

NOTE 1: This is a closed book examination. For example, class text, copies of overhead slides and printed notes may not be used. There are 11 pages. The last page, only, may be separated and used as an easy reference for the second question. Answer all questions in the space provided.

NOTE 2: Some of the questions in this examination are open ended; however, they can be answered to the level discussed in class by short organized answers. It is recommended that you spend part of your time organizing your answer, rather than writing down ideas in the order they occur to you. The conciseness and organization of your answers will be taken into consideration in the grading.

NOTE 3: There are 100 marks in total. As a guide to managing your time, the marks awarded for each question are indicated in parenthesis at the start of the question.

NOTE 4: You are also being tested on your ability to understand the questions. In the case of a perceived ambiguity, state a clear assumption and proceed to answer the question.

NOTE 5: Cheating is an academic offence. Your signature on the signup sheet indicates that you understand and agree to the University's policies regarding cheating on exams.

I. (31 marks; continued on next two pages) General questions about database systems, the relational model and about the SQL language. Answer each of the following using no more than a few sentences in each case.

(a) Explain how physical data dependencies can increase the cost of maintaining an information system.

(b) What are the two general problems relating to data that are addressed by adopting the use of database technology when implementing an information system?

(c) Explain each of the following terms.

1. relational completeness

2. atomicity

3. domain dependence

(d) Describe the primary means of defining external views in the SQL language.

- (e) Assume relation R has numeric attributes {A, B, C}, with C as the primary key. Express the following query in the range restricted fragment of the relational calculus.

```
select distinct r1.A from R r1, R r2
where r1.C = r2.A
and not exists ( select * from R r3
                 where r3.b = r1.B or r3.B = R2.B )
```

- (f) What are three features of the SQL query language that make it more expressive than the relational calculus?

II.(45 marks; continued on next three pages) Consider the following relational database schema for maintaining customer rental information for a hypothetical car rental agency. (NOTE: this schema is reproduced on the final page of the exam.)

```
CREATE TABLE customer
(  cnum          INTEGER NOT NULL,
   cname         VARCHAR(20) NOT NULL,
   city          VARCHAR(20) NOT NULL,
   PRIMARY KEY (cnum) );
```

```
CREATE TABLE car
(  license       INTEGER NOT NULL,
   make          VARCHAR(20) NOT NULL,
   model         VARCHAR(20) NOT NULL,
   year          INTEGER NOT NULL,
   PRIMAY KEY (licence) );
```

```
CREATE TABLE pickup
(  rnum          INTEGER NOT NULL,
   cnum          INTEGER NOT NULL,
   license       INTEGER NOT NULL,
   fee           INTEGER NOT NULL,
   PRIMARY KEY (rnum),
   FOREIGN KEY (cnum) REFERENCES customer,
   FOREIGN KEY (license) REFERENCES car );
```

```
CREATE TABLE dropoff
(  rnum          INTEGER NOT NULL,
   PRIMARY KEY (rnum),
   FOREIGN KEY (rnum) REFERENCES pickup );
```

The database schema reflects two main events relating to car rentals: (a) a customer takes possession of a car at the start of a rental agreement (in this case, a tuple is added to the **pickup** table), and (b) a customer returns the car at the end of a rental agreement and pays the agreed rental fee (in this case, a tuple is added to the **dropoff** table). Thus, the database records information about both past and ongoing car rentals.

For each of the following two queries, *indicate if the query is a conjunctive query*, and then translate the query to *both the relational calculus and SQL*.

- (a) The number and name of each customer who has rented the same car at least twice in the past.

(b) The license and fee of cars that are currently rented by some customer in Waterloo.

Translate the following query to the SQL DDL only.

- (c) The make and models of cars that have generated the lowest received revenue, that is, for which the total rental fees for past rentals of cars of these makes and models is among the lowest.

III.(24 marks; continued on next page) Questions on the SQL standard, application development and ER modelling. Answer each of the following using no more than a few sentences in each case.

(a) What is the view update problem?

(b) Is it possible for an information system using a *call level interface* (CLI) to a SQL database to factor the overhead of compiling SQL queries? Justify your answer.

(c) Explain the main purpose of each of the following.

1. a host variable

2. a preprocessor for embedded SQL

3. the COMMIT WORK command

(d) Explain how many-to-one relationships can be captured by using general cardinality constraints.

The following is reproduced from Question II. You may detach this page to assist composing your answers to Question II. NOTE: do not write on this page.

```
CREATE TABLE customer
(  cnum          INTEGER NOT NULL,
   cname         VARCHAR(20) NOT NULL,
   city         VARCHAR(20) NOT NULL,
   PRIMARY KEY (cnum) );
```

```
CREATE TABLE car
(  license       INTEGER NOT NULL,
   make         VARCHAR(20) NOT NULL,
   model        VARCHAR(20) NOT NULL,
   year         INTEGER NOT NULL,
   PRIMAY KEY (licence) );
```

```
CREATE TABLE pickup
(  rnum          INTEGER NOT NULL,
   cnum          INTEGER NOT NULL,
   license       INTEGER NOT NULL,
   fee          INTEGER NOT NULL,
   PRIMARY KEY (rnum),
   FOREIGN KEY (cnum) REFERENCES customer,
   FOREIGN KEY (license) REFERENCES car );
```

```
CREATE TABLE dropoff
(  rnum          INTEGER NOT NULL,
   PRIMARY KEY (rnum),
   FOREIGN KEY (rnum) REFERENCES pickup );
```

The database schema reflects two main events relating to car rentals: (a) a customer takes possession of a car at the start of a rental agreement (in this case, a tuple is added to the pickup table), and (b) a customer returns the car at the end of a rental agreement and pays the agreed rental fee (in this case, a tuple is added to the dropoff table). Thus, the database records information about both past and ongoing car rentals.