OLLSCOIL NA GAILLIMHE
UNIVERSITY OF GALWAY

Semester 1 Examinations 2022-2023

Course Instance Code(s)	3BCT1
Exam(s)	Third B.Sc. Computer Science and IT
Module Code(s)	CT3532
Module(s)	Database Systems 2
Paper No.	1
External Examiner(s)	Dr. R. Trestian
Internal Examiner(s)	Professor M. Madden *Dr. C. O'Riordan

Instructions: Answer any 3 questions. All questions are equally weighted.

Duration	2 hours
No. of Pages	3
Discipline(s)	Computer Science
Course Co-ordinator(s)	Dr. C. O'Riordan

Requirements:

Release in Exam Venue	Yes
Handout	None
Statistical/ Log Tables	None
Cambridge Tables	None
Graph Paper	None
Log Graph Paper	None
Other Materials	None

PTO

CT3532 Database Systems 2

Exam Duration: 2 Hours

Question 1

(a) Outline an approach to generating a minimal cover set. Illustrate your approach by generating a minimal cover set for **F**, where **F** is:

$$\mathbf{F} = \{A \rightarrow B, ABCD \rightarrow E, EK \rightarrow GH, ABCK \rightarrow EG\}$$

(8)

(b) Given a relation R and a set of functional dependencies, outline an algorithm to generate a schema such that all the resulting relations satisfy BCNF. (8)

(c) Explain what is meant by the term *denormalisation*. Outline, with the aid of examples, two separate forms of denormalisation. Discuss scenarios where denormalisation can be used. (9)

Question 2

(a) In the context of concurrency control, explain with a suitable example what is meant by the *temporary update problem.* (5)

(b) Outline the two phase locking approach to concurrency control and show that it guarantees conflict serializability. (10)

(c) With respect to database recovery, what is meant by a *commit point*? In distributed databases, the database items are distributed across a number of sites with some items replicated across a number of site. Outline an approach that could be used to commit transactions in such a system. (10)

Question 3

(a) Outline an efficient algorithm for performing a join between two relations. Discuss the efficiency of your algorithm. (8)

(b) Explain how the join algorithm in (a) can be improved given a parallel architecture. Specify the improvement in efficiency. (8)

(c) Explain, with the use of suitable example, an efficient means to jointly index a number of attributes. Discuss the efficiency of your approach. (9)

Question 4

(a) A B+tree is a commonly used data structure used for efficient access to data. Given a B+tree built on some attribute a_i, write pseudo-code to return all values of a_i in some defined range. (7)

(b) In the context of parallel databases, compare round-robin and range partitioning techniques. Discuss the relative merits of these approaches for handling range queries. (8)

(c) Discuss the motivations for adopting a dynamic hashing approach. Describe, with the aid of an example, any approach to hashing to a dynamic file. (10)

END