國立交通大學試題紙

一百零九學年度第二學期
博士班資格考

科目：作業系統 A　　　日期：110年7月14日　第1頁 共3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. (3%) Which of the following was NOT a design goal of the Internet?
   (a) Provide reliable communication.
   (b) Support the interconnection of different physical networks.
   (c) Use routers to store and forward packets.
   (d) Have decentralized control of the network.

2. (3%) A port number in IP protocols is used in the:
   (a) Data Link layer.
   (b) Network layer.
   (c) Transport layer.
   (d) Session layer.

3. (3%) The TCP layer does not deal with:
   (a) Reliable data transfer.
   (b) In-order data transfer.
   (c) Flow control.
   (d) Secure data delivery.

4. (3%) If event A has a Lamport timestamp of 2 and event B has a Lamport timestamp of 3, you are certain that:
   (a) A happened before B if the events took place on the same process.
   (b) A happened before B regardless of the processes on which A and B took place.
   (c) A and B are concurrent.
   (d) B happened before A, regardless of the processes on which A and B took Place

5. (3%) Which mutual exclusion algorithm requires, on average, the fewest messages in a group of 100 machines?
   (a) Centralized.
   (b) Token-ring.
   (c) Lamport's.
   (d) Ricart & Agrawala

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

6. (3%) The Ricart & Agrawala mutual exclusion algorithm:
   (a) Does not depend on time stamps in messages while Lamport's does.
   (b) Cannot handle the case where two or more processes request the same resource at the same time.
   (c) Does not require a process to send messages to the entire group while Lamport's does.
   (d) Requires fewer messages than Lamport's algorithm.

7. (3%) A practical way to handle network partitioning is to:
   (a) Have each process start an election.
   (b) Elect a leader only if a majority of nodes can be reached.
   (c) Run a normal election and multicast the result to the subset of processes that can be reached.
   (d) Divide and conquer: elect multiple coordinators.

8. (3%) Which file system supports a whole file cache model?
   (a) NFS
   (b) AFS
   (c) SMB
   (d) GFS (Google File System)

9. (3%) The wait-die algorithm:
   (a) Ensures that a circular wait can never develop.
   (b) Has a process give up after waiting for a resource beyond a specific time.
   (c) Requires a process to kill any process that is using resources it needs.
   (d) Uses a coordinator to detect and kill any process that is waiting for a resource

10. (3%) The following is not a property of transactions:
    (a) Permanent: once committed, the results of a transaction are made permanent.
    (b) Serializable: the result of concurrent transactions must be the same as if they were run in serial order.
    (c) Indivisible: the transaction appears as an indivisible action.
    (d) Bounded: a transaction must have an upper bound on the time it takes to complete

11. (3%) A bully election algorithm picks the process that:
    (a) First started the election.
    (b) Is nominated by the most processes.
    (c) Has the highest numbered Lamport timestamp of all election messages.
    (d) Has the highest numbered process ID.

12. (3%) Which disk scheduling algorithm is most vulnerable to starvation?
    (a) SCAN.
    (b) Shortest Seek Time First (SSTF).
    (c) LOOK.
    (d) First Come, First Served (FCFS)

13. (3%) Page fault frequency is a way of:
    (a) Measuring the effectiveness (hit ratio) of a TLB.
    (b) Identifying the specific pages that are best targets for eviction.
    (c) Measuring the efficiency of page fault processing within an operating system.
    (d) Ensuring that all processes have their working sets in memory.

14. (3%) Thrashing in a virtual memory system is caused by:
    (a) A process making too many requests for disk I/O.
    (b) Multiple processes requesting disk I/O simultaneously.
    (c) Processes not having their working set resident in memory.
    (d) Slow operating system response to processing a page fault.

15. (5%) A system has the following blocks queued for writing: 8000, 3000, 5000, 2000. The most recently written block was 4500. The block read before that was 2500.
    (a) What sequence of writes will a C-SCAN algorithm generate?
    (b) What sequence of writes will a SCAN algorithm generate?
    (c) What sequence of writes will a FCFS algorithm generate?
    (d) What sequence of writes will a SSTF algorithm generate?
    (e) What sequence of writes will a LOOK algorithm generate?

16. (3%) An inode-based file system uses 4 Kbyte blocks and 4-byte block numbers. What is the largest file size that the file system can handle if an inode has 12 direct blocks, 1 indirect block, and 1 double indirect block?

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) At runtime, Program X uses a single process with multiple threads.
   A. If we want to refactor the program to use a single process with **a single thread**, how can we interleave the code executions from different threads? (i.e., would it be sufficient to finish the execution of one thread and begin the execution of the next thread?)
   B. Suppose we want to refactor the program to a multi-process program, in which each process contains just one thread. What modifications will be needed with respect to the synchronization code in the original program code (i.e., the program that uses a single process with multiple threads)?

2. (10 pts) Assume Program X and Program Y run in different processes. At some time point, the two programs need to exchange data with each other.
   A. Can the data exchange be made as efficient as if the two programs were run by different threads in the same process? If so, explain how this can be done. If not, please describe the technical issues.
   B. If Y needs to wait for X, explain how this can be achieved with a semaphore.

3. (10 pts) Compare the following real-time scheduling algorithms by listing their advantages and disadvantages, respectively.

	Rate Monotonic Scheduling	Earliest Deadline First Scheduling
Advantages		
Disadvantages		

4. (10 pts) Has the Banker's algorithm for deadlock detection been widely used on general-purpose operating systems such as Windows or Linux? If yes, give an example showing how it is used on general-purpose operating systems. If not, explain why it is not widely in use.

5. (10 pts) For the Dining-Philosophers Problem, if a philosopher always picks the chopstick on the right-hand side first and then picks the chopstick on the left-hand side, would a deadlock situation still occur? If so, please give an example. If not, please provide the reason.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國立交通大學試題紙

一百零九學年度第二學期
博士班資格考

科目：計算理論 A　　　　日期：110 年 7 月 13 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. (15%) Design a deterministic Turing machine with no more than 5 states to accept the language $L = \{a^n b^m : n > m \geq 0\}$.
2. (20%) Show that the set $L = \{\langle M_1, M_2 \rangle : L(M_1) \cap L(M_2) \neq \emptyset\}$ is recursively enumerable, where M_i's are Turing machines and $\langle \cdot \rangle$ is the standard encoding function of Turing machines.
3. (15%) The Hamiltonian circle problem (HC) is, for an input of undirected graph $G = (V, E)$, to find a circle on G that passes each node in V exactly once, where $V = \{v_1, v_2, \ldots, v_n\}$ and $E \subseteq V \times V$. Give a nondeterministic polynomial-time algorithm to solve the HC problem and explain that the algorithm is indeed polynomial-time nondeterministically.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國 立 交 通 大 學 試 題 紙

一百零九學年度第二學期
博士班資格考

科目：計算理論 B

日期： 110 年 7 月 13 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%) Prove or disprove that if A and B are both regular languages, $A \cup B$ is a regular language.
2. (10%) Prove or disprove that if A and B are both nonregular languages, $A \cup B$ is a nonregular language.
3. (10%) Prove or disprove that $L = \{ \mathtt{a}^i\mathtt{b}^i\mathtt{c}^j \mid i \neq j \}$ is context-free.
4. (10%) Prove or disprove that $L = \{ \mathtt{a}^i\mathtt{b}^i\mathtt{c}^j \mid i, j \geq 0 \}$ is context-free.
5. (10%) Prove or disprove that if $L_1 = \{ \mathtt{a}^i\mathtt{b}^i\mathtt{c}^j \mid i, j \geq 0 \}$ and $L_2 = \{ \mathtt{a}^i\mathtt{b}^j\mathtt{c}^j \mid i, j \geq 0 \}$, $L_1 \cap L_2$ is context-free.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國 立 交 通 大 學 試 題 紙

一百零九學年度第二學期
博士班資格考

科目：計算機架構 A　　　　日期： 110 年 7 月 13 日　第 1 頁 共 2 頁

* 請 "✓" 明 Please tick ✓　　✓不可看書 Closed-book　　可看書 Open-book

* 請將答案依題號順序寫入答案卷 Write your answers in the answer sheets in problem and sub-problem order.

* 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%) About the term "***Computer Architecture***,"
   (a) Why do the authors of the textbook think that referring "***Computer Architecture***" only to "***instruction set design***" was incorrect?
   (b) A suggested myopic (short-sighted 近視的) view of "Computer Architecture" is "***Instruction Set Architecture*** (***ISA***)." What does ***ISA*** mean?
   (c) There are typically seven dimensions of an ***ISA***. List these seven, with necessary, very concise (精簡) explanations to each. (Four of the seven are enough to earn you all points.)
   (d) What is the authors' view of genuine "***Computer Architecture***?"
   (e) Computer architecture design is strongly influenced by supporting implementation technologies. What are the five crucial and rapid changing technologies according to the textbook (at its time)? (Only names, no explanations needed. Listing three correctly is good enough.)

2. (10%) In talking about using computer(s), major concerns include ***Power***, ***Energy***, ***Dependability***, and ***Performance***.
   (a) How does ***Power*** affect computer design (i.e., what design considerations must be faced)?
   (b) Why is ***Energy*** an important factor? Will lower energy always comes with a lower power design?
   (c) While ***Dependability*** is hard to quantify, we can do this, alternatively: For a module, show its ***MTTF***, ***FIT***, ***MTTR***, ***MTBF***, and ***Availability*** values. Define these five terms (verbally or mathematically).
   (d) In examining performance of a computer system, why is using "instructions per second" a pitfall? What metric shall we use indeed?

3. (10%) Given a memory system containing the following hierarchies: a ***virtual L1*** cache memory, a ***physical L2*** cache memory, a ***main memory***, and a ***hard disk***. For a program under execution, let the virtual memory space size be 2^{32} ***Bytes***, the useable physical memory space size be 2^{24} ***Bytes***, and the system has a ***translation-lookaside buffer***. The sizes of the hierarchies are: ***L1 is 32 KBytes***, ***L2 is 256 KBytes***, ***main memory is 128 Mbytes***, and ***hard disk is 1 TBytes***. For the following questions, assume that the computer is the simplest, uni-tasked, and ignore any system-management related information such as PID (process identification number), which is beyond the scope of this program itself.
   (a) When the program reads an instruction,
      (1) How many bits does the instruction memory address coming out of the CPU contain?
      (2) Is there any address handling involved before this address is sent to L1, and how many address bits are sent? (Clearly and precisely state any due address handlings if any. Do the same for the following (3) to (5).)
      (3) If L1 misses, is there any address handling involved before this address is sent to L2, and how many address bits are sent?
      (4) If L2 misses too, is there any address handling involved before this address is sent to main memory, and how many address bits are sent?
      (5) If even main memory misses (which we generally call page fault), is there any address handling involved before this address is sent to hard disk, and if yes, EXACTLY WHEN does this take place? Again, how many address bits are sent?
   (b) Repeat (a) (1)~(5) for another case when the program loads or stores a piece of data.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

4. (10%) The Appendix B lists six ***basic cache optimizations***. The purpose of each of them may be ***reducing the miss rate***, ***reducing the miss penalty***, ***reducing the hit time***, or even more than one of the above three. While some optimizations may require use of extra hardware pieces (by "extra" I mean very different hardware, not just changing sizes of the already-existing hardware), and may bring negative influences to some of the purposes. Complete the following table by filling each blank box with an integer -2 ~ +2 (-2 meaning very negative and +2 meaning very positive), and for the "***extra hardware pieces required***" column, it should not get a negative score, and you need to list the extra hardware pieces required <u>*very clearly and precisely*</u> in the accompanying blank(s). (An empty box will be regarded as a "0".)

Optimization	Extra hardware pieces required	Extra hardware pieces list	Reducing miss rate	Reducing miss penalty	Reducing hit time
Larger block size	(-2~+2 here)	(Names, e.g., address buffer)	(-2~+2 here)	(-2~+2 here)	(-2~+2 here)
Larger cache size					
Higher associativity					
Multilevel caches					
Read priority over writes					
Avoiding address translation during cache indexing					

5. (5%) We study the three types of cache misses mentioned in Section 2.1 of the textbook (but remember that there is yet another type, which has to do with multi-processing):
   (a) List the three, each with sufficient explanation of what it is.
   (b) On the log of cache misses in executing a program, state how you can identify each of the misses as being of what cache-miss type.

6. (5%) Parallelism at different levels is now the driving force of computer design.
   (a) Two basic kinds of parallelism exist in applications. What are they?
   (b) Two types of parallel computers are designed: the SIMD and MIMD, to take advantage of those two in (a), respectively. Point out which of the two computer types is to take advantage of which of the two parallelism kinds. Also, briefly explain how each of the two computer types can take advantage of that respective parallelism kind.
   (c) Disregard that the two parallel computer types are designed for different purposes, overall, what are SIMD's advantages over MIMD?

國立交通大學試題紙

一百零九學年度第二學期
博士班資格考

科目：計算機架構 B　　　　日期： 110 年 2

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. (4%) Describe the characteristics of the instruction set architecture (ISA) for a reduced instruction set computer (RISC)

2. (8%) Write the code sequence to compute C = A + B followed by D = A – B for an ***accumulator architecture*** and a ***register-register*** (***load-store***) ***architecture.*** Determine the total code size and the number of bytes of instructions and data moved to or from memory, i.e., the instruction bytes fetched and the memory-data bytes transferred, for your code. Assume that the opcode is always 1 byte, all memory addresses are 2 bytes, all data operands are 4 bytes, all register specifiers are 4 bits long, **all instructions are an integral number of bytes in length**, and the variables A, B, C, and D are initially in memory. (You should minimize your code size as possible as you can.)

For *accumulator* architecture:

Code sequence	Instr bytes fetched	Data bytes transferred
…		
Total:		

For *load-store* architecture:

Code sequence	Instr bytes fetched	Data bytes transferred
Total:		

3. (6%) Describe the following schemes of dynamic branch prediction: ***1-bit prediction***, ***2-bit prediction***, ***correlating branch prediction***, and ***tournament branch prediction***.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國 立 交 通 大 學 試 題 紙

一百零九學年度第二學期
博士班資格考

科目：計算機架構 B　　　　　　　　日期： 110 年

4. (6%) Given a computer implemented in **single-cycle** implementation, assume that it has a clock cycle time of 18 ns.
   (a) (3%) Design the computer as a **5-stage pipelined** implementation. After the stages are split by functionality, the measured times were IF, 4 ns; ID, 3 ns; EX, 3.5 ns; MEM, 5 ns; and WB, 2.5 ns. The pipeline register delay is 1 ns. What is the clock cycle time of the 5-stage pipelined machine? If there was no hazard in this pipeline, what is the speedup of the pipelined machine over the single cycle machine?
   (b) (3%) Assume that there will be a stall due to a data hazard every 5 instructions, and branches constitute 25% of the instructions, the misprediction rate is 10%, and the branch misprediction penalty is 2 cycles for the **5-stage pipelined** machine. Calculate the CPI of the pipelined machine and the speedup of the pipelined machine over the single cycle machine.

5. (6%) Describe the goal of ***loop unrolling*** and the decisions and transformations made for obtaining the efficient unrolled code of a loop.

6. (8%) Given the structure of **Tomasulo algorithm with reorder buffer** (**ROB**),
   (a) (4%) Describe the key idea of this Tomasulo algorithm and the main purposes of using ***reservation stations*** and ***ROB*** in this structure.
   (b) (4%) Describe the following four steps involved in instruction execution of this algorithm: ***Issue***, ***Execute***, ***Write result***, ***Commit***.

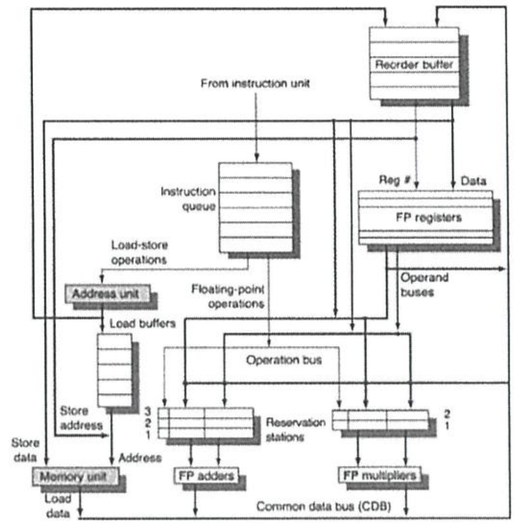

7. (6%) Describe the characteristics, one main advantage, and one main disadvantage of the following two classes of **MIMD multiprocessors**, and draw the block diagram of the basic structure for each of them: ***centralized shared-memory multiprocessor*** and ***distributed shared-memory multiprocessor***.

8. (6%) For a **centralized shared-memory multiprocessor**, which one of the following protocols, ***snooping*** protocol and ***directory-based*** protocol, should be applied for cache coherence? Describe the characteristics of the suitable protocol and explain the reasons.

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國立交通大學試題紙

一百零九學年度第二學期
博士班資格考

科目：演算法 A　　　　日期：110 年 7 月 14 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. (a) Use the least significant digit first radix sort to sort the following 3-digit numbers, 187, 26, 321, 231, 9, 1, 783,436, 91, 266. 4%

   (b) Show that least significant digit first sorting algorithm correctly sort n k-digit numbers in linear time if k is a constant. 6%

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`. If we have a binary search tree of n nodes, show that there must be $n+1$ unused pointer fields, regardless the tree looks like. 7%

3. Linear time algorithm for selection problem, given n numbers, find the k'th largest in linear time: Typical implementation starts with partitioning the n numbers into $\lceil n/5 \rceil$ groups and each group contains 5 numbers or less (the last group). My question is, how about we partition the numbers into 3 in a group or 7 in a group? Please explain. 8%

4. Binary Tree: Root is level 1. Show that there are at most $2^k - 1$ nodes in a level k Binary tree. 6%

5. We can implement a queue, `Q`, using two stacks `S1` and `S2`.

   (a) Use pseudo code to describe the `inQueue` method (insert an item into `Q`). 3%

   (b) Use pseudo code to describe the `deQueue` method (delete the first one from `Q`). 5%

   (c) What is the worst case cost required for `deQueue`. 5%

   (d) What is the amortized cost for `deQueue`. 6%

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

國立交通大學試題紙

一百零九學年度第二學期
博士班資格考

科目：演算法 B　　　　日期： 110 年 7 月 14 日　第 1 頁 共 2 頁

請"✓"明　✓不可看書　可看書

1. 14%

(1) Use Dynamic programming technique to solve the set-partition problem. See the problem definition below.
(Define the object function, write the recursive relation, give the initial condition, illustrate the table, indicate where the answer is, and estimate the time complexity and space complexity.)

(2) Is the algorithm in part(1) a polynomial time algorithm? Explain.

Set-partition problem:
Given a set S of n positive integers, $S = \{s_1, s_2, s_3, \ldots, s_n\}$,
is there a partition of S into two disjoint subsets A and $S - A$ such that

$$\sum_{s_i \in A} s_i = \sum_{s_i \in S-A} s_i \quad ?$$

2. 10%

(1) Give a Huffman binary coding tree for the following data: ABCATDABEBAAAAAABEB.
. Compute the total number of bits to send this message.

(2) Give a Huffman ternary coding tree for the data in part (1).

3. 14%

(1)2% Describe the Bellman-Ford algorithm for solving the shortest path problem.

(2)1% State the time complexity. Is it a polynomial time algorithm?

(3)4% Consider the graphs containing negative weight cycles. Can Bellman-Ford algorithm find a shortest path between two nodes? Explain or prove.

(4)3% Is there a polynomial time algorithm to find a negative weight cycle? Explain or prove.

(5)4% Is there a polynomial time algorithm to find the largest negative weight cycle? Explain or prove.
(A negative weight cycle such that the absolute value of its total weight is maximum..)

4. 12%

Assume only the basic knowledge of NP- complete problems; namely, Satisfiability problem, Vertex Cover problem, Hamiltonian Cycle problem, and the Subset-Sum problem are NP-complete.
And assume $P \neq NP$.

Consider the maximum programs stored on one disk problem:
Suppose that there are n programs, let l_i be the amount of storage needed to store the ith program,

◎請用深黑色鋼筆或原子筆出題

命題老師簽名：

i=1 to n. Let L be the storage capacity of the disk. Suppose that there is only one disk with storage capacity L.

For each of the following two problems, decide whether it is a P-problem or an NP-hard problem. If it is P, describe a polynomial time algorithm to solve it. If it is NP-hard, prove it.

(1) Determine the maximum total length of these n programs that can be stored on the disk.

(2) Determine the maximum number of these n programs that can be stored on the disk.

Note: The rule here to store a program on a tape is the following: A program has to be stored consecutively as a whole. One program cannot be split into several small parts and stored in different places, nor can be stored only part of it.