【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 2 分(答錯兩個選項倒扣 4 分，最多扣至該題 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 3-2=1 分)

1. Which one(s) of the following statements about modern operating systems are correct?
   (a) Microkernels can suffer from performance decreases due to overhead of user space to kernel space communication
   (b) To create and use shared memory, several system calls have to be invoked.
   (c) Direct Memory Access (DMA) can reduce the number of interrupts for data transmission.
   (d) None of the above statement is correct.

2. Which one(s) of the following statements about processes and threads are correct?
   (a) Threads created by a process do not share their program counters.
   (b) After `exec()` system call is applied, the process id is not changed.
   (c) When a process issues a blocking I/O request, it will be moved to the ready queue.
   (d) A process can be removed forcibly from the CPU (running state), as a result of an interrupt.

3. Assume that a system issues multiple producer and consumer processes for a shared buffer (size **n**=5) in the following pseudo codes with semaphore **empty**, **mutex** and **full**.

   A producer process:

```
do {
   ...
   //produce an item in nextp
   ... //Section PA
   wait(empty);
   ... //Section PB
   wait(mutex);
   ... //Section PC
   //add nextp to buffer
   ... //Section PD
   signal(mutex);
   ... //Section PE
   signal(full);
}while(TRUE);
```

   A consumer process:

```
do {
   wait(full);
   ... //Section CA
   wait(mutex);
   ... //Section CB
   //remove an item from buffer to nextc
   ... //Section CC
   signal(mutex);
```

```
        ... //Section CD
        signal(empty);
        ... //Section CE
        //consume the item in nextc
        ...
      }while(TRUE);
```

Which of the following statements are correct?

(a) The semaphore **empty** should be initialized to 0.

(b) Up to one consumer process can enter "Section CA" (between `wait(full)` and `wait(mutex)`) at a time.

(c) When there are only one producer process and one consumer process, the binary semaphore **mutex** is redundant.

(d) None of the above statement is correct.

4. Consider a system with four processes P_1~P_4 and three resource type A, B, C. Suppose at time T_0, the resource allocation state is:

	Allocation			MAX			Available		
	A	B	C	A	B	C	A	B	C
P_1	2	0	1	2	0	3	1	0	2
P_2	0	1	1	1	2	1			
P_3	1	1	0	1	1	3			
P_4	1	0	2	2	1	2			
P_5	1	1	0	1	2	0			

Which of the following statements are correct?

(a) $<P_4, P_1, P_2, P_3, P_5>$ is a safe sequence

(b) $<P_1, P_3, P_5, P_4, P_2>$ is a safe sequence

(c) Operating systems have to estimate the MAX value when a process enters the ready queue.

(d) None of the above statement is correct.

5. Given the following codes with pthread,

```
...//include header files

void func ( void *ptr );
sem_t mutex;
int m;

int main(){
    int i;
    int n;
    int p[3];
    pthread_t thrd[3];

    m = 0; n = 0;
    p[0] = 0; /* argument to threads */
    p[1] = 1;
    p[2] = 2;
    sem_init(&mutex, 0, 1); /* initialize mutex to 1 */

    for(i=0; i<3; i++ )
      pthread_create (&thrd[i], NULL, (void *) &func, (void *) &p[i]);
```

```
    for(i=0; i<3; i++)
       pthread_join(thrd[i], NULL);

    sem_destroy(&mutex); /* destroy semaphore */
    exit(0);
} /* main() */

void func ( void *ptr ){
    int a, n = 0;
    a = *((int *) ptr);
    sem_wait(&mutex);
    m++;
    printf("m = %d\n", m);
    sem_post(&mutex);

    n=a;
    printf("n= %d\n", n);

    pthread_exit(0);
} /* func() */
```

Which of the following statements about executing this program are correct?
(a) Race condition does not occur.
(b) The sum of all output numbers is less than 7.
(c) The executed program outputs the number "2" at most once.
(d) All performed "func" share the content of variable "m".

6. Which one(s) of the following statements about stack and heap are correct?
(a) Memory blocks on the stack can be freed at any time.
(b) Memory blocks on the heap can be freed at any time.
(c) Typically, each thread has its own stack.
(d) Typically, each thread has its own heap.

7. Given 4 hard disks with equal capacity, if we want to assemble a RAID (Redundant Array of Inexpensive Disks) to tolerate single disk failure, which of the following configurations would suffice?
(a) RAID 0.
(b) RAID 1.
(c) RAID 1+0.
(d) RAID 5.

8. Which of the following are true about the use of page table in virtual memory management?
(a) Each user process has its own virtual address space by default.
(b) The page table of a Linux process is managed by the C runtime library in the process (e.g., libc-x.xx.so)
(c) Use of shared memory can reduce the number of page table entries in a page table
(d) For the unused regions in the virtual address space, the space overhead of the corresponding page table entries can be negligible.

9. Assume that a program performs the following system calls to create file `f` under directory `d` and write the string "hello" to file `f`:

```
fd = create("d/f", 0666);
write(fd, "hello", 5);
```

Which of the following file system operations will be involved?
(a) Get a free i-node for `f`.
(b) Update the content of directory `d` to include the i-node mapping of `f`.
(c) Get free data block(s) for `f` and link the data blocks with the i-node of `f`.
(d) Write "hello" to the data block(s).

10. Assume that you use Google Chrome to browse Facebook.com via HTTPS (HyperText Transfer Protocol Secure) connections. Which of the following statements are correct? (assume there exist no security vulnerabilities)
(a) The ISP (Internet Service Provider) cannot tell which website you are browsing by looking at the network traffic.
(b) The ISP cannot know the contents of the Facebook pages you are browsing by looking at the network traffic.
(c) It is impossible to know the contents of the Facebook pages you are browsing by looking at the memory content of the Google Chrome process.
(d) It is impossible to know the contents of the Facebook pages you are browsing by looking at the disk files (including executables, data files, cached web pages, and etc.) of Google Chrome.

11. Which of the following statements are *correct*?
(a) Instructions *addi* and *lb* (load byte) both need to perform sign extension.
(b) The range of numbers represented by n-bit 2's complement is: $-(2^{n-1})+1 \sim 0 \sim (2^{n-1})$.
(c) Consider an addition of two 32-bit signed integers with a 32-bit ripple-carry adder, which consists of 32 1-bit full adders connected in series: if the carry-in and the carry-out of the most-significant-bit (MSB) full adder have different logic values, then there is an overflow due to the addition.
(d) The operation of adding one very large positive integer and one very large negative integer is likely to produce an overflow.

12. Which of the following 32-bit *hexadecimal* data will have identical storage sequence in the (byte-addressable) memory no matter the machine is big-endian or little-endian?
(a) 56788765_{hex}
(b) 56787856_{hex}
(c) 56666656_{hex}
(d) 56655665_{hex}
(e) 55665566_{hex}

13. Which of the following statements are *correct*?
(a) The number of pipeline stages affects latency, not throughput; thus pipelining

improves the performance of a processor by decreasing the latency of a job (i.e., an instruction) to be done.

(b) For a given program, its average cycles per instruction (CPI) is mainly affected by the instruction set architecture, but not affected by the programming language used.

(c) By reducing the clock frequency of a processor from 2.5 GHz to 2 GHz, and also reducing its supply voltage from 1.2 Volt to 0.9 Volt, the overall power consumption of this processor will be reduced by 55% theoretically.

(d) According to Amdahl's law, it is theoretically impossible to achieve a 2X speedup on executing a program which is 56% parallelizable, by using an 8-core multiprocessor.

14. Which of the following MIPS instructions are *not* R-type?
    (a) addi $t1, $t2, 3 # add-immediate
    (b) bne $t1, $t2, 3 # branch on not-equal
    (c) sll $t1, $t2, 3 # shift left logical
    (d) jr $ra # jump-return

15. Which of the following statements are *correct*?
    (a) A branch delay slot is defined as a sequence of instructions with neither embedded branch (except at the end) nor branch target (except at the beginning); a compiler identifies branch delay slots for code optimization.
    (b) An issue slot is defined as a group of instructions to be issued together in the context of multiple-issue processors.
    (c) $0.3125*2^{130}$ can be represented by the IEEE 754 single-precision floating-point standard format without any loss of accuracy.
    (d) Although IEEE 754 double-precision floating-point representation (denoted by *double*) is more precise than the single-precision counterpart (denoted by *float*), it is possible that *float*-to-*double* casting (conversion) causes loss of precision.

16. Below shows a MIPS datapath design with control and the control signal setting of the ALU control unit for R-type, lw/sw and branch-equal instructions. Which components in (a) and (b) are required during the operation of Branch-on-Equal instruction? The 4 output signals of ALU control unit are the 4 input control signals (*AInvert*, *BInvert* and *Operation*) for ALU cells (Fig. 1-2). Which terms in (c) and (d) should be involved in the optimal design of control signal *BInvert*?

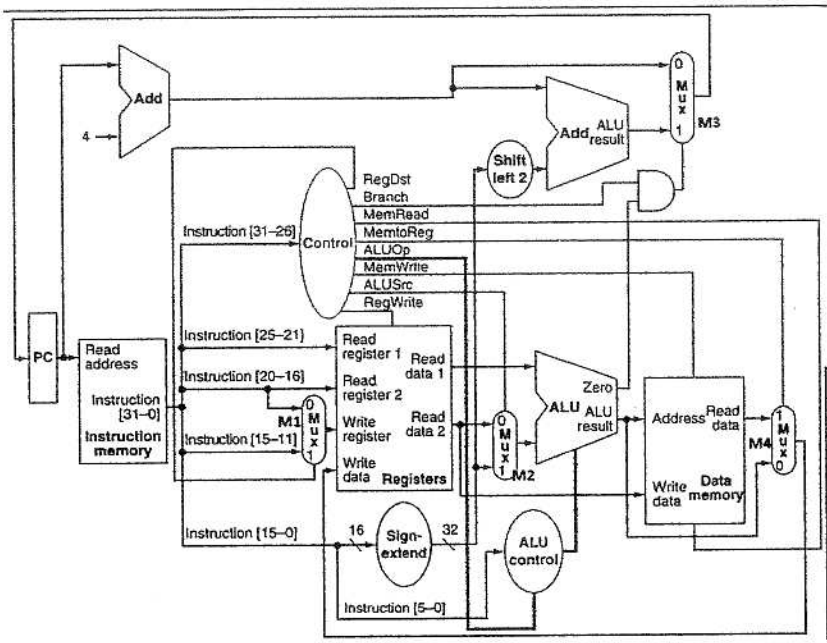


Fig. 1-1

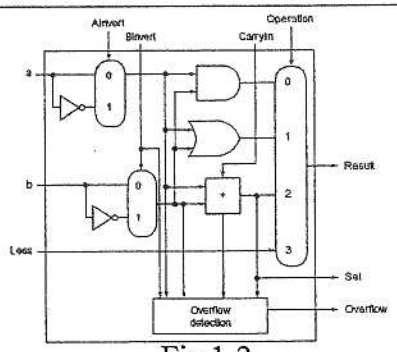


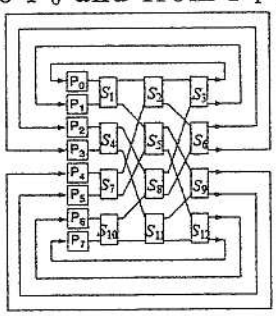
Fig. 1-2

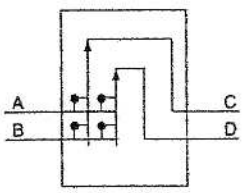
ALUOp		Funct field						Control output
ALUOp1	ALUOp0	F5	F4	F3	F2	F1	F0	Ainv, Binv, Operation
0	0	X	X	X	X	X	X	0, 0, 10 (lw/sw)
X	1	X	X	X	X	X	X	0, 1, 10 (beq)
1	X	X	X	0	0	0	0	0, 0, 10 (add)
1	X	X	X	0	0	1	0	0, 1, 10 (sub)
1	X	X	X	0	1	0	0	0, 0, 00 (AND)
1	X	X	X	0	1	0	1	0, 0, 01 (OR)
1	X	X	X	1	0	1	0	0, 1, 11 (slt)

Fig. 1-3

(a) The multiplexer M1.
(b) Signed-extend unit.
(c) ($ALUOp1$,$ALUOp0$,$F5$,$F4$,$F3$,$F2$,$F1$,$F0$)=(1,X,X,X,X,X,1,X).
(d) ($ALUOp1$,$ALUOp0$,$F5$,$F4$,$F3$,$F2$,$F1$,$F0$)=(X,0,X,X,0,X,X,X).

17. Consider the parallel processor and computing design, which of the following statements are **true**?
(a) As compared to conventional processors, vector processors demand less bandwidth for accessing instructions since vector instructions replace a series of repeated instructions on a set of vectorized data.
(b) As compared to conventional processors, the execution of two adjacent vector instructions with load-use data hazard between them can save much time for stalling because vector processors can avoid data hazard however conventional processors need to stall for every execution of two instructions with load-use data hazard.
(c) The cost to establish warehouse-scale computers (WSC) is very expensive and disks rather than processors are the main maintenance issue for WSC.
(d) The following figure shows an Omega network for multi-core interconnection network topology and the design of a switch box. We can achieve the data transfer from P_0 to P_6 and from P_1 to P_4 concurrently.





18. Consider the control unit and data dependency issues in MIPS pipelined CPU design.

Which of the following statements are **wrong**?

(a) The control signals are generated in the register-read stage and then each control signal is sent to the stage that needs the control signal through pipeline registers. Since the control signal *RegWrite* is used by the registers in the register-read stage, *RegWrite* can be sent to the registers directly and does not need to be sent to pipeline registers.

(b) For double data hazard issue, like the code sequence (add \$1,\$1,\$2; add \$1,\$1,\$3; add \$1,\$1,\$4), the following statement can detect one of double hazard cases: if (MEM/WB.RegWrite and (MEM/WB.RegisterRd ≠ 0) and not (EX/MEM.RegWrite and (EX/MEM.RegisterRd ≠ 0) and (EX/MEM.RegisterRd = ID/EX.RegisterRs)) and (MEM/WB.RegisterRd = ID/EX.RegisterRs))

(c) Consider the code sequence (*lw* \$2,12(\$1); *or* \$5, \$4, \$2; *and* \$6, \$4, \$2; *and* \$8, \$2, \$5) and assume the first instruction enters the pipeline CPU at the first clock cycle. The third instruction (*and*) enters the pipeline CPU at the third clock cycle, and will stay at the second stage (register-read) for two clock cycles since the data hazard between the first and the second instructions cannot be resolved by forwarding scheme and a bubble is required for this situation.

(d) Consider the dynamic multiple-issue processor. The reservation station is used to keep the required input data from functional units or commit unit for coming instructions. As the input data for one instruction is available in a reservation station, the instruction is loaded into the associated reservation station for execution.

19. About the memory design, which of the following statements are **true**?

(a) Temporal locality includes the access to induction variables.

(b) A direct-mapped cache has 32 blocks, each of which has 32 bytes. The addresses 0XB50 and 0XB54 are both mapped to block 26 in the cache.

(c) For a fixed-size cache, a design of a block with too large size tends to have a side-effect that a cache read loads unnecessary data into a block, causing degradation in cache performance. This situation usually happens in the matrix operations that perform on a subset of a matrix rather than on the entire matrix, such as on a subset of 256×256 elements in a matrix of 1024×1024 elements.

(d) For a CPU running at 4GHz with a 2-level cache, the base CPI is 1.0 and the average L1-cache miss rate per instruction is 2%. The average access time in the L2-cache and main memory are 3*ns* and 75*ns*, respectively, and the average cache miss rate is 0.4%, which means the average rate that needs to access main memory is 0.4%. Then the effective CPI is 2.44.

20. For the schemes to improve the throughput of MIPS pipelined CPU, which of the following statements are **true**?

(a) For speculation scheme, it is not necessary to speculate two instructions with a store preceding a load that reference the same memory location, where this kind of speculation allows to perform a load before a store, since sometimes we cannot determine if two locations referenced by a store and a load are the same due to the pointer aliasing problem.

(b) The next three sub-problems consider a static dual-issue pipelined CPU with one issue in ALU/branch instructions and another issue in store/load instructions. The use latency of the code sequence (*lw* \$3, 20(\$1); *add* \$5, \$4, \$3; *sw* \$3,

12($2)) is 1 and is the same as that of the same code sequence running in a single-issue pipelined CPU. There is no performance gain in this code sequence using static dual-issue pipelined CPU.

(c) Loop unrolling technique is applied to make two copies of the code sequence (Loop: *lw* $t0, 0($s1); *lw* $t1, 4($s1); *addu* $t0, $t0, $t1; *sw* $t0, 0($s1); *addi* $s1, $s1,–4; *bne* $s1, $zero, Loop) and the unrolling result is shown in the following table. *A*=0; *B*=12; *C*=4;

Label	ALU/branch	Load/Store	Cycle
Loop:	*addi* $s1, $s1, -8	*lw* $t0, *A*($s1)	1
		lw $t1, *B*($s1)	2
			3
	addu $t0, $t0, $t1	*lw* $t2, *C*($s1)	4 // rename t0 as t2
		sw $t0, *D*($s1)	5
		lw $t3, *E*($s1)	6 // rename t1 as t3
	addu $t0, $t2, $t3		7
	bne $s1, $zero, Loop	*sw* $t0, *F*($s1)	8

(d) Follow previous sub-problem (c), *D*=8; *E*=8; *F*=4;

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 21~24，題組 B 包含題號 25~27 ，題組 C 包含題號 28~30，題組 D 包含題號 31~33。

A. Consider the following set of processes performed on a single-core processor. These processes arrive at the following arrival time, with the length of the CPU burst given in the following table:

Process	Arrival Time	CPU Burst Time
P_1	0	6
P_2	2	8
P_3	3	5
P_4	5	3
P_5	7	1

Please calculate the waiting time W under the *preemptive shortest-remaining-time first* scheduling method. Assume $W = W_3*4^3 + W_2*4^2 + W_1*4^1 + W_0$ ($0 \leq W_0, W_1, W_2, W_3 < 4$). Please choose the correct values of W_3, W_2, W_1, and W_0.

21. W_3=? (a) 3, (b) 2, (c) 1, (d) 0
22. W_2=? (a) 3, (b) 2, (c) 1, (d) 0
23. W_1=? (a) 3, (b) 2, (c) 1, (d) 0
24. W_0=? (a) 3, (b) 2, (c) 1, (d) 0

B. Consider the following memory maps of a `cat` process running on a Linux system in the NCTU CS computing center:

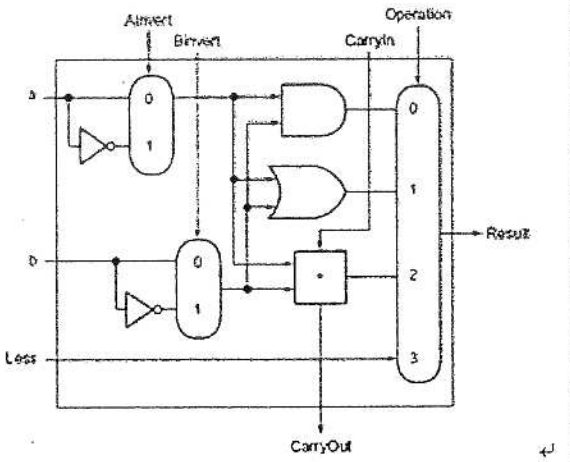
```
[wclin@linux1 ~]$ cat /proc/self/maps
00400000-0040b000 r-xp 00000000 fd:00 33634378                           /usr/bin/cat
0060b000-0060c000 r--p 0000b000 fd:00 33634378                           /usr/bin/cat
0060c000-0060d000 rw-p 0000c000 fd:00 33634378                           /usr/bin/cat
00d87000-00da8000 rw-p 00000000 00:00 0                                  [heap]
7f11d3b8c000-7f11da0b6000 r--p 00000000 fd:00 100737152                   /usr/lib/locale/locale-archive
7f11da0b6000-7f11da279000 r-xp 00000000 fd:00 67216364                    /usr/lib64/libc-2.17.so
7f11da279000-7f11da479000 ---p 001c3000 fd:00 67216364                    /usr/lib64/libc-2.17.so
7f11da479000-7f11da47d000 r--p 001c3000 fd:00 67216364                    /usr/lib64/libc-2.17.so
7f11da47d000-7f11da47f000 rw-p 001c7000 fd:00 67216364                    /usr/lib64/libc-2.17.so
7f11da47f000-7f11da484000 rw-p 00000000 00:00 0
7f11da484000-7f11da4a6000 r-xp 00000000 fd:00 67778326                    /usr/lib64/ld-2.17.so
7f11da690000-7f11da693000 rw-p 00000000 00:00 0
7f11da6a4000-7f11da6a5000 rw-p 00000000 00:00 0
7f11da6a5000-7f11da6a6000 r--p 00021000 fd:00 67778326                    /usr/lib64/ld-2.17.so
7f11da6a6000-7f11da6a7000 rw-p 00022000 fd:00 67778326                    /usr/lib64/ld-2.17.so
7f11da6a7000-7f11da6a8000 rw-p 00000000 00:00 0
7ffed373a000-7ffed375c000 rw-p 00000000 00:00 0                          [stack]
7ffed37b1000-7ffed37b3000 r-xp 00000000 00:00 0                          [vdso]
ffffffffff600000-ffffffffff601000 r-xp 00000000 00:00 0                   [vsyscall]
```

25. Which is the memory region that contains the executable code of the `cat` program?
    (a) 00400000-0040b000
    (b) 0060b000-0060c000
    (c) 0060c000-0060d000
    (d) None of the above

26. Which is the physical address range containing the executable code of the C library?
    (a) 7f11da0b6000-7f11da279000
    (b) 7f11da479000-7f11da47d000
    (c) 7f11da47d000-7f11da47f000
    (d) None of the above

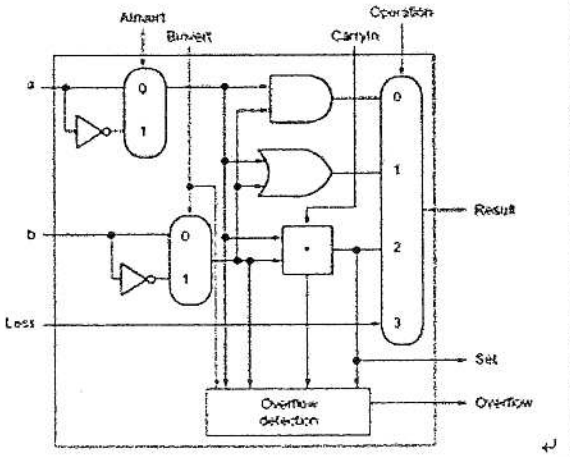
27. When memory space becomes tight, we may need to swap out memory pages to the disk. Which of the following statement about swapping out memory pages is correct?
    (a) We can skip swapping out memory pages that are backed by a memory-mapped file.
    (b) We can skip swapping out memory pages that are dirty.
    (c) We can skip swapping out anonymous pages (pages that are not backed by memory-mapped file).
    (d) The number of anonymous pages that can be swapped out is bounded by the size of the swap file (or the swap partition) on the disk.

C. Consider a 32-bit ALU, consisting of 32 1-bit ALUs connected in series (shown in the following) where:
NOT gate (inverter) has a propagation delay of 1 unit,
AND gate has a propagation delay of 1 unit,
OR gate has a propagation delay of 1 unit,
2-to-1 mux has a propagation delay of 2 unit,
4-to-1 mux has a propagation delay of 4 unit,
full adder (square) has a propagation delay of 3 units,
overflow detection (rectangle) has a propagation delay of 3 units.



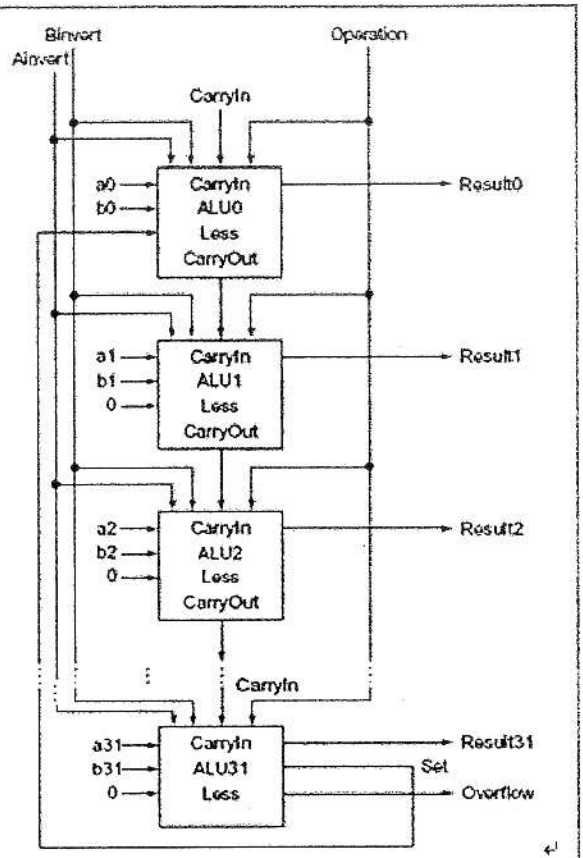
1-bit ALU

(Instance of **ALU0 ~ ALU30** on the right)



1-bit ALU (most significant bit)

(Instance of **ALU31** on the right)



32-bit ALU

28. What is the worst-case (i.e., maximum) cumulative propagation delay for a 32-bit bitwise NOR operation/instruction?
    (a) 7 units
    (b) 8 units
    (c) 32 units
    (d) 96 units

29. What is the worst-case (i.e., maximum) cumulative propagation delay for a 32-bit ADD operation/instruction??
    (a) 9 units
    (b) 10 units
    (c) 100 units
    (d) 102 units

30. What is the worst-case (i.e., maximum) cumulative propagation delay for an SLT (set

on less-than) operation/instruction??
(a) 99 units
(b) 100 units
(c) 102 units
(d) 103 units

D. Consider the (7, 4) Hamming code. Each code word contains seven bits, c_1, c_2, d_1, c_3, d_2, d_3, d_4, with satisfying the following three equations:
$c_1 \oplus d_1 \oplus d_2 \oplus d_4 = 0$; $c_2 \oplus d_1 \oplus d_3 \oplus d_4 = 0$; $c_3 \oplus d_2 \oplus d_3 \oplus d_4 = 0$.
Answer the first sub-problem for this (7, 4) Hamming code.
As we add an additional parity bit, c_4, in the end, the code can detect two errors. One additional equation holds for this need: $c_1 \oplus c_2 \oplus d_1 \oplus c_3 \oplus d_2 \oplus d_3 \oplus d_4 \oplus c_4 = 0$. Answer the second and the last sub-problems with the new code.

31. As receiving a 7-bit data (1 1 0 1 1 0 0),
(a) The received data is correct.
(b) The received data has one bit error in c_1.
(c) The received data has one bit error in c_2.
(d) The received data has one bit error in d_2.

32. As receiving a 8-bit data (1 0 1 1 0 1 0 1)
(a) The received data is correct.
(b) The received data has an error in bit d_1.
(c) The received data has an error in bit c_2.
(d) The received data has an error in bit c_4.

33. As receiving a 8-bit data (1 0 0 0 1 0 1 1)
(a) The received data has double errors.
(b) The received data is correct.
(c) The received data has an error in d_3.
(d) The received data has an error in c_4.