考試科目	計算機系統	814所別	資訊科學系碩士班	考試時間	3月17日 星期六 第二節

一、Computer Organization and Design

1. **[ADC]** The successive approximation converter is one of the most widely used types of ADC. Given the following simplified block diagram:
    *   Explain how the control logic works using a flowchart. (6%)
    *   Assume V_A=10.4V, use a simple four-bit converter with a step size of 1V to illustrate the conversion process. (4%)
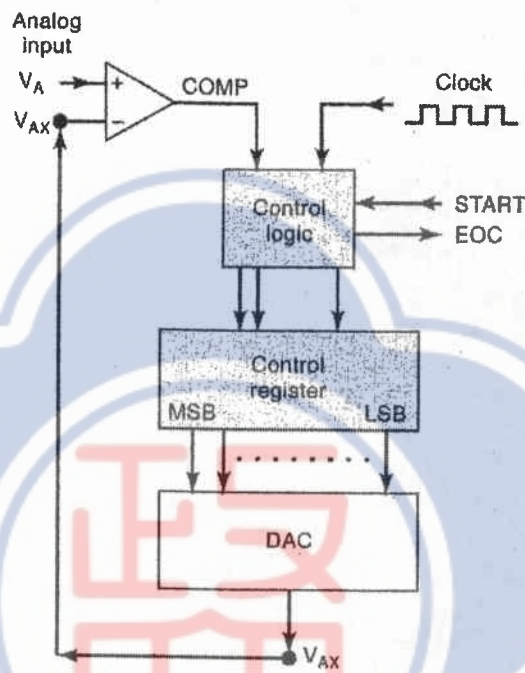
2. **[Instruction Set Architecture, Cache, Performance]**
    *   One of the differences between RISC architectures and CISC architecture is supposed to be the reduced types of instructions available. A student thinks it would be a good idea to simplify the instruction set even more to remove the special case instructions that take immediate operands such as "li", "addi", etc. Explain to him/her why this might not be such a good idea. (3%)
    *   Explain how a memory system that pages to secondary storage depends on locality of reference for efficient operation. (3%)
    *   Program A consists of 2000 consecutive add instructions, while program B consists of a loop that executes a single add instruction 2000 times. You run both programs on a certain machine and find that program B consistently executes faster. Explain. (4%)

備 考	試題隨卷繳交

命題委員： 　　　　　063　（簽章）96年3月5日

3. **[Adder]** A half-adder takes two input bits A and B to produce sum (S) and carry (C_{out}) outputs.

   - Use basic logic gates to construct the circuits for a half adder. (6%)
   - Use exactly two half-adders and one OR gate to construct a full adder. (4%)

4. **[Pipelining]** Refer to the following figure, if the time for an ALU operation can be shortened by 25%; (a) Will it affect the speedup obtained by pipelining? If yes, by how much? Otherwise, why? (5%) (b) What if the ALU now takes 25% more time? (5%)
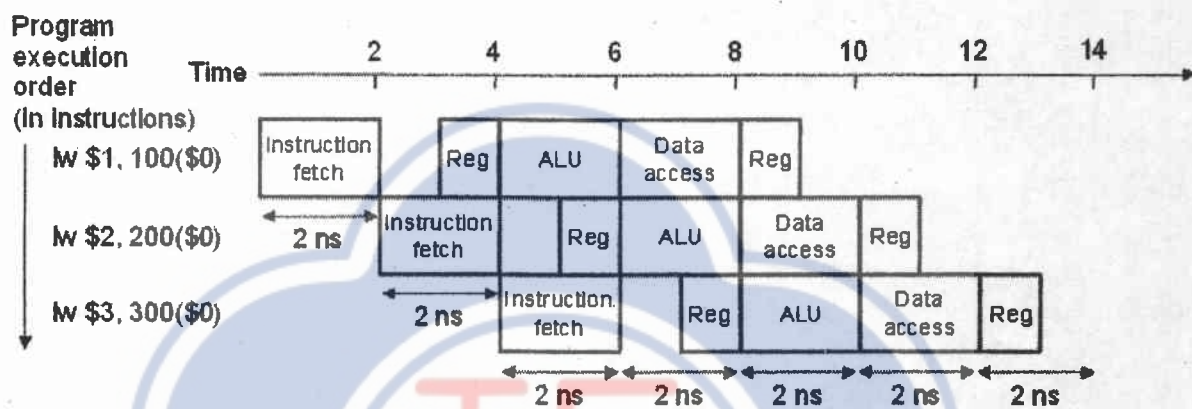
5. **[Memory]** Give a computer system that features:

   - a single processor
   - 32-bit virtual addresses
   - a cache of 2^{10} sets that are four-way set-associative and have 8-byte blocks
   - a main memory of 2^{26} bytes;
   - a page size of 2^{12} bytes.

   (1) Does this system cache virtual or physical addresses? (2%)

   (2) How many bytes of data from memory can the cache hold? (excluding tags) (2%)

   (3) In the cache, each block of data must have a tag associated with it. How many bits long are these tags? (2%)

   (4) How many comparators are needed to build this cache while allowing single cycle access? (2%)

   (5) At any one time, what is the greatest number of page-table entries that can have their valid bit set to 1? (2%)

備 考	試 題 隨 卷 繳 交

命題紙使用說明:1.試題將用原件印製,敬請使用黑色墨水正楷書寫或打字(紅色不能製版請勿使用)。
　　　　　　　2.書寫時請勿超出格外,以免印製不清。
　　　　　　　3.試題由郵寄遞者請以掛號寄出,以免遺失而示慎重。

二、Operating Systems

1. (5%) [Context switching] For each of the specified operations (a~e) below, indicate whether they must be performed (A) during a context switch between two processes, (B) during a context switch between two threads in the same process (address space), (C) during both kinds of context switches, and (D) during neither type of context switch. Assume that the threads are *kernel* managed (i.e., visible to the operating system).

(a) (1%) Modify the memory management unit to load the new process/threads's memory mappings.

(b) (1%) Save the registers of the process/thread being swapped out to its process/thread descriptor (control block).

(c) (1%) Save the top stack frame of the process/thread being swapped out to its process/thread descriptor (control block).

(d) (1%) Select the next process/thread to execute, based on priority information and other kernel internal state.

(e) (1%) Reset the "return register" (e.g., r0) to zero to indicate that the context switch occurred without any errors.

2. (9%) [CPU Scheduling] Consider a *multi-level feedback queue* in a single-CPU system. The first level (queue 0) is given a quantum of 8 ms, the second one a quantum of 16 ms, the third is scheduled FCFS. Assume six jobs (J1 ~ J6) arrive all at time zero with the following job times (in ms): 4, 7, 12, 20, 25 and 30. Show the *Gantt chart* for this system (3%) and compute the average waiting (3%) and turnaround time (3%).

3. (5%) [Deadlock] There are four conditions that are necessary for deadlock to occur: (A) mutual exclusion, (B) hold and wait, (C) no preemption, and (D) circular wait. With *deadlock prevention*, the system ensures that deadlock does not occur by preventing one of these conditions from holding. Match each of the following techniques with the one deadlock condition (A, B, C, or D) that it prevents.

(a) (1%) Impose a total ordering (or ranking) on how resources are acquired

(b) (1%) When a process requests a resource that is already held, force the process holding the resource to release it

(c) (1%) Only allow a process to request a resource when the process has none

(d) (1%) Allow all processes to access the resource simultaneously

(e) (1%) Require each process to grab all desired resources at once

4. (11%) [Miscellaneous topics]

(a) (4%) Assume that an inode-based file system has 2^{16} data blocks. For some reasons, the system designers have to trim down the size of an inode by reducing the number of direct pointers to 5 and keep only the single and double indirect pointers. The block size is also shrunk to 2Kbytes. What is the *maximum* file size in the trimmed system?

(b) (3%) Suppose each process spends 40% of its time in an I/O state. *How many* such processes are needed to bring the system CPU utilization to higher than 95%? Show your calculations. You will receive no credit if you only provide a number.

(c) (4%) Give *two specific factors* that we do NOT need to consider for the selection of disk scheduling algorithm when we decide to use RAM disk.

5. (10%) [Synchronization and Deadlock] A family has three children, each of whom owns one piece of ski equipment (Bobby has boots, Suzy has skis, and Joey has poles). A child needs all three pieces of equipment to ski, and that the parents have a complete set of equipment that they lend to their children. Each day, the parents randomly choose one child to let ski, and place the pieces of equipment that they need outside the condo's front door. Each day, each child attempts to ski. We model this scenario with by using *semaphores* to pass signals between *four threads*: PARENTS, BOBBY, SUZY, and JOEY.

```
semaphore skis = 0, poles = 0, boots = 0, done = 0;
PARENTS: for (;;) { //thread
   who = random(BOBBY,SUZY,JOEY);//select one randomly
   if (who == BOBBY)      { V(skis); V(poles); }
   else if (who == SUZY)  { V(poles); V(boots); }
   else /* who == JOEY */ { V(boots); V(skis); }
   P(done); // Wait for child to finish
}
BOBBY: //thread        SUZY: //thread        JOEY: //thread
for (;;){              for (;;){             for (;;){
  P(skis);              P(poles);             P(boots);
  P(poles);            P(boots);             P(skis);
  ski();               ski();                ski();
  V(done);}            V(done); }            V(done); }
```

(a) (5%) The program above is not deadlock free. Show a scenario where deadlock occurs using a step-by-step execution sequence table for the four threads.

PARENTS	BOBBY	SUZY	JOEY

time

(b) (5%) Show that two children can safely ski at the same (i.e., mutual exclusion is guaranteed.) using a step-by-step execution sequence table like the above one.

6. (10%) [Memory management] Consider the page replacement policies of OPT (Optimal algorithm), FIFO, and LRU. Which of the following statements (A~J) are true? Be careful to notice when the phrase states "better than or equal to" versus "strictly better than".

A. OPT always performs better than or equal to LRU.
B. OPT always performs strictly better than LRU.
C. LRU always performs better than or equal to FIFO.
D. LRU always performs strictly better than FIFO.
E. OPT with $n+1$ pages of physical memory always performs better than or equal to OPT with n pages.
F. OPT with $n+1$ pages of physical memory always performs strictly better than OPT with n pages.
G. FIFO with $n+1$ pages of physical memory always performs better than or equal to FIFO with n pages.
H. FIFO with $n+1$ pages of physical memory always performs strictly better than FIFO with n pages.
I. LRU with $n+1$ pages of physical memory always performs better than or equal to LRU with n pages.
J. LRU with $n+1$ pages of physical memory always performs strictly better than LRU with n pages.

備 考	試 題 隨 卷 繳 交

命題委員： 　　　　066　　（簽章） 96 年 3 月 1 日