

題號： 360

國立臺灣大學 111 學年度碩士班招生考試試題

科目： 計算機結構與作業系統(B)

題號： 360

節次： 2

共 10 頁之第 1 頁

答題注意事項/Instructions to answer the questions :

- 1) 考題分為【單選題】與【複選題】。各題的配分不完全相同，分別標示。
There are two types of questions: **Multiple Choices** and **Multiple Choices with at least one correct answer**. Each question has its individual score.
- 2) 【單選題】：只有一個選項為正確答案。不倒扣。
Multiple Choices: only one among the four or five choices is the correct answer. No penalty for incorrect answer.
- 3) 【複選題】：至少有一個選項為正確答案。每一個選項分別計分。錯誤選項為零分，不倒扣。

Multiple Choices with one or more correct answers: at least one among the four or five choices are the correct answer. Each answer will be graded individually. No penalty for incorrect selection.

見背面

【單選題】：只有一個選項為正確答案。不倒扣。

Multiple Choices: only one among the four or five choices is the correct answer. No penalty for incorrect answer.

1. (5 pts) Given the logical address of a memory and its page tables where the page size is 4 bytes, the number of entries on each page is 4, and the size of physical memory is 32 bytes. Please provide the physical address of variable "g" and "o" shown in below figure.

Logical Address	Variable
0	a
1	b
2	c
3	d
4	e
5	f
6	g
7	h
8	i
9	j
10	k
11	l
12	m
13	n
14	o
15	p

Logical Page	Physical
0	10
1	7
2	10
3	10

- (A) 26 and 10
 (B) 26 and 7
 (C) 22 and 10
 (D) 10 and 26
 (E) 7 and 10
2. (5 pts) Please provide the page offset and the size of page table according to the following parameters. Consider a system with a 24-bit logical address, each entry is 1 byte, and each page is 4 KiB.
- (A) 9 and 32 KiB
 (B) 10 and 16 KiB
 (C) 11 and 8 KiB
 (D) 12 and 4 KiB
 (E) 13 and 1 KiB
3. (5 pts) The following line opens a file for write only.


```

...
int fd = open(file_name, O_CREAT | O_EXCL | O_WRONLY,
new_file_mode);
    
```

```

if (fd == -1) {
    /* Handle Error */
}
...
    
```

Please indicate the sequence of executing the above line if fd is equal to -1, using the symbol in the circle.



- (A) A -> B -> C -> D
- (B) E -> B
- (C) A -> D
- (D) A -> E
- (E) B -> C -> A

4. (5 pts) Compilers could affect program execution efficiency. We have two compiler options, C1 and C2. Given a program A, compiler C1 results in a dynamic instruction count of 1.0×10^9 and has an execution time of 1.1 second, while compiler C2 results in a dynamic instruction count of 1.2×10^9 and an execution time of 1.5 second. Assume the processor has a clock cycle time of 1ns. What is the average CPI for that program compiled with C1 and C2, respectively?

- (A) C1= 1.3 ; C2=1.35
- (B) C1= 1.2 ; C2=1.30
- (C) C1= 1.1 ; C2=1.25
- (D) C1= 1.0 ; C2=1.20
- (E) None of the above

5. (5 pts) Consider the fragment of RISC-V assembly codes below:

#1	sd x29, 12(x16)
#2	ld x29, 8(x16)

見背面

#3	sub x17, x29, x14
#4	add x15, x17, x14
#5	sub x15, x30, x14

Assume that we have a standard 5-stage pipeline (IF, ID, EX, MEM, WB) without data forwarding, and a register can be written and read back in the same cycle without introducing a hazard. Which CPU cycles do instruction #3 and instruction #4 complete assuming a perfect cache?

- (A) 7 ; 9
- (B) 8 ; 9
- (C) 8 ; 10
- (D) 9 ; 11
- (E) 9 ; 12

6. (5 pts) Below is the RISC-V assembly program generated for the body of function f : " $int f(int a, int b, int c)$ ". The program follows RISC-V's convention and uses registers x10-x18 to pass parameters and return values in function calls.

```
addi x2, x2, -16
sd x1, 0(x2)
slli x5, x12, 5
sd x5, 8(x2)
jal x1, m
ld x11, 8(x2)
lui x12, 0x11232
addi x12, 0x334
ld x1, 0(x2)
addi x2, x2, 16
jal x0, g
```

Which of the following is the respective C code? Note that the addi instruction in RISC-V is sign-extended for immediate values.

- (A) return g(m(a, b*5), 0x11232334)
- (B) return g(m(a, b), c*32, 0x33411232)
- (C) return g(m(a, b), 0x11232334, c*32)
- (D) return g(m(a, c), b*32, 0x11232334)
- (E) return g(m(a, b), c*32, 0x11232334)

7. (5 pts) Consider the RISC-V code below were executed concurrently on two cores.

```
[x2] = 1
[x3] = 1
-----
ld    x1, 0(x2)
addi  x1, x1, 1
sd    x1, 0(x2)
ld    x4, 0(x3)
add   x4, x1, x4
sd    x4, 0(x3)
```

Memory is shared by the two cores. Before the executing the code, assuming registers x2 and x3 of each of the cores contain different values ($x2_{\text{core0}} \neq x3_{\text{core0}}$ and $x2_{\text{core1}} \neq x3_{\text{core1}}$). The same register x2 and x3 of the two cores respectively contain the same value ($x2_{\text{core0}} = x2_{\text{core1}}$ and $x3_{\text{core0}} = x3_{\text{core1}}$). The memory location [x2] and [x3] each contains the integer value 1 before executing the code. Assuming no instruction reordering is possible. What are all the possible values of [x2] and [x3] after code execution has completed?

- (A) [x2] = 2, [x3] = 3 or 4 or 5
- (B) [x2] = 2 or 3, [x3] = 3 or 4 or 5 or 6
- (C) [x2] = 2 or 3, [x3] = 3 or 4 or 5 or 6 or 7
- (D) [x2] = 1 or 2 or 3, [x3] = 2 or 3 or 4 or 5
- (E) [x2] = 3, [x3] = 4 or 5 or 6 or 7

【複選題】：至少有一個選項為正確答案。每一個選項分別計分。錯誤選項為零分，不倒扣。

Multiple Choices with one or more correct answers: at least one among the four or five choices are the correct answer. Each answer will be graded individually. No penalty for incorrect selection.

8. (4 pts) The UNIX system call **wait()** returns the process identifier of a terminated child process. Select correct description(s).
- (A) Cascading termination means that "if a process terminates, then all its children must also be terminated."
 - (B) The process identifier of a zombie process is released even if its parent has not yet invoked **wait()**.
 - (C) If there is no child process when **wait()** is invoked, then it is as if no **wait()** is there.
 - (D) If there is an orphan process, the system can assign a process as the new parent which periodically invokes **wait()**.
9. (4 pts) Regarding threads, select correct description(s).
- (A) Implicit threading transfers the creation and management of threading from application developers to compilers and run-time libraries.
 - (B) Thread pools limit the number of threads that exist at any time, which is advantageous to resource-constrained systems.
 - (C) A system with a single computing core cannot support concurrent computation.
 - (D) POSIX thread (Pthread) cancellation occurs only when a thread reaches a cancellation point, which is an asynchronous cancellation.

10. (4 pts) Given a set of 3 **periodic** processes as shown below:

Process	Arrival Time	Period	Burst Time
P ₁	0	30	10
P ₂	10	70	10
P ₃	20	100	20

Perform the **rate-monotonic scheduling**. Select correct description(s).

- (A) The worst-case turnaround time (from the arrival time to the time of the completion of a process) of P₂ is 10.
- (B) The worst-case turnaround time of P₃ is 30.

- (C) If we increase the burst time of P_3 by X time units, the worst-case turnaround time of P_3 will also be increased by X time units.
- (D) If we change the arrival time of a process, the worst-case turnaround times of all processes will be the same.

11. (5 pts) Which of the following are correct for Virtual Memory?

- (A) Each process has its own virtual space
- (B) Allow more processes to run concurrently
- (C) Allow address spaces to be shared by several processes
- (D) Use SEVERAL bases
- (E) Use ONE base

12. (5 pts) Which of the following are correct for consistency semantics for file systems?

- (A) Consistent semantic specify when the ownership of a lock will be transferred to another process.
- (B) Consistent semantics specify when modification of data by one process will be observable by other processes.
- (C) UNIX semantic is supported by modern distributed file systems.
- (D) Session semantic requires one file to be associated with a single physical image that is accessed as an exclusive resource.
- (E) Consistent semantic specify how multiple processes of a system are to access a shared file.

13. (5 pts) Regarding the critical-section problem and synchronization, select correct description(s).

- (A) A resource-allocation graph has a cycle if and only if there is a deadlock.
- (B) The "progress" requirement implies the "bounded waiting" requirement.
- (C) The "bounded waiting" requirement implies the "progress" requirement.
- (D) The "freedom from deadlock" property implies the "freedom from starvation" property.
- (E) The "freedom from starvation" property implies the "freedom from deadlock" property.

見背面

14. (5 pts) Assume a virtual memory system with an 4KiB page size. The cache capacity is 16KiB, and the block size is 8-word (1 word = 4 bytes). Which cache architecture could be implemented as a virtually-indexed, physically tagged cache:
- (A) Direct-mapped cache
 - (B) 2-way associative cache
 - (C) 4-way associative cache
 - (D) 8-way associative cache
 - (E) 16-way associative cache
15. (5 pts) Assume the CPI for INT instructions, FP instructions, L/S instructions, and branch instructions for a given processor P is 1, 2, 4, and 2, respectively. Consider a program F that includes 12×10^6 INT instructions, 15×10^6 FP instructions, 50×10^6 L/S instructions, and 13×10^6 branch instructions. Assume the program executes on a machine with a single processor P that has a 3.2 GHz clock rate. Which of the following are correct?
- (A) If one can optimize the program F by removing 80% of the INT instructions, the execution time of F can be reduced by more than 3%.
 - (B) It is possible to make program F run 2 times faster by reducing the CPI of the FP instructions.
 - (C) One can optimize program F to run more than 2 times faster by reducing the CPI of the L/S instructions to 1.2.
 - (D) Suppose a new type of INT instruction called INT1 is added to the instruction set of processor P to replace the existing INT instructions. On average, through the use of the new INT instruction, we can reduce the number of INT instructions needed to execute a program by 50%, while increasing the clock cycle time by only 10%. If we replace the INT instructions in program F with the INT1 instructions, program F will run faster on processor P.
 - (E) If 30×10^6 system call instructions (each has CPI of 10) are added to the program F, F will run more than 2 times slower after the addition.

16. (5 pts) Recent hardware architectural extensions include the support for Trusted Execution Environment (TEE). Hardware isolates the software running in the TEE from the outside attackers. Intel provides the hardware support called Software Guard Extensions (SGX) that allows programs to run in its TEE called enclave in userspace. Intel SGX encrypts the enclave's code and data in CPU registers and RAM memory. Attackers that control another application or privileged systems software such as an operating system can only read or write encrypted code and data of the enclave. Which of the following are true?

- (A) Intel SGX ensures that the software running in the TEE can never crash.
- (B) As the OS kernel is involved when loading the application's binary into the enclave, Intel SGX supports remote attestation to authenticate the binary to ensure the application's safety.
- (C) It is possible to run an entire OS kernel such as Linux in the SGX enclave.
- (D) Although the enclave is encrypted by the Intel SGX hardware, an attacker residing in the OS kernel could still modify the enclave's page tables to affect the enclave's execution.
- (E) To support encrypting the enclave's secrets for persistent storage, the SGX enclave could use a persistent hardware-based encryption key to securely encrypt and store its sensitive data in a way that ensures the data can be retrieved only when the trusted environment is restored.

17. (5 pts) Which of the following are true about RISC-V instructions about the registers in the register file and instruction size?

- (A) If the registers were halved, I-type instructions could have 2 more immediate bits.
- (B) If the registers were halved, there must be less I-type instructions.
- (C) If the registers were halved, shift amounts would change to 0-63.
- (D) Increasing the size of each bit field to support more registers potentially makes each instruction longer.
- (E) Increasing the number of registers must increase the code size of a RISC-V assembly program.

見背面

18. (8 pts) Given a set of n processes, the arrival time of the i -th process is $(i-1)$, and the burst time of the i -th process is $(2n-2i+2)$, as shown below:

Process	Arrival Time	Burst Time
P_1	0	$2n$
P_2	1	$2n-2$
P_3	2	$2n-4$
...
P_n	$n-1$	2

Perform **preemptive scheduling** to minimize the **average turnaround time** (from the arrival time to the time of the completion of a process) of all processes. Assume that the minimized average turnaround time of all processes is $Wn^3 + Xn^2 + Yn + Z$. Select correct description(s).

- (A) $W + X + Y + Z = 1$.
- (B) $X - Y + Z = -1$.
- (C) $2XY = 1$.
- (D) $W^2 + 3Z = 2$.

19. (10 pts) Which of the following statements are correct?

- (A) ARM big.LITTLE is a multi-core architecture designed for energy-efficiency.
- (B) You are asked to deliver 90x speedup from a 100-processor system. The target application can only have at most 5% of codes that can't be serialized.
- (C) A RISC processor can always deliver better performance than a CISC processor.
- (D) A larger block size could potentially lead to more false sharing misses.
- (E) A larger block size could potentially help reduce compulsory misses.

(以下空白)

試題隨卷繳回