Where an algorithm is asked for, you may write in any suitable pseudocode. Correct syntax for any computer language is not expected.

Answer 3 questions.

1. '*RPC is intrinsically inferior to distributed object-based communication.*' Discuss. [33]

2.
   a)
      i) What forms of name are common in distributed systems? [3]

      ii) Say what location mechanisms might be appropriate for each and outline the advantages and disadvantages of each of the approaches you have described. [6]

   b) In a particular distributed system, objects are named using UIDs. These objects may be saved onto floppy disks which can then be extracted from the system and replaced some time later in a machine anywhere in the world and restored onto that machine.

      i) Assuming that I only have the UID of an object, devise a scheme by which I can locate an object wherever it happens to be in the world (assuming, of course, that it is resident on some machine). [17]

      ii) What are the overheads of your scheme in terms of the number of messages exchanged and the amount of storage used? [2]

      iii) How fault tolerant is your scheme? [5]

3.
   a) Briefly define the terms Secret Key Cryptography and Public Key Cryptography, explaining the important characteristics of each. [8]

   b) Show how secret key cryptography can be used to establish secure communication channel between a client and a server in two different security domains. [16]

   c) State the security threats your solution is intended to address and explain how this is accomplished with your solution. [9]

[TURN OVER]

4. UniCoL is a medium sized London-based company, which uses computers quite extensively, but employs outside contractors to write and maintain software for them. They have just taken over two other companies in different UK cities and would like to have distributed stock control and order processing applications in order to increase efficiency.

To this end, they have just had a presentation from the ACME consultancy sales team. The substance of the presentation was that, although distributed systems are complex, buying the ACME transactional system will make programming as simple as it is in a centralised system.

The MD of UniCoL has a few questions for you, an independent consultant:

a) The presentation stated that distributed systems are complex, but failed to explain why. What is the cause of the complexity? [10]

b) What precisely is a transaction? [4]

c) To what extent can a transactional system really simplify distributed systems programming and to what extent are the problems fundamental? [12]

d) The links between the different companies are known to be flaky and disconnections of up to 3 hours are relatively common. What problems, if any, could this cause if UniCoL adopted a transactional system, given that they would like the system to continue to work insofar as is possible even when there is a link down. [7]

[CONTINUED]

5.

a) '*Distributed systems are more reliable than centralised systems.*' Under what conditions is this statement true? Illustrate your answer with examples. [8]

b) In the following configuration, all the nodes have a single replica and are equal with respect to machine type, processing power, and loading. Would one therefore assign a single vote to each replica? Justify your answer. [6]

c)

i) Explain how the voting structure should be modified after a failure or partition to ensure that the availability of the system is retained. [6]

ii) Say why the approach you have chosen works. You should demonstrate that if any single node fails, then all processing could continue in the partition containing the other replicas. [9]

iii) What happens when a partition is repaired or a node recovers? [4]

[END OF PAPER]