

Points missed: \_\_\_\_\_

Student's Name: \_\_\_\_\_

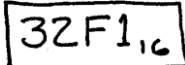
Total score: \_\_\_\_\_/100 points

East Tennessee State University  
Department of Computer and Information Sciences  
CSCI 4717 – Computer Architecture  
TEST 3 for Fall Semester, 2006

**Read this before starting!**

- The total possible score for this test is 100 points.
- This test is *closed book and closed notes*
- *Please turn off all cell phones & pagers during the test.*
- You may use one sheet of scrap paper that you will turn in with your test.
- **When possible, indicate final answers by drawing a box around them. This is to aid the grader (who might not be me!) Failure to do so might result in no credit for answer.**

**Example:**

A handwritten example of a boxed answer, showing the text "32F1,6" enclosed in a hand-drawn rectangular box.

- **1 point will be deducted** per answer for missing or incorrect units when required. **No** assumptions will be made for hexadecimal versus decimal, so you should always include the base in your answer.
- If you perform work on the back of a page in this test, indicate that you have done so in case the need arises for partial credit to be determined.

“Fine print”

Academic Misconduct:

Section 5.7 "Academic Misconduct" of the East Tennessee State University Faculty Handbook, October 21, 2005:

"Academic misconduct will be subject to disciplinary action. Any act of dishonesty in academic work constitutes academic misconduct. This includes plagiarism, the changing of falsifying of any academic documents or materials, cheating, and the giving or receiving of unauthorized aid in tests, examinations, or other assigned school work. Penalties for academic misconduct will vary with the seriousness of the offense and may include, but are not limited to: a grade of 'F' on the work in question, a grade of 'F' of the course, reprimand, probation, suspension, and expulsion. For a second academic offense the penalty is permanent expulsion."

- True or **false**: A page fault always results in a page being removed from memory to make room for a new page. (2 points)
- True** or false: In a system that uses paging where processes are divided into pages, all pages except for the last page of a process are equal sizes. (2 points)
- True** or false: In a system that uses paging, the process page size is equal to the memory page frame size. (2 points)
- True or **false**: In a system that uses paging, the pages of a process must be in order in memory. (2 points)
- True or **false**: In a system that uses paging, the pages of a process must be in contiguous (adjacent or unbroken) memory. (2 points)
- In a system that uses paging, \_\_\_\_\_ maintains a list of the free frames. (2 points)
  - the applications or processes
  - a hardware mechanism in main memory
  - a hardware mechanism in the cache
  - d.) the operating system**
  - none of the above
- What problem is typically caused by large virtual memory page sizes? (2 points)


Wasted space

- Using the page table shown to the right representing a specific process, calculate the physical address from the logical address  $079A_{16}$ . Assume a page size of  $2^8 = 256$ . (3 points)
  - $9A_{16}$
  - $79A_{16}$
  - $449A_{16}$
  - $FE9A_{16}$
  - $6C9A_{16}$
  - $9A9A_{16}$
  - g.) Not enough of the page table given to calculate**

Start of page table	page number
44	0
9A	1
2C	2
54	3
4A	4
FE	5
6C	6
??	7


Since the size of a page is 256 bytes, then the last 8 bits of a physical address ( $2^8 = 256$ ) is the offset into a page. This means that the last 8 bits of both the physical and logical addresses represent offset. Note that this automatically requires that the last byte of the physical address must also be 9A. For the logical address, the bits above the lower 8 bits represent the page number in the page table. The frame address contained in the page table should replace the page number to give us the physical address. Since there is no page 7 however, we do not have enough information.

**LOGICAL ADDRESS**

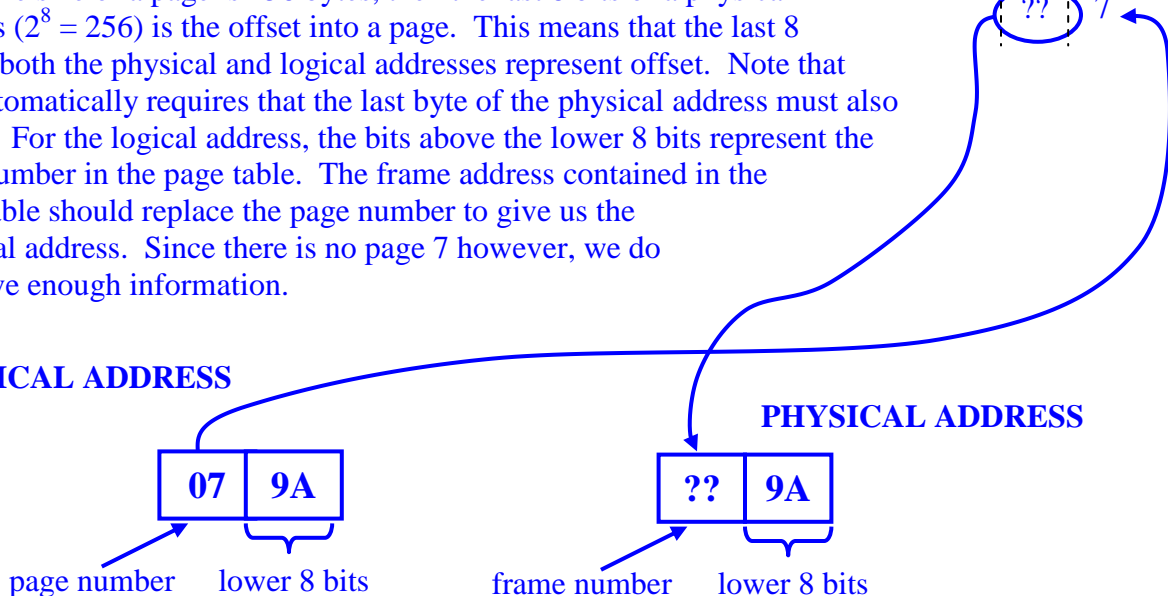


page number      lower 8 bits

**PHYSICAL ADDRESS**




frame number      lower 8 bits




9. Assume the page size of the previous problem changes to  $2^9 = 512$ , but everything else remains the same. What is the physical address from the logical address  $079A_{16}$ . (3 points)

Now with 9 bits used for the page offset, the page number is taken from the bits above the lower nine bits. (Note that  $079A_{16} = 0000\ 0111\ 1001\ 1010_2$ .)



This gives us a page number of  $0000011_2 = 3_{10}$ . Going to entry 3 of our page table gives us a frame number of  $54_{16} = 01010100_2$ . Substituting the frame number  $01010100_2$  for our page number should give us the physical address.



Put the frame number with the offset and get  $01010100110011010_2$  which equals  $A99A_{16}$ .

- a.)  $39A_{16}$       b.)  $79A_{16}$       c.)  $559A_{16}$       d.)  $549A_{16}$       e.)  $6C9A_{16}$   
f.)  $A99A_{16}$       g.) Not enough of the page table given to calculate      h.) None of the above

10. A(n) \_\_\_\_\_ saves the memory required for a page table by maintaining a table of only the pages represented by real memory and *not* all of the pages for a process. (2 points)

- a.) translation lookaside buffer      b.) virtual page table      c.) segmentation table  
d.) partition table      e.) inverted page table      f.) none of the above

11. The single instruction below is a three-operand instruction. In the table below, write three short programs that do exactly the same thing, one with two-operand instructions, one with one-operand instructions, and one with zero-operand instructions. If it fits the need of the instruction, feel free to use register names R1, R2, etc. (5 points)

MULT A, B, C ; A = B × C

<i>Two operand instructions</i>	<i>One operand instructions</i>	<i>Zero operand instructions</i>
MOV A, B MULT A, C	LOAD B MULT C STORE A	PUSH B PUSH C MULT POP A

12. Which of the following operations is performed by the zero-operand assembly language code shown to the right? (2 points)

- a.)  $Y = ((B \times C) + D) \times (A - E)$
- b.)  $Y = A \times (B \times (C + D) - E)$
- c.)  $Y = (A \times (B + C)) \times (D - E)$
- d.)  $Y = ((A \times B) + C) \times (D - E)$
- e.)  $Y = (A + (B \times C) - D) \times E$
- f.) None of the above

PUSH A  
 PUSH B  
 MULT  
 PUSH C  
 ADD  
 PUSH D  
 PUSH E  
 SUB  
 MULT  
 POP Y

The first operation is to put A and B on the top of the stack. This means that when the MULT operation is performed, it multiplies A and B and puts  $A \times B$  on the top of the stack. This automatically eliminates answers a, b, c, and e just because of the fact that  $A \times B$  is not a “first operation” of any of them. Let’s now see if answer d is still viable.

Pushing C onto the stack then doing an ADD adds  $A \times B$  to C giving us  $(A \times B) + C$ . Answer d is still in the running.

Next, D and E are pushed giving us the top three items on the stack as  $(A \times B) + C$ , D, and then E. The following SUB operation subtracts E from D then places the result  $D - E$  on the stack right above  $(A \times B) + C$ .

The MULT operation multiplies  $D - E$  by  $(A \times B) + C$  giving us  $((A \times B) + C) \times (D - E)$ . The popped value is then stored in Y. Therefore, the answer is d.

13. From the list below of stages of a 6-stage pipeline, identify the stage at which the outcome of a conditional branch has been determined. Circle one. (2 points)

- a.) Fetch instruction (FI)
- b.) Decode instruction (DI)
- c.) Calculate operands (CO)
- d.) Fetch operands (FO)
- e.) Execute Instruction (EI)
- f.) Write Operand (WO)

14. From the list below of stages of a 6-stage pipeline, identify all of the stage(s) that may vary considerably in duration depending on the instruction or operand. Circle all that apply. (2 points)

- a.) Fetch instruction (FI)
- b.) Decode instruction (DI)
- c.) Calculate operands (CO)
- d.) Fetch operands (FO)
- e.) Execute Instruction (EI)
- f.) Write Operand (WO)

15. From the list below of stages of a 6-stage pipeline, identify all of the stages that are *always* used in the execution of all instructions. Circle all that apply. (2 points)

- a.) Fetch instruction (FI)
- b.) Decode instruction (DI)
- c.) Calculate operands (CO)
- d.) Fetch operands (FO)
- e.) Execute Instruction (EI)
- f.) Write Operand (WO)

16. From the list below of stages of a 6-stage pipeline, identify the stage after which the processor check for interrupts? Circle one. (2 points)

- a.) Fetch instruction (FI)
- b.) Decode instruction (DI)
- c.) Calculate operands (CO)
- d.) Fetch operands (FO)
- e.) Execute Instruction (EI)
- f.) Write Operand (WO)

17. An instruction prefetch architecture is essentially a \_\_\_\_\_ - stage pipeline. (2 points)  
 a.) 1      **(b.) 2**      c.) 3      d.) 4      e.) 5      f.) 6      g.) 7
18. In an ideal implementation, what is the speed up of a processor with a  $k$ -stage pipeline over a non-pipelined processor if the duration of each stage is  $\tau$ ? Don't consider pipeline flushes or delays incurred between stages. (2 points)  
 a.)  $k \cdot \tau$     b.)  $k-2$     c.)  $k-1$     **(d.)  $k$**     e.)  $(k+1) \cdot \tau$     f.)  $k+\tau-1$     g.) none of the above
19. How many bits are required for each conditional branch in a branch history table in order to remember the past 2 branch outcomes for a specific instruction? (2 points)  
**(a.) 2**      b.) 3      c.) 4      d.) 5      e.) 6      f.) 7      g.) 8
20. List two of the three causes discussed in class of disrupting a pipeline. (3 points)

There were actually many more than three discussed in class. They include:

- Resource conflict
- Data dependency
- Conditional branch
- Interrupts
- Varied instruction time or time added to fetch multiple operands

*For problems 21, 22, and 23, consider the following section of code.*

```
for (i=0; i<2; i++)
{
  for (j=0; j<5; j++)
  {
    for (k=0; k<4; k++)
      sum += array_val[i, j, k];
  }
}
```

21. Once compiled, how many conditional jumps would be contained in the machine code resulting from the above section of code, i.e., static occurrence? (2 points)
22. After fully executing the above section of code, how many conditional jumps would the CPU have encountered, i.e., dynamic occurrence? (2 points)

There would be one conditional jump at the end of each for-loop: **3 conditional jumps**.

First of all, let us figure out for each for-loop how many times their associated conditional jump would be encountered.

- $k$ -loop: one loop for each  $k=0, k=1, k=2,$  and  $k=3$ . That makes four times the  $k$ -loop's conditional branch is encountered.
- $j$ -loop: one loop for each  $j=0, j=1, j=2, j=3,$  and  $j=4$ . That makes five times the  $j$ -loop's conditional branch is encountered.
- $i$ -loop: one loop for each  $i=0$  and  $i=1$ . That makes two times the  $k$ -loop's conditional branch is encountered.

Since the  $j$ -loop is executed once for every time the  $i$ -loop is executed, then the  $j$ -loop loops  $2 \times 5$  times or 10 times. Since the  $k$ -loop is executed once for every time the  $j$ -loop is executed, then the  $k$ -loop loops  $10 \times 4$  times or 40 times. Therefore, the number of times a conditional jump is encountered is the number of times a conditional jump is encountered for each of the loops.

$$\begin{aligned}\text{Total encounters with conditional jump} &= \text{times for } i\text{-loop} + \text{times for } j\text{-loop} + \text{times for } k\text{-loop} \\ &= 2 + 10 + 40 \\ &= 52 \text{ conditional jumps}\end{aligned}$$

23. Using the static branch prediction algorithm “branch always,” how many of the conditional jumps calculated in the previous problem would have been predicted *incorrectly*? (2 points)

“Branch always” predicts *wrong* once at the last execution of a loop. That happens once for the  $i$ -loop, twice for the  $j$ -loop, and ten times for the  $k$ -loop. This means there will be 13 incorrect predictions.

24. There are 3 types of static branch prediction methods: predict always taken, predict never taken, and predict based on opcode (fill in the blank) (2 points)

25. For the six architectural characteristics listed below, identify whether the statement more closely identifies a CISC architecture or a RISC architecture. (6 points)

CISC    RISC

- |                                     |                                     |  |
|-------------------------------------|-------------------------------------|--|
| <input checked="" type="checkbox"/> | <input type="checkbox"/>            | Assembly language is closer to high-level language                       |
| <input type="checkbox"/>            | <input checked="" type="checkbox"/> | Tends to have fewer stages in pipeline                                   |
| <input type="checkbox"/>            | <input checked="" type="checkbox"/> | Tends to have fewer addressing modes                                     |
| <input checked="" type="checkbox"/> | <input type="checkbox"/>            | Allows for indirect addressing   |
| <input checked="" type="checkbox"/> | <input type="checkbox"/>            | Registers tend to be more specialized than general purpose               |
| <input checked="" type="checkbox"/> | <input type="checkbox"/>            | Allows the operands of an arithmetic instruction to be memory references |

26. What is the absolute minimum number of registers required to execute the code below. (2 points)

- a.) 2      b.) 3      **c.) 4**      d.) 5      e.) 6      f.) 7      g.) 8

```
int done = 0;
int user_input << cin;
while (!done)
{
    int calc = 0;
    switch (user_input)
    {
        case 0:
            for (int i=0; i<5; i++) calc = (calc << 1) + i;
            break;
        case 1:
            for (int j=0; j<10; j++) calc = (calc << 1) + j;
            break;
        default:
            calc = 7;
            done = 1;
            break;
    }
    for (int k=0; k<calc; k++) cout << ".";
}
```

27. **True** or false: The purpose behind the delayed branch is to avoid flushing the pipeline. (2 points)

28. True or **false**: The purpose behind the delayed load is to avoid flushing the pipeline. (2 points)

29. After which line in the code below does a NOP need to be inserted to provide a **delayed load** to avoid a problem with a data dependency? (2 points)

- a.) L1      **b.) L2**      c.) L3      d.) L4      e.) L5      f.) L6

```
L1:    mov    al,[1000h]    ;Load al w/value from memory
L2:    mov    ah,[1001h]    ;Load ah w/value from memory
L3:    add    ah,al        ;ah += al
L4:    dec    al          ;al--
L5:    bne    L3          ;Branch if al is not equal to 0
L6:    mov    [1002h],ah    ;Store ah to memory
```

30. Identify the write-read, write-write, and read-write dependencies in the instruction sequence below by entering each line pair with a dependency in the correct column of the table to the right. For example, if L1 and L4 had a write-write dependency (which they don't), you would enter L1-L4 in the column labeled "write-write". (4 points)

L1: R3 = R2 + R5  
 L2: R3 = R3 + 16  
 L3: R3 = R3 + R5  
 L4: R5 = R1 - R2  
 L5: R1 = R3 + 32

write-read	write-write	read-write
L1 - L2*	L1 - L2*	L1 - L4*
L1 - L3	L1 - L3	L2 - L3*
L2 - L3*	L2 - L3*	L4 - L5*
L1 - L5		L3 - L4*
L2 - L5		
L3 - L5*		

\* - indicates required pairs. The other pairs are implied by the required pairs.

For problems 31, 32, and 33, use the figure to the right which represents the execution of 6 instructions on an "in-order-issue/in-order-completion" machine.

Assume that there is only one data dependency in the sequence, a true data dependency (write-read) between I1 and I4, i.e., I4 depends on the completion of I1 before it can execute. As for resource dependencies, assume that any instructions that share a column in the execute stage require the same resource.

Decode		Execute			Write		Cycle
I1	I2						1
	I2	I1					2
I3	I4	I1					3
I3	I4	I2					4
I5	I6		I4	I3	I1	I2	5
I5	I6			I3			6
		I5	I6		I3	I4	7
					I5	I6	8

31. In an out-of-order-issue/out-of-order-completion architecture, what is the earliest cycle I4 could enter the execute stage? (2 points)

a.) 1    b.) 2    c.) 3    **d.) 4**    e.) 5    f.) 6    f.) 7    g.) 8

I4 depends on I1 for its data, but the only reason it is waiting for I2 in the above diagram is that for in-order-issue/in-order-completion, all instructions from the previous decode cycle must be in the write pipes before the next instructions can enter the execute pipes. Therefore, I4 can be executed with I2 (cycle 4) in an out-of-order completion architecture. Since I3 doesn't have any dependencies, it could enter the pipe as soon as I3.

32. In an out-of-order-issue/out-of-order-completion architecture, what is the earliest cycle I5 could enter the execute stage? (2 points)

a.) 1    b.) 2    c.) 3    d.) 4    **e.) 5**    f.) 6    f.) 7    g.) 8

In an out-order-issue/out-of-order completion architecture, I5 (without any dependencies) requires two things: it must have entered through the decode operation and it's execution resource must be available, i.e., all previous instructions that used I5's execution pipe (I1 and I2) must be finished with it. With four instructions in front of it, I5 isn't brought into the decode pipe until cycle 3. After that, it must wait for I1 and I2 to be done with its execution pipe. That happens at cycle 5. Therefore, I5 can enter the pipe at cycle 5.



33. In an out-of-order-issue/out-of-order-completion architecture, what is the earliest cycle I3 could enter the execute stage? (2 points)
- a.) 1    b.) 2     c.) 3    d.) 4    e.) 5    f.) 6    g.) 7    h.) 8

Just like I5, in an out-of-order-issue/out-of-order completion architecture, I3 must have finished with the decode operation and its execution resource must be available. No other instructions are shown to be sharing the same execution pipe with I3, so all we have to do is get I3 from the decode pipe. With two instructions in front of it, I3 isn't brought into the decode pipe until cycle 2. Therefore, I3 can enter the pipe at cycle 3, i.e., right after it finishes being decoded.

34. Multiple Instruction/Multiple Data Stream architectures are referred to as tightly coupled when: (2 points)

- a.) they communicate over a local area network (LAN)  
b.) the network they are on is dedicated to the system, i.e., there is no other traffic  
c.) they share a single hard drive or RAID for storage across a network  
 d.) they communicate through a shared memory  
e.) none of the above

35. Which Multiple Instruction/Multiple Data Stream architecture is better suited to processes that require a high level of interaction, loosely coupled or tightly coupled? (2 points)

- a.) loosely coupled     b.) tightly coupled

36. A cluster is an example of: (2 points)

- a.) a single instruction/single data stream architecture  
b.) a single instruction/multiple data stream architecture  
c.) a multiple instruction/single data stream architecture  
d.) a tightly coupled multiple instruction/multiple data stream architecture  
 e.) a loosely coupled multiple instruction/multiple data stream architecture

37. Assume a multiprocessor system uses the MESI protocol. If the current state of a line in processor A's cache is *modified*, do valid copies of the data exist in other caches? (2 points)

- a.) yes     b.) no    c.) cannot tell

38. Assume a multiprocessor system uses the MESI protocol. If the current state of a line in processor A's cache is *exclusive*, do valid copies of the data exist in other caches? (2 points)

- a.) yes     b.) no    c.) cannot tell

39. Assume a multiprocessor system uses the MESI protocol. If the current state of a line in processor A's cache is exclusive and processor B loads the same line into its cache, what is the new state of that line in processor A's cache? (2 points)

- a.) modified    b.) exclusive     c.) shared    d.) invalid    e.) cannot be determined

40. Which SMP bus configuration is simplest due to its architecture being closest to the single-processor architecture? (2 points)

- a.) time-shared bus                      b.) multiport memory                      c.) central controller

41. Which SMP bus configuration is most reliable due to the fact that there is no central point of failure? (2 points)

- a.) time-shared bus                      b.) multiport memory                      c.) central controller

42. The snoopy protocol is more suited to the \_\_\_\_\_ interconnection method for symmetric multiprocessors. (2 points)

- a.) time-shared bus                      b.) multiport memory                      c.) central controller

43. For the four characteristics listed below, identify whether the statement more closely identifies an SMP system or a cluster. (4 points)

SMP    Cluster

- |                                     |                                     |  |
|-------------------------------------|-------------------------------------|--|
| <input checked="" type="checkbox"/> | <input type="checkbox"/>            | Operation is closer to that of a single processor system |
| <input type="checkbox"/>            | <input checked="" type="checkbox"/> | Easier to upgrade incrementally                          |
| <input checked="" type="checkbox"/> | <input type="checkbox"/>            | Lower overall power consumption                          |
| <input checked="" type="checkbox"/> | <input type="checkbox"/>            | Older, more established track record                     |