

Initials \_\_\_\_\_

Signature: \_\_\_\_\_

I.D.Number: \_\_\_\_\_

Printed Name: \_\_\_\_\_

**McGill University**  
**Department of Electrical and Computer Engineering**

**Course ECSE-322 -- Computer Engineering**

**MidTerm Test**

**EXAMPLE**

**PLEASE NOTE CAREFULLY:**

Sign this paper, fill in your student ID number and print your name at the top of the first page and on the mark sense sheets - if you do not do this the exam may not be marked. **Initial each page of the exam paper in case the sheets should get separated.** Make sure the signed paper in its entirety is handed in at the end of the examination.

**INSTRUCTIONS:**

*This exam consists of two parts; part 1 consists of a set of 16 multiple choice questions, part 2 consists of a set of 2 questions with short answers - YOU SHOULD CHOOSE ONE OF THE TWO QUESTIONS IN PART 2. The answers to the questions in Part 1 (the first 16 questions) should be entered on the computer marked sheets, the answers to Part 2 should be written on this question paper in the space provided. DO NOT USE ANY OTHER EXAM BOOKS FOR ANSWERS TO BE MARKED.*

*Unless otherwise stated on the exam paper, the scoring method for this examination will assign 1 mark for each correct answer, 0 marks for a blank or wrong answer. All multiple choice questions will be weighted equally in scoring.*

*This is a closed book exam. However, the candidates are allowed to bring in one sheet of letter size paper which may have handwritten notes on both sides.*

*Please sign this paper at the top of this page, write your name legibly, and read the important notice above.*

**ENTER THE NUMBER OF THE SHORT ANSWER QUESTION IN PART 2 TO BE MARKED HERE:**

Initials \_\_\_\_\_

**TEST QUESTIONS**

*PART 1 - This part consists of 16 multiple choice questions. Indicate the correct answer to each question on the computer readable sheets provided.*

1. It is necessary to provide a mechanism for enabling and disabling interrupts because:
- a) Programmed Input and Output cannot be implemented if interrupts are enabled.
  - b) The device may not be able to respond to a CPU request fast enough.
  - c) The CPU can be paralysed by spurious interrupts being received.
  - d) Interrupts are not needed for synchronous data transfers.
  - e) All priority mechanisms require interrupts to be disabled.

Answer c

2. One key difference between the set of instructions used in an interrupt service routine (ISR) and those used for programmed input and output (PIO) is:
- a) In PIO but not in an ISR, the address of the device status register must be known
  - b) In an ISR, there is no need to wait for the device to be ready
  - c) The ISR does not use the device data register
  - d) There is no vectoring capability for PIO
  - e) There is no difference between the two sets of instructions

Answer b

3. Which of the following pieces of information is NOT explicitly represented in a composite video signal for a television picture?
- a) The dot clock
  - b) The horizontal retrace period
  - c) The vertical blanking period
  - d) The vertical synch pulse
  - e) All of the above are explicitly represented.

Answer a

4. Which of the following statements is TRUE?
- a) I/O is impossible without interrupts
  - b) A computer needs as many interrupt levels as there are I/O devices
  - c) Priority interrupts cause the device which has waited the longest to be handled next

Initials \_\_\_\_\_

- d) Interrupt service routines may themselves be interrupted
- e) Polled I/O is preferred over interrupt-driven I/O on a multi-user system.

Answer d

5. All computer graphics displays function with a dot clock. The dot clock frequency is dependent on the resolution of the display. In general (for a PC display) the dot clock frequency is of the order of:
- a) 1 - 1000 Hz
  - b) 1000 - 100000 Hz
  - c) 100000 - 1000000 Hz
  - d) 1000000 - 100000000 Hz
  - e) Greater than 100000000 Hz

Answer d

6. A Byte-Addressable machine is one in which
- a) The Least Significant Bit is addressed first
  - b) The Most Significant Bit is addressed first
  - c) Words and Long Words are not supported
  - d) Memory is a one-dimensional structure
  - e) Each byte can be individually retrieved.

Answer e

7. The XON/XOFF protocol is used
- a) Only with a full RS-232 communications link
  - b) To avoid losing data through overflowing the buffers
  - c) To minimize the number of bits that need to be sent
  - d) To achieve higher throughput than can be achieved with RS-232
  - e) To remove the wait loop that can occur in programmed input and output processes.

Answer b

8. Which of the following input devices does NOT return a relative position?
- a) A track ball
  - b) A mouse
  - c) A keyboard using arrow keys
  - d) The pen used with the display in classroom 2000
  - e) A joy stick

Initials \_\_\_\_\_

Answer d

9. The frame buffer in a graphics system is used to:
- Prevent the graphics display from overflowing and thus losing data
  - Store graphical data before it is presented to the display device
  - Provide an animation capability for a graphics display
  - Store the horizontal and vertical synchronization information
  - Allow a graphics system to be controlled over a link using an RS-232 protocol.

Answer b

10. Which of the following is NOT a requirement for an interface designed to use an RS-232 protocol to communicate with a peripheral device?
- Accept or send serial representations of characters
  - Recognize appropriate bus select signals
  - Provide a clock signal to synchronize CPU/interface communications
  - Place a character on, or remove a character from, the bus data lines
  - Provide status/control information to both devices connected to it.

Answer c

11. In a polled interrupt system, the concept of *device priority* is implemented by
- The address of the device interface on the bus
  - The use of an interrupt acknowledge line
  - The address of the interrupt service routine in memory
  - The sequence in which devices are checked for a set interrupt
  - The bus line used to signal the interrupt.

Answer d

12. Which of the following statements is TRUE when describing an array?
- An array is only addressed by a single integer
  - All the data items in an array are of the same type
  - An array is always a simple list of items
  - Adding two arrays is an  $O(N^2)$  process
  - Searching for data in an array is an  $O(N)$  process

Answer b

Initials \_\_\_\_\_

13. An unordered data set is to be sorted by either an *Exchange Sort* or a *Quick Sort*. If the data set contains  $N$  items, what is the ratio of the time taken, on average, between the Quick and Exchange Sorts?
- a)  $N$
  - b)  $\log_2(N)$
  - c)  $N/\log_2(N)$
  - d)  $\log_2(N)/N$
  - e) They will take the same time.

Answer d

14. The worst case performance of a hashing storage system (i.e. when a collision occurs on each access) is
- a)  $O(1)$
  - b)  $O(\log_2 N)$
  - c)  $O(N)$
  - d)  $O(N \log_2 N)$
  - e)  $O(N^2)$

Answer e

15. An ordered data set is to be sorted by either an *Exchange Sort* or a *Quick Sort*. If the data set contains  $N$  items, what is the ratio of the time taken, on average, between the Quick and Exchange Sorts? Note that the Quick Sort is to use an algorithm that requires no extra storage space beyond the original array?
- a)  $N$
  - b)  $\log_2(N)$
  - c)  $N/\log_2(N)$
  - d)  $\log_2(N)/N$
  - e) They will take the same time.

Answer a

16. Which of the following statements is TRUE when describing a vectored multi-dimensional array?
- a) The time to evaluate the mapping polynomial is increased by using vectoring
  - b) The use of base addresses makes the storage of sparse data easier
  - c) Computing an array index only requires addition operations
  - d) Computing an array index only requires multiplication operations
  - e) The storage space needed for the base address arrays is approximately equal to that needed for the original array.

Initials \_\_\_\_\_

Answer c

**PART 2 of this exam starts on the next page. Please turn over the page and begin Part 2 NOW!**

Initials \_\_\_\_\_

PART 2 - This part consists of 2 short answer questions. YOU SHOULD ANSWER ONLY ONE OF THE QUESTIONS - IF YOU ANSWER BOTH, ONLY ONE WILL BE MARKED. The question is worth 10 marks (Each question consists of 5 sections, each worth 2 marks). Write your answers to the question in the space provided on this exam paper. Do not use extra paper - any answers written outside of the allowed space will not be marked.

17. (a) To implement an effective interrupt capability for a device, two bits are needed in the Device Status Register. These two bits describe four possible states. One of the states can never occur. Describe each of these states and indicate which one is (theoretically) impossible. (2 marks)

State 1: \_\_\_\_\_

State 2: \_\_\_\_\_

State 3: \_\_\_\_\_


State 4: \_\_\_\_\_

Answers:

*Interrupt not enabled, Interrupt not set*  
*Interrupt enabled, Interrupt not set*  
*Interrupt not enabled, Interrupt set – illegal*  
*Interrupt enabled, Interrupt set*

Initials \_\_\_\_\_

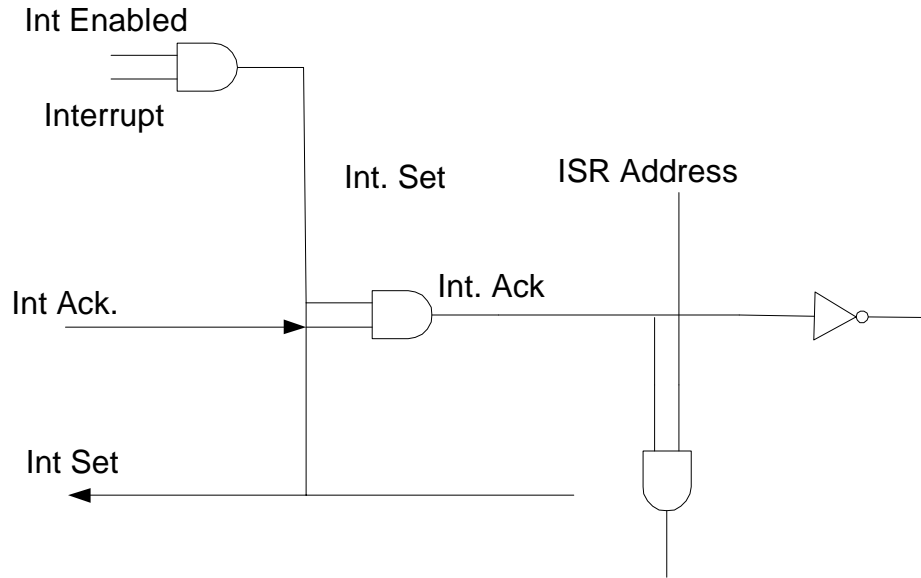
- (b) Draw a flow diagram that shows the sequence of events that take place when a vectored interrupt occurs starting from the time at which the interrupt is enabled and ending when the interrupt service routine begins execution. (2 marks)





Initials \_\_\_\_\_

(c) Sketch the circuitry needed on an interface to implement a vectored interrupt protocol in particular showing how the bits described in (a) are used in the protocol. Show only the logic related to the control lines. (2 marks)



Initials \_\_\_\_\_

- (d) A printer is to be configured to serve 10 computers connected on a network. The printer is capable of printing 4 pages per minute but each computer can output a page of text to the printer in 5 seconds, providing that the printer is available (i.e. not receiving data from another computer). A complete page of text comprises 48 lines of text (on average) and each line contains 90 characters (on average) and each computer only prints complete pages.

On average, each computer prints 20 pages an hour. If a collision is detected, (i.e. two computers try to print at exactly the same time) one computer backs off and sends its print job immediately the other has finished provided there is room in the buffer of the printer for the page. The printer implements an XON/XOFF protocol.

Assuming worst case conditions, calculate the size of the memory buffer needed in the printer so that the waiting time for any computer to access the printer is minimized. (2 marks)

*4 pages per minute = 15 seconds per page  
1 page of text = 48 \* 90 = 4320 characters*

*200 pages per hour  
200 pages arrive in 1000 seconds = 200 \* 4320 characters in 1000 seconds  
= 864000 characters in 1000 seconds*

*In 1000 seconds printer can print 1000/15 pages  
= 66.7 pages = 66.7 \* 4320 = 288000 characters*

*Buffer size needed = 864000 – 288000 = 576000 bytes (approx) = 562 kB (approx)*

- (e) If the total pages to be printed are uniformly distributed over an hour, would a buffer still be needed on the printer? If so, how big should it be? (2 marks)

*Yes – because one page arrives in 5 seconds but takes 15 seconds to print. Thus a buffer of 2/3 of a page is needed:*

$$0.667 * 4320 = 2882 \text{ bytes}$$

Initials \_\_\_\_\_

18. (a) In dealing with data in a computer system, or anywhere else, sorting the data is always a slower process than searching for data. Explain why this is so. (2 marks)

*Because each item to be sorted has to be checked and then sorted. The actual storage of an item is equivalent to a search operation. i.e. for each item to be stored, find the storage location and put it in the place.*

- (b) Given the statement in part (a) of this question, under what condition is it worth sorting the data before searching for particular items? Describe the condition and give an algebraic proof of the condition you have stated. (2 marks)

*If data in a particular data set is to be accessed a large number of times. The number would have to be greater than  $N$  where  $N$  is the number of data items..*

Proof:

*Accessing an item in an unordered set takes  $O(N)$  operations, i.e.  $= N$   
If  $S$  items are to be found, the total time is  $S.N$*

*If the data set is ordered,  $S$  items can be found in  $S \log_2 N$  operations  
The number of operations taken to sort the data – using a Quick sort to get the best average performance – is  $N \log_2 N$ .*

*Thus the total time with an ordered data set to find  $S$  items is  
 $(N+S) \log_2 N$*

*To make the set worth sorting, we have*

$$\begin{aligned} S.N &= (N+S) \log_2 N \\ S(N - \log_2 N) &= N \log_2 N \\ S &= N \log_2 N / (N - \log_2 N) \end{aligned}$$

Initials \_\_\_\_\_

- (c) Give two operations that are usually defined for any Abstract Data Type and describe what they do. (2 marks)

(1) \_\_\_\_\_ **Store** – places a data item in a data structure (allow words like Add, Insert...)

(2) \_\_\_\_\_ **Retrieve** – accesses an item within a data structure (Allow words like Delete, Remove..)

- (d) A particular application requires the storage of an array with 1000 columns and 1000 rows where less than 10,000 of the array values are non-zero. Describe an implementation for the array that would be more space efficient than a standard two-dimensional array implementation requiring one million positions. In addition, if the array consists of integers, what is the ratio of memory used by your storage scheme to that used by the standard implementation? (Assume that the standard implementation is not vectored). (2 marks)

*To minimize the storage, 2 sets of pointers are required – one to indicate the row of each non-zero and one to indicate the column. In this case, only the non zeros will be stored.*

*Thus there are 10,000 integer values plus 10,000 integer row pointers plus 10,000 column pointers*

*Total = 30,000 integers.*

*The standard implementation would use 1000000 integers*

*The ratio is  $30000/1000000 = 3/100$  or 3%*

***An Alternate solution:***

*Instead of the second array of pointers for the columns, just keep a pointer for the start of each column. This would then only need 1000 integers*

*The total number of integers now needed would be 10,000 values plus 10,000 row pointers plus 1000 column offsets*

*Total = 21000*

*The ratio is  $21000/1000000 = 2.1\%$*

Initials \_\_\_\_\_

- (e) If the array in part (d) is vectored when stored with a standard storage scheme what would the ratio become? Also draw the flow chart for adding two sparse arrays stored in the format you have specified in part (d). Use the next page of this booklet for drawing your flowchart. (2 marks)

*If vectoring is used, then base address arrays are held for the offset of each column so 1000 extra integers would be needed in the standard scheme:*

$$\text{Total used} = 1001000$$

*Ratio:*

$$30000/1001000 = 2.98\% \text{ or } 21000/1001000 = 2.09\%$$

-

Initials \_\_\_\_\_

Flowchart:

*This is not a properly drawn flowchart...*

Assume we have 2 input arrays A and B, each has row and column arrays (AR, BR, AC, BC) A is of length Y, B is of length X

The output is in C with CR and CC.

Set L = 1                    'L is the pointer in C  
 Set D = 1                    'D is the pointer in A  
 Set E = 1                    'E is the pointer in B

```

For I = 1 To M
For J = 1 To N
  IF D <> Y THEN
    IF AR(D) = I and AC(D) = J
    THEN
      IF E <> X THEN
        IF BR(E) = I and BC(D) = J
        THEN
          CR(L) = AR(D)
          CC(L) = AC(D)
          C(L) = A(D) + B(E)
          D = D+1
          E = E+1
          L = L+1
        ELSE
          CR(L) = AR(D)
          CC(L) = AC(D)
          C(L) = A(D)
          D = D+1
          L = L+1
        END IF
      ELSE
        IF BR(E) = I and BC(D) = J
        THEN
          CR(L) = BR(D)
          CC(L) = BC(D)
          C(L) = B(E)
          E = E+1
          L = L+1
        END IF
      END IF
    END IF
  NEXT N

```

Initials \_\_\_\_\_

*NEXT M*

Initials \_\_\_\_\_

