

COMPUTER SCIENCE TRIPOS Part IA – 2023 – Paper 3

2 Databases (djg11)

- (a) Is the relational database join operator associative? Describe an example where different associations might result in significant execution time differences. [4 marks]
- (b) Give an example relation where there is more than one potential key. Ensure your example has at least two reasons why some of these candidate keys are not suitable for use as the primary key. Explain the reasons. [4 marks]
- (c) In this question, an XML document is a tree with named nodes, called elements, whose leaves are character strings. In addition, an element has an unordered list of string pairs where the strings are an attribute name and its value.

```
xml_t = | LEAF of string
        | ELEMENT of string * (string * string) ulist * xml_t list
```

- (i) A document may be stored in XML in various ways, varying from rigidly structured to loosely structured. Describe when this can be useful. How can variations in structure be tolerated or reported? [5 marks]
- (ii) How might all the data held in a relational DBMS sensibly be exported into a single XML document? [4 marks]
- (d) What is the purpose of the project operator in the relational algebra? Would there be an equivalent operator in a document or graph database? [3 marks]