

2003 Paper 4 Question 3

Data Structures and Algorithms

An application requires a hash table to hold up to 10^6 key–value pairs where both the keys and the values are integers. Carefully describe the two possible implementations: (a) *open hashing* using linked lists of key–value pairs outside the table and (b) *closed hashing* in which all key–value pairs are held within the hash table. [10 marks]

Discuss how you would decide which method is most suitable for the given application. [4 marks]

Assuming the table has exactly 10^6 entries and that keys, values and list pointers are all of size 4 bytes, and that you have allocated a total of 16 million bytes for the table (and hash chains), estimate the expected number of key comparisons necessary to locate an existing entry for (a) the open hash table and (b) the closed hash table. [6 marks]