國立陽明交通大學 試 題 紙

一一四學年度第二學期
博士班資格考

科目：作業系統 A　　日期： 115 年 7 月 1 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. (5 pts) How many times does the following code print "hello"?

```
main() {
   if (fork() == 0) {
      if (fork() <> 0) printf("hello\n");
   }; printf("hello\n");
}
```

2. (6 pts) List two advantages of kernel-level threads over user-level threads.

3. (6 pts) Explain Belady's Anomaly in the context of page replacement algorithms. Which replacement algorithm has this anomaly?

4. (6 pts) What is the difference between a preemptive scheduler and a non-preemptive scheduler? Provide an example scenario where preemptive scheduling is beneficial.

5. (4 pts) Describe the concept of "demand paging" as a virtual memory implementation. Why is it commonly used?

6. (5 pts) In a system with a TLB hit ratio of 85%, if TLB access time is 15ns and memory access time is 100ns, calculate the effective memory access time (EAT).

7. (12 pts) List three attributes that are shared among threads in the same process, and three that are not shared.

8. (6 pts) In a UNIX inode with 12 direct blocks and 1 single indirect block, assuming a block size of 8KB and address size of 4 bytes, calculate the maximum file size.

國立陽明交通大學 試 題 紙

一一四學年度第二學期
博士班資格考

科目：作業系統 B

日期： 115 年 7 月 1 日

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. [5pts] Modern x86 64-bits systems adopt 4-level page tables (including the PGD, PUD, PMD and PT). Please elaborate how a CPU walks the 4-level page table. (Hint: draw a picture to explain how CPU gets each page table's physical address)

2. [5pts] Please describe the interrupt procedure between a hardware device and the CPU. (Shall include how CPU gets the ISR's starting address. Hint: 3 key steps)

3. [5pts] When an user process runs a system call, please explain how the user process notifies the kernel and how the kernel handles the system call.

4. [5pts] Modern x86 systems usually adopt multi-level page tables. What is the minimum levels are there for the system where the page size, virtual address space size and page table entry size are set to 32KB, 2^{58}B, and 64B, respectively.
(Please show the calculation process. No points without the calculation process.)

5. [30pts] Please answer **True** or **False** for each statement below:
   a. Once a user process being interrupted, the CPU switches to kernel mode to run corresponding OS code to handle the interrupt.
   b. In Linux, the fork() system call creates a new process with a separate PID by duplicating the parent's address space (using copy-on-write), while sharing most kernel data structures
   c. When an OS adopts Copy-on-Write (CoW) to fork a process, memory pages will be shared between the parent and child processes until one of them writes to the page.
   d. Interrupt handlers are split into top and bottom halves. Top halves indicate high-priority interrupts and bottom halves indicate low-priority interrupts.
   e. If a required data is not in DRAM, the CPU will raise a page fault, and the user process will run the page fault handler to transfer the data from storage to memory, possibly using DMA for efficiency.
   f. Virtual memory allows each process to use identical virtual address ranges independently, ensuring isolation even if physical memory is shared.
   g. Physical memory refers to the actual memory location of data, while virtual memory represents the data address in the storage devices.
   h. On the operating system adopting paging (e.g., Linux), processes can share page table with others to save memory consumption.

i. Physical memory indicates the location where the data actually stored on the memory. Virtual memory indicates the data address in the CPU cache.
j. When an x86 CPU cannot find an address translation in the TLB, the CPU's MMU automatically checks the page table to find the physical address, instead of requiring Linux to do the translation.
k. Renaming a file in the Ext-based file system is realized by directly replacing the old filename by a new one, and the inode will not be changed.
l. Linux systems cache file data and heap data inside the page cache to minimize the frequency of accessing storage devices.
m. In the Ext-based file system, filenames are maintained in the corresponding directory block not in the inode.
n. In the same file system, two different files in different directories can have the same file name with the same file descriptors.
o. Using the standard C library, a user program can open and read files without the help from the operating system.

a.	b.	c.	d.	e.	f.	g.	h.
i.	j.	k.	l.	m.	n.	o.	

國立陽明交通大學 試 題 紙

一一四學年度第二學期
博士班資格考

科目：計算理論 A

日期：115 年 6 月 30 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 that a language is decidable if and only if it is Turing-recognizable (r.e.) and co-Turing-recognizable (co-r.e.).

2. (10\%) A 2cnf-formula is an AND of clauses, where each clause is an OR of at most two literals. Let 2SAT={<F> | F is a satisfiable 2cnf-formula}. Show that 2SAT can be solved in polynomial time

3. (10%) We say graphs G and H isomorphic if the nodes of G may be reordered so that it is identical to H. Let ISO={<G, H>| G and H are isomorphic}. Show that ISO is in NP.

4. (20%) Consider the following problem: Let S be a set of integers. Can we partition S into two parts: A and B such that the sum of integers in A is equal to the sum of integers in B? Prove that this problem is NP-complete. You can use the problem SUBSET-SUM for reduction.

國立陽明交通大學 試 題 紙

一一四學年度第二學期
博士班資格考

科目：計算理論 B

日期： 115 年 6 月 30 日

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 L_1 and L_2 are both regular languages, the union of L_1 and L_2 is also regular.

2. (10%) Prove or disprove that if the union of L_1 and L_2 is regular, L_1 and L_2 are both regular.

3. (10%) Prove or disprove that if the union of L_1 and L_2 is context-free, L_1 and L_2 are both context-free.

4. (10%) Prove or disprove that $L = \{ w \mid w = w^R, w \in \{ 0, 1 \}^* \}$ is a context-free language, where w^R is the reverse of w.

5. (10%) Prove or disprove that $L = \{ w \mid w \neq w^R, w \in \{ 0, 1 \}^* \}$ is a context-free language, where w^R is the reverse of w.

國立陽明交通大學 試 題 紙

一一四學年度第二學期
博士班資格考

科目：計算機架構 A　　日期：115 年 6 月 30 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. Briefly answer each question. (4pts each)
   i. A processor reduces its clock period by 20% through deeper pipelining. Using the Iron Law of Performance, explain why execution time may not improve proportionally. Give two possible reasons.

   ii. A fixed-length ISA reduces its architected registers from 32 to 16. Suppose the freed encoding bits are used either to double the immediate field size or to add new instructions. Under what workload characteristics might each choice be preferable?

   iii. Consider an 8-issue in-order processor and a 4-issue out-of-order processor implemented with similar transistor budgets. Why might the narrower processor achieve higher IPC?

   iv. A branch predictor improves from 99% to 99.5% accuracy. Under what processor and workload characteristics would this small improvement produce a large performance gain?

   v. Register renaming removes WAR and WAW dependencies. Why can a processor still experience stalls even if all false dependencies are eliminated?

   vi. Modern processors often increase instruction window size rather than issue width. Under what workload conditions would increasing issue width be more beneficial?

   vii. Intel Itanium exposed more scheduling decisions to the compiler than conventional superscalar processors. Explain one advantage and two disadvantages of this approach.

   viii. When increasing SMT threads from 2 to 4, identify one processor structure that typically scales with thread count and one shared structure that may become a performance bottleneck?

   ix. A processor implements MESI and maintains coherence perfectly. Give an example of a memory behavior that may still violate programmer expectations and explain why a fence could be required.

   x. Is cache coherence sufficient to implement sequential consistency? Justify your answer.

2. True or False. Briefly justify your answer in one or two sentences. (2pts each)

   i. If software compatibility were not a concern, many features of the x86 ISA would likely not appear in a newly designed general-purpose ISA today.
   ii. Increasing the reorder buffer size can improve performance even when issue width remains unchanged.
   iii. Register renaming increases the number of architectural registers visible to software.
   iv. Two processors implementing the same ISA may require different numbers of cycles to execute the same instruction sequence.
   v. If branch prediction were perfect, an out-of-order processor would no longer require register renaming.

國立陽明交通大學 試 題 紙

一一四學年度第二學期
博士班資格考

科目：計算機架構 B

日期： 115 年 6 月 30 日

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. There is a processor P1. P1 has a clock cycle time of 0.4 ns and CPI of 2.
   (a) What is the clock rate of the processor P1? (2 pt)
   (b) If the processor P1 executes a program in 20 seconds, please find the number of instructions executed by the processor P1. (3 pt)
   (c) We reduce the execution time from 20 seconds to 10 seconds, but this leads to an increase of 25% in the CPI. What clock rate should we have to get this time reduction? (3 pt)
2. Assume that main memory accesses take 160 ns and that 40% of all instructions access data memory. The following table shows data for the L1 cache attached to the processor P1.

	L1 size	L1 Miss Rate	L1 Hit Time
P1	8KiB	5%	0.5 ns

   (a) What is the average memory access time (AMAT) for P1 (in cycles)? (3 pt)
   (b) Assume a base CPI of 1.0 without any memory stalls, what is the total CPI for P1? (3 pt)
   (c) What is the AMAT for the processor P1 with the addition of an L2 cache? (3 pt)

L2 size	L2 Miss Rate	L2 Hit Time
1MiB	95%	5 ns

   (d) What would the L2 miss rate need to be in order for P1 with an L2 cache to be faster than P2 without an L2 cache? (3 pt)
3. The following list provides parameters of a virtual memory system

Virtual address (bits)	Physical DRAM	Page Size	PTE Size (byte)
46	32 GiB	4 KiB	4

   (a) For a single-level page table, how many page table entries (PTEs) are needed? (3 pt)
   (b) For a single-level page table, how much physical memory is needed for storing the page table? (3 pt)
   (c) How many levels of page tables are needed if the segments are limited to the 4KiB page size? (3 pt)
4. Consider a compute with a memory system with 16 bits for physical address, 64 bits for virtual address, page size of 4 KB, 16 bit words, and word (16 bit) addressable memory. The computer system contains a 32 KB data cache that is virtually indexed and physically tagged. The data cache is 8-way set associative and the cache line size is 16 words.
   (a) How many virtual pages does this system have? (2pt)
   (b) Assume each page table entry (PTE) has 3 state bits in addition to the address translation. How many bits does each PTE consume? (3 pt)
   (c) What is the total number of sets in this 8-way set associative cache? (3 pt)
   (d) If the architect decreases the cache associativity to 4-way, what would the cache index bits be? (3 pt)

5. The bus-based dual-core multi-processor is shown in below. The Figure below represents a symmetric shared memory architecture. Each processor has an L1 write-back private cache and use the MSI write-invalidate snooping cache coherence protocol. Each cache is direct-mapped with four blocks, and each block holds two words. The tag contains the full address in hexadecimal and the data is shown in decimal. Please state the resulting state, tag, and value of the caches and memory after the bus transaction (if any) in a given operation below.
[Hint: P0 read address 0x220, P0 places Read Miss on the bus at the address 0x220, Transfer block at 0x220 in memory to P0, Processor P0 read block 0 (S, 0x220, 6)]
   (a) P0 write the value 54 to address 0x220. Please write down the state and data value of (state, 0x220, data value) in the P0. (3 pt)
   (b) P0 reads address 0x228. Please write down the state and data value of (state, 0x228, data value) in the P0 after the bus transaction through MSI cache coherence protocol. (3 pt)
   [Hint: P0 write-back block 1, which replaced: (M, 0x208, 32)]
   (c) P1 write address the data value 23 to the address 0x218. Please write down the state and data value in P0 and P1, respectively. What is the (state, 0x218, data value) in P0? What is the (state, 0x128, data value) in P1? (4pt)

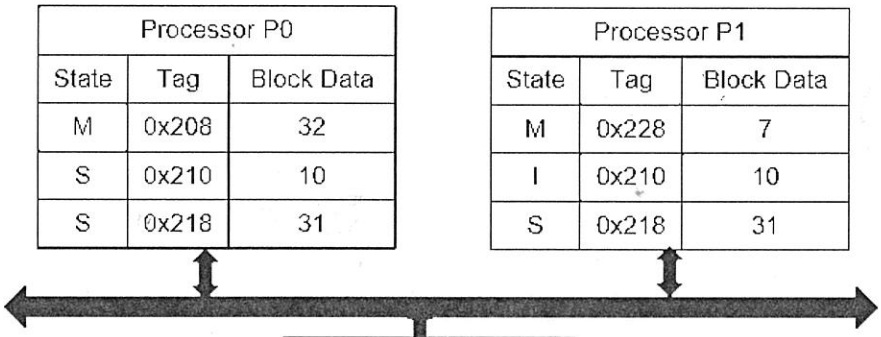

Processor P0		
State	Tag	Block Data
M	0x208	32
S	0x210	10
S	0x218	31

Processor P1		
State	Tag	Block Data
M	0x228	7
I	0x210	10
S	0x218	31

Memory	
Addr	Block Data
0x208	32
0x210	10
0x218	31
0x220	6
0x228	19

國立陽明交通大學 試 題 紙

一一四學年度第二學期
博士班資格考

科目：演算法 A　　日期： 115 年 7 月 1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. Implement a queue by using two stacks S_1 and S_2.

10% (a) Describe the **enqueue** (insertion to the rear) and **dequeue** (deletion from the front) operations. You should define the stack operations first.

10% (b) Show that a **dequege** needs worst case $O(n)$ time but $O(1)$ amortized time.

10% 2. Consider a tree constructed in this way: T_0 is a node. T_k is constructed by merging two $T_{(k-1)}$: The root of one $T_{(k-1)}$ is the leftmost child of the root of the other $T_{(k-1)}$. Figure 1 shows a T_3. If T_k has n nodes, show that the height of T_k is $O(\log n)$.

10% 3. Union/Find operation for sets representation. A set is a rooted tree. Every node in the tree has a pointer that points to its parent. The root's pointer points to itself. The union of two sets is to merge two trees by making the pointer of the root of one tree pointing to the root of the other tree. The union's weighting rule is to make the root of the smaller tree a child of the root of the larger tree. A smaller tree means it has fewer nodes. Show that a tree with n nodes constructed in this way has height $O(\log n)$. An example of T_3 is shown in Figure 2.

10% 4. T is a binary tree. If every non-leaf node (internal node) has two children, show that the number of leaf-nodes is always 1 more than the number of internal nodes. Figure 3 shows an example.

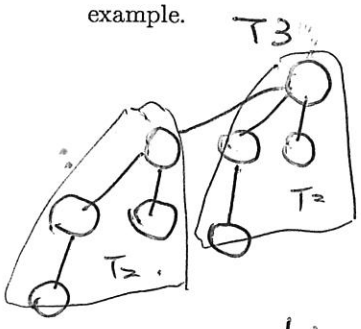

Figure 1.

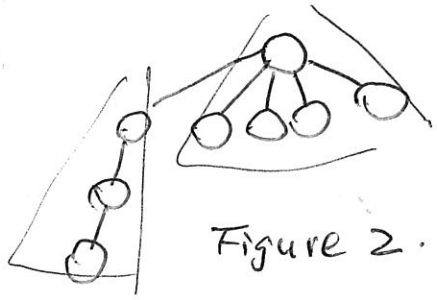
Figure 2.

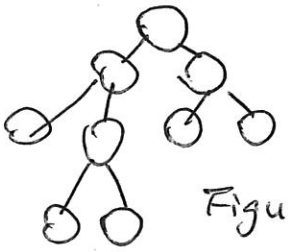
Figure 3.

9 nodes.
4 internal nodes.
5 external nodes (leaves).

國立陽明交通大學 試 題 紙

一一四學年度第二學期
博士班資格考

科目：演算法 B　　日期：115 年 7 月 1 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. Given the following two grids with shaded obstacles, you are asked to count the number of paths which start from the lower-left square, end in the upper-right square, and step either upward or rightward, avoiding the obstacles. For the smaller grid (of 4-by-4 squares with 2 obstacles) on the top, there are 7 paths, as enumerated on its right. How many paths are there, for the larger grid (of 6-by-6 squares with 3 obstacles) on the bottom.

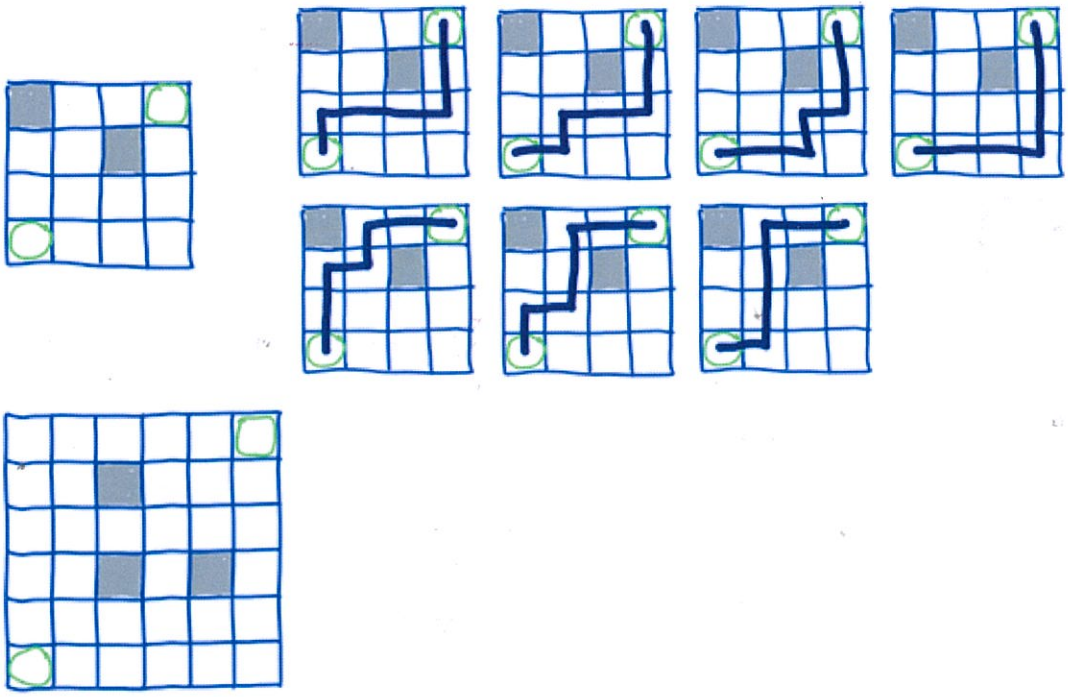

***** DO THIS *****

(a) [6 points] "Derive" (rather than "count") the number of paths in the bottom grid of 6-by-6 squares. Show your derivation with concise explanation, and finally, clearly show the "number" (of paths).

(b) [9 points] Write the pseudo-code implementation of your methodology used in **(a)**, but for a general-size grid of m-by-n squares, denoted by $grid[1 \sim m][1 \sim n]$, with obstacles indicated by $obstacles[1 \sim m][1 \sim n]$ (i.e., $obstacles[i][j] = 1$ if $grid[i][j]$ is an obstacle). Note that, in your pseudo-code, we are particularly looking for something like "initialization" and "recursive formulation" for grading.

2. Kruskal's algorithm is a greedy approach used to find a Minimum Spanning Tree (MST) for a connected, weighted, undirected graph $\boldsymbol{G} = (\boldsymbol{V}, \boldsymbol{E})$. The algorithm proceeds as follows:
   (1) Sort all edges in non-descending order of their weights.
   (2) Start with an empty set of edges $\boldsymbol{T}$.
   (3) For each edge (u, v) in the sorted list, if adding (u, v) to $\boldsymbol{T}$ does not create a cycle, then add it to $\boldsymbol{T}$.
   (4) Repeat until $\boldsymbol{T}$ contains $|\boldsymbol{V}| - 1$ edges.

   ***** DO THIS *****

   (a) [6 points] Prove that there exists at least one Minimum Spanning Tree (MST) of a graph $\boldsymbol{G}$ which contains the edge with the absolute smallest weight in $\boldsymbol{G}$.
   HINT: Suppose an MST $\boldsymbol{T}$ does not include the smallest-weight edge e, and show that you can swap it with another edge to form a new MST.

   (b) [9 points] Prove the correctness of Kruskal's algorithm.
   HINT: Using mathematical induction is recommended.

3. **[10 points]** Derive the satisfiability of each of the following 2-SAT problems (in conjunctive normal form, CNF). If it is satisfiable, answer **SAT**; otherwise (if it is un-satisfiable), answer **UNSAT**.

   (a) $F = (a + b)\ (a + \neg b)\ (\neg a + b)\ (\neg a + \neg b)$
   (b) $F = (a + b)\ (a + \neg b)\ (\neg a + c)\ (\neg a + \neg c)$
   (c) $F = (a + b)\ (\neg a + b)\ (\neg b + c)\ (\neg b + d)\ (\neg c + \neg d)$
   (d) $F = (a + b)\ (\neg a + b)\ (\neg b + \neg c)\ (\neg b + \neg d)\ (c + d)$
   (e) $F = (a + b)\ (\neg a + b)\ (\neg b + \neg c)\ (\neg b + \neg d)\ (c + \neg d)$

4. **[10 points]** Given CNF: $F = (a + b)\ (\neg a + b)\ (\neg b + c)\ (\neg b + d)\ (\neg c + \neg d)$

   (a) Draw F's implication graph. **HINT: There are 8 vertices and 10 edges in the graph.**
   (b) Find SCC(s) in F's implication graph.
   (c) Explain the satisfiability of F based on your result of finding SCC(s).