COMPUTER SCIENCE TRIPOS Part IA – 2024 – Paper 3

2 Databases (djg11)

(*a*) It is reckoned to be impossible to simultaneously provide the CAP trio in DBMS design: these are C---------y, A----------y and P-------- t-------e. Complete and define these terms and explain why. [4 marks]

(*b*) What is one significant difference between a reflexive relation in discrete maths and a database relation between a domain and itself? Give simple examples of both. Can either form of relation usefully have a one-to-one cardinality? [6 marks]

(*c*) An rDMBS holds tables with these four schemas:

R1:(A, B, C, D), R2:(A, B, C), R3:(A, B, D, E, F) and R4:(A, D).

You are told that values of F are always predictable from values of E, but it might be costly to make that prediction. Also, database updates might be much rarer than reads. What rearrangment of the schemas might be good and why? [5 marks]

(*d*) An `is_a` relation between entity types is defining a two-level hierarchy: for instance, lions and whales are both mammals. There are three basically different ways that a two-level `is_a` hierarchy can be modelled using relational tables. What are they? Multi-level `is_a` relationships commonly arise: which way then might be best? [5 marks]