 試 題 紙

一百一零學年度第二學期
博士班資格考

科目：計算機架構 A

日期： 111 年 7 月 5 日 第 1 頁 共 2 頁

請 "✓" 明 ✓不可看書 可看書

* 請將答案依題號順序寫入答案卷

答題時字跡需工整，否則不予計分。Write your answers legibly; otherwise you will get zero score.

1. (7%) In classifying computer architecture, according to our reference textbook, what is the most basic differentiation?
   (a) (1%) Name it clearly.
   (b) (3%) What options do we have regarding this differentiation?
   (c) (3%) Well explain why this feature is so basic and important.

2. (4%) Assume that operations supported by most ISAs (Instruction Set Architectures) can be classified into SIMPLE, MEDIUM, and COMPLEX, based on the amount of execution time needed by individual operations; of course the more time an operation requires, the more work can be done. And we know that the processor design trend is to make the instruction execution not only multi-cycle, but also pipelined. Now if you are to design a new general-purpose processor with good execution speed, what operation types should you include in your processor's ISA? (All SIMPLE, all COMPLEX, mixture of SIMPLE and MEDIUM, etc.) Well support your answer.

3. (8%) Pipelining a processor creates three major pipeline hazards. Name these, each with
   (a) (3%) What this hazard is about, and
   (b) (3%) A well-known technique that we can use to eliminate or at least alleviate 減輕 harm of that hazard, with proper explanation.
   Other than the pipeline hazards, there are also costs that we must pay whatsoever.
   (c) (2%) List them, each with proper explanations.

4. (8%) In dealing with exceptions,
   (a) (2%) What is synchronous exception, and what is asynchronous exception?
   (b) (3%) Give an example exception after which the program execution should resume, and another after which the program should terminate. Support you answers.
   (c) (3%) In pipeline processor design, if exception is restartable, what capabilities should the processor have?

5. (10%) Given C-like code "for (i=0; i<10000; i=i+1) x[i] = x[i] + y;" where x is an array of 64-bit floating-point numbers, R0 now contains address of x[0], R1 contains address right after x[9999], and y is a 64-bit floating-point scalar now contained in register F0.
   (a) (2%) Write a MIPS-like code for this loop.
   (b) (2%) If considering pipeline features given below, indicate the loop execution pattern cycle-by-cycle the sequence of the instructions and necessary stalls are issued.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

(c) (3%) Pipeline-schedule the loop to minimize the stall cycles, and then show only the cycle-by-cycle loop execution pattern, including necessary stalls.

(d) (3%) Unroll the loop the minimum number of times to eliminate any stall. Show your final MIPS-like code. (Do not need to worry about the prolog and epilog code.)

Instruction producing result	Instruction using result	Latency in clock cycles
FP ALU op	Another FP ALU op	3
FP ALU op	Store double	2
Load double	FP ALU op	1
Load double	Store double	0

6. (4%) We study the Tomasulo’s algorithm. Answer the following questions, each with sufficient explanation.

(a) (2%) What difficulties has it overcome? And how?

(b) (2%) Does it support precise interrupt? Why or why not?

7. (9%) Multiprocessors provide a higher-level parallelism, a common trick we should be familiar with.

(a) (3%) Define UMA, and list its other well-known name(s).

(b) (3%) Repeat (a) for NUMA.

(c) (3%) For cache memory, there are three types of misses called the 3 Cs. With multiprocessing, there comes yet another C. Explain what this last C is about.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國立陽明交通大學 試 題 紙

一百一零學年度第二學期
博士班資格考

科目：計算機架構 B　　日期：111 年 7 月 5 日　第 1 頁 共 2 頁

請"✓"明　✓不可看書　可看書

* 請將答案依題號順序寫入答案卷

答題時字跡需工整，否則不予計分。Write your answers legibly; otherwise you will get zero score.

1. (7%)
   (a) (3%) Describe the three main concerns from the system perspective of **power** and **energy**.
   (b) (4%) Define the following two main measures of **dependability**: ***module reliability*** and ***module availability***. Write the equations with respect to *mean time to failure* (MTTF) and *mean time to repair* (MTTR) for measuring these two quantities of a non-redundant system with repair.

2. (8%) A 5-GHz processor was used to execute a benchmark program with the instruction mix and clock cycle counts shown in the following table.
   (a) (3%) Assume that the total number of instructions executed is 8×10^9, determine the *effective CPI*, and *execution time* of this program.
   (b) (3%) Assume that an enhancement proposal is to reduce the CPI of the floating-point operations to 8 with 20% lengthening of the clock cycle time. Calculate the *effective CPI* and the *speedup* of the enhancement.
   (c) (2%) Calculate the *fraction of the time* for executing load and store operations of this program in the original processor. Assume that we build an optimizing compiler to discard two thirds (2/3) of the Load/Store operations from the original instruction mix, determine the *speedup* of the enhancement to the original design by applying ***Amdahl's Law***.

Instruction type	Frequency	CPI
Integer ALU ops	40%	1
Floating-point ops	15%	20
Loads/Stores	25%	4.8
Branches	20%	2

3. (8%)
   (a) (3%) Describe the following three **block placement** approaches for **cache** and the **block identification** method corresponding to each approach: ***direct mapped***, ***set associative***, and ***fully associative***.
   (b) (3%) Assume that the CPU address is N bits and byte addressable, the cache has 2^i blocks, the block size is 2^j bytes, and the cache has 2^k ways if it is set-associative. An exemplar address division diagram is given below. Determine the number of bits of the ***tag***, ***index***, and ***block offset*** fields for each of the three block placement approaches mentioned in (a).
   (c) (2%) Describe the following two **write options** for **cache**: ***write through*** and ***write back***.

Block address		Block offset
Tag	Index	

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

4. (8%)
   (a) (2%) Describe the main characteristic of program access, including instructions and data, that make ***memory hierarchy*** design feasible and efficient.
   (b) (6%) Explain the main reason that make the strategies adopted for the design of ***cache*** and ***virtual memory*** different, and describe the policies adopted for ***block placement***, ***block identification***, ***block replacement***, and ***write strategy*** of a **paged virtual memory**.

5. (8%) Typically, *miss rate*, *miss penalty*, *hit time*, and *bandwidth* are the main metrics to evaluate the performance of a cache memory. Describe each of the following cache optimizations and indicate its positive impact to the performance metrics:
   (a) (2%) ***multilevel cache***
   (b) (2%) ***nonblocking cache***
   (c) (2%) ***virtual cache*** (***virtually addressed cache***)
   (d) (2%) ***compiler-controlled prefetching***

6. (11%)
   (a) (2%) Describe two advantages of ***SIMD*** (single instruction, multiple data) architecture versus ***MIMD*** (multiple instructions, multiple data) architecture.
   (b) (3%) Describe the following three variations of SIMD: ***vector architectures***, ***multimedia SIMD instruction set extensions***, and ***graphics processing units*** (***GPUs***).
   (c) (6%) Describe each of the following optimizations of **vector architecture** and the problem it solves:
   i. ***vector-length register*** ii. ***predicate register*** iii. ***stride***

國立陽明交通大學 試 題 紙

一百一零學年度第二學期
博士班資格考

科目：演算法 A

日期：111 年 7 月 6 日 第 1 頁 共 1 頁

請 "✓" 明 ✓不可看書 可看書

* 請將答案依題號順序寫入答案卷

答題時字跡需工整，否則不予計分。Write your answers legibly; otherwise you will get zero score.

1. Given a function, $f(n) = a^{\log_c n}$ where a and c are non-negative constant and $a > c$. Now the question, is $f(n)$ polynomial or exponential? Show me the reason(s) or give me a proof. You can assume n is exact power of c. 10%

2. When we implement a binary search tree, we need to define a class or a struct for `node`. In a `node`, we have three fields, one for `data`, and two pointers `leftChild` and `rightChild`. Let call an unused pointer field a leaf. So if there is only one key in a tree T, then T has a `node` as the root, the `data` field stores the key, and two usused pointers are the two leaves. That is, T has an internal node and two leaves.

   (a) If we have a binary search tree storing n keys, show that there must be $n+1$ leaves (unused pointers), regardless the shape of the tree. 5%

   (b) We use $|T|$ to denote the number of leaves in tree T. Now the question: Given a tree T, T_L and T_R are the left and right subtrees of T, we have
   $$\frac{1}{3}|T| \leq |T_L| \leq \frac{1}{3}|T|$$
   and
   $$\frac{1}{3}|T| \leq |T_R| \leq \frac{1}{3}|T|.$$
   Furthermore, both T_L and T_R meet the above property. What is the minimum and maximum height of T. Please give me a "more accurate" answer, don't use the big O notation. 10%

3. We can implement a queue, `Q`, using two stacks `S1` and `S2`. Standard stack operations are `S.push(item)` (push `item` into `S`), `S.pop()` (return `item` on top of `S`, and remove `item` from `S`), and `S.empty()` (answer whether `S` is empty). We add one more operation `S.multiPop(`k`)`, pop out k items from S if there are more than k items in S, otherwise pop all the items.

   (a) Use pseudo code to describe the `inQueue` method (insert an item into `Q`). 2%

   (b) Use pseudo code to describe the `deQueue` method (delete the first one from `Q`). 3%

   (c) What is the amortized cost for `inQueue` and `deQueue`. 10%

4. Consider the problem, if we have two binary heaps (the heap we used in heap sort), we have to merge these two binary heaps into a binary heap. What is the best algorithm (the most efficient algorithm) to carry out the merging? Describe your algorithm and prove the time bound. 10%

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國立陽明交通大學 試 題 紙

一百一零學年度第二學期
博士班資格考

科目：演算法 B

日期： 111 年 7 月 6 日 第 1 頁 共 1 頁

請 "✓" 明 ✓不可看書 可看書

* 請將答案依題號順序寫入答案卷

答題時字跡需工整，否則不予計分。Write your answers legibly; otherwise you will get zero score.

1. **(5pt.)** Prove or disprove that if any NP-complete is solved in polynomial time, then P = NP.
2. **(10pt.)** A minimum spanning tree (MST) of an edge weighted and connected graph is a spanning tree whose weight (the sum of the weights of its edges) is no larger than the weight of any other spanning tree. Show an efficient greedy algorithm to find a minimum spanning tree for an edge weighted and connected graph. And discuss the time complexity of it.
3. **(10pt.)** Let G be a weighted directed connected graph. Show an efficient algorithm to find a directed cycle of negative weight in G, if such a directed cycle exists.
4. **(10pt.)** Show that the hamiltonian-path problem can be solved in polynomial time on directed acyclic graphs by giving an efficient algorithm. Hint: topological sort.
5. **(15pt.)** Consider a 2-CNF formula φ with n variables and m clauses. Show that 2-CNF SAT by giving an efficient polynomial algorithm. Note $(w \vee x) \wedge (\neg x \vee \neg y) \wedge (y \vee z) \wedge (z \vee x) \wedge (\neg w \vee z)$ is a 2-CNF formula with 4 variables and 5 clauses.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國立陽明交通大學 試 題 紙

一百一零學年度第二學期
博士班資格考

科目：作業系統 A　　　　日期： 111年7月6日　第 1 頁 共 1頁

請 "✓" 明　　✓不可看書　　可看書

* 請將答案依題號順序寫入答案卷

答題時字跡需工整，否則不予計分。Write your answers legibly; otherwise you will get zero score.

1. [10 pts] List a benefit and a drawback of using multilevel page tables.
2. [10 pts] LRU (Least-Recently Used) and LFU (Least-Frequently Used) are two representative page replacement algorithms. Which one is the better choice for the following workloads? Why?
   a) Scanning a very large item list and checking whether each scanned item is present in a small read-only database.
   b) Running 10 programs sequentially. Each program runs for 3 minutes.
3. [10 pts] What is a dangling pointer in file systems? Show a solution to the dangling pointer problem.
4. [10 pts] Free-space management for file systems is often based on an allocation bitmap or on a free-block list. Which one is more prone to free-space fragmentation? Why?
5. [10 pts] First-Come First Serve (FCFS), Shortest-Job First (SJF), and SCAN are typical disk scheduling algorithms. Which one is prone to the starvation problem the most? Why?

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國立陽明交通大學 試 題 紙

一百一零學年度第二學期
博士班資格考

科目：作業系統 B

日期： 111 年 7 月 6 日　第 1 頁 共 2 頁

請 "✓" 明　✓不可看書　可看書

* 請將答案依題號順序寫入答案卷

答題時字跡需工整，否則不予計分。Write your answers legibly; otherwise you will get zero score.

1. Consider using semaphores for the first readers-writers problem, which allows either one writer or at most *n* readers in the critical session at the same time (no reader be kept waiting unless a writer has already obtained the permission to use the shared object). We use an integer `read_count` to count the number of readers that are currently reading the data. In addition, we use two binary semaphores: `rw_mutex` and `mutex`. The former is common to both reader and writer processes. The `mutex` semaphore is used to ensure mutual exclusion when the variable `read_count` is updated. The reader's code is shown on the right-hand side. (a) What is the initial value of `mutex`? (3%) (b) What should be in ⑤? (3%) (c) What should be in ⑥? (4%)

```
sem_wait(&①);
read_count++;
if (read_count == ②)
  sem_wait(&③);
sem_post(&④);
(read the data)
sem_wait(&⑤);
read_count--;
if (read_count == ⑥)
  sem_post(&⑦);
sem_post(&⑧);
```

2. Consider the following snapshot of a system consisting of 5 processes (P_0~ P_4) and four types of resource (A, B, C, D):

	Allocation				Max				Available			
	A	B	C	D	A	B	C	D	A	B	C	D
P_0	1	2	0	1	1	5	4	2	7	1	6	1
P_1	2	2	1	0	7	2	2	4				
P_2	0	0	1	3	7	1	6	3				
P_3	4	1	3	0	4	6	3	5				
P_4	3	0	1	6	9	6	7	6				

Answer the following questions using the banker's algorithm:

(a) What are the contents of the matrix *Need*? (5%)

(b) Is the system in a safe state? If yes, show one safe sequence. If not, state the reason. (5%)

(c) If another request from process P_1 arrives for (0,0,1,2), can the request be granted immediately? Why or why not? (5%)

3. In case of preemptive CPU scheduling, CPU scheduling takes place under conditions (1), (2), (3), and (4) in the following state transition diagram. In case of non-preemptive CPU scheduling, only under which conditions does CPU scheduling take place? (5%)

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國立陽明交通大學 試 題 紙

一百一零學年度第二學期
博士班資格考

科目：作業系統B

日期： 111年7月6日 第2頁 共2頁

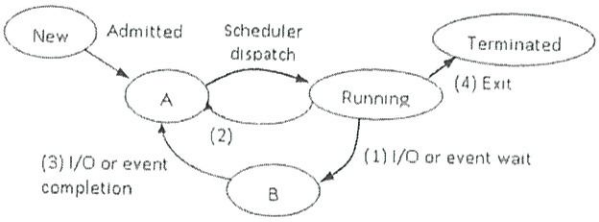

4. (a) What is the main disadvantage of spinlock in a system with a single-core CUP? (5%) (b) How do you evaluate spinlock in a multiprocessor system? (5%)

5. Consider the following C code. List all the possible outputs (an output looks like "ABCDE"). (10%)

```
int main() {
  printf("A");
  if (fork() == 0) {
    /* wait a random time */
    printf("B");
    if (fork() == 0) {
   /* wait a random time */
      printf("C");
    }
   else {
      wait(NULL);
      printf("D");
    }
  }
  else {
    /* wait a random time */
    printf("E");
    wait(NULL);
  }
  return 0;
}
```

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：