

COMPUTER SCIENCE TRIPOS Part IA – 2026 – Paper 3

1 Databases (djg11)

- (a) Dozens of computers around the globe coordinate with each other to implement a distributed database.
- (i) Name two advantages and two disadvantages resulting from a distributed implementation. [4 marks]
 - (ii) Various policies are possible for deciding at which site(s) a data item is stored. Discuss two extreme possibilities and intermediate approaches. [4 marks]
 - (iii) How might law or security affect such distribution policies and would or should this influence the data model (i.e. the logical schema). [2 marks]
- (b) Consider a graph database that is subject to updates once a day. It stores data as nodes with attributes connected by arcs that possess further attributes.
- (i) Discuss whether the distinctive feature of a graph database is the underlying data model (i.e. the logical schema) or its query language, or anything else. [5 marks]
 - (ii) Often there are queries that must compare today's data with that from a month ago and compute metrics on the intermediate changes. Discuss how this might normally be solved and whether a graph database is intrinsically better for this sort of query. Consider using a special form of graph edge. [5 marks]