2005 Paper 6 Question 8

Databases

(*a*) OLAP and OLTP.

    (*i*) What is *on-line transaction processing* (OLTP)? [2 marks]

    (*ii*) What is *on-line analytic processing* (OLAP)? [2 marks]

    (*iii*) If you were designing a relational database system, how would your approach to schema design differ for OLTP and OLAP systems?

    [3 marks]

    (*iv*) In OLAP, what is the meaning of the terms *drill down*, *roll up*, and *slice*?

    [3 marks]

    (*v*) What is a *star schema*? [1 mark]

(*b*) Suppose we have the following relational schema,

        Supplier(sid:integer, name:string, postcode:string)
        Parts(pid:integer, name:string, description:string)
        SuppliedBy(sid:integer, pid:integer, weight:integer)

where the underlined attributes represent the primary keys of the associated relation. The table *SuppliedBy* implements a relationship between suppliers and parts — indicating which parts are supplied by which supplier — using foreign keys pointing into the *Parts* and *Supplier* tables. The *weight* attribute is the parts weight in grams.

Write an SQL query that will return a list, without duplicates, of all postcodes associated with suppliers of parts less than one kilogram in weight. [5 marks]

(*c*) Define and explain the ACID properties of database transactions. [4 marks]