

McMaster University
SFWRENG 3DB3 Fall 2024
Assignment 3
Due: December 4, 2024 at 10:00pm

November 15, 2024

I. Database Design (44 marks)

Question 1 (14 marks) Finding Keys

- (a) (12 marks) Consider a relation schema $R(A, B, C, D, E)$, and the set of functional dependencies

$$\mathbf{F} = \{A \rightarrow BC \\ CD \rightarrow E \\ B \rightarrow D \\ E \rightarrow A\}$$

Find all candidate keys (i.e., minimal keys) of relation R . Show all the steps you took to derive each key, and clearly state which of Armstrong's axioms are used in each step.

- (b) (2 marks) Is the FD $AB \rightarrow C$ entailed by \mathbf{F} ? Show your work that supports your answer.

Question 2 (10 marks) Minimal Cover

Given the relational schema $T(A, B, C, D)$ and the FDs: $F = \{ABC \rightarrow D, CD \rightarrow A, \text{ and } CA \rightarrow B, AD \rightarrow C, CD \rightarrow B\}$. Compute the minimal cover F' of F . Show all your work (derivation) to compute F' .

Question 3 (10 marks) Armstrong's Axioms

Prove the following using Armstrong's axioms (using only the axioms presented in class). Show the steps of your proof, and indicate which of Armstrong's axioms is applied in each step.

- a) (5 marks) Given the relational schema $R(A, B, C, D, E, F)$ and the FDs $F_1: \{AB \rightarrow C, A \rightarrow D, CD \rightarrow EF\}$. Show that $AB \rightarrow F$.
- b) (5 marks) Given the relational schema $R(A, B, C, D, E, F)$ and the FDs: $C \rightarrow D, BE \rightarrow A$, and $BEF \rightarrow C$. Show that BEF is a key.

Question 4 (10 marks) 3NF, BCNF

You have been hired to design the database for the *3DB3 Company*. After meeting with your client, you have agreed upon the schema below, consisting of three relations.

- *Company* (*companyID*, *companyName*, *cityName*, *country*, *assets*)
These represent, respectively, the ID of company, its name, the city and country where it is located, and the assets of the company.
- *Department* (*deptID*, *deptName*, *companyID*, *cityName*, *country*, *deptMgrID*)
These represent, respectively, the ID of the department, its name, the *companyID* it belongs to, the city and country it's located in, and the ID of the department manager.
- *City* (*cityID*, *cityName*, *country*)
The ID of the city, its name, and the country where it is located.

However, your design is not yet complete, since you must still represent the following information about *3DB3 Company*:

- Two different companies cannot have the same *company ID*.
- Two different departments cannot have the same *deptID*.
- Two different cities cannot have the same *cityID*.
- Two different cities in the same country cannot have the same name.
- The company name and the city it's located in determine the *company ID*.
- Two departments in the same company must have different names.
- One manager cannot run two different departments.

To complete your job, answer the following questions:

1. What are the functional dependencies (and keys) that can be inferred from this additional information? List all of them.
2. Your client will be satisfied if your design is a good one, i.e., the schema satisfies either the BCNF or the 3NF.
Is the design of your schema with the functional dependencies from part (1) above a good one? Justify your answer. If the design is not a good one, provide a better one, using one of the decomposition algorithms discussed in class, and explain how it is better.

II. Transactions and Concurrency (20 marks)

Question 5 (9 marks) Schedules

Consider schedules S_1 , S_2 , S_3 below. State which of the following properties hold (or not) for each schedule: strict, avoids cascading aborts, recoverability. Provide a brief justification for each answer.

(a) (3 marks)

S_1 : r1(X); r2(Z); r1(Z); r3(X); r3(Y); w1(X); c1; w3(Y); c3; r2(Y); w2(Z); w2(Y); c2

(b) (3 marks)

S_2 : r1(X); r2(Z); r1(Z); r3(X); r3(Y); w1(X); w3(Y); r2(Y); w2(Z); w2(Y); c1; c2; c3;

(c) (3 marks)

S_3 : r1(X); r2(Z); r3(X); r1(Z); r2(Y); r3(Y); w1(X); w2(Z); w3(Y); w2(Y); c3; c1; c2;

Question 7 (6 marks) Serializability

Which of the following schedules is (conflict) serializable? For each serializable schedule, find the equivalent serial schedules.

(a) (3 marks) r1(X); r3(X); w1(X); r2(X); w3(X)

(b) (3 marks) r3(X); r2(X); w3(X); r1(X); w1(X)

Question 8 (5 marks) Locking

Consider the following locking protocol: Before a transaction T writes a data object A, T has to obtain an exclusive lock on A. For a transaction T, we hold these exclusive locks until the end of the transaction. If a transaction T reads a data object A, no lock on A is obtained. State which of the following properties are ensured by this locking protocol: serializability, conflict-serializability, recoverability, avoids cascading aborts, avoids deadlock. Explain and justify your answer for each property.

Grading

This assignment is worth 14% towards your final grade.

Submission

All files are to be submitted using the Avenue system. Please ensure your answers are typed and clearly readable. Include your name and student ID number in the file. Submit your solutions to all questions in a file called **asg3.pdf**.