2002 Paper 12 Question 1

Data Structures and Algorithms

You have available a 20 Gbyte disc on which you need to hold an indexed sequential file consisting of variable length records each having a 20 byte key. Records, including the key, are typically 500 bytes long but never exceed 1000 bytes. The total size of all the records is never more than 10 Gbytes.

(a) Suggest, in detail, how you would represent this file on the disc. You should choose an organisation that allows

(i) efficient insertion of new records,

(ii) efficient updating of existing records identified by key, and

(iii) efficient inspection of all records in key order.

[14 marks]

(b) If the total size of the database is 10 Gbytes, estimate, for your organisation of the file, how many disc transfers would be needed to access a record with a given key, and estimate how many transfers would be required to read the entire database in sequential order. [6 marks]