

# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)  
系所班別：資訊聯招

考試日期：105 年 2 月 3 日第 3 節

第 1 頁，共 10 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！！

請使用答案卡作答

複選題：該題全部答對得 4 分，只錯一個答案可得 2 分，錯兩個或兩個以上答案不給分，整題未作答不給分

1. Which of the following statements are true?
  - (a) Threads within the same process share the same heap and stack.
  - (b) Bootstrap program is loaded at power-up or reboot and typically stored in RRAM.
  - (c) A two phase locking protocol for atomic transactions has two phases. One phase obtains locks but not release any locks and another release locks but not obtain any new locks.
  - (d) Caching is an important structure that performed at many levels in a computer.
  - (e) A process is a program in execution. It is a unit of work within the system. Program is an active entity, process is a passive entity.
2. Which of the following statements are true?
  - (a) Computer system can be divided into four components including hardware, operating system, application programs and users.
  - (b) Most operating systems have three general methods used to pass parameters including passing in registers, storing in a block and pushing onto a stack.
  - (c) A CPU-bound process spends more time doing computations and consists of many short CPU bursts.
  - (d) The goal of multi-threads is to utilize CPU in maximum. Hence, we should create threads as much as possible.
  - (e) It is important that the short-term scheduler select a good process mix of I/O-bound and CPU-bound processes.
3. Which of the following statements are true?
  - (a) An I/O-bound process spends more time doing I/O than computations and consists of a few long CPU bursts.
  - (b) The long-term scheduler selects processes among the processes that are ready to execute and allocates the CPU to one of them.
  - (c) In UNIX, the fork() system call is used to load a new program into the process' memory space.
  - (d) Remote procedure call (RPC) abstracts procedure calls between processes on networked systems.
  - (e) Main memory is a volatile storage device that loses its contents when power is turned off.
4. Which of the following statements are true?
  - (a) Dual-mode operation allows OS to protect itself and other system components. If OS doesn't have dual-mode operation, every user can modify the kernel system without permissions.
  - (b) DMA is used for high-speed I/O devices to transfer blocks of data from buffer storage directly to main memory without CPU intervention.
  - (c) The user program can deal with both of logical addresses and physical addresses.
  - (d) Message passing is the only mechanism of inter-process communication.
  - (e) If a system does not employ either Deadlock-prevention or Deadlock-avoidance, there is no other method to figure out the deadlock problem.



# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)

系所班別：資訊聯招

考試日期：105 年 2 月 3 日 第 3 節

第 2 頁，共 10 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!!

5. Which of the following statements are true?
  - (a) A multithreaded process has a program counter, pointing to the next instruction to execute for all threads.
  - (b) The short-term scheduler controls the degree of multiprogramming.
  - (c) In UNIX, the `exec()` system call is used to create a new process.
  - (d) Blocking is considered synchronous. Non-blocking is considered asynchronous.
  - (e) A ready process cannot do I/O or anything else that might block it.
6. Which of the following are true about SAN (storage area network) vs. NAS (network attached storage)?
  - (a) Both are accessed via network.
  - (b) NAS operates at block level while SAN operates at file level.
  - (c) CIFS/SMB and NFS are examples of NAS protocols.
  - (d) For databases, SAN is often the preferred choice over NAS.
  - (e) Almost any machine running Microsoft Windows with LAN connectivity can be configured to access a NAS.
7. A UNIX file system uses 4KB (4096 bytes) blocks and 4-byte disk address. Each inode contains 12 direct entries, 1 singly-direct entry, and 1 doubly-indirect entry. The maximum file size would be about X MB (1MB = 1,048,576 bytes). Let  $R_0 = X \bmod 5$ ,  $R_1 = X \bmod 4$ , and  $R_2 = X \bmod 3$ . Select the values of  $R_0$ ,  $R_1$ , and  $R_2$  from following. (for instance, if  $X=256$ ,  $R_0 = 1$ ,  $R_1 = 0$ , and  $R_2 = 1$ , you should select a and b)
  - (a) 0
  - (b) 1
  - (c) 2
  - (d) 3
  - (e) 4
8. Which of the following RAID configurations improve disk I/O performance during normal operation and also provide fault tolerance for single disk failure?
  - (a) RAID 0
  - (b) RAID 1
  - (c) RAID 1+0
  - (d) RAID 0+1
  - (e) RAID 5
9. Which of the following are true about using contiguous allocation in filesystem:
  - (a) Filesystem is prone to internal fragmentation.
  - (b) The space overhead of file allocation record is minimal. (only need to track a file's starting location and its length)
  - (c) Very easy to increase the size of a file dynamically.
  - (d) Sequential access to file is efficient.
  - (e) Random access to file is efficient.
10. It is now a common practice to execute potentially vulnerable code in a child process. For instance, a web browser may spawn a child process for running flash player. Which of the following could be the motivations?
  - (a) Separating the memory address space so that malicious code would not be able to corrupt the memory content of the main program.
  - (b) Preventing malicious code from accessing the system resources (e.g. opened files, sockets) used by the main program.
  - (c) Unnecessary privileges can be stripped away from the child process so that malicious code would not be able to invoke certain system calls.



# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)

系所班別：資訊聯招

考試日期：105 年 2 月 3 日 第 3 節

第 3 頁，共 10 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!!

- (d) Child process can run with a different user ID so that access control can be applied to prevent malicious code from accessing critical files on the system.
- (e) Reduce the memory usage of the main program so that denial-of-service attack will be much harder to carry out.
11. Which of the following statement are correct?
- (a) Writing to programs with a set of powerful instructions is the shortcut to yield high performance.
- (b) The number of pipeline stages affects latency, not throughput; thus pipeline improves the performance of a processor by decreasing the latency of a job (i.e. an instruction) to be done.
- (c) For a given program, its average cycles per instruction (CPI) is affected not only by the instruction set architecture, but also by the compiler used.
- (d) By reducing the clock frequency of a processor from 2 GHz to 1.5 GHz, and also reducing its supply voltage from 1.25 Volt to 1 Volt, the overall power consumption of this processor will be reduced by 40% theoretically.
- (e) Hexadecimal integer value 0xABCCBBA has identical storage sequence in the (byte-addressable) memory no matter the machine is big-endian or little-endian.
12. Consider a byte-addressable memory hierarchy with 44-bit addresses. The memory hierarchy adopts a 4-way set associative cache of 64 KB with every block in a set containing 64 bytes. Which of the following statements are *correct*?
- (a) Among an address of 44 bits, the least-significant 6 bits is the "offset".
- (b) Among an address of 44 bits, the most-significant 30 bits is the "tag".
- (c) The total number of bits required to store the entire cache (including valid bits, tags, and data) is  $(1+30+64 \times 4 \times 8) \times 2^8 = 532,224$  (bits).
- (d) By increasing the block size while using the same size of storage for data, the total number of bit required to storage the entire cache will decrease.
- (e) Given the following sequence of access addresses:  $(0E1B01AA050)_{16}$ ,  $(0E1B01AA073)_{16}$ ,  $(0E1B2FE3057)_{16}$ ,  $(0E1B4FFD85F)_{16}$ ,  $(0E1B01AA04E)_{16}$ , assuming that the cache is initially empty, there will be *three* set containing exactly *one* referenced block at the end of the sequence.
13. Which of the following statements (about virtual memory and page table) are *correct*?
- (a) For virtual memory, write-through is more practical than write-back.
- (b) For virtual memory, full associativity is typically used for minimizing page fault rate.
- (c) Given a 32-bit virtual address space with 2 KB per page and 4 bytes per page table entry, the total page table size is 8 MB.
- (d) It is possible to miss in cache and translation look-aside buffer (TLB), but hit in page table.
- (e) It is possible to miss in TLB, but hit in cache and page table.
14. Which/of the following statements are *correct*?
- (a) It is possible that increasing block size could increase the overall miss rate.
- (b) It is possible to decrease the occurrence of capacity misses by increasing block size.
- (c) It is possible to decrease the occurrence of conflict misses by increasing associativity.
- (d) It is possible to "eliminate" the occurrence of conflict misses by further increasing associativity.
- (e) It is possible that increasing associativity could increase the overall miss rate.

# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)

系所班別：資訊聯招

考試日期：105 年 2 月 3 日第 3 節

第 4 頁，共 10 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！！

15. Which of the following statements are *correct*?


- (a) RAID (redundant array of inexpensive/independent disks) not only enhances the reliability/availability of data storage, but also boosts the performance of data access.
- (b) RAID 0 provides fault tolerance by replicating data to mirror disks.
- (c) The advantage of RAID 4 over RAID 3 is on the byte-level striping/interleaving of data across disks, so as to allow more parallel data accesses.
- (d) The advantage of RAID 5 over RAID 4 is on the distribution of parity blocks across disks, so as to avoid a single parity disk from being the bottleneck.
- (e) The advantage of RAID 6 over RAID 5 is on the higher degree of fault tolerance through more disk redundancy.

ARM recently built up a new CPU design center at Hsinchu in Taiwan, the first CPU design center in Asia. You are the principal engineer in a team in charge of developing a brand new CPU, called NCTU (Next-generation Compute Terabit Unit).

The following what—if questions refer to some assumptions. The latencies for logic blocks in the figure are listed, while the latency of other blocks is almost negligible:

| Main control | I-mem | Adder | Registers | ALU  | Sign-extend | D-mem |
|--------------|-------|-------|-----------|------|-------------|-------|
| 50ps         | 100ps | 50ps  | 40ps      | 50ps | 10ps        | 100ps |

The NCTU1, the basic standard 5-stage pipeline processor, is organized as the following diagram. All instructions executed by a pipelined processor are broken down as ALU: 40%, beq: 30%, lw: 25% and sw: 5%.





# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)

系所班別：資訊聯招

考試日期：105 年 2 月 3 日第 3 節

第 5 頁，共 10 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!!

16. Which of following statement can be true for ideal NCTU1 (assume there is no stalls, where all instructions are executed ideally)?
- (a) The best clock rate of NCTU1 is 10GHz/
  - (b) The MIPS (million instructions per second) can be reached up to 10000.
  - (c) All pipeline registers between stages have the same width (number of bits).
  - (d) The pipeline clock cycle time is equal to the average of all stage latencies.
  - (e) The longest instruction will be lw instruction. Reducing number of lw instructions in a program can improve the NCTU1 system throughput (e.g. CPI or instructions executed per second).
17. Ideal execution does not exist at all. Which of following statements can be true for improving NCTU1?
- (a) The pipe stages look unbalanced. The ideal CPI can be reduced to less than 1 if we move around the building blocks for better balance.
  - (b) The overall performance can be always improved by making the pipeline deeper as the cycles are shorter.
  - (c) Solve the load-use data hazard (load instructions followed by a dependent instruction) by inserting a hazard detection unit at the ID stage.
  - (d) Increase the register file with more read ports (three reads and a write) to avoid instruction conflicts.
  - (e) If ALU control decoder (marked by dashed circle) needs to take 50ps, moving the ALU control to the ID Stage can improve NCTU1 for allowing more concurrent activities.
18. For the statements on right (i is \$r8 and A is \$r10), you solve the pipeline dependency of NCTU1 as the textbook. How many cycles totally are needed for NCTU1 to execute them all? (you may need to insert NOP if necessary.)
- ```
// C or C++
A[i] += 10;
A[i] *= 2;

// MIPS
add    $r1, $r8, $r10
lw     $r2, 4($r1)
addi   $r2, $r2, 10
multi  $r2, $r2, 2
sw     $r2, 12($r1)
```
- (a) Total 9 cycles ideally if we don't consider hazard problems.
  - (b) Total 19 cycles if there is no forwarding unit and no hazard detection unit, and we use NOPs.
  - (c) Total 11 cycles if there is a forwarding unit and no hazard detection unit for load-use.
  - (d) Total 10 cycles if there are forwarding unit and hazard detection unit for load-use.
  - (e) The forwarding unit will not be used in the lw instruction (load-use scenario) when hazard detection unit is triggered.
19. Now we hope to design a reduced low cost version using fewer stages, a new 4-pipe CPU called NCTU2. We can remove MEM stage and put the memory access in parallel with the ALU. Which of following statements will be true for NCTU2?
- (a) It is feasible except all load/store instructions with nonzero offsets, whereas we have to convert them into register-based only (no offset). Specifically, convert
- ```
lw $r3, 30($r5)
```

# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)

考試日期：105 年 2 月 3 日第 3 節

系所班別：資訊聯招

第 6 頁，共 10 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！！

into

```
addi $r1, $r5, 30
lw $r3, ($r1)
```

- (b) The NCTU2 has better overall system throughput, as instructions take less cycles to complete.
- (c) The NCTU2 will have a worse clock rate, as its cycle time is longer than that of NCTU1.
- (d) NCTU2 still needs a stall hazard detection unit for load-use condition.
- (e) The data forwarding unit is still needed for NCTU2.

Branches share a significant portion of program execution. In current NCTU1, branch outcomes are determined at MEM stage. It's one of the key problems.

20. Which of following statements can be true for NCTU1 by considering branches only?
- (a) The CPI with all branch stalls is 1.9 for NCTU1 (assume no other stalls).
  - (b) If the NCTU1 is designed with deeper pipelined stages, the branch control hazard becomes harder to solve.
  - (c) Assuming branches are not taken can continue execution down the sequential instruction stream and thus improve the NCTU1 without changing hardware.
  - (d) If the branch handling is moved to the EXE stage, it can improve CPI without increasing the cycle time.
  - (e) If more hardware resources such as forwarding comparator are provided, the branch resolving can be perfectly moved to the IF stage while still maintaining the normal pipelining mechanism.



# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)

系所班別：資訊聯招

考試日期：105 年 2 月 3 日第 3 節

第 7 頁，共 10 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!!

題組：各題組下的小題為單選，該題組內的小題完全答對得 5 分，任一小題答錯倒扣 0 分(該題組以零分計)，各小題均未作答不給分

## 題組 A

Now we have a new NCTU3 derived from NCTU1, where the branch handling is moved to the ID stage.

21. Which of following statements will be true?

- (a) It needs more hardware costs such as extra adder for subtraction and a forwarding unit.
- (b) The branch target address adder has to be replicated.
- (c) The good thing is that it improves CPI without increasing cycle time.
- (d) The branch only CPI of new NCTU3 will become 1.1.

22. Even with NCTU3, we hope to improve the brunch problem further. Which of following statements is NOT true?

- (a) Delayed branch is a mechanism to delay the effect of branch. Branch delay slot in NCTU3 is 1. It also needs no extra hardware resources in NCTU3.
- (b) Compiler can help in filling delay slots. If compiler can find any safe instruction for 50% of branches, the CPI can be reduced to 1.15.
- (c) Predicting branch as not-taken is a good mechanism and it needs no extra hardware resources in NCTU3.
- (d) Branch prediction buffer is good to predict the branch outcome, but it does not help in predicting the branch target.

## 題組 B

Given a 16-byte cache (byte-addressable, initially empty) and a sequence of access addresses:  $(2)_{10}$ ,  $(14)_{10}$ ,  $(0)_{10}$ ,  $(17)_{10}$ ,  $(12)_{10}$ ,  $(3)_{10}$ , please derive the corresponding hit/miss sequence for each of the following cache designs. There is one question for each design. You may use the attached tables to find the answers to the questions; filling out the tables is for your own convenience and is not the answers to be submitted.

23. 8 bytes per block, direct-mapped

| Address | Cache index | Hit or miss? |
|---------|-------------|--------------|
| 2       |             |              |
| 14      |             |              |
| 0       |             |              |
| 17      |             |              |
| 12      |             |              |
| 3       |             |              |

How many of these six accesses are "hit"?

- (a) 0
- (b) 1
- (c) 2
- (d) 3
- (e) 4

# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)

系所班別：資訊聯招

考試日期：105 年 2 月 3 日第 3 節

第 8 頁，共 10 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！！

24. 4 bytes per block, 2-way set associative

| Address | Cache index | Hit or miss? |
|---------|-------------|--------------|
| 2       |             |              |
| 14      |             |              |
| 0       |             |              |
| 17      |             |              |
| 12      |             |              |
| 3       |             |              |

How many of these six accesses are "hit"?

- (a) 0
- (b) 1
- (c) 2
- (d) 3
- (e) 4

題組 C

Considering the banker's algorithm, there is a system with 5 processes (P0-P4) and 3 types of resources (A-C). Resource A has 10 instances, resource B has 5 instances and resource C has 7 instances. At the beginning, the following snapshot of the system has been taken:

| PID | Allocation | Max   | Available |
|-----|------------|-------|-----------|
|     | A B C      | A B C | A B C     |
| P0  | 0 1 0      | 7 5 3 | 3 3 2     |
| P1  | 2 0 0      | 3 2 2 |           |
| P2  | 3 0 2      | 9 0 2 |           |
| P3  | 2 1 1      | 2 2 2 |           |
| P4  | 0 0 2      | 4 3 3 |           |

25. Please list a sequence of processes satisfies the safety criteria.

- (a) <P0, P2, P1, P4, P3>
- (b) <P1, P4, P2, P3, P0>
- (c) <P2, P0, P3, P1, P4>
- (d) <P3, P1, P4, P0, P2>
- (e) <P4, P3, P0, P2, P1>

26. Suppose that the process P1 request 1 additional resource A and 2 addition resources C. Can this request be immediately granted?

- (a) Yes
- (b) No
- (c) Maybe

27. Suppose that the process P4 request 2 additional resources B. Can this request be immediately granted?

- (a) Yes
- (b) No
- (c) Maybe



# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)  
系所班別：資訊聯招

考試日期：105 年 2 月 3 日第 3 節

第 9 頁，共 10 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!!

題組 D

Assume you are the system administrator (root) of a UNIX workstation. There are four users nancy, john, frank, and charlene on the workstation. User nancy and john belong to the student group, while user frank and charlene belong to the faculty group.

```
[root@localhost ~]#  
[root@localhost ~]# id frank  
uid=1004(frank) gid=1004(frank) groups=1004(frank), 1005(faculty)  
[root@localhost ~]# id charlene  
uid=1001(charlene) gid=1001(charlene) groups=1001(charlene), 1005  
(faculty)  
[root@localhost ~]# id john  
uid=1002 (john) gid=1002(john) groups=1002(john), 1006 (student)  
[root@localhost ~]# id nancy  
uid=1003(nancy) gid=1003(nancy) groups=1003(nancy), 1006(student)  
[root@localhost ~]#  
[root@localhost ~]# umask -S  
u=rwx, g=rx, o=rx  
[root@localhost ~]#
```

User account info and umask setting of the UNIX workstation

28. We would like to create a directory /project for nancy and john to store their files for a course project. They both need to be able to create new files under project directory. Specifically, we **only** allow the administrator and users of the student group to access files in the project directory. To achieve this, which of the following command would suffice?

- (a) chown -R root /project; chgrp -R student /project; chmod -R 770 /project
- (b) chown -R root /project; chgrp -R faculty /project; chmod -R 770 /project
- (c) chown -R frank /project; chgrp -R student /project; chmod -R 770 /project
- (d) chown -R root /project; chgrp -R student /project; chmod -R 700 /project
- (e) chown -R student /project; chgrp -R student /project; chmod -R 770 /project

29. User nancy adds a file quicksort.cpp under the project directory. The project directory now looks like the following:

```
[nancy@localhost project]$ ls -al  
total 12  
drwxrwx---, 2 root student 4096 Dec 19 17:00  
dr-xr-xr-x, 19 root root 4096 Dec 19 16:15  
-rw-rw-r--, 1 nancy nancy 5 Dec 19 16:32 quicksort.cpp  
[nancy@localhost project]$
```

# 國立交通大學 105 學年度碩士班考試入學招生試題

科目：計算機系統(1003)

系所班別：資訊聯招

考試日期：105 年 2 月 3 日第 3 節

第 10 頁，共 10 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！

User nancy does not want other users in the student group to have read or write access to the file. Which of the following command can be used by nancy to achieve this?

- (a) `chmod 666 quicksort.cpp`
- (b) `chmod 600 quicksort.cpp`
- (c) `chgrp root quicksort.cpp`
- (d) This is not possible because the project directory is associated with the studnet group.
- (e) This is not possible because the nancy is not the system administrator. She cannot use the `chmod / chgrp` commands.