國立陽明交通大學 110 學年度碩士班考試入學招生試題

科目：計算機系統(1103)　　考試日期：110 年 2 月 3 日 第 3 節

系所班別：資訊聯招　　第 1 頁, 共 9 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

一、複選題(80%)，共二十題，每題全對得 4 分，答對一個選項得 1 分(答對兩個選項得 2 分，以此類推)，答錯一個選項倒扣 1 分(答錯兩個選項倒扣 2 分，最多扣至本科目計算機系統 0 分為止)，整題未作答不給分。(舉例每題有 a、b、c、d 四個選項，某題答案為 a、b，而作答時選 a、b、c，則三個選項正確一個錯誤，該題得 3-1=2 分)

1. For different directory structures, which of following statements are true?
   (a) A tree-structured directory prohibits the sharing of files or subdirectories by symbolic links.
   (b) General graph directory structures may suffer from the problem of having dangling pointers to non-existent files.
   (c) Acyclic-graph directory structures may have the problem that file's reference count is not 0, but it cannot be accessed.
   (d) Garbage collection is to recover unused disk space. It is necessary only when there are possible cycles in the directory structures.
2. Assume there are four frames in the system. Given the memory reference patterns as follows (in terms of frame number): 50, 40, 80, 20, 50, 10, 80, 40. Which of following statements are true?
   (a) There are 5 page faults if LRU (least recently used) replacement algorithm is used.
   (b) There are 5 page faults if optimal page-replacement algorithm is used.
   (c) LRU can never exhibit Belady's anomaly.
   (d) A process is thrashing if it spends more time in paging than execution. The effects of thrashing can be limited by using a global replacement algorithm.
3. Assume both index block and data block use the same block size of 1K bytes and assume each entry of the index block consists of 4 bytes. Which of following statements are true?
   (a) The maximum file size that can be supported for single-level indexed allocation is 2^{18} bytes.
   (b) Two-level indexed allocation uses a first-level index block to point to a set of second-level index blocks. The maximum file size that can be supported by a two-level indexed allocation is 2^{26} bytes.
   (c) Both linked allocation and indexed allocation do not suffer from external fragmentation.
   (d) For an application requiring direct access, it is better to use linked allocation than indexed allocation.
4. Buddy system and Slab allocation are two techniques for allocating memory for kernel processes. Which of following statements are true?
   (a) With the buddy system, memory requests can be satisfied very quickly because objects are created in advance.
   (b) In Slab allocation scheme, a slab may consist of both used and free objects.
   (c) Using the buddy system, a request for 29KB memory will be satisfied with a 36KB segment of memory.
   (d) Using the Slab allocation scheme, no memory is wasted due to fragmentation.
5. Consider a disk queue with requests for blocks on cylinders in the listed order:
   120, 25, 5, 11, 80, 185, 90, 85.
   If the disk head is initially at cylinder 88 and is moving upward to next service

cylinder 90. Which of following statements are true?
(a) If SSTF (Shortest seek-time first) is used, the requests will be serviced in the order of 90, 85, 80, 120, 185, 25, 11, 5.
(b) If C-SCAN (Circular SCN) is used, the requests will be serviced in the order of 90, 120, 185, 85, 80, 25, 11, 5.
(c) SSTF may suffer from the problem of starvation
(d) C-SCAN is a variant of SCAN, designed to reduce the disk seek time needed by SCAN.

6. System calls are the interface between application programs and kernel. What descriptions about system calls are correct?
(a) Just like hardware devices that generate interrupts, system calls also generate interrupts and jump to the kernel.
(b) Just like hardware devices, there are also vector entries for system calls and they point to interrupt service routines for both hardware and software interrupts.
(c) If there are many parameters to a system call, the user program can store all parameters in a table and pass the address of the table as a parameter to the system call.
(d) If memory copy between kernel and user program is heavy, it can be offloaded to DMA.

7. State of a process switches between ready, running, and waiting, when device interrupt, clock interrupt, or system call happens. Suppose now a process is in the running state.
(a) When a device interrupt happens, the process enters the waiting state first.
(b) When a clock interrupt happens, the process enters the ready state first.
(c) When an I/O system call is invoked, the process enters the waiting state first.
(d) When a non-I/O system call is invoked, the process enters the ready state first.

8. In message passing for inter-process communication, with send() and receive() being two system calls, if send() in sender is non-blocking but receive() in receiver is blocking.
(a) If receive() happens before send(), the receiver will not get the message.
(b) If receive() happens after send(), the receiver will get the message.
(c) If now reversed, send() in sender is blocking and receive() in receiver is non-blocking. If receive() happens before send(), the receiver will get the message.
(d) Continued from (c), if receive() happens after send(), the receiver will not get the message.

9. For a user thread there can be a corresponding kernel thread. Do they both maintain stack and PCB (process control block)? Separately or jointly? What possibilities below are valid?
(a) User thread has its stack and PCB, but the corresponding kernel thread does not have.
(b) User thread and kernel thread share the same stack and PCB.
(c) User thread and kernel thread have their own stack separately but PCB jointly.
(d) User thread and kernel thread have their own stack and PCB separately.

10. Modern operating systems adjust priority in scheduling processes. What strategies are commonly used?
(a) Increase priority when a process uses up its quantum.
(b) Increase priority when a process has frequent I/O.
(c) Increase priority when a process wakes up from a long sleep.
(d) Increase priority for a foreground process.

11. Consider the CPU performance, which of the following statements are correct?
    (a) If you want to decrease the response time of a task, you can promote your PC by replacing the old single-core CPU with a new multi-core CPU and using a new operating system that can partition a task into several load-balanced sub-tasks and execute all sub-tasks on all cores concurrently. Assume the performance of the single-core CPU is the same as that of each core of new multi-core CPU.
    (b) If we run two programs on two CPUs of different ISAs to perform the same task, the CPU with smaller CPI must have less CPU time than the other CPU since smaller CPI also implies that the required number of cycles to complete a task is also less than that of the other CPU.
    (c) The CPU time of performing a program on a CPU is determined by the instruction count, the CPI, and the clock rate of the CPU.
    (d) Consider using the metric of Millions of Instructions per Second (MIPS) to measure the CPU performance. Two programs are run on single-core CPUs C_A and C_B of different ISAs to solve the same task. C_A requires less time than C_B to complete the task while C_A has less MIPS than C_B. The abovementioned facts implies that C_A is likely a CISC-based CPU and C_B is a RISC-based CPU.
12. Consider the instruction set of 32-bit MIPS CPU, which of the following statements are correct?

```
fib:
addi $sp,$sp, -12
sw   $ra, 0($sp)
sw   $s1, 4($sp)
sw   $a0, 8($sp)
slti $t0, $a0, 2
beq  $t0,$0,L1
addi $v0,$a0,0
j    EXIT
```

```
L1:
addi $a0,$a0, -1
jal  fib
addi $s1,$v0,0
addi $a0,$a0,-1
jal  fib
add  $v0,$v0,$s1
```

```
EXIT:
lw   $ra,0($sp)
lw   $a0,8($sp)
lw   $s1,4($sp)
addi $sp,$sp, 12
jr   $ra
```

    (a) One of the reason why there are only 32 registers in the CPU is that the use of more registers will decrease the number of instructions that can be encoded in an instruction.
    (b) As generating a MIPS instruction, a compiler generally assigns a register to the most frequently used variable to diminish memory access time.
    (c) The reason why MIPS does not support the instructions of *blt* (branch less than) and *bgt* (branch greater than) is that the hardware design of "greater than" or "less than" is more complex than "equal" and "not equal", which may make the design of supporting *blt* and *bgt* have less clock rate.
    (d) The above figure shows a MIPS assembly program. In the procedure prologue, three registers were saved: $*ra*, $*s*1, and $*a*0. In fact, only $*ra* and $*s*1need to be stored and $*a*0 does not.
13. Consider the 32-bit MIPS CPU, which of the following statements are correct?
    (a) Floating-point instructions have different register set from integer instructions.
    (b) The code sequence (*lui* $*t*1, 288; *ori* $*t*1, $*t*1, 26) performs the operation of loading a constant 0x 0120 001A into register $*t*1.
    (c) Consider the addition of three floating-point numbers ($a+b+c$) on a parallel program. The addition results of $(a+b)+c$ and $a+(b+c)$ are always the same, i.e., floating-point number addition holds the associativity.
    (d) Assume the value of register PC is 0xAB05FF84, the highest address of

instruction *j* can go to is 0xAFFF FFFC.

14. Consider the 32-bit MIPS CPU. Which of the following statements are correct?
    (a) The reason why sometimes the assembler replaces the following instructions (*beq* $*s*2, $*s*3, *L*1; *add* $*t*0, $*s*1, $*s*2) with the following code sequence (*bne* $*s*2, $*s*3, *L*2; *j L*1; *L*2: *add* $*t*0, $*s*1, $*s*2) is that the hardware to realize *beq* is slower than that to realize *bne* and *j* is also a faster instruction than *beq* due to no need to access registers.
    (b) Consider the 32-bit floating-point numbers using IEEE Std. 754, the smallest positive single-precision denormalized number is 1.0×2^{-149}.
    (c) The concept of sticky bit is to use two bits to record the number of 1s in the already rounded bits.
    (d) Consider the following code sequence (*try*: *add* $*t*0, $*zero*, $*s*0; *ll* $*t*1, 0($*s*2); *sc* $*t*0, 0($*s*2); *beq* $*t*0, $*zero*, *try*; *add* $*s*0, $*zero*, $*t*1;) running on one core of a multi-core CPU. The atomic pair of instructions (*ll*, *sc*) in the code sequence protects the data at memory 0($*s*2) from memory write of another core during the data swap of register $*s*0 and memory 0($*s*2).
15. Consider the arithmetic operation unit in the following figure for 32-bit MIPS CPU. Which of the following statement are correct?
    (a) As the operation unit performs a multiplication, the *Divisor* register acts as a *multiplicand* register and the *Remainder* register acts as a *Product* register. The lower word of *Remainder* register accommodates the dividend and the higher word of *Product* register accommodates the multiplier.
    (b) The performance gain of Booth's algorithm comes from consecutive 0s or 1s in the multiplier.
    (c) Using the non-restoring division algorithm (NRDA) and restoring division algorithm (RDA) to perform a division of 7 divided by 2 that is represented by a signed 4-bit binary number. NRDA requires 5 subtractions and 3 additions while RDA requires 2 subtractions and 3 additions.
    (d) If we want to improve the performance of 32-bit ALU, we can use carry look-ahead adder to avoid waiting for carry propagation. To exploit the hierarchical design concept, we also can apply group propagate and group generate to design a multi-level carry look-ahead adder. The multi-level design owns the benefit of using systematic way to deal with the adder of large bit number in a simple way without the other benefit such as delay.
16. Using MIPS processor design and implementation for example, which of the following statements are correct?
    (a) A control signal generally called a "clock" is used widely to control what most combinational components (such as ALU, multiplexer, etc.) should do.
    (b) If every instruction must be executed in one clock cycle, the implementation must include at least two independently operable memory units.
    (c) If every instruction can be executed using multiple clock cycles and the execution is not pipelined, this can result in the cheapest-possible hardware design (even considering the may-be-required stage latches).
    (d) For pipelined processor design, at least two independently operable memory units are a must.
17. Given following figures and tables for a few MIPS sample instructions, which of the following statements are correct?

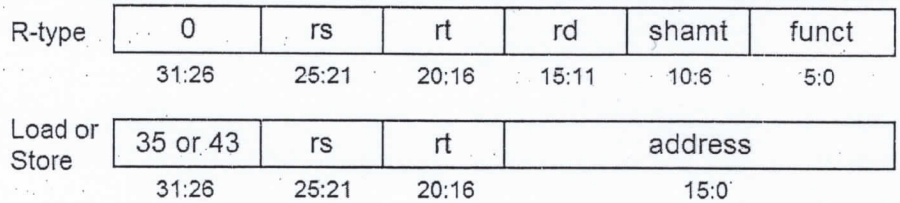

opcode	ALUOp	Operation	funct	ALU function	ALU control
lw	00	load word	XXXXXX	add	0010
sw	00	store word	XXXXXX	add	0010
beq	01	branch equal	XXXXXX	subtract	0110
R-type	10	add	100000	add	0010
		subtract	100010	subtract	0110
		AND	100100	AND	0000
		OR	100101	OR	0001
		set-on-less-than	101010	set-on-less-than	0111

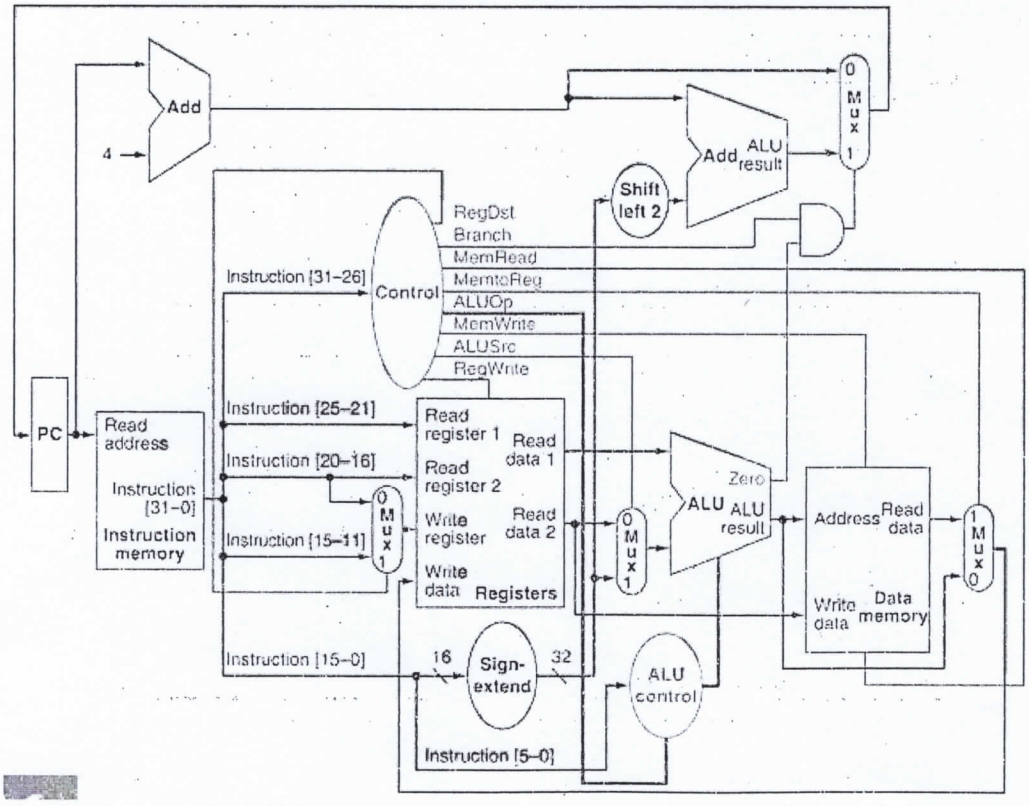

(a) The binary machine instruction encoding for “and $r4, $r6, $r8” ($r4=$r6and$r8) is 0000 0000 1100 1000 0010 0000 0010 0100

(b) The binary machine instruction encoding for “sw $r5, -7($r9)” is 1010 1101 0010 0101 1000 0000 0000 0111

(c) The instruction “beq ...” can result in a next PC (program counter) value of [(the address of “beq ...”)+(the address field in 2’s complement form)].

(d) The instruction R-type “sub ...” has a Branch=0, RegDst=1, ALUSrc=0, “ALU

control" output=0110, and MemtoReg=0.

18. We study the MIPS processor hazards using pipelined designs for a few typical instructions, which of the following statements are correct?
    (a) If there are structural hazards, any contended hardware must be duplicated many copies to satisfy any possible needs, or else the pipelined processor cannot work properly.
    (b) For sequentially fetched instructions, with a single pipeline, in-order issue and in-order completion of instructions are both guaranteed.
    (c) For true-data dependence (read-after-write), data dependence detection and forwarding can effectively alleviate the hazards, while stalls are many times still required.
    (d) For sequentially fetched instructions with a single pipeline, in-order issue can be designed to never cause anti-dependencies (write-after-read), while in doing loop unrolling, one must rename registers to remove other name-dependencies.
19. Memory design is typically hierarchical. Which of the following statements are correct?
    (a) The purpose of such design is to make memory look both large and fast. Any application running with a hierarchical memory system will have faster average memory accesses compared with a flat, non-hierarchical memory system.
    (b) Let the L1 cache L1$ be the closest-to-processor hierarchy of the memory system. In considering its role and performance, the hit time is a more important factor than the hit rate.
    (c) A processor-issued memory address will be partitioned by the L1$ into three fields: tag, index, and byte offset. Only one of these fields, the tag, needs to be stored in a proper cache line (or block) together with a valid bit, the corresponding data, and other helpful information such as a dirty bit.
    (d) Since processor accesses to L1$ will miss, the average memory access time of a program is expressed as [(L1$ hit time)+(next lower hierarchy hit time)*(L1$ miss rate)].
20. In a computer system, cache memory and main memory are two intermediate memory hierarchies. Which of the following statements are not correct?
    (a) Both hierarchies are designed based on the same principle of typical applications: The Principle of Locality.
    (b) The locality in a cache design is based on spatial locality and temporal locality; while the locality in the main memory design is based on only spatial locality, excluding temporal locality.
    (c) For typical computer systems, the cache controller and virtual memory manager are both software to better lower their miss rates.
    (d) Upon cache misses, to better utilize available hardware resources, CPU usually switches to execute other ready and waiting jobs.

二、題組(20%)，共四個題組，每一題組內所有小題完全答對得 5 分，答錯任一小題或未作答得 0 分。題組 A 包含題號 21~23，題組 B 包含題號 24~26 ，題組 C 包含題號 27~29，題組 D 包含題號 30~33。

A. Consider the system with a 32-bit logical address space, a page size of 4K bytes, and physical memory of size 256M bytes. For all the questions below, assume each entry in the page table consists of 4 bytes.

21. Assuming that the single-level page-table scheme is used, for a logical address 0xabcdef12, the value used to index the page table is _____ (a) 0xf12, (b) 0xabc, (c) 0xabcde, (d) 0xabcd.
22. Assuming that the two-level page-table scheme is used and the outer page table has 256 entries, for a logical address 0xabcdef12, the values used to index the outer and inner page tables are __A__ and __B__ , respectively. (a) A=0xab, B=0xcdef12, (b) A =0xab , B =0xcde, (c) A=0xabcd, B=0xef12, (d) A=0xab, B=0xcd.
23. Assuming that the two-level page-table scheme is used and the outer page table has 256 entries, each process may need up to ____ bytes of memory space for the page table? (a) $2^{10} + 2^{22}$, (b) $2^{8} + 2^{12}$, (c) $2^{10} + 2^{14}$, (d) $2^{13} + 2^{25}$.

B. Consider the following monitor solution to the 5 dining philosophers problem where there are 5 chopsticks between 5 philosophers who sit on a round table. A philosopher needs to get both chopsticks on the right and left sides to eat. Suppose the index *i* is incremented counterclockwise. With the monitor defined below, each philosopher *i* invokes the operations pickup() and putdown() in the following sequence:

```
DiningPhilosophers.pickup(i);
    EAT
DiningPhilosophers.putdown(i);
```

```
monitor DiningPhilosophers
{
    enum { THINKING; HUNGRY, EATING) state [5] ;
    condition self [5];
    void pickup (int i) {
            state[i] = HUNGRY;
            test(i);
            if (state[i] != HUNGRY) self[i].wait();
}
    void putdown (int i) {
            state[i] = THINKING;
             test((i + 4) % 5);
             test((i + 1) % 5);
}
void test (int i) {
             if ((state[(i + 4) % 5] != EATING) &&
             (state[i] == HUNGRY) &&
             (state[(i + 1) % 5] != EATING) ) {
                    state[i] = EATING ;
                    self[i].signal() ;
             }
    }
        initialization_code() {
            for (int i = 0; i < 5; i++)
            state[i] = THINKING;
        }
}
```

24. One statement in the above program is wrong. Name which statement.
    (a) state[i] = HUNGRY,
    (b) if (state[i] != HUNGRY) self[i].wait()
    (c) state[i] = EATING
    (d) self.signal()
25. What is the purpose of the function test(i)?
    (a) Check whether both right and left chopsticks are available.
    (b) Check whether both neighbors are not EATING, and if both not, eat.
    (c) Check whether both right neighbors are HUNGRY, and if both not, eat.
    (d) Check whether either right neighbor or left neighbor is HUNGRY, and if yes, wake it up to eat.
26. What is the purpose of the function test((i+4)%5)?
    (a) Check whether my left neighbor is HUNGRY and its both neighbors are not EATING, and if both not EATING, wake it up to eat.
    (b) Check whether my right neighbor is HUNGRY and its both neighbors are not EATING, and if both not EATING, wake it up to eat.
    (c) Check whether my left neighbor is HUNGRY and its both neighbors are not HUNGRY, and if both not HUNGRY, wake it up to eat.
    (d) Check whether my right neighbor is HUNGRY and its both neighbors are not HUNGRY, and if both not HUNGRY, wake it up to eat.

C. Consider two one-bit ALU designs and a 32-bit ALU design for 32-bit MIPS CPU design.

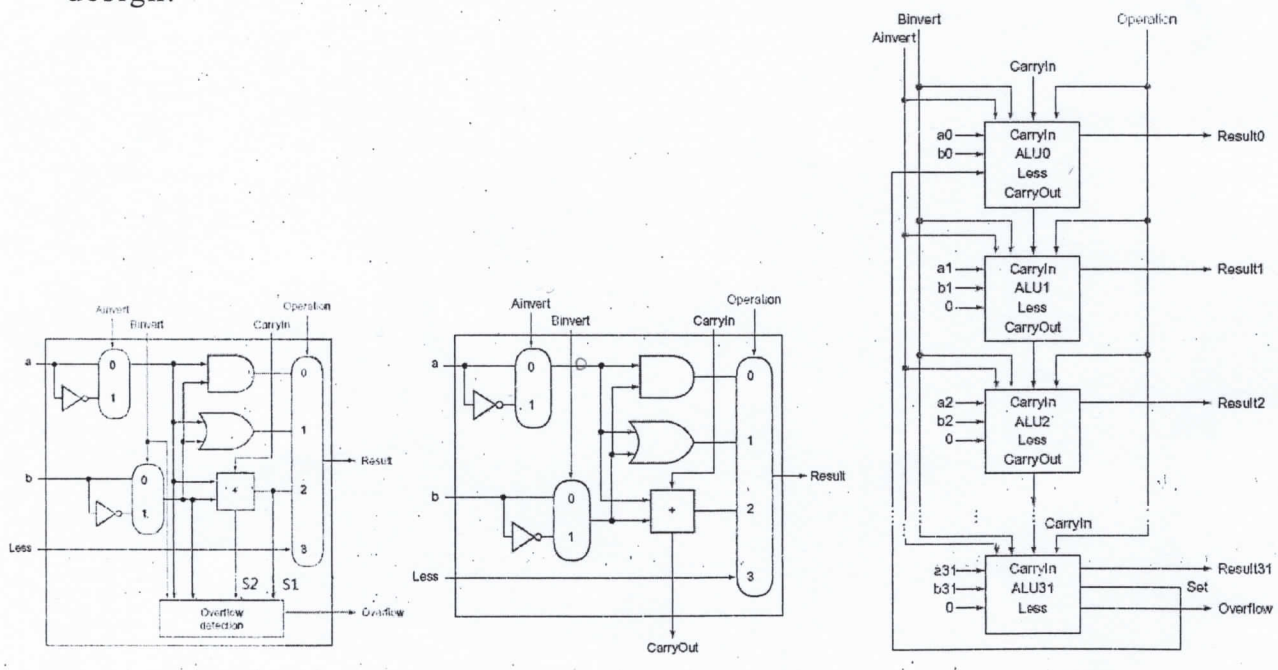

27. Which of the following statements is correct?
    (a) The left figure is the most significant bit (MSB) ALU and the *set* signal comes from *S*2.
    (b) The middle figure is the MSB ALU and the *set* signal comes from *CarryOut*.
    (c) The concept to realize *slt* is that a subtraction is performed first and then the *CarryOut* of the MSB ALU is the *set* signal that is connected to the least significant bit ALU.
    (d) The *set* signal comes from the *S*1 of the ALU in the left figure.

28. Which of the following statements is correct?
    (a) The control signal setting to realize a *slt* instruction is: *AInvert*=1; *BInvert*=1; Operation=10; *CarryIn*=1.
    (b) The control signal setting to realize a *slt* instruction is: *AInvert*=0; *BInvert*=1; Operation=11; *CarryIn*=1.
    (c) The control signal setting to realize a *sub* instruction is: *AInvert*=0; *BInvert*=1; Operation=00; *CarryIn*=1.
    (d) The control signal setting to realize a *sub* instruction is: *AInvert*=0; *BInvert*=1; Operation=10; *CarryIn*=0.
29. Which of the following statements is correct?
    (a) The control signal setting to realize a *bne* instruction is: *AInvert*=0; *BInvert*=1; Operation=10; *CarryIn*=1.
    (b) The control signal setting to realize a *nor* instruction is: *AInvert*=1; *BInvert*=1; Operation=01; *CarryIn*=0.
    (c) The control signal setting to realize a *lw* instruction is: *AInvert*=1; *BInvert*=1; Operation=10; *CarryIn*=1.
    (d) The control signal setting to realize a *sw* instruction is: *AInvert*=1; *BInvert*=1; Operation=00; *CarryIn*=0.

D. Given that (1) memory space is 32-bit byte addressable, (2) main memory size is 2^{mm} bytes, (3) page frame size is 2^{pf} bytes, (4) Cache size is 2^{cache} bytes, and (5) cache block size is 2^{cb} bytes:

30. The memory space available to a program is (a) 2^{32}, (b) 2^{31}, (c) 2^{mm}, (d) 2^{mm-1}, bytes
31. The number of virtual pages is (a) 2^{mm}, (b) 2^{32}-2^{mm}, (c) $2^{32-cache}$, (d) 2^{32-pf}.
32. The page size is (a) $2^{32}/2^{mm}$, (b) $2^{mm}/2^{pf}$, (c) 2^{pf}, (d) $2^{pf}/2^{cb}$.
33. The size of the cache block tag in bits is (choose the most proper one) (a) [<(32-cb)], (b) 32-cache, (c) (32-cache)~[<(32-cb)], (d)$2^{(32-cache)}$~$2^{[<(32-cb)]}$.