

V.S.B. ENGINEERING COLLEGE, KARUR
Department of Computer Science and Engineering
Academic Year: 2018-2019 (EVEN Semester)
Assignment Question for Students

Class: II Year / IV Semester Computer Science and Engineering
Name of Subject: Software Engineering

Sl. No.	Assignment Topic
1	To help counter terrorism, many countries are planning the development of computer systems that track large numbers of their citizens and their actions. Clearly this has privacy implications. Discuss the ethics of developing this type of system. Design the system to find the illegal activities from already stored citizen's database.
2	A flood warning system is to be procured which will give early warning of possible flood dangers to sites that are threatened by floods. The system will include a set of sensors to monitor the rate of change of river levels, links to a meteorological system giving weather forecasts, links to the communication systems of emergency services (police, coastguard, etc.), video monitors installed at selected locations, and a control room equipped with operator consoles and video monitors. Controllers can access database information and switch video displays. The system database includes information about the sensors, the location of sites at risk and the threat conditions for these sites (e.g., high tide, southwesterly winds), tide tables for coastal sites, the inventory and location of flood control equipment, contact details for emergency services, local radio stations, and so on. Draw a block diagram of a possible architecture for such a system. You should identify the principal sub-systems and the links between them.
3	You are an engineer involved in the development of a financial system. During installation, you discover that this system will make a significant number of people redundant. The people in the environment deny you access to essential information to complete the system installation. To what extent should you, as a systems engineer, become involved in this? Is it your professional responsibility to complete the installation as contracted? Should you simply abandon the work until the procuring organization has sorted out the problem?
4	Is it ethical for an engineer to agree to deliver a software system with known faults to a customer? Does it make any difference if the customer is told of the existence of these faults in advance? Would it be reasonable to make claims about the reliability of the software in such circumstances? Draw the milestones for each module to rectify the problems.
5	As an expert in computer security, you have been approached by an organization that campaigns for the rights of torture victims and have been asked to help the organization gain unauthorized access to the computer systems of an American company. This will help them confirm or deny that this company is selling equipment that is used directly in the torture of political prisoners. Discuss the ethical dilemmas that this request raises and how you would react to this request.
6	You are asked by your manager to deliver software to a schedule that you know can only be met by asking your project team to work unpaid overtime. All team members have young children. Discuss whether you should accept this demand from your manager or whether you should persuade your team to give their time to the organization rather than to their families. What factors might be significant in your decision?
7	Discover ambiguities or omissions in the following statement of requirements for part of a ticket-issuing system. An automated ticket-issuing system sells rail tickets. Users select their destination and input a credit card and a personal identification number. The rail ticket is issued and their credit card account charged. When the user presses the start button, a menu display of potential destinations is activated, along with a message to the user to select a destination. Once a destination has been selected, users are requested to input their credit card. Its validity is checked and the user is then requested to input a personal identifier. When the credit transaction has been validated, the ticket is issued. Draw a sequence diagram showing the actions performed in the ticket-issuing system. You may make any reasonable assumptions about the system. Pay particular attention to specifying user errors.
8	Using the technique suggested here, where natural language is presented in a standard way, write plausible user requirements for the following functions: <ul style="list-style-type: none"> • The cash-dispensing function in a bank ATM • The spelling-check and correcting function in a word processor • An unattended petrol (gas) pump system that includes a credit card reader. <p style="text-align: center;">The customer swipes the card through the reader and then specifies the amount of fuel required. The fuel is delivered and the customer's account debited.</p>

Sl. No.	Assignment Topic
9	You have taken a job with a software user who has contracted your previous employer to develop a system for them. You discover that your company's interpretation of the requirements is different from the interpretation taken by your previous employer. Discuss what you should do in such a situation. You know that the costs to your current employer will increase if the ambiguities are not resolved. You have also a responsibility of confidentiality to your previous employer.
10	A software system is to be developed to manage the records of patients who enter a clinic for treatment. The records include records of all regular patient monitoring (temperature, blood pressure, etc.), treatments given, patient reactions and so on. After treatment, the records of their stay are sent to the patient's doctor who maintains their complete medical record. Identify the principal viewpoints which might be taken into account in the specification of this system and organize these using a viewpoint hierarchy diagram.
11	Draw a context model for a patient information system in a hospital. You may make any reasonable assumptions about the other hospital systems that are available, but your model must include a patient admissions system and an image storage system for X-rays, as well as other diagnostic records.
12	Draw state machine models of the control software for an automatic washing machine that has different programs for different types of clothes.
13	Draw state machine models of the control software for the software for a DVD player.
14	Draw state machine models of the control software for a telephone answering system that records incoming messages and displays the number of accepted messages on an LED. The system should allow the telephone customer to dial in from any location, type a sequence of numbers (identified as tones) and play the recorded messages.
15	<p>Suggest appropriate reliability metrics for the following classes of software system. Give reasons for your choice of metric. Predict the usage of these systems and suggest appropriate values for the reliability metrics:</p> <ul style="list-style-type: none"> • A system that monitors patients in a hospital intensive care unit • A word processor • An automated vending machine control system • A system to control braking in a car • A system to control a refrigeration unit • A management report generator.
16	Bank teller machines rely on using information on the user's card giving the bank identifier, the account number and the user's personal identifier. They also derive account information from a central database and update that database on completion of a transaction. Using your knowledge of ATM operation, write Z schemas defining the state of the system, card validation (where the user's identifier is checked) and cash withdrawal.
17	Your customer wants to develop a system for stock information where dealers can access information about companies and can evaluate various investment scenarios using a simulation system. Each dealer uses this simulation in a different way, according to his or her experience and the type of stocks in question. Suggest a client-server architecture for this system that shows where functionality is located. Justify the client-server system model that you have chosen.
18	Using the basic application types introduced the systems and explains a point-of-sale system in a supermarket.
19	Using the basic application types introduced the systems and explains a system that sends out reminders that magazine subscriptions are due to be paid.
20	Using the basic application types introduced the systems and explains a photo album system that provides some facilities for restoring old photographs.
21	<p>Based on an input-process-output model, expand the Compute salary function and draw a data-flow diagram that shows the computations carried out in that function. You need the following information to do this:</p> <ul style="list-style-type: none"> • The employee record identifies the grade of an employee. This grade is then used to look up the table of pay rates. • Employees below a particular grade may be paid overtime at the same rate as their normal hourly pay rate. The extra hours for which they are to be paid are indicated in their employee record. • The amount of tax deducted depends on the employee's tax code (indicated in the record) and their annual salary. Monthly deductions for each code and a standard salary are indicated in the tax tables. These are scaled up or down depending on the relationship between the actual salary and the

Sl. No.	Assignment Topic
	standard salary used.
22	Using the UML graphical notation for object classes, design the following object classes identifying attributes and operations. Use your own experience to decide on the attributes and operations that should be associated with a telephone system. Find out the modules for this system.
23	Using the UML graphical notation for object classes, design the following object classes identifying attributes and operations. Use your own experience to decide on the attributes and operations that should be associated with a printer for a personal computer system. Find out the modules for this system.
24	Using the UML graphical notation for object classes, design the following object classes identifying attributes and operations. Use your own experience to decide on the attributes and operations that should be associated with a library catalogue system. Find out the modules for this system.
25	Identify possible objects in the following systems and develop an object-oriented design for them. You may make any reasonable assumptions about the systems when deriving the design. A group diary and time management system is intended to support the timetabling of meetings and appointments across a group of coworkers. When an appointment is to be made that involves a number of people, the system finds a common slot in each of their diaries and arranges the appointment for that time. If no common slots are available, it interacts with the user to rearrange his or her personal diary to make room for the appointment.
26	Identify possible objects in the following systems and develop an object-oriented design for them. You may make any reasonable assumptions about the systems when deriving the design. A petrol (gas) station is to be set up for fully automated operation. Drivers swipe their credit card through a reader connected to the pump; the card is verified by communication with a credit company computer; and a fuel limit is established. The driver may then take the fuel required. When fuel delivery is complete and the pump hose is returned to its holster, the driver's credit card account is debited with the cost of the fuel taken. The credit card is returned after debiting. If the card is invalid, the pump returns it before fuel is dispensed.
27	A charity has asked you to prototype a system that keeps track of all donations they have received. This system has to maintain the names and addresses of donors, their particular interests, the amount donated and when the donation was made. If the donation is over a certain amount, the donor may attach conditions to the donation (e.g., it must be spent on a particular project), and the system must keep track of these and how the donation was spent. Discuss how you would prototype this system, bearing in mind that the charity has a mixture of paid workers and volunteers. Many of the volunteers are retirees who have had little or no computer experience.
28	Design the interfaces of components that might be used in a system in an emergency control room. You should design interfaces for a call-logging component that records calls made, and a vehicle-discovery component that, given a post code and an incident type, finds the nearest suitable vehicle to be dispatched to the incident.
29	<p>The door lock control mechanism in a nuclear waste storage facility is designed for safe operation. It ensures that entry to the storeroom is only permitted when radiation shields are in place or when the radiation level in the room falls below some given value (danger Level). That is:</p> <ol style="list-style-type: none"> i. If remotely controlled radiation shields are in place within a room, the door may be opened by an authorized operator. ii. If the radiation level in a room is below a specified value, the door may be opened by an authorized operator. iii. An authorized operator is identified by the input of an authorized door entry code.
30	A software manager is in charge of the development of a safety-critical software system that is designed to control a radiotherapy machine to treat patients suffering from cancer. This system is embedded in the machine and must run on a special-purpose processor with a fixed amount of memory (8 Mbytes). The machine communicates with a patient database system to obtain the details of the patient and, after treatment, automatically records the radiation dose delivered and other treatment details in the database. The COCOMO method is used to estimate the effort required to develop this system and an estimate of 26 person-months is computed. All cost driver multipliers were set to 1 when making this estimate. Explain why this estimate should be adjusted to take project, personnel, product and organizational factors into account. Suggest four factors that might have significant effects on the initial COCOMO estimate and propose possible values for these factors. Justify why you have included each factor.
31	Explain why incremental development is the most effective approach for developing business software systems. Why is this model less appropriate for real-time systems engineering?
32	Giving reasons for your answer based on the type of system being developed, suggest the most appropriate generic software process model that might be used as a basis for managing the development of the following systems:

Sl. No.	Assignment Topic
	<ul style="list-style-type: none"> i. A system to control anti-lock braking in a car ii. A virtual reality system to support software maintenance iii. A university accounting system that replaces an existing system iv. An interactive travel planning system that helps users plan journeys with the lowest environmental impact
33	To reduce costs and the environmental impact of commuting, your company decides to close a number of offices and to provide support for staff to work from home. However, the senior management who introduce the policy are unaware that software is developed using agile methods, which rely on close team working and pair programming. Discuss the difficulties that this new policy might cause and how you might get around these problems.
34	You have taken a job with a software user who has contracted your previous employer to develop a system for them. You discover that your company's interpretation of the requirements is different from the interpretation taken by your previous employer. Discuss what you should do in such a situation. You know that the costs to your current employer will increase if the ambiguities are not resolved. However, you have also a responsibility of confidentiality to your previous employer.
35	Based on your experience with a bank ATM, draw an activity diagram that models the data processing involved when a customer withdraws cash from the machine.
36	<p>Draw state diagrams of the control software for:</p> <ul style="list-style-type: none"> i. An automatic washing machine that has different programs for different types of clothes. ii. The software for a DVD player. iii. A telephone answering system that records incoming messages and displays the number of accepted messages on an LED. The system should allow the telephone customer to dial in from any location, type a sequence of numbers (identified as tones), and play any recorded messages.
37	You work for a very small software organization—only 11 people are involved in developing software. Is SPI for you? Explain your answer
38	<p>Suggest appropriate reliability metrics for the following classes of software system. Give reasons for your choice of metric. Predict the usage of these systems and suggest appropriate values for the reliability metrics:</p> <ul style="list-style-type: none"> i. A system that monitors patients in a hospital intensive care unit ii. A word processor iii. An automated vending machine control system iv. A system to control braking in a car v. A system to control a refrigeration unit vi. A management report generator.
39	Bank teller machines rely on using information on the user's card giving the bank identifier, the account number and the user's personal identifier. They also derive account information from a central database and update that database on completion of a transaction. Using your knowledge of ATM operation, write Z schemas defining the state of the system, card validation (where the user's identifier is checked) and cash withdrawal.
40	Your customer wants to develop a system for stock information where dealers can access information about companies and can evaluate various investment scenarios using a simulation system. Each dealer uses this simulation in a different way, according to his or her experience and the type of stocks in question. Suggest a client-server architecture for this system that shows where functionality is located. Justify the client-server system model that you have chosen.
41	Is unit testing possible or even desirable in all circumstance = provide examples to justify your answer.
42	<p>Perform requirements model engineering and specify a set of essential and real use cases describing the functionality that your business solution/system will exhibit. To start with, using URN GRL diagrams, identify goals/subgoals, actors and their dependencies, and high-level tasks. Perform scenario modeling using URN UCM diagrams and group scenarios as needed to identify essential use cases. Using essential/real Use Case diagrams identify the following:</p> <ul style="list-style-type: none"> ✓ Actors – who/what interacts with the business solution/system? ✓ Use Cases – what behaviors does the business solution/system exhibit?

Sl. No.	Assignment Topic
43	Draw diagrams for weather mapping system using the following and explain them. (a) Subsystem (b) Sequence model (c) State machine model
44	To help counter terrorism, many countries are planning the development of computer systems that track large numbers of their citizens and their actions. Clearly this has privacy implications. Discuss the ethics of developing this type of system. Design the system to find the illegal activities from already stored citizen's database.
45	Given a set of numbers 'n', the function FindPrime(a[],n) print a number-if it is a prime number. Draw a control flow graph, calculate the cyclomatic complexity and enumerate all paths. State how many test cases are needed to adequately cover the code in terms of branches, decisions and statement? Develop the necessary test cases using sample values 'a' and 'n'.
46	Consider an online railway reservation system, which allows the user to select route, book / cancel tickets using net banking/ credit/debit cards. The site also maintains the history of the passengers. For the above system, list and draw the use case scenario, sequential diagram and collaboration diagram.
47	Tamil Nadu Electricity Board (TNEB) would like to automate its billing process. Customers apply for a connection (domestic/commercial). EB staffs take readings and update the system. Each customer is required to pay charges bi-monthly according to the rates set for the type of connection. Customers can choose to pay either by cash/card. A bill is generated on payment. Monthly reports are provided to the EB manager. (i) Draw the level-0 and level-1 (context flow diagram) (ii) Draw the level-2
48	Suppose you have a budgeted cost of a project as Rs. 9,00,000. The project is to be completed in 9 months. After a month, you have completed 10 percent of the project at a total expense of Rs 1,00,000. The planned completion should have been 15 percent. You need to determine whether the project is on-time and on-budget? Use Earned Value analysis approach and interpret.
49	Illustrate the COCOMO model for software cost estimation. Use it to estimate the effort required to build software for a simple ATM that produce 12 screens, 10 reports and has 80 software components. Assume average complexity and average developer maturity. Use application composition model with object points.
50	Compare function point value for a project with the following information domain characteristics: No. of external inputs – 30 No. of external outputs – 52 No. of external inquiries – 22 No. of log files- 12 No. of external interface files-2 Assume complexity adjustment values for above are average (4,5,4,10,7 respectively).
51	Suppose you have a budgeted cost of a project as Rs. 9,00,000. The project is to be completed in 9 months. After a month, you have completed 10 percent of the project at a total expense of Rs 1,00,000. The planned completion should have been 15 percent. You need to determine whether the project is on-time and on-budget? Use Earned Value analysis approach and interpret.
52	An application has the following: 10 low external inputs, 8 high external outputs, 13 low internal logical files, 17 high external interface files, 11 average external inquiries and complexity adjustment factor of 1.10. What are the unadjusted and adjusted functional point counts?
53	Given a set of 'n' numbers, write an algorithm that finds whether a given number is positive, negative, zero, even or odd. Finally, the total number in each category is also printed. Draw the flow graph and enumerate paths for testing. Determine the number of independent paths using cyclomatic complexity. Consider the pseudo code for simple subtraction given below. 1. Program 'simple subtraction' 2. Input (x,y) 3. Output(x) 4. Output(y) 5. if x>y then DO 6. x-y=z

Sl. No.	Assignment Topic
	7. Else $y-x=z$ 8. Endif 9. Output (z) 10. Output "End program" Perform basis path testing and generate test cases.
54	Draw use case and data flow diagrams for library management system.

Class: II Year / IV Semester Computer Science and Engineering

Name of Subject: Operating systems

Sl.No.	Assignment Topic
1	Show that, if the wait() and signal() semaphore operations are not executed atomically, then mutual exclusion may be violated.
2	Original versions of Apple's mobile iOS operating system provided no means of concurrent processing. Discuss three major complications that on current processing adds to an operating system.
3	Discuss, with examples, how the problem of maintaining coherence of cached data manifests itself in the following processing environments: a. Single-processor systems c. Distributed systems
4	Some computer systems do not provide a privileged mode of operation in hardware. Is it possible to construct a secure operating system for these computer systems? Give arguments both that it is and that it is not possible.
5	Describe the differences among short-term, medium-term, and long term scheduling.
6	Describe the actions taken by a kernel to context-switch between processes.
7	Consider both the system level and the programmer level. a. Synchronous and asynchronous communication b. Automatic and explicit buffering c. Send by copy and send by reference d. Fixed-sized and variable-sized messages What resources are used when a thread is created How do they differ from those used when a process is created?
8	How does an interrupt differ from a trap? Can traps be generated intentionally by a user program? If so, for what purpose?
9	Consider the hierarchical paging scheme used by the VAX architecture. How many memory operations are performed when a user program executes a memory-load operation?
10	Can a multithreaded solution using multiple user-level threads achieve better performance on a multiprocessor system than on a singleprocessor system? Explain
11	Consider the following set of processes, with the length of the CPU-burst time given in milliseconds: Process Burst Time Priority 1. P1 10 3

	<p>2. P2 29 1</p> <p>3. P3 3 4</p> <p>4. P4 7 5</p> <p>5. P5 12 2</p> <p>The processes are assumed to have arrived in the order P1, P2, P3, P4, P5, all at time 0</p> <p>a. Draw four Gantt charts illustrating the execution of these processes using FCFS, SJF, A non preemptive priority (a smaller priority number implies a higher priority), and RR (quantum = 1) scheduling.</p> <p>b. What is the turnaround time of each process for each of the scheduling algorithms in part a?</p> <p>c. What is the waiting time of each process for each of the scheduling algorithms in Part a?</p> <p>d. Which of the schedules in part a results in the minimal average waiting time (over all processes)?</p>
12	<p>Consider the following page reference string: 1, 2, 3, 4, 2, 1, 5, 6, 2, 1, 2, 3, 7, 6, 3, 2, 1, 2, 3, 6.</p> <p>How many page faults would occur for the following replacement algorithms, assuming one, two, three, four, five, six, or seven frames? Remember all frames are initially empty, so your first unique pages will all cost one fault each.</p> <p>LRU replacement, FIFO replacement, Optimal replacement</p>
13	<p>Consider the following page reference string: 1,2,3,4,2,1,5,6,1,2,3,7,6,3,2,1,2,3,6 How many page faults would occur for the LRU, FIFO, LFU and optimal page replacement algorithms assuming two and five frames?</p>
14	<p>Consider the following page reference string: 1,2,3,4,2,1,5,6,1,2,3,7,6,3,2,1,2,3,6 How many page faults would occur for the LRU, FIFO, LFU and optimal page replacement algorithms assuming two and five frames?</p>
15	<p>Google's Chrome browser and its practice of opening each new website in a separate process. Would the same benefits have been achieved if instead Chrome had been designed to open each new website in a separate thread? Explain.</p>
16	<p>Consider a multicore system and a multithreaded program written using the many-to-many threading model. Let the number of user-level threads in the program be greater than the number of processing cores in the system. Discuss the performance implications of the following scenarios.</p> <p>a. The number of kernel threads allocated to the program is less than the number of processing cores.</p> <p>b. The number of kernel threads allocated to the program is equal to the number of processing cores.</p> <p>c. The number of kernel threads allocated to the program is greater than the number of processing cores but less than the number of user-level threads.</p>
17	<p>Explain why Windows, Linux, and Solaris implement multiple locking mechanisms. Describe the circumstances under which they use spinlocks, mutex locks, semaphores, adaptive mutex locks, and condition variables. In each case, explain why the mechanism is needed.</p>
18	<p>What is the meaning of the term <i>busy waiting</i>? What other kinds of waiting are there in an operating system? Can busy waiting be avoided altogether? Explain your answer.</p>
19	<p>What are two differences between user-level threads and kernel-level threads? Under what circumstances is one type better than the other?</p>
20	<p>What is the effect of allowing two entries in a page table to point to the same page frame in memory? Explain how this effect could be used to decrease the amount of time needed to copy a large amount of memory from one place to another. What effect would updating some byte on the one page have on the other page?</p>
21	<p>Describe a mechanism by which one segment could belong to the address space of two different processes.</p>
22	<p>Compare the memory organization schemes of contiguous memory allocation, pure segmentation, and pure paging with respect to the following issues:</p> <p>a. External fragmentation</p> <p>b. Internal fragmentation</p> <p>c. Ability to share code across processes</p>
23	<p>Explain why mobile operating systems such as iOS and Android do not support swapping.</p>

24	Although Android does not support swapping on its boot disk, it is possible to set up a swap space using a separate SD nonvolatile memory card. Why would Android disallow swapping on its boot disk yet allow it on a secondary disk?
25	Consider the following page reference string: 1,2,3,4,2,1,5,6,1,2,3,7,6,3,2,1,2,3,6 How many page faults would occur for the LRU, FIFO, LFU and optimal page replacement algorithms assuming three and four frames?
26	Consider the following page reference string: 1,2,3,4,2,1,5,6,1,2,3,7,6,3,2,1,2,3,6 How many page faults would occur for the LRU, FIFO, LFU and optimal page replacement algorithms assuming two and six frames?
27	Consider the following page reference string: 1,2,3,4,2,1,5,6,1,2,3,7,6,3,2,1,2,3,6 How many page faults would occur for the LRU, FIFO, LFU and optimal page replacement algorithms assuming seven frames?
28	Explain why sharing a reentrant module is easier when segmentation is used than when pure paging is used.
29	Under what circumstances does a multithreaded solution using multiple kernel threads provide better performance than a single-threaded solution on a single-processor system?
30	Under what circumstances do page faults occur? Describe the actions taken by the operating system when a page fault occurs.
31	Discuss the hardware support required to support demand paging.
32	Two segment-replacement algorithms, one based on the FIFO pagereplacement scheme and the other on the LRU page-replacement scheme. Remember that since segments are not the same size, the segment that is chosen for replacement may be too small to leave enough consecutive locations for the needed segment. Consider strategies for systems where segments cannot be relocated and strategies for systems where they can.
33	What is the copy-on-write feature, and under what circumstances is its use beneficial? What hardware support is required to implement this feature?
34	Is it possible for a process to have two working sets, one representing data and another representing code? Explain.
35	Is disk scheduling, other than FCFS scheduling, useful in a single-user environment? Explain your answer.
36	Suppose that a disk drive has 5,000 cylinders, numbered 0 to 4,999. The drive is currently serving a request at cylinder 2,150, and the previous request was at cylinder 1,805. The queue of pending requests, in FIFO order, is: 2,069, 1,212, 2,296, 2,800, 544, 1,618, 356, 1,523, 4,965, 3681 Starting from the current head position, what is the total distance (in cylinders) that the disk arm moves to satisfy all the pending requests for each of the following disk-scheduling algorithms? a. FCFS b. SSTF c. SCAN d. LOOK e. C-SCAN f. C-LOOK
37	Consider the following page reference string: 7,5,7,8,6,0,1,8,9,7,5,4,3,2,1. How many page faults would occur for the following replacement algorithms, assuming one, two, three, four, five, six, or seven frames? Remember all frames are initially empty, so your first unique pages will all cost one fault each. LRU replacement , FIFO replacement , Optimal replacement

38	<p>Consider the following page reference string: 8,5,6,4,7,2,1,9,7,2,0,1,6,8,9,0,1. How many page faults would occur for the following replacement algorithms, assuming one, two, three, four, five, six, or seven frames? Remember all frames are initially empty, so your first unique pages will all cost one fault each. LRU replacement , FIFO replacement , Optimal replacement</p>																												
39	<p>Compare the performance of C-SCAN and SCAN scheduling, assuming a uniform distribution of requests. Consider the average response time (the time between the arrival of a request and the completion of that request's service), the variation in response time, and the effective.</p>																												
40	<p>Some systems automatically delete all user files when a user logs off or a job terminates, unless the user explicitly requests that they be kept. Other systems keep all files unless the user explicitly deletes them. Discuss the relative merits of each approach.</p>																												
41	<p>Why do some systems keep track of the type of a file, while others leave it to the user and others simply do not implement multiple file types? Which system is "better"?</p>																												
42	<p>Similarly, some systems support many types of structures for a file's data, while others simply support a stream of bytes. What are the advantages and disadvantages of each approach?</p>																												
43	<p>Give an example of an application that could benefit from operating system support for random access to indexed files.</p>																												
44	<p>Why is it advantageous to the user for an operating system to dynamically allocate its internal tables? What are the penalties to the operating system for doing so?</p>																												
45	<p>Contrast the performance of the three techniques for allocating disk blocks (contiguous, linked, and indexed) for both sequential and random file access.</p>																												
46	<p>How does DMA increase system concurrency? How does it complicate hardware design?</p>																												
47	<p>Dynamically loadable kernel modules give flexibility when drivers are added to a system, but do they have disadvantages too? Under what circumstances would a kernel be compiled into a single binary file, and when would it be better to keep it split into modules? Explain your answer.</p>																												
48	<p>Consider the following snapshot of a system:</p> <table style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th></th> <th style="text-align: center;"><u>Allocation</u></th> <th style="text-align: center;"><u>Max</u></th> <th style="text-align: center;"><u>Available</u></th> </tr> <tr> <th></th> <th style="text-align: center;">A B C D</th> <th style="text-align: center;">A B C D</th> <th style="text-align: center;">A B C D</th> </tr> </thead> <tbody> <tr> <td>P_0</td> <td style="text-align: center;">0 0 1 2</td> <td style="text-align: center;">0 0 1 2</td> <td style="text-align: center;">1 5 2 0</td> </tr> <tr> <td>P_1</td> <td style="text-align: center;">1 0 0 0</td> <td style="text-align: center;">1 7 5 0</td> <td></td> </tr> <tr> <td>P_2</td> <td style="text-align: center;">1 3 5 4</td> <td style="text-align: center;">2 3 5 6</td> <td></td> </tr> <tr> <td>P_3</td> <td style="text-align: center;">0 6 3 2</td> <td style="text-align: center;">0 6 5 2</td> <td></td> </tr> <tr> <td>P_4</td> <td style="text-align: center;">0 0 1 4</td> <td style="text-align: center;">0 6 5 6</td> <td></td> </tr> </tbody> </table> <p>Answer the following questions using the banker's algorithm:</p> <ol style="list-style-type: none"> What is the content of the matrix <i>Need</i>? Is the system in a safe state? If a request from process P_1 arrives for (0,4,2,0), can the request be granted immediately 		<u>Allocation</u>	<u>Max</u>	<u>Available</u>		A B C D	A B C D	A B C D	P_0	0 0 1 2	0 0 1 2	1 5 2 0	P_1	1 0 0 0	1 7 5 0		P_2	1 3 5 4	2 3 5 6		P_3	0 6 3 2	0 6 5 2		P_4	0 0 1 4	0 6 5 6	
	<u>Allocation</u>	<u>Max</u>	<u>Available</u>																										
	A B C D	A B C D	A B C D																										
P_0	0 0 1 2	0 0 1 2	1 5 2 0																										
P_1	1 0 0 0	1 7 5 0																											
P_2	1 3 5 4	2 3 5 6																											
P_3	0 6 3 2	0 6 5 2																											
P_4	0 0 1 4	0 6 5 6																											
49	<p>Consider the following page reference string:</p>																												

	7,2,4,3,8,9,4,1,6,7,9,1,4,5,3,8,7,6,4,2,1,2. How many page faults would occur for the LRU, FIFO, LFU and optimal page replacement algorithms assuming three and five frames?																												
50	Suppose a stack is to be used by the processor to manage procedure calls and returns. Can the program counter be eliminated by using the top of the stack as a program counter? Explain.																												
51	Assume a multithreaded application uses only reader–writer locks for synchronization. Applying the four necessary conditions for deadlock, is deadlock still possible if multiple reader–writer locks are used?																												
52	<p>Consider the following snapshot of a system:</p> <table style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th></th> <th style="text-align: center;"><u>Allocation</u></th> <th style="text-align: center;"><u>Max</u></th> <th style="text-align: center;"><u>Available</u></th> </tr> <tr> <th></th> <th style="text-align: center;">A B C D</th> <th style="text-align: center;">A B C D</th> <th style="text-align: center;">A B C D</th> </tr> </thead> <tbody> <tr> <td>P_0</td> <td style="text-align: center;">2 0 0 1</td> <td style="text-align: center;">4 2 1 2</td> <td style="text-align: center;">3 3 2 1</td> </tr> <tr> <td>P_1</td> <td style="text-align: center;">3 1 2 1</td> <td style="text-align: center;">5 2 5 2</td> <td></td> </tr> <tr> <td>P_2</td> <td style="text-align: center;">2 1 0 3</td> <td style="text-align: center;">2 3 1 6</td> <td></td> </tr> <tr> <td>P_3</td> <td style="text-align: center;">1 3 1 2</td> <td style="text-align: center;">1 4 2 4</td> <td></td> </tr> <tr> <td>P_4</td> <td style="text-align: center;">1 4 3 2</td> <td style="text-align: center;">3 6 6 5</td> <td></td> </tr> </tbody> </table> <p>Answer the following questions using the banker’s algorithm:</p> <ol style="list-style-type: none"> Illustrate that the system is in a safe state by demonstrating an order in which the processes may complete. If a request from process P_1 arrives for (1, 1, 0, 0), can the request be granted immediately? If a request from process P_4 arrives for (0, 0, 2, 0), can the request be granted immediately? 		<u>Allocation</u>	<u>Max</u>	<u>Available</u>		A B C D	A B C D	A B C D	P_0	2 0 0 1	4 2 1 2	3 3 2 1	P_1	3 1 2 1	5 2 5 2		P_2	2 1 0 3	2 3 1 6		P_3	1 3 1 2	1 4 2 4		P_4	1 4 3 2	3 6 6 5	
	<u>Allocation</u>	<u>Max</u>	<u>Available</u>																										
	A B C D	A B C D	A B C D																										
P_0	2 0 0 1	4 2 1 2	3 3 2 1																										
P_1	3 1 2 1	5 2 5 2																											
P_2	2 1 0 3	2 3 1 6																											
P_3	1 3 1 2	1 4 2 4																											
P_4	1 4 3 2	3 6 6 5																											
53	Consider the following page reference string: 1,2,3,4,2,1,5,6,1,2,3,7,6,3,2,1,2,3,6 How many page faults would occur for the LRU, FIFO, LFU and optimal page replacement algorithms assuming seven frames?																												
54	Suppose that a disk drive has 5,000 cylinders, numbered 0 to 4,999. The drive is currently serving a request at cylinder 2,150, and the previous request was at cylinder 1,805. The queue of pending requests, in FIFO order, is: 2,069, 1,212, 2,296, 2,800, 544, 1,618, 356, 1,523, 4,965, 3681 Starting from the current head position, what is the total distance (in cylinders) that the disk arm moves to satisfy all the pending requests for each of the following disk-scheduling algorithms? a. FCFS b. SSTF c. SCAN d. LOOK e. C-SCAN f. C-LOOK																												

Class: II Year / IV Semester Computer Science and Engineering**Name of Subject: Computer Architecture**

Sl.No.	Assignment Topic
1	Let $A = A(1), A(2), \dots, A(1000)$ and $B = B(1), B(2), \dots, B(1000)$ be two vectors (one-dimensional arrays) comprising 1000 numbers each that are to be added to form an array C such that $C(I) = A(I) + B(I)$ for $I = 1, 2, \dots, 1000$. Using the IAS instruction set, write a program for this problem. Ignore the fact that the IAS was designed to have only 1000 words of storage.
2	a. On the IAS, what would the machine code instruction look like to load the contents of memory address 2? b. How many trips to memory does the CPU need to make to complete this instruction during the instruction cycle?
3	On the IAS, describe in English the process that the CPU must undertake to read a value from memory and to write a value to memory in terms of what is put into the MAR, MBR, address bus, data bus, and control bus.
4	Given the memory contents of the IAS computer shown below, Address Contents 08A 010FA210FB 08B 010FA0F08D 08C 020FA210FB show the assembly language code for the program, starting at address 08A. Explain what this program does.
5	In the IBM 360 Models 65 and 75, addresses are staggered in two separate main memory units (e.g., all even-numbered words in one unit and all odd-numbered words in another). What might be the purpose of this technique?
6	Draw the typical block diagram of a DMA controller and explain how it is used for direct data transfer between memory and peripherals.
7	While browsing at Billy Bob's computer store, you overhear a customer asking Billy Bob what is the fastest computer in the store that he can buy. Billy Bob replies, "You're looking at our Macintoshes. The fastest Mac we have runs at a clock speed of 1.2 gigahertz. If you really want the fastest machine, you should buy our 2.4-gigahertz Intel Pentium IV instead." Is Billy Bob correct? What would you say to help this customer?
8	The ENIAC was a decimal machine, where a register was represented by a ring of 10 vacuum tubes. At any time, only one vacuum tube was in the ON state, representing one of the 10 digits. Assuming that ENIAC had the capability to have multiple vacuum tubes in the ON and OFF state simultaneously, why is this representation "wasteful" and what range of integer values could we represent using the 10 vacuum tubes?
9	A benchmark program is run on a 40 MHz processor. The executed program consists of 100,000 instruction executions, with the following instruction mix and clock cycle count: Instruction Type Instruction Count Cycles per Instruction Integer arithmetic 45000 1, Data transfer 32000 2, Floating point 15000 2 Control transfer 8000 2 Determine the effective CPI, MIPS rate, and execution time for this program.
10	A processor accesses main memory with an average access time of T_2 . A smaller cache memory is interposed between the processor and main memory. The cache has a significantly faster access time of T_1 . The cache holds, at any time, copies of some main memory words and is designed so that the words more likely to be accessed in the near future are in the cache. Assume that the probability that the next word accessed by the processor is in the cache is H , known as the hit ratio. a. For any single memory access, what is the theoretical speedup of accessing the word in the cache rather than in main memory? b. Let T be the average access time. Express T as a function of T_1 , T_2 , and H . What is the overall speedup as a function of H ?

11	Draw the typical block diagram of a DMA controller and explain how it is used for direct data transfer between memory and peripherals.
12	Show the micro-operations and control signals in the same fashion for the processor in for the following instructions: • Load Accumulator, • Store Accumulator • Add to Accumulator • AND to Accumulator • Jump • Jump if • Complement Accumulator
13	Assume that propagation delay along the bus and through the ALU are 20 and 100 ns, respectively. The time required for a register to copy data from the bus is 10 ns. What is the time that must be allowed for a. transferring data from one register to another? b. incrementing the program counter?
14	Write the sequence of micro-operations required for the bus structure of to add a number to the AC when the number is a. an immediate operand b. a direct-address operand c. an indirect-address operand
15	A stack is implemented Show the sequence of micro operations for a. popping b. pushing the stack. Let be the percentage of program code that can be executed simultaneously by n processors in a computer system. Assume that the remaining code must be executed sequentially by a single processor. Each processor has an execution rate of x MIPS. a. Derive an expression for the effective MIPS rate when using the system for exclusive execution of this program, in terms of n, and x.
16	Can you foresee any problem with the write-once cache approach on bus-based multiprocessors? If so, suggest a solution.
17	Consider a situation in which two processors in an SMP configuration, over time, require access to the same line of data from main memory. Both processors have a cache and use the MESI protocol. Initially, both caches have an invalid copy of the line. Depicts the consequence of a read of line x by Processor P1. If this is the start of a sequence of accesses, draw the subsequent figures for the following sequence: 1. P2 reads x. 2. P1 writes to x (for clarity, label the line in P1's cache 3. P1 writes to x (label the line in P1's cache 4. P2 reads x.
18	shows the state diagrams of two possible cache coherence protocols. Deduce and explain each protocol, and compare each to MESI.
19	Consider an SMP with both L1 and L2 caches using the MESI protocol. As explained, one of four states is associated with each line in the L2 cache. Are all four states also needed for each line in the L1 cache? If so, why? If not, explain which state or states can be eliminated.
20	(i) Explain how a processor works with simple instruction sets. (RISC). (ii) How can you improve the performance of a processor with complex
21	Consider a situation in which two processors in an SMP configuration, over time, require access to the same line of data from main memory. Both processors have a cache and use the MESI protocol. Initially, both caches have an invalid copy of the line. Depicts the consequence of a read of line x by Processor P1. If this is the start of a sequence of accesses, draw the subsequent figures for the following sequence: 1. P2 reads x. 2. P1 writes to x (for clarity, label the line in P1's cache 3. P1 writes to x (label the line in P1's cache 4. P2 reads x.
22	(i) Explain how a read operation can be performed with data transfer signals in a PCI bus. (ii) Narrate the sequence of events caused by a SCSI controller.
23	(i). Explain in detail about interface standards used in the design of a processor bus. (ii). Discuss in detail about the objectives and technical details of USB.
24	Draw the typical block diagram of a DMA controller and explain how it is used for direct data transfer between memory and peripherals.
25	(i). Explain about division algorithms. What is the need of division algorithms in ALU. (ii). How multiplication operation is performed by ALU.
26	(i) Explain the objectives and advantages of nanoprogramming. (ii) Narrate the sequence of events for an unconditional branch instruction.
27	(i) Explain Handshake protocol. Depict clearly how it controls data transfer during an input operation. (ii) Explain the data transfer over a synchronous bus.
28	(i) Explain how a processor works with simple instruction sets. (RISC). (ii) How can you improve the performance of a processor with complex
29	(i) With neat diagram explain the basic organization of a micro-programmed control unit with a typical set of micro instructions.

	(ii) How can you structure Microinstructions? Explain with the help of an example.
30	Given the memory contents of the IAS computer shown below, Address Contents 08A 010FA210FB 08B 010FA0F08D 08C 020FA210FB show the assembly language code for the program, starting at address 08A. Explain what this program does.
31	In the IBM 360 Models 65 and 75, addresses are staggered in two separate main memory units (e.g., all even-numbered words in one unit and all odd-numbered words in another). What might be the purpose of this technique?
32	(i) Explain how a processor works with simple instruction sets. (RISC). (ii) How can you improve the performance of a processor with complex
33	On the IAS, describe in English the process that the CPU must undertake to read a value from memory and to write a value to memory in terms of what is put into the MAR, MBR, address bus, data bus, and control bus.
34	Consider a situation in which two processors in an SMP configuration, over time, require access to the same line of data from main memory. Both processors have a cache and use the MESI protocol. Initially, both caches have an invalid copy of the line. Depicts the consequence of a read of line x by Processor P1. If this is the start of a sequence of accesses, draw the subsequent figures for the following sequence: 1. P2 reads x. 2. P1 writes to x (for clarity, label the line in P1's cache 3. P1 writes to x (label the line in P1's cache 4. P2 reads x.
35	Write the sequence of micro-operations required for the bus structure of to add a number to the AC when the number is a. an immediate operand b. a direct-address operand c. an indirect-address operand
36	(i) Explain how a processor works with simple instruction sets. (RISC). (ii) How can you improve the performance of a processor with complex
37	Consider an SMP with both L1 and L2 caches using the MESI protocol. As explained, one of four states is associated with each line in the L2 cache. Are all four states also needed for each line in the L1 cache? If so, why? If not, explain which state or states can be eliminated.
38	Assume that propagation delay along the bus and through the ALU are 20 and 100 ns, respectively. The time required for a register to copy data from the bus is 10 ns. What is the time that must be allowed for a. transferring data from one register to another? b. incrementing the program counter?
39	Explain the various mapping techniques associated with cache memories. Explain how the virtual address is converted into real address in a paged virtual memory system.
40	Explain in detail about interface standards used in the design of a processor bus. Discuss in detail about the objectives and technical details of USB.
41	Write the sequence of micro-operations required for the bus structure of to add a number to the AC when the number is a. an immediate operand b. a direct-address operand c. an indirect-address operand
42	While browsing at Billy Bob's computer store, you overhear a customer asking Billy Bob what is the fastest computer in the store that he can buy. Billy Bob replies, "You're looking at our Macintoshes. The fastest Mac we have runs at a clock speed of 1.2 gigahertz. If you really want the fastest machine, you should buy our 2.4-gigahertz Intel Pentium IV instead." Is Billy Bob correct? What would you say to help this customer?
43	A processor accesses main memory with an average access time of T_2 . A smaller cache memory is interposed between the processor and main memory. The cache has a significantly faster access time of T_1 . The cache holds, at any time, copies of some main memory words and is designed so that the words more likely to be accessed in the near future are in the cache. Assume that the probability that the next word accessed by the processor is in the cache is H , known as the hit ratio. a. For any single memory access, what is the theoretical speedup of accessing the word in the cache rather than in main memory?

	b. Let T be the average access time. Express T as a function of T1, T2, and H. What is the overall speedup as a function of H?
44	Shows the state diagrams of two possible cache coherence protocols. Deduce and explain each protocol, and compare each to MESI.
45	a. On the IAS, what would the machine code instruction look like to load the contents of memory address 2? b. How many trips to memory does the CPU need to make to complete this instruction during the instruction cycle?
46	Let $A = A(1), A(2), \dots, A(1000)$ and $B = B(1), B(2), \dots, B(1000)$ be two vectors (one-dimensional arrays) comprising 1000 numbers each that are to be added to form an array C such that $C(I) = A(I) + B(I)$ for $I = 1, 2, \dots, 1000$. Using the IAS instruction set, write a program for this problem. Ignore the fact that the IAS was designed to have only 1000 words of storage.
47	Explain the single bus and three bus organization of the data path inside a processor with necessary diagrams. Explain the various design methods of hardwired control unit.
48	Can you foresee any problem with the write-once cache approach on bus-based multiprocessors? If so, suggest a solution.
49	(i) Explain the objectives and advantages of nanoprogramming. (ii) Narrate the sequence of events for an unconditional branch instruction.
50	(i) Explain how a read operation can be performed with data transfer signals in a PCI bus. (ii) Narrate the sequence of events caused by a SCSI controller.
51	Given the memory contents of the IAS computer shown below, Address Contents 08A 010FA210FB 08B 010FA0F08D 08C 020FA210FB show the assembly language code for the program, starting at address 08A. Explain what this program does.
52	Draw the typical block diagram of a DMA controller and explain how it is used for direct data transfer between memory and peripherals.
53	shows the state diagrams of two possible cache coherence protocols. Deduce and explain each protocol, and compare each to MESI.
54	Assume that propagation delay along the bus and through the ALU of Figure 15.6 are 20 and 100 ns, respectively. The time required for a register to copy data from the bus is 10 ns. What is the time that must be allowed for a. transferring data from one register to another? b. incrementing the program counter?

Class: II Year / IV Semester Computer Science and Engineering

Name of Subject: Database Management Systems

S.No	Assignment Topic
1.	What is the difference between a false positive and a false drop? If it is essential that no relevant information be missed by an information retrieval query, is it acceptable to have either false positives or false drops? Why?
2.	Explain the difference between the three storage types—volatile, nonvolatile, and stable—in terms of I/O cost.
3.	Explain why logical undo logging is used widely, whereas logical redo logging (other than physiological redo logging) is rarely used.

4.	List the physical storage media available on the computers you use routinely. Give the speed with which data can be accessed on each medium.
5.	Consider the relations <i>section</i> and <i>takes</i> . Give an example instance of these two relations, with three sections, each of which has five students. Give a file structure of these relations that uses multitable clustering.
6.	Using the university example, write relational-algebra queries to find the course sections taught by more than one instructor in the following ways: a. Using an aggregate function. b. Without using any aggregate functions.
7.	Physiological redo logging can reduce logging overheads significantly, especially with a slotted page record organization. Explain why.
8.	Give a tuple-relational-calculus expression to find the maximum value in relation $r(A)$.
9.	<p><u><i>employee</i> (person name, street, city)</u> <u><i>works</i> (person name, company name, salary)</u> <u><i>company</i> (company name, city)</u></p> <p>Consider the given relational database.</p> <p>Give an expression in the relational algebra to express each of the following queries:</p> <p>a. Find the names of all employees who work for “First Bank Corporation”.</p> <p>b. Find the names and cities of residence of all employees who work for “First Bank Corporation”.</p> <p>c. Find the names, street address, and cities of residence of all employees who work for “First Bank Corporation” and earn more than \$10,000.</p>
10.	Construct a decision-tree classifier with binary splits at each node, using tuples in relation $r(A, B, C)$ shown below as training data; attribute C denotes the class. Show the final tree, and with each node show the best split for each attribute along with its information gain value. (1, 2, a), (2, 1, a), (2, 5, b), (3, 3, b), (3, 6, b), (4, 5, b), (5, 5, c), (6, 3, b), (6, 7, c)
11.	List at least two reasons why database systems support data manipulation using a declarative query language such as SQL, instead of just providing a library of C or C++ functions to carry out data manipulation.
12.	Consider the following expressions, which use the result of a relational algebra operation as the input to another operation. For each expression, explain in words what the expression does.

	<p>a. $\sigma_{year \geq 2009}(takes) \bowtie student$</p> <p>b. $\sigma_{year \geq 2009}(takes \bowtie student)$</p> <p>c. $\Pi_{ID.name.course_id}(student \bowtie takes)$</p>
13.	<p>List two reasons why null values might be introduced into the database.</p> <p>Discuss the relative merits of procedural and nonprocedural languages.</p>
14.	<p>List some benefits and drawbacks of an anticipatory standard compared to a reactionary standard.</p>
15.	<p>Explain the reason for the use of degree-two consistency. What disadvantages does this approach have?</p>
16.	<p>Explain the difference between two-tier and three-tier architectures. Which is better suited for Web applications? Why?</p>
17.	<p>Explain what problems are caused by the design of the table?</p>
18.	<p>Suppose a JDO database had an object <i>A</i>, which references object <i>B</i>, which in turn references object <i>C</i>. Assume all objects are on disk initially. Suppose a program first dereferences <i>A</i>, then dereferences <i>B</i> by following the reference from <i>A</i>, and then finally dereferences <i>C</i>. Show the objects that are represented in memory after each dereference, along with their state (hollow or filled, and values in their reference fields).</p>
19.	<p>Describe how the ideas behind the RAID organization can be used in a broadcast-data environment, where there may occasionally be noise that prevents reception of part of the data being transmitted.</p>
20.	<p>a. What are the three broad levels at which a database system can be tuned to improve performance?</p> <p>b. Give two examples of how tuning can be done for each of the levels.</p>
21.	<p>Define a model of repeatedly broadcast data in which the broadcast medium is modeled as a virtual disk. Describe how access time and data-transfer rate for this virtual disk differ from the corresponding values for a typical hard disk</p>
22.	<p>Consider two-dimensional vector data where the data items do not overlap. Is it possible to convert such vector data to raster data? If so, what are the drawbacks of storing raster data obtained by such conversion, instead of the original vector data?</p>
23.	<p>Consider a database schema with a relation <i>Emp</i> whose attributes are as shown below, with types specified for multivalued attributes.</p> <p><i>Emp</i> = (<i>ename</i>, <i>ChildrenSet multiset(Children)</i>, <i>SkillSet multiset(Skills)</i>)</p> <p><i>Children</i> = (<i>name</i>, <i>birthday</i>)</p> <p><i>Skills</i> = (<i>type</i>, <i>ExamSet setof(Exams)</i>)</p>

	<p>$Exams = (year, city)$</p> <p>a. Define the above schema in SQL, with appropriate types for each attribute.</p> <p>b. Using the above schema, write the following queries in SQL.</p> <p>i. Find the names of all employees who have a child born on or after January 1, 2000.</p> <p>ii. Find those employees who took an examination for the skill type “typing” in the city “Dayton”.</p> <p>iii. List all skill types in the relation <i>Emp</i>.</p>
24.	Explain how a TP monitor manages memory and processor resources more effectively than a typical operating system.
25.	Compare the use of embedded SQL with the use in SQL of functions defined in a general-purpose programming language. Under what circumstances would you use each of these features?
26.	Describe at least 3 tables that might be used to store information in a social-networking system such as Facebook.
27.	Suppose there are two relations r and s , such that the foreign key B of r references the primary key A of s . Describe how the trigger mechanism can be used to implement the on delete cascade option, when a tuple is deleted from s .
28.	Explain why the nested-loops join algorithm would work poorly on database stored in a column-oriented manner. Describe an alternative algorithm that would work better and explain why your solution is better.
29.	The execution of a trigger can cause another action to be triggered. Most database systems place a limit on how deep the nesting can be. Explain why they might place such a limit.
30.	Explain the concept of physical data independence, and its importance in database systems.
31.	Explain how a TP monitor manages memory and processor resources more effectively than a typical operating system.
32.	Instead of storing shared structures in shared memory, an alternative architecture would be to store them in the local memory of a special process, and access the shared data by interprocess communication with the process. What would be the drawback of such architecture?
33.	Will functional dependencies be preserved if a relation is converted to a temporal relation by adding a time attribute? How is the problem handled in a temporal database?
34.	Explain why it may be impractical to require serializability for longduration transactions.
35.	Discuss the modifications that need to be made in each of the recovery schemes covered in Chapter 16 if we allow nested transactions. Also, explain any differences

	that result if we allow multilevel transactions.
36.	Explain how to use a histogram to estimate the size of a selection of the form $A \leq v(r)$.
37.	Create a small example of four small documents, each with a PageRank, and create inverted lists for the documents sorted by the PageRank. You do not need to compute PageRank, just assume some values for each page.
38.	Consider two-dimensional vector data where the data items do not overlap. Is it possible to convert such vector data to raster data? If so, what are the drawbacks of storing raster data obtained by such conversion, instead of the original vector data?
39.	What is lock de-escalation, and under what conditions is it required? Why is it not required if the unit of data shipping is an item?
40.	<p>A car-rental company maintains a database for all vehicles in its current fleet. For all vehicles, it includes the vehicle identification number, license number, manufacturer, model, date of purchase, and color. Special data are included for certain types of vehicles:</p> <ul style="list-style-type: none"> • Trucks: cargo capacity. • Sports cars: horsepower, renter age requirement. • Vans: number of passengers. • Off-road vehicles: ground clearance, drivetrain (four- or two-wheel drive). <p>Construct an SQL schema definition for this database. Use inheritance where appropriate.</p>
41.	Memory systems can be divided into multiple modules, each of which can be serving a separate request at a given time. What impact would such a memory architecture have on the number of processors that can be supported in a shared-memory system?
42.	Explain the notions of transparency and autonomy. Why are these notions desirable from a human-factors standpoint?
43.	<p><u>branch(branch name, branch city, assets)</u></p> <p><u>customer (customer name, customer street, customer city)</u></p> <p><u>loan (loan number, branch name, amount)</u></p> <p><u>borrower (customer name, loan number)</u></p> <p><u>account (account number, branch name, balance)</u></p> <p><u>depositor (customer name, account number)</u></p> <p>a. Find all loan numbers with a loan value greater than \$10,000.</p> <p>b. Find the names of all depositors who have an account with a value greater than \$6,000.</p> <p>c. Find the names of all depositors who have an account with a value greater than \$6,000 at the “Uptown” branch.</p>

44.	Study and summarize the facilities that the database system you are using provides for dealing with inconsistent states that can be reached with lazy propagation of updates.
45.	Describe how LDAP can be used to provide multiple hierarchical views of data, without replicating the base-level data.
46.	How does the concept of an object in the object-oriented model differ from the concept of an entity in the entity-relationship model?
47.	Explain the distinction between a type x and a reference type $\mathbf{ref}(x)$. Under what circumstances would you choose to use a reference type?
48.	<p>Suppose that you have been hired as a consultant to choose a database system for your client's application. For each of the following applications, state what type of database system (relational, persistent programming language-based OODB, object relational; do not specify a commercial product) you would recommend. Justify your recommendation.</p> <p>a. A computer-aided design system for a manufacturer of airplanes. b. A system to track contributions made to candidates for public office. c. An information system to support the making of movies.</p>
49.	List five responsibilities of a database-management system. For each responsibility, explain the problems that would arise if the responsibility were not discharged.
50.	<p><u>branch(branch name, branch city, assets)</u> <u>customer (customer name, customer street, customer city)</u> <u>loan (loan number, branch name, amount)</u> <u>borrower (customer name, loan number)</u> <u>account (account number, branch name, balance)</u> <u>depositor (customer name, account number)</u></p> <p>Consider the given bank database.</p> <p>Give an expression in the relational algebra for each of the following queries.</p> <p>a. Find the names of all branches located in "Chicago". b. Find the names of all borrowers who have a loan in branch "Downtown".</p>
51.	The Google search engine provides a feature whereby Web sites can display advertisements supplied by Google. The advertisements supplied are based on the contents of the page. Suggest how Google might choose which advertisements to supply for a page, given the page contents.
52.	What benefit does rigorous two-phase locking provide? How does it compare with other forms of two-phase locking?

53.	Show that the two-phase locking protocol ensures conflict serializability, and that transactions can be serialized according to their lock points.
54.	Show that there are schedules that are possible under the two-phase locking protocol, but are not possible under the timestamp protocol, and vice versa.

Class: II Year / IV Semester Computer Science and Engineering
Name of Subject: Probability and Queuing Theory

Sl. No.	Assignment Topic
1.	Two random variables have the joint pdf $f(x, y) = \begin{cases} \frac{1}{3}(x + y); & 0 \leq x \leq 1 \\ 0 & \text{elsewhere} \end{cases}$. Find the correlation coefficient and regression lines.
2.	The number of monthly breakdowns of a computer is a random variable having a poisson distribution with mean equal to 1.8. Find the probability that this computer will function for a month (1) without a breakdown (2) with only one breakdown.
3.	Two chips numbered 1 to 10 are mixed in a bowl. Two chips are drawn from the bowl successively and without replacement. What is the probability that their sum is 10?
4.	A committee of 5 persons is to be selected randomly from a group of 5 men and 10 women. (1) Find the probability that the committee consists of 2 men and 3 women. (2) Find the probability that the committee consists of all women.
5.	Let X be the random variable denotes the outcome of the roll of a fair die. Compute the mean and variance of X .
6.	Two fair dice are tossed. Find the probability of each of the following events: (1) The sum of the outcomes of the two dice is equal to 7. (2) The sum of the outcomes of the two dice is equal to 7 or 11 (3) The outcome of second die is greater than the outcome of the first die. Both dice come up with even numbers.
7.	If the probability mass function of a random variable X is given by $P(X = x) = kx^3, x = 1, 2, 3, 4$, find the value of $k, P((1/2 < X < 5/2) / X < 1)$, mean and variance of X .
8.	The joint probability mass function of (X, Y) is given by $p(x, y) = \frac{1}{72}(2x + 3y), x = 0, 1, 2$ and $y = 1, 2, 3$. Find all the marginal and conditional probability functions of X and Y .
9.	A random variable X is uniformly distributed over $(0, 10)$. Find (1) $P(X < 3), P(X > 7)$ and $P(X < 5)$ (2) $P(X = 7)$.
10.	The joint pdf of the random variable (X, Y) is $f(x, y) = 3(x + y), 0 < x < 1, 0 < y < 1, x + y \leq 1$. Find $Cov(X, Y)$.
11.	Let X be the random variable denotes the outcome of the roll of a fair die. Compute the mean and variance of X .

12.	The pdf of a random variable X is given by $f(x) = \begin{cases} 2x, & 0 \leq x \leq b \\ 0, & \text{otherwise} \end{cases}$. For what value of b is $f(x)$ a valid pdf? Also find the cdf of the random variable X with the above pdf.
13.	A continuous random variable has PDF $f(x) = Kx^2e^{-x}$, $x \geq 0$. Find K , r^{th} moment, mean, and variance.
14.	The two lines of regression are $8X - 10Y + 66 = 0$, $40X - 18Y - 214 = 0$. The variance of X is 9. Find the mean values of X and Y . Also find the coefficient of correlation between the variables X and Y .
15.	The pdf of a random variable X is given by $f(x) = \begin{cases} 2x, & 0 \leq x \leq b \\ 0, & \text{otherwise} \end{cases}$. For what value of b is $f(x)$ a valid pdf? Also find the cdf of the random variable X with the above pdf.
16.	Out of 1000 balls, 50 are red and the rest white. If 60 balls are picked at random, what is the probability of picking up (1) 3 red balls (2) not more than 3 red balls in the sample?
17.	The length of time a person speaks over phone follows exponential distribution with mean 6 mins. What is the probability that the person will talk for (1) more than 8 mins, (2) between 4 and 8 mins.
18.	Out of 800 families with 4 children each, how many families would be expected to have (i) Two boys and 2 girls (ii) atleast 1 boy (iii) atmost 2 girls (iv) children of both genders. Assume equal probability for both genders.
19.	Find the MGF of a random variable X whose density function is given by $f(x) = \lambda e^{-\lambda(x-a)}$, $x \geq a$. Hence find its mean and variance.
20.	A random variable X has the pdf $f(x) = \begin{cases} 2e^{-2x}, & x \geq 0 \\ 0, & x < 0 \end{cases}$ obtain the mgf and first four moments about the origin. Find the mean and variance of the same.
21.	The time in hours required to repair a machine is exponentially distributed with perimeter $\lambda = \frac{1}{2}$. (1) What is the probability that the repair time exceeds 2h. (2) What is the conditional probability that a repair takes atleast 10h given that its duration exceeds 9h?
22.	In a partially destroyed laboratory record only the lines of regressions and variance of x are available. The regression equations are $8x - 10y + 66 = 0$ and $40x - 18y = 214$ and variance of $x = 9$. Find i. The correlation coefficient between x and y ii. Mean values of x and y
23.	If the joint density function of X and Y is given by $f(x, y) = \begin{cases} x^2 + \frac{xy}{8}; & 0 < x < 1; 0 < y < 2 \\ 0 & \text{otherwise} \end{cases}$. Find (i) $P(X > \frac{1}{2})$ (ii) $P(Y < X)$ (iii) $P(X + Y \geq 1)$ and (iv) Find the conditional density functions.
24.	The number of typing mistakes that a typist makes on a given page has a Poisson distribution with a mean of 3 mistakes. What is the probability that she makes (1) Exactly 2 mistakes (2) Fewer than 4 mistakes (3) no mistakes on a given page.

25.	If the probability that a target is destroyed on any shot is 0.5. What is the probability that it would be destroyed on 6 th attempt?
26.	Three balls are drawn at random without replacement from a box containing 2 white, 3 red and 4 black balls. If x denotes the number of white balls drawn and y denotes the number of red balls drawn, find the joint probability distribution of (x, y) .
27.	The joint density of X and Y is given by $p(x, y) = k(2x + 3y)$, $x = 0, 1, 2, y = 1, 2, 3$. Find all the marginal and conditional distribution. Also find the probability distribution of $X + Y$.
28.	A coin is biased so that a head is twice as likely to appear as a tail. If the coin is tossed 6 times, find the probabilities of getting (i) exactly two heads (ii) at least 3 heads (iii) at most 4 heads.
29.	The joint pdf of (X, Y) is $f(x, y) = e^{-(x+y)}$, $x, y \geq 0$. Are X and Y independent?
30.	Bill and Jim like to shoot at targets. Jim can hit a target with a probability of 0.8, while Bill can hit a target with a probability of 0.7. If both fire at the target at the same time, what is the probability that the target is hit at least once.
31.	Two fair dice are tossed. Find the probability of each of the following events: (1) The sum of the outcomes of the two dice is equal to 7. (2) The sum of the outcomes of the two dice is equal to 7 or 11 (3) The outcome of second die is greater than the outcome of the first die. Both dice come up with even numbers.
32.	A radar system has a probability of 0.1 of detecting a certain target during a single scan. Find the probability that the target will be detected (1) at least two times in four consecutive scans (2) at least once in 20 consecutive scans.
33.	Out of 1000 balls, 50 are red and the rest white. If 60 balls are picked at random, what is the probability of picking up (1) 3 red balls (2) not more than 3 red balls in the sample? 1.173
34.	Four persons are chosen at random from a group containing 3 men, 2 women and 4 children. Show that the chance that exactly two of them will be children is $\frac{10}{21}$.
35.	If X and Y are independent random variables each normally distributed with mean zero and variance σ^2 , find the pdf of $R = \sqrt{X^2 + Y^2}$ and $\theta = \tan^{-1}\left(\frac{Y}{X}\right)$.
36.	A car hire firm has 2 cars. The number of demands for a car on each day is distributed as poisson variate with mean 0.5. Calculate the proportion of days on which (1) neither car is used (2) some demand is refused.
37.	If the probability that a target is destroyed on any shot is 0.5. What is the probability that it would be destroyed on 6 th attempt?
38.	Three balls are drawn at random without replacement from a box containing 2 white, 3 red and 4 black balls. If x denotes the number of white balls drawn and y denotes the number of red balls drawn, find the joint probability distribution of (x, y) .
39.	The joint density of X and Y is given by $p(x, y) = k(2x + 3y)$, $x = 0, 1, 2, y = 1, 2, 3$. Find all the marginal and conditional distribution. Also find the probability distribution of $X + Y$.
40.	A coin is biased so that a head is twice as likely to appear as a tail. If the coin is tossed 6 times, find the probabilities of getting (i) exactly two heads (ii) at least 3 heads (iii) at most 4 heads.
41.	The joint pdf of (X, Y) is $f(x, y) = e^{-(x+y)}$, $x, y \geq 0$. Are X and Y independent?
42.	A and B shoot independently until each has his own target. The probability of their hitting the target at each shot is $\frac{3}{5}$ and $\frac{5}{7}$ respectively? find the probability that B will require more shots than A.

43.	The joint probability mass function of (X, Y) is given by $p(x, y) = \frac{1}{27}(x + 7y)$, $x = 0, 1, 2$ and $y = 1, 2, 3$. Obtain $f(y/x)$ and $f(x/y = 1)$.																						
44.	The joint PDF of (X, Y) is given by $f(x, y) = \begin{cases} 24xy; & x > 0, y > 0, x + y \leq 1 \\ 0 & \text{elsewhere} \end{cases}$. Find the conditional mean and variance of Y given X .																						
45.	The joint pdf of the random variable (X, Y) iff. $f(x, y) = (x + y)$, $0 \leq x \leq 1, 0 \leq y \leq 1$, find $cov(X, Y)$.																						
46.	In a certain city, the daily consumption of electric power in millions of Kilowatt – hours can be considered as a random variable following gamma distribution with parameters $\lambda = \frac{1}{2}$ and $\alpha = 3$. If the power plant in this city has a daily capacity of 12 million Kilowatt hours, what is the probability that this supply of power will be insufficient on any given day?																						
47.	Marks obtained by 10 students in Mathematics (x) and Statistics (y) are given below. <table border="1" data-bbox="311 491 1482 647"> <tbody> <tr> <td>x</td> <td>60</td> <td>34</td> <td>40</td> <td>50</td> <td>45</td> <td>40</td> <td>22</td> <td>43</td> <td>42</td> <td>64</td> </tr> <tr> <td>y</td> <td>75</td> <td>32</td> <td>33</td> <td>40</td> <td>45</td> <td>33</td> <td>12</td> <td>30</td> <td>34</td> <td>51</td> </tr> </tbody> </table> <p>Also find y when $x = 55$</p>	x	60	34	40	50	45	40	22	43	42	64	y	75	32	33	40	45	33	12	30	34	51
x	60	34	40	50	45	40	22	43	42	64													
y	75	32	33	40	45	33	12	30	34	51													
48.	The life time of particular variety of electric bulbs may be considered as a random variable with mean 1200 hours and standard deviation 250 hours. Using central limit theorem find the probability that the average life time of 60 bulbs exceed 1250 hours.																						
49.	The number of monthly breakdowns of a computer is a random variable having a poisson distribution with mean equal to 1.8. Find the probability that this computer will function for a month (1) without a breakdown (2) with only one breakdown.																						
50.	Find the MGF of a random variable X whose density function is given by $f(x) = \lambda e^{-\lambda(x-a)}$, $x \geq a$. Hence find its mean and variance.																						
51.	Let X be the random variable denotes the outcome of the roll of a fair die. Compute the mean and variance of X .																						
52.	If the probability that a target is destroyed on any shot is 0.5. What is the probability that it would be destroyed on 6 th attempt?																						
53.	The joint density of X and Y is given by $p(x, y) = k(2x + 3y)$, $x = 0, 1, 2, y = 1, 2, 3$. Find all the marginal and conditional distribution. Also find the probability distribution of $X + Y$.																						
54.	A coin is biased so that a head is twice as likely to appear as a tail. If the coin is tossed 6 times, find the probabilities of getting (i) exactly two heads (ii) atleast 3 heads (iii) atmost 4 heads.																						

Class: II Year / IV Semester Computer Science and Engineering**Name of Subject: Design and Analysis of Algorithm**

Sl. No.	Assignment Questions
1	Explain how merge sort works? Sort the following sequence: { 17, 28, 31, 35, 65, 25, 42, 86, 45, 52 }. Write the Merge Sort Algorithm. Analyze the time complexity .
2	<p>Locker doors:</p> <p>There are n lockers in a hallway numbered sequentially from 1 to n. Initially, all the locker doors are closed. You make n passes by the lockers, each time starting with locker #1. On the ith pass, $i = 1, 2, \dots, n$, you toggle the door of every ith locker: if the door is closed, you open it, if it is open, you close it.</p> <p>For example, after the first pass every door is open; on the second pass you only toggle the even-numbered lockers (#2, #4, ...) so that after the second pass the even doors are closed and the odd ones are opened; the third time through you close the door of locker #3 (opened from the first pass), open the door of locker #6 (closed from the second pass), and so on. After the last pass, which locker doors are open and which are closed? How many of them are open?</p>
3	Consider the problem of scheduling n jobs of known durations t_1, \dots, t_n for execution by a single processor. The jobs can be executed in any order, one job at a time. You want to find a schedule that minimizes the total time spent by all the jobs in the system. (The time spent by one job in the system is the sum of the time spent by this job in waiting plus the time spent on its execution.)
4	Design a dynamic programming algorithm for the change-making problem: given an amount n and unlimited quantities of coins of each of the denominations d_1, d_2, \dots, d_m , find the smallest number of coins that add up to n or indicate that the problem does not have a solution.
5	How we can apply branch and bound technique in Flow-shop scheduling problems?
6	Write a report on the college admission problem (residents-hospitals assignment) that generalizes the stable marriage problem.
7	Consider the problem of the roommates, which is related to but more difficult than the stable marriage problem: “An even number of boys wish to divide up into pairs of roommates. A set of pairings is called stable if under it there are no two boys who are not roommates and who prefer each other to their actual roommates. Give an instance of this problem that does not have a stable pairing
8	Find a stable-marriage matching for the instance defined by the following ranking matrix:

	<table style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th></th> <th><i>A</i></th> <th><i>B</i></th> <th><i>C</i></th> <th><i>D</i></th> </tr> </thead> <tbody> <tr> <td>α</td> <td>1,3</td> <td>2,3</td> <td>3,2</td> <td>4,3</td> </tr> <tr> <td>β</td> <td>1,4</td> <td>4,1</td> <td>3,4</td> <td>2,2</td> </tr> <tr> <td>γ</td> <td>2,2</td> <td>1,4</td> <td>3,3</td> <td>4,1</td> </tr> <tr> <td>δ</td> <td>4,1</td> <td>2,2</td> <td>3,1</td> <td>1,4</td> </tr> </tbody> </table>		<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	α	1,3	2,3	3,2	4,3	β	1,4	4,1	3,4	2,2	γ	2,2	1,4	3,3	4,1	δ	4,1	2,2	3,1	1,4
	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>																						
α	1,3	2,3	3,2	4,3																						
β	1,4	4,1	3,4	2,2																						
γ	2,2	1,4	3,3	4,1																						
δ	4,1	2,2	3,1	1,4																						
9	Card guessing Design a strategy that minimizes the expected number of questions asked in the following game. You have a deck of cards that consists of one ace of spades, two deuces of spades, three threes, and on up to nine nines, making 45 cards in all. Someone draws a card from the shuffled deck, which you have to identify by asking questions answerable with yes or no. Explain it with Huffman tree.																									
10	Discuss about Optimal Binary Search Trees with help of an example? Let $n=4$, and $(a_1, a_2, a_3, a_4) = (\text{do}, \text{if}, \text{int}, \text{while})$, $p(a:4)=(0.15,0.10,0.05,0.10,0.20)$ and $q(0:4) = 0.05,0.10,0.05,0.05,0.10)$, generate the OBST.																									
11	Write down driving directions for going from your school to your home with the precision required by Travelling salesman algorithm.																									
12	Glove selection: There are 22 gloves in a drawer: 5 pairs of red gloves,4 pairs of yellow, and 2 pairs of green. You select the gloves in the dark and can check them only after a selection has been made. What is the smallest number of gloves you need to select to have at least one matching pair in the best case? What is in the worst case?																									
13	Missing socks: Imagine that after washing 5 distinct pairs of socks, you discover that two socks are missing. Of course, you would like to have the largest number of complete pairs remaining. Thus, you are left with 4 complete pairs in the best-case scenario and with 3 complete pairs in the worst case. Assuming that the probability of disappearance for each of the 10 socks is the same, find the probability of the best-case scenario; the probability of the worst-case scenario; the number of pairs you should expect in the average case.																									
14	Explain how hashing can be applied to check whether all elements of a list are distinct. What is the time efficiency of this application?																									
15	A top-down 2-3-4 tree is a B-tree of order 4 with the following modification of the insert operation. Whenever a search for a leaf for a new key encounters a full node (i.e., a node with three keys), the node is split into two nodes by sending its middle key to the node's parent (or, if the full node happens to be the root, the new root for the middle key is created). Construct a top-down 2-3-4 tree by inserting the following list of keys in the initially empty tree: 10, 6, 15, 31, 20, 27, 50, 44, 18.																									
16	Explain Agricultural Land allocation using Dynamic Programming.																									
17	Estimate how many times faster it will be to find gcd (31415, 14142) by Euclid's algorithm compared with the algorithm based on checking consecutive integers from $\min\{m, n\}$ down to gcd(m, n).																									
18	Use the most appropriate notation among O , Θ , and Ω to indicate the time efficiency class of sequential search a. in the worst case. b. in the best case. c. in the average case.																									

19	Explain how to apply the local area network connection in rural areas?
20	Explain how pile of adverts can be selected to be played on air using garden city radio. Which technique is used?
21	Suggest how any sorting algorithm can be augmented in a way to make the best-case count of its key comparisons equal to just $n - 1$ (n is a list's size, of course). Do you think it would be a worthwhile addition to any sorting algorithm?
22	Invention of chess: According to a well-known legend, the game of chess was invented many centuries ago in northwestern India by a sage named B. When he took his invention to his king, the king liked the game so much that he offered the inventor any reward he wanted. B asked for some grain to be obtained as follows: just a single grain of wheat was to be placed on the first square of the chess board, two on the second, four on the third, eight on the fourth, and so on, until all 64 squares had been filled. What would the ultimate result of this algorithm have been?
23	Consider the scenario of a central information server in the realm of public railroad transport on wide-area networks. Such a system has to process a large number of on-line queries for optimal travel connections in real time. Which algorithm is applied for this method?
24	How Lossless compression is achieved using Huffman Tree?
25	Improve the implementation of the matrix multiplication algorithm by reducing the number of additions made by the algorithm. What effect will this change have on the algorithm's efficiency?
26	Door in a wall: You are facing a wall that stretches infinitely in both directions. There is a door in the wall, but you know neither how far away nor in which direction. You can see the door only when you are right next to it. Design an algorithm that enables you to reach the door by walking at most $O(n)$ steps where n is the (unknown to you) number of steps between your initial position and the door.
27	Consider the problem of counting, in a given text, the number of substrings that start with an A and end with a B. (For example, there are four such substrings in CABAAXBYA.) (a) Design a brute-force algorithm for this problem and determine its efficiency class (b) Design a more efficient algorithm for this problem.
28	Design a recursive algorithm for computing 2^n for any nonnegative integer n that is based on the formula: $2^n = 2^{n-1} + 2^{n-1}$.
29	Apply mergesort to sort the list E, X, A, M, P, L, E in alphabetical order.
30	One can implement mergesort without a recursion by starting with merging adjacent elements of a given array, then merging sorted pairs, and so on. Implement this bottom-up version of mergesort in the language of your choice.
31	Puzzle: A tromino is an L-shaped tile formed by adjacent 1-by-1 squares. The problem is to cover any $2n$ -by- $2n$ chessboard with one missing square (anywhere on the board) with trominoes. Trominoes should cover all the squares of the board except the missing one with no overlaps. Design a divide-and-conquer algorithm for this problem
32	How we can apply public Transport Planning System using Dijkstra algorithm?

33 Odd pie fight: There are $n \geq 3$ people positioned in a field (Euclidean plane) so that each has a unique nearest neighbor. Each person has a cream pie. At a signal, everybody hurls his or her pie at the nearest neighbor. Assuming that n is odd and that nobody can miss his or her target, true or false: There always remains at least one person not hit by a pie?

34 What is Travelling Salesperson Problem? Write about the branch and bound solution for TSP. Solve the following TSP:

Fig: Cost Adjacency Matrix

35 The Dutch flag problem is to rearrange an array of characters R, W, and B (red, white, and blue are the colors of the Dutch national flag) so that all the R's come first, the W's come next, and the B's come last. Design a linear in-place algorithm for this problem.

36 Nuts and bolts: You are given a collection of n bolts of different widths and n corresponding nuts. You are allowed to try a nut and bolt together, from which you can determine whether the nut is larger than the bolt, smaller than the bolt, or matches the bolt exactly. However, there is no way to compare two nuts together or two bolts together. The problem is to match each bolt to its nut. Design an algorithm for this problem with average-case efficiency in $\Theta(n \log n)$.

37 Give an example of an array of n elements for which the sentinel mentioned in the text is actually needed. What should be its value? Also explain why a single sentinel suffices for any input.

38 Apply quick sort to sort the list E, X, A, M, P, L, E in alphabetical order. Draw the tree of the recursive calls made.

39 Traverse the following binary tree
 a. in preorder. b. in inorder. c. in postorder.

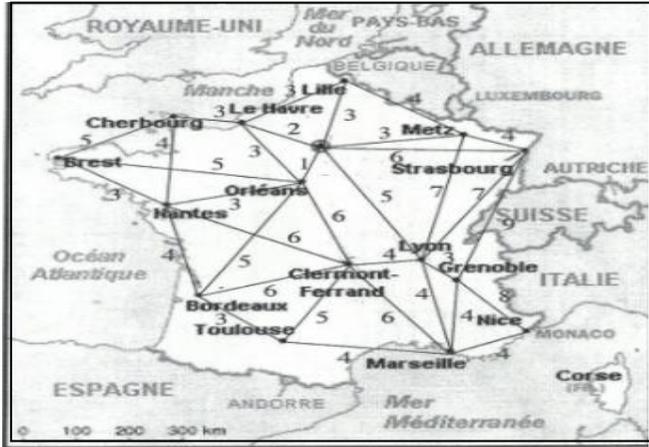
40 Chocolate bar puzzle: Given an n -by- m chocolate bar, you need to break it into nm 1-by-1 pieces. You can break a bar only in a straight line, and only one bar can be broken at a time. Design an algorithm that solves the problem with the minimum number of bar breaks. What is this minimum number? Justify your answer by using properties of a binary tree.

41 Shortest path around There is a fenced area in the two-dimensional Euclidean plane in the shape of a convex polygon with vertices at points $P_1(x_1, y_1)$, $P_2(x_2, y_2)$, ..., $P_n(x_n, y_n)$ (not necessarily in this order). There are two more points, $A(x_A, y_A)$ and $B(x_B, y_B)$, such that $x_A < \min\{x_1, x_2, \dots, x_n\}$ and $x_B > \max\{x_1, x_2, \dots, x_n\}$. Design a reasonably efficient algorithm for computing the length of the shortest path between A and B.

42 Alternating glasses: There are $2n$ glasses standing next to each other in a row, the first n of them filled with a soda drink while the remaining n glasses are empty. Make the glasses alternate in a filled-empty-filled-empty pattern in the minimum number of glass moves.

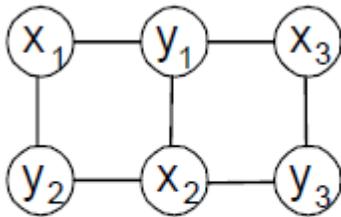
43 Apply insertion sort to sort the list E, X, A, M, P, L, E in alphabetical order.

44 How to connect the main cities of France in order to minimize the cost of project using Prim's algorithm.

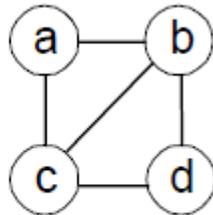


45 Apply shellsort to the list
S, H, E, L, L, S, O, R, T, I, S, U, S, E, F, U, L

46 A graph is said to be bipartite if all its vertices can be partitioned into two disjoint subsets X and Y so that every edge connects a vertex in X with a vertex in Y . (We can also say that a graph is bipartite if its vertices can be colored in two colors so that every edge has its vertices colored in different colors; such graphs are also called 2-colorable). For example, graph (i) is bipartite while graph (ii) is not.



(i)



(ii)

	<ul style="list-style-type: none"> a. Design a DFS-based algorithm for checking whether a graph is bipartite. b. Design a BFS-based algorithm for checking whether a graph is bipartite.
47	Celebrity problem: A celebrity among a group of n people is a person who knows nobody but is known by everybody else. The task is to identify a celebrity by only asking questions to people of the form: "Do you know him/her?" Design an efficient algorithm to identify a celebrity or determine that the group has no such person. How many questions does your algorithm need in the worst case?
48	Construct a 2-3 tree for the list C, O, M, P, U, T, I, N, G.
49	The birthday paradox asks how many people should be in a room so that the chances are better than even that two of them will have the same birthday (month and day). Find the quite unexpected answer to this problem. What implication for hashing does this result have?
50	Puzzle: A tromino is an L-shaped tile formed by adjacent 1-by-1 squares. The problem is to cover any $2n$ -by- $2n$ chessboard with one missing square (anywhere on the board) with trominoes. Trominoes should cover all the squares of the board except the missing one with no overlaps. Design a divide-and-conquer algorithm for this problem