



VERSION A

COMPUTER ENGINEERING

ECSE-322A

6 December 2005, 0900 - 1200

Examiner: *Professor D.A.Lowther*

Co-Examiner: *Professor T.Arbel*

Signature:

Signature:

STUDENT NAME AND SIGNATURE

MCGILL I.D. NUMBER

INSTRUCTIONS:

- This is a **CLOSED BOOK** examination. Only one (1) hand-written (not photo-copied) double sided **CRIB SHEET** on letter size (8.5 by 11 inch) paper permitted.
- This is a **MULTIPLE CHOICE** examination. It will be subject to the Harp Hogan Program. This is a program that will be run by the Computing Centre to check for the possibility of cheating.
- **MARK YOUR ANSWERS TO THE MULTIPLE CHOICE QUESTIONS ON THE ENCLOSED COMPUTER SHEET USING PENCIL ONLY.** Each question is worth 1 mark, a wrong answer will receive 0.
- Space is provided on the exam in which to answer the short answer questions (46 to 49). **DO NOT USE ANY OTHER EXAM BOOKS FOR ANSWERS TO BE MARKED.**
- You are permitted **REGULAR AND TRANSLATION** dictionaries.
- **FACULTY STANDARD CALCULATOR** permitted **ONLY**.
- *Student answers in the multiple choice section will be subject to analysis by the McGill Exam Security Computer Monitoring Program in order to identify unusually similar answer patterns.*
- This examination consists of two parts; part 1 consists of **45** multiple choice questions, part 2 consists of a set of **4** short answer questions **OF WHICH YOU SHOULD ATTEMPT ONLY 3**. The examination consists of a total of pages including this cover page. Please ensure that you have a complete examination paper **BEFORE** starting.
- **Please sign this paper at the top of this page, write your name legibly, and read the instructions above. Initial each page of the exam paper in case the sheets should get separated.**

THIS EXAMINATION PAPER MUST BE RETURNED.

LIST QUESTIONS ATTEMPTED IN PART 2 HERE:

| | |
|--|--|
| | |
| | |
| | |

Initials _____

TEST QUESTIONS

PART 1 - This part consists of 45 multiple choice questions. Indicate the correct answer to each question on the computer readable sheets provided.

1. In a particular sorting process the first iteration places 1 data item in the correct position; the second places a further 2; by the fourth pass 15 items are correctly placed. The *time complexity* of this process is

- (a) $O(\text{constant})$
- (b) $O(\log_2 N)$
- (c) $O(N)$
- (d) $O(N \log_2 N)$
- (e) $O(N^2)$

Answer (d).

2. In a tree data structure a *leaf of the tree* is identified by

- (a) it is the only node containing data
- (b) it does not have two children
- (c) all the children are nulls
- (d) some of the children are nulls
- (e) none of the above

Answer (c)

3. An *inorder traversal* of a binary tree structure visits the nodes in the sequence

- (a) root, left, right
- (b) root, right, left
- (c) right root, left
- (d) left, root, right
- (e) left, right, root

Answer (a).

4. In a hashed data structure of size M and a data set of size N ($N < m$) the maximum number of collisions that could occur is

- (a) $O(M)$
- (b) $O(M^2)$
- (c) $O(MN)$
- (d) $O(N)$

Initials _____

(e) $O(N^2)$.

Answer (d)

5. *Vectoring* is often used in multi-dimensional arrays because

- (a) It provides a way to save memory
- (b) It reduces the time needed to access an array element
- (c) It removes the need to perform an address calculation
- (d) It simplifies the calculation of a vector-vector product
- (e) It is never used.

Answer (b)

6. A *queue* is often referred to as a

- (a) Last in first out structure
- (b) First in last out structure
- (c) First in first out structure
- (d) Last in most important out structure
- (e) Single in multiple out structure.

Answer (b)

7. The effect of limiting the number of bits used to represent the *exponent* of a floating point number is to

- (a) limit the resolution with which numbers can be represented
- (b) limit the dynamic range over which numbers can be represented
- (c) remove the ability to represent negative exponents
- (d) remove the ability to represent negative number
- (e) allow the use of hidden bit normalization.

Answer (b)

8. The effect of limiting the number of bits used to represent the *mantissa* of a floating point number is to

- (a) limit the resolution with which numbers can be represented
- (b) limit the dynamic range over which numbers can be represented
- (c) remove the ability to represent negative exponents
- (d) remove the ability to represent negative number
- (e) allow the use of hidden bit normalization.

answer (a)

Initials _____

9. Which of the following does NOT need to be represented by a *dynamic data structure*?

- (a) a library
- (b) a book
- (c) a bus line up
- (d) a pile of dishes
- (e) a list of the students in ECE at McGill.

Answer (b)

10. The *Read Only Memory* in the printer discussed in class is used to

- (a) Store characters to be printed as they are received
- (b) Convert characters to ASCII codes
- (c) Convert ASCII Codes to control signals for the printer
- (d) Remove the need for hardwired logic in the printer
- (e) Handle the serial line communications protocol.

Answer (c)

11. The *Address Decoder* on an interface is used to

- (a) convert the address on the bus to a location in memory
- (b) store the address in memory for a DMA transfer
- (c) provide a clock signal for incrementing the byte count in a DMA transfer
- (d) provide a load signal for the Device Data Register
- (e) provide control information to the device.

Answer (d)

12. An *Interrupt Service Routine* is BEST described as

- (a) An additional control line on the bus to allow an interface to signal to the CPU
- (b) A piece of software stored at a specific location in the memory of a processor
- (c) A subroutine accessed at the end of a “fetch and execute” cycle in the CPU
- (d) A subroutine which controls the movement of data to and from an interface and is initiated by a signal on the bus.
- (e) A routine in the operating system designed to move data to and from an interface.

Answer (d)

13. The *XON/XOFF* protocol was developed to

- (a) Provide synchronization between two ends of a serial line communication system

Initials _____

- (b) Prevent buffer overflow happening in a peripheral device
- (c) Maximize the speed at which a graphics device can operate
- (d) Reduce the overhead in serial line data transmission
- (e) Minimize the number of connections needed for an RS-232 communication system.

Answer (b)

14. The *serial line character display* discussed in class is faster than a bit mapped raster display because

- (a) It only displays characters
- (b) It does not need to use a raster display system
- (c) It uses hardwired digital logic
- (d) It is not directly connected to the computer bus
- (e) It uses a dedicated CPU.

Answer (c)

15. Which of the following devices provides an *absolute* position?

- (a) Glide pad
- (b) Track ball
- (c) Joystick
- (d) Pen
- (e) They all do

Answer (d)

16. The main reason for implementing *DMA* in a computer system is

- (a) To free off the CPU to do other work
- (b) To simplify the writing of interrupt service routines
- (c) To remove the need for a buffer on the device interface
- (d) To reduce the number of logic circuits needed in the interface.
- (e) To speed up the transfer of data between an interface and the main memory.

Answer (e)

17. Which of the following signals is NOT embedded in a *Composite Video Signal*?

- (a) Clock
- (b) Horizontal synchronization
- (c) Vertical synchronization
- (d) Vertical blanking period

Initials _____

- (e) Image (video) data.

Answer (a)

18. In *Cycle Stealing DMA*

- (a) The device takes the bus only when it has data to transfer
- (b) The device takes and holds the bus during the entire DMA process
- (c) The CPU only grants the bus to the device when it does not need it
- (d) The BUS GRANT line is never set
- (e) The CPU has to enable BUS REQUEST before each DMA transfer.

Answer (a)

19. Which of the following is NOT a characteristic of a *magnetic tape*?

- (a) It is removable
- (b) It is non-volatile
- (c) It provides direct access to data
- (d) The cost/byte is extremely low
- (e) It can provide a large amount of secondary storage.

Answer (c)

20. A *floppy disk* rotates slower than a *hard disk* because

- (a) It is likely to “stretch” due to the centrifugal forces at high speed and thus the data would not be readable
- (b) If it is rotated faster, the data rate would be too high for the interface
- (c) DMA cannot be implemented for its interface
- (d) The friction between the head and the disk would destroy both
- (e) Too much power would be needed to rotate it faster.

Answer (d)

21. Which of the following pieces of information is NOT needed for find data on a *floppy disk*?

- (a) Platter
- (b) Surface
- (c) Track
- (d) Sector
- (e) All are needed.

Answer(a)

Initials _____

22. The *maximum throughput* of a bus is

- (a) dependent on the number of interfaces connected to it
- (b) dependent on the address space of the computer
- (c) decreases as more devices are added to it
- (d) related to the bus protocol implemented
- (e) related to the data transfer rate of the fastest device on the bus.

Answer (d)

23. In a *semi-synchronous* protocol

- (a) there is no clock signal
- (b) devices can respond on any clock edge
- (c) a device may take several clock cycles to respond
- (d) a device must respond within one clock cycle
- (e) a devices responds to changes in control signals

Answer (c)

24. Which of the following control signals is NOT needed to implement an *asynchronous protocol*?

- (a) clock
- (b) address strobe
- (c) data strobe
- (d) read
- (e) write

Answer (a)

25. A *Non-interlocked protocol* is one in which

- (a) Signals can be read on any edge of the clock signal
- (b) Devices on the bus check all changes in state of the control lines
- (c) Devices only check and respond to activation of control lines
- (d) Devices do not wait for a strobe (synchronization) signal before reading the data or address lines
- (e) Devices do not have to wait for the Bus Grant signal before placing data on the bus.

Answer (c)

26. The MAIN role of an *Arbiter* in allocating the bus is to

Initials _____

- (a) Allow only one device to access the bus at a time
- (b) Allow only one device to control the bus at a time
- (c) Ensure that the highest priority device always gets the bus
- (d) Provide a prioritized but fair access to the bus
- (e) Ensure that a device in control of the bus cannot be interrupted.

Answer (d)

27. Which of the following control signals is NOT used in the *bus arbitration process*?

- (a) Bus Request
- (b) Bus Grant
- (c) Bus Busy
- (d) Bus Release
- (e) Bus Acknowledge

Answer (e)

28. In the event that the bus is being used and a higher priority device has requested it, which one of the following signals would be activated?

- (a) Bus Request
- (b) Bus Grant
- (c) Bus Busy
- (d) Bus Release
- (e) Bus Acknowledge

Answer (a)

29. In an *Asynchronous Protocol* performing a **read** operation, when would the Slave Synchronization signal be set?

- (a) As soon as the data is stable
- (b) As soon as the address is stable
- (c) As soon as the master synchronization signal is set
- (d) As soon as the read signal is set
- (e) It would not be set since it is not needed in this operation.

Answer (a)

30. Which of the following best describes a *process*?

- (a) A user program
- (b) A routine in the operating system

Initials _____

- (c) A device driver
- (d) A piece of software that performs an action and can be controlled
- (e) A piece of software that accesses a resource.

Answer (d)

31. In the context of an *operating system*, which of the following is the best description of *concurrency*?

- (a) Allowing several hardware devices to operate simultaneously
- (b) Allowing several processes to be active simultaneously
- (c) Allowing several processors to access memory simultaneously
- (d) Allowing one device to request a resource while another is using it
- (e) Allowing several users to access the processor simultaneously.

Answer (b)

32. When a *process* has requested a *resource* and the resource is not available, the process is said to be

- (a) running
- (b) ready
- (c) sleeping
- (d) blocked
- (e) completed

Answer (d)

33. A *privileged instruction* is one which

- (a) Accesses a non-shareable resource
- (b) Can only be executed in kernel mode
- (c) Can only be executed by processes having a high enough priority
- (d) Can be executed when the system is deadlocked
- (e) Is used during process swapping.

Answer (b)

34. A *semaphore* is

- (a) A signal used by a resource to indicate that it is busy
- (b) A software interrupt used to request service
- (c) A non-negative integer used for control
- (d) A variable whose value is incremented by a wait operation
- (e) A small piece of software used to synchronize processes.

Initials _____

Answer (c)

35. Which of the following is NOT a condition for *deadlock*?

- (a) All the resources are non-shareable
- (b) All processes need one more resource
- (c) All processes hold onto resources while waiting for new ones
- (d) Resources cannot be relinquished when they are being used
- (e) Each process holds resources currently being requested by the next process in the chain.

Answer (b)

36. The *volatile environment* of a process is

- (a) The main memory used by the process when executing
- (b) The devices, other than disks and tapes, that it uses when executing
- (c) The set of instructions being executed by a process
- (d) A subset of modifiable shared resources used by the process, e.g. registers and program counter
- (e) The set of non-shareable resources used by the process and protected by the operating system.

Answer (d)

37. The *directory* on a device is

- (a) The area where files are stored on the device
- (b) The area containing the addresses of files on the device
- (c) The area on the device containing information about the device formatting
- (d) The area on the device used for decoding file addresses
- (e) The area on the device containing boot information for the operating system.

Answer (b)

38. The *head to disk spacing* in a CD-ROM is of the range of

- (a) 1 to 9 micro-inches
- (b) 10 to 99 micro-inches
- (c) 100 to 999 micro-inches
- (d) 1000 to 9999 micro-inches
- (e) 10,000 to 99,999 micro-inches

Answer (e)

Initials _____

39. The *physical law* that states that a current flowing in a conductor will create a magnetic field is known as

- (a) Ampere's Law
- (b) Faraday's Law
- (c) Ohm's Law
- (d) The Kerr Effect
- (e) Moore's Law

Answer (a)

40. *Burst feeding* is a system by which

- (a) Graphics data is loaded into a frame buffer in the vertical blanking period
- (b) Many bytes of data are transmitted to the buffer on a device through a single interrupt
- (c) A device interface moves data directly to or from memory
- (d) The need for a buffer on a device is removed
- (e) Several bytes of data can be sent to several devices simultaneously.

Answer (b)

41. The most common (and important) *disadvantage* of a *linked list* structure for storing *files* on disk is

- (a) That a file can be spread out over the entire disk leading to long access times
- (b) That a large number of unusable spaces are created on the disk as files are edited
- (c) That the directory **MUST** be stored in random access memory
- (d) That a file can be completely lost if a single link is destroyed
- (e) That files must be interleaved on the disk.

Answer (a)

42. *Tri-state buffers* are used to place address and data signals on the bus because

- (a) They have a smaller logic delay than normal buffers
- (b) They allow all devices to read data from the bus simultaneously
- (c) They allow write access to the bus to be controlled by the bus arbiter
- (d) They simplify the design of cycle stealing direct memory access
- (e) They can be used to auto-increment the address and word count.

Answer (c)

Initials _____

43. When a computer is first turned on, it executes a *bootstrap* to begin loading the operating system. The first part of the bootstrap is found
- (a) In the random access memory
 - (b) In the first sector of the first track of a floppy disk
 - (c) In the first sector of the first track of a hard disk
 - (d) In a read only memory
 - (e) It is built into the processor.

Answer (d)

44. *Sector Interleaving* is a technique by which
- (a) The number of heads needed on a disk drive is minimized
 - (b) More information can be stored on a given disk
 - (c) The fracturing of a disk is minimized
 - (d) The interface/disk drive interaction is optimized
 - (e) DMA transfers are simplified.

Answer (d)

45. The *address space* of a computer is determined by
- (a) The amount of RAM installed
 - (b) The word width of the processor
 - (c) The number of address lines
 - (d) The number of data lines on the bus
 - (e) The operating system

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 4 short answer questions. YOU SHOULD ANSWER ONLY 3 (THREE) OF THE QUESTIONS - IF YOU ANSWER 4, ONLY 3 WILL BE MARKED. Each question is worth 12 marks (Each question consists of 4 sections). 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.

46. (a) Give brief descriptions of the following terms (each is worth 1 mark):

(i) Polled interrupt system:

When the CPU receives an Interrupt Request, it jumps to a global interrupt service routine. Within the routine, each device is checked in turn to determine which one has set an interrupt. The first device found in the list with an interrupt request bit in the status register set is serviced.

(ii) Device Data Register:

A single register (memory location) on the device interface which is used for passing data between the interface and the bus. It connects to the bus data lines.

(iii) Cursor:

An imaginary pointer which indicates the current editing/drawing position in a graphical display frame buffer.

Initials _____

(b) A computer is to be connected to a printer via a serial line consisting of three wires – one of which is a common signal ground.

- (i) Describe (either by words or a diagram) how reliable transmission of ASCII characters over this link could be achieved assuming that both ends of the system have the same clock speed. (2 marks)

The solution is to describe the use of Start/Stop bits added to the character. Just saying Start/Stop will only generate half the marks. The solution should be stated in words:

Before transmitting a character, the link is held in an agreed idle state (state 1). One clock period before the character arrives, the line is dropped to state 0. The character is then transmitted. At the end of the ASCII character (i.e. 8 bits) the line is returned to the idle state for a minimum of 1 or 2 clock periods.

- (ii) If the printer is equipped with a memory buffer, describe a process which might be used to avoid buffer overflow. Give this process as a set of steps. (1 mark)

When the printer is ready to receive characters, i.e. the buffer is almost empty, it sends a message to the computer interface. The message is a single character (XON)

The interface receives the character and enables data transmission to the printer

Data is sent as a continuous stream to the printer

Initials _____

When the printer buffer reaches about 80% full, the printer sends a transmit off character (XOFF) to the interface

The interface then disables data transmission to the printer.

(c) A printer buffer is designed for a printer which can print c characters per second. A page of text consists, on average, of d characters and the connection between the interface and the printer can run at a maximum speed of t bytes per second. Characters are transmitted on a serial link using ASCII coding and come in bursts, i.e. a page is transmitted at the maximum rate of the link and then the transmission will pause until the next page is ready..

- (i) What is the maximum number of pages per hour that the printer can expect to receive from the computer? (1 mark)

In one hour, transmitting continuously, the link can send $3600t$ characters.

Thus the number of pages is $3600t/d$

- (ii) What is the minimum buffer size that will be needed (if t is greater than c) and what is the limit on the number of pages to be printed per hour, p , if this buffer is not to overflow? (1 mark)

*The total number of bytes being sent from the computer is $p*d$ per hour but the characters come in bursts at t characters/second. Each page takes d/t seconds to transmit*

*Thus, for a single page, a buffer of size $b = d - (d/t)*c = d(1 - c/t)$ is needed.*

This buffer will be emptied in d/c seconds. Thus if p is less than $3600c/d$ the single page buffer is sufficient.

- (iii) Sketch a graph showing the buffer size as a function of the number of pages per hour assuming that the only limit to the number of pages to be printed per

Initials _____

hour is set by the transmission speed. Assume that the printer will work for a maximum period of one hour. (1 mark)

The graph should be level at $b = d(1 - c/t)$ for p up to $3600c/d$ (this point should be shown on the graph).


The maximum buffer size needed will occur when the pages are printed as fast as the connection permits so the max value of p is given by $3600t/d$ (this should be shown on the graph). In an hour the system will have sent $3600t$ characters and $3600c$ will have been printed so the maximum buffer size will be $3600(t - c)$ bytes.

(d) The following sequence of operations describes an interrupt process. Draw the timing diagram for the process and design the logic circuit that will implement it.

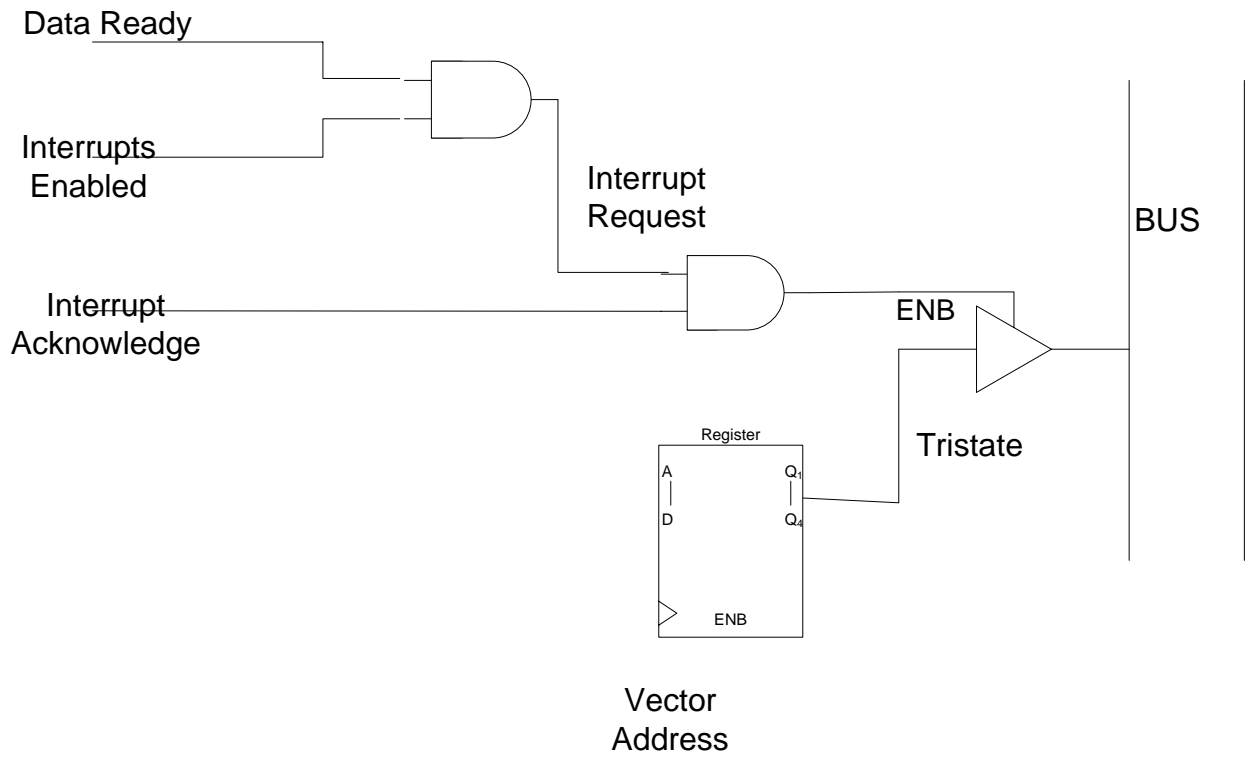
Device has data ready to send
Interrupts are enabled
An interrupt is generated
An interrupt acknowledge signal is sent back
The Interrupt Service Routine address is placed on the bus

(ii) Timing Diagram: (1.5 marks)

Initials _____



(iii) Circuit: (1.5 marks)



Initials _____

47. (a) Explain the following terms (maximum one sentence for each) (1 mark each)

(i) Logic Delay:

The basic delay through a digital logic gate

(ii) Capacitance Delay:


The delay due to the capacitance of the output of a gate

(iii) Propagation Delay:

The delay due to the finite time signals take to move down a wire

Initials _____

(b) (i) Draw the timing diagram for a write operation in the semi-synchronous bus protocol. (2 marks)



(ii) Give one advantage of a semi-synchronous protocol when compared to a synchronous protocol (0.5 marks)

Can accommodate devices of various speeds thus slow speed devices can be handled within the protocol.

(iii) Give one disadvantage of the semi-synchronous protocol when compared to an asynchronous protocol (0.5 marks)

There is a limit to the maximum clock rate, the asynchronous protocol can run as fast as the fastest devices on the bus.

Initials _____

(c) An interlocked asynchronous bus using only Master and Slave Synchronization signals has a set up time for both the address and data by the Master of t nanoseconds, the master (or slave) takes two gate delays (g nanoseconds per delay) to respond to a signal; the propagation speed of signals along the bus is v meters/nanosecond and the bus is L meters long. If the data path is 4 bytes wide, derive an expression for the maximum data transfer rate, R , (in bytes per second) that can be achieved for a **WRITE** operation. (3 marks)

Propagation time = L/v nanoseconds, for write address and data set up at same time

$$\text{Time} = t + (L/v) + 2g + (L/v) + 2g = (t + 4g + (2L/v))$$

$$\text{Maximum rate, } R = (4v * 10^9) / (vt + 4gv + 2L) \text{ bytes per second}$$

(Note subtract 1 if eqn for prop time wrong, subtract 1 if return signals not been allowed for. Get 1 for recognizing effect of t and $2g$, get 1 for final result. Lose 0.5 if 10^{**9} missing.


Lose a mark if needed 3 transfers – one of these is overlapped with the setup time..

Lose a mark if 4 assumed series rather than parallel

Accept 2 or 4 delays.)

Initials _____

(d) The diagram below relates to a sequence of operations on an arbitration bus.



(i) Describe the operations that are taking place (1 mark)

This is the sequence of events that happens when a low priority device requests and gets the bus and then a higher priority device requests the bus..

(ii) Name the signals in the diagram (2 marks)

Initials _____

Signal 1:
__B_RQST0_____

Signal 2:
__B_GRNT0_____

Signal 3:
__B_RQST1_____

Signal 4:
__B_GRNT1_____

Signal 5:
__B_RLS_____

Signal 6:
__B_BUSY_____

Initials _____

48. (a) Explain what is meant by each of the following terms. (1 mark each)

(i) Sector:

An angular area on the surface of a magnetic disk – it is the smallest physical data entity on the disk. Each track is made up of the same number of sectors

(ii) Cylinder:

A section through a hard disk drive obtained by taking the same track on every surface.

(iii) Seek Time:

The time taken for the head in a magnetic disk to move to the correct track.

Initials _____

(b) (i) Explain why Sector Interleaving is often implemented on a hard disk. (1 mark)

Sector interleaving is used to optimize the data transfer rate between the disk and the interface. If the interface cannot move data onto the bus at the rate at which it is received from the disk, then it has to stop accepting data from the disk. This could involve waiting for an entire revolution of the disk. By reducing the AVERAGE transfer rate from the disk by skipping sectors, it is possible to match the interface and bus performance. Thus a file may not be stored on consecutive sectors but several files may be interleaved.

- (ii) It is desired to find an optimum sector interleave factor for a particular disk/interface pair. Describe, using pseudo-code or a flow chart, an algorithm which would allow this to be determined (2 marks).

The solution should be something of the form:

For a sector interleave factor from 1 to N

Write a file to the disk with the current interleave factor

Start a timer

Read the file back from the disk

Stop the timer

Record the time

If the time is greater than the previous time, stop

Else repeat for the next interleave factor

The desired interleave factor is the next to last factor tried above

Initials _____

(c) A hard disk drive is designed to have p platters and to have t tracks per inch. The outer radius of the disk is R inches. Data is fed to the disk at B bytes per second. The maximum areal recording density is D bits per square inch.

- (i) What should be the radius of the innermost track if the amount of data to be stored on the disk is to be maximized? (1 mark)

The linear density of recording is $D^{0.5}$ bits per inch. If the inner track is at distance r , then the amount of data stored on a track is $2\pi r D^{0.5}$. The total number of tracks is $t(R-r)$ and each track stored the same amount of data. Thus the total data on the disk is given by:

$$\text{Total data} = T = t(R-r) * 2\pi r D^{0.5}$$

T maximize this:

$$dT/dr = 2\pi D^{0.5} (R-2r) = 0$$

$$\text{Thus } r = R/2$$

- (ii) If the inner radius is set to provide the maximum storage capability, what should the rotational speed of the disk be to be able to accept data at the rate the interface is sending it? (Assume the seek time is zero). (0.5 marks)

The interface is sending data at $8B$ bits per second. The data rate for a track on the disk is the total data on the track multiplied by the rotational speed of the disk:

$$8B = \pi R D^{0.5} x, \text{ where } x \text{ is the number of revolutions per second.}$$

Thus

$$x = 8B / (\pi R D^{0.5})$$

- (iii) What would happen to the total data stored on the disk and the disk performance if the rotational speed was increased above the speed from part (ii)? (0.5 marks)

Initials _____

The number of bits that can be stored per track is reduced so the total data on the disk will reduce.

In reality there would also be more track to track head moves for the same amount of data

- (iv) If the disk speed was reduced by 20% from that in part (ii) but all the other parameters stayed constant? What is the total amount of data that can be stored on the disk if one surface is used for timing?(1 mark)

As the disk slows down but the data rate stays the same, the physical size of the bits on the disk (at the innermost recorded track) reduces. This cannot happen in the model considered in part (ii) since the inner track bits are already at the minimum size for this disk. Assuming all the parameters stay the same, the inner recording radius will have to increase:

If $y = 0.8x$, then the bits per revolution must be $8B/0.8x$

The bits per track = $2\pi*r*D^{0.5}$*

Thus $8B/0.8x = 2\pi*r*D^{0.5}$*

*$r = 10B/(x*2*\pi*D^{0.5})$*

$r = 10B/(16B / R)$

$r = 10R/16 = 5R/8 = 0.675R$

The total data on a surface is then $2\pi*0.675R*D^{0.5}$*

$= 13.5\pi*R*D^{0.5}$ bits $= (13.5/8)*\pi*R*D^{0.5} = 1.687*\pi*R*D^{0.5}$ bytes*

*And the total data on the disk = $(p-1)*1.687*\pi*R*D^{0.5}$ bytes*

Initials _____

(d) Describe, either in the form of a flowchart or as a sequence of operations, the steps taken during a data transfer using a cycle stealing Direct Memory Access system. Start from the point at which it is determined that input is needed from a device and end when all the data has been transferred. In particular, indicate the changes in state of the control lines.

The diagram should be in the form of two columns – one for the CPU, the other for the interface and the communication between the two should be indicated.(3 marks)

*CPU**Interface*

*Input request
Load the interface with
the starting memory
address and number of
bytes to be transferred
Enable DMA (-0.5 if missing)*

*Load data in the input register
Set Bus Request when ready*

*Set Bus Grant to
relinquish the bus*

*Set the “Bus Busy” line
Place address on the address lines
Place data on the data lines*

Memory latches address

Issue “Memory Write”

Memory latches data

*Reset the “Bus Request”
Increment MAR
Decrement Byte Count*

Reset “Bus Grant”

0.25 marks off if the repeat not clearly stated

-1.5 if not cycle stealing.

-0.5 if loop to zero bytes not shown

Initials _____

49. (a) (i) Give four functions of the Nucleus (Kernel) of an Operating System. (0.5 marks each)

Process management

(1) _____

Memory management

(2) _____

I/O Management

(3) _____

Security

(4) _____

(ii) Name the two parts an Operating System is often divided into. (0.5 marks each)

System nucleus (or kernel)

(1) _____

System utilities

(2) _____

(b) There are three main problems which can occur as a result of Concurrency. Name each of these and give a brief explanation of the problem (maximum 20 words) (1 mark each)

Exclusion – this is due to the fact that some system resources are not shareable.

Problem 1: _____

Initials _____

Synchronization – this occurs when two or more processes are interdependent, i.e. each needs information from the other to operate.


Problem 2: _____

Deadlock – this occurs when several processes block each other and none can execute because of a lack of access to resources.

Problem 3: _____

Initials _____

(c) (i) The diagram below shows the state transitions of a process. Fill in all the blank states indicated in the figure. (2 marks)



(ii) Explain how a process may become “blocked” and how it can move out of this state.

A Process becomes blocked when _____

It needs a resource which is not available

(0.5 marks)

A Process can move out of a blocked state when _____

The resource it is waiting on becomes available

(0.5 marks)

Initials _____

(d) (i) Three (3) processes share four (4) resources which can be reserved and released only one at a time. Each process needs a maximum of two (2) resources to be able to complete (it does not matter which ones). Explain why *deadlock* occurs and show that it cannot occur in this system. (2 marks)

Deadlock requires that all resources are allocated but no process has enough to run to completion.

In this problem, on start up, each process reserves 1 resource.

1 process reserves the free resource and runs to completion

It then frees of 2 resources. Each of the other processes can then complete.

Hence deadlock cannot occur.

(ii) In any operating system one process always exists. Name the process and explain why it can never be in a “blocked” state.

Null process

Process: _____ (0.5 marks)

Explanation: _____

Needs no resources

_____ (0.5 marks)