

Obligatory Exercise 4

Due date: October 6, 2009 at 0800 (8am)

1. Given is a relational database schema on ocean-going shipping. The schema consists of a single relation R, with attributes as follows:

$R(\text{Ship_ID}, \text{Ship_Name}, \text{Ship_Type}, \text{Ship_Country}, \text{Voyage_ID}, \text{Start_Date}, \text{End_Date}, \text{Cargo}, \text{Port}, \text{Dock_Date})$

The following functional dependencies hold.

$\{\text{Ship_ID}\} \rightarrow \{\text{Ship_Name}, \text{Ship_Type}, \text{Ship_Country}\}.$

$\{\text{Ship_Name}, \text{Ship_Country}\} \rightarrow \{\text{Ship_ID}\}.$

$\{\text{Voyage_ID}\} \rightarrow \{\text{Ship_ID}, \text{Cargo}, \text{Start_Date}, \text{End_Date}\}$

$\{\text{Ship_ID}, \text{Dock_Date}\} \rightarrow \{\text{Voyage_ID}, \text{Port}\}.$

- (a) Identify the candidate keys of this schema.
 - (b) Construct a lossless, dependency-preserving decomposition of this schema into 3NF. It is not necessary to use Algorithm 11.4 of the text (Algorithm 15.4 in the third edition), but it is essential to provide a clear justification as to why the answer which you provide is in 3NF.
 - (c) Construct a lossless decomposition of this schema into BCNF.
 - (d) Explain why there is no lossless, dependency-preserving decomposition of this schema into BCNF.
2. Given is the relational database schema $R(A,B,C,D,E,F)$, constrained by the functional dependency set $\mathcal{F} = \{AD \rightarrow F, CD \rightarrow EF, F \rightarrow AD, D \rightarrow AE\}$. Answer the following:

- (a) Determine a minimal cover for the set \mathcal{F} . It is not necessary to follow a formal algorithm, but a clear justification of why the answer is correct must be given.
- (b) Provide a 3NF decomposition, using Algorithm 11.2 of the textbook (15.1 in the third edition).
- (c) Indicate whether or not the decomposition of (b) is lossless. Justify your answer.
- (d) If your answer to (c) is no, then apply Algorithm 11.4 of the textbook (15.4 in the third edition) to obtain a lossless, dependency-preserving 3NF decomposition.
- (e) Indicate whether or not the decomposition of (d) is a BCNF decomposition. Justify your answer. (If the answer to (c) is yes, so no new decomposition is constructed for (d), then use the decomposition of part (b)).

5DV021, Obligatory Exercise 4, page 2

3. Given is a relational database schema $R(A, B, C, D)$ constrained by the functional dependencies in $\{A \rightarrow BC, C \rightarrow B, D \rightarrow B\}$. For each of the decompositions $\{R[ABC], R[AD]\}$, $\{R[AB], R[BC], R[CD]\}$, $\{R[AB], R[BCD]\}$, $\{R[AB], R[AC], R[DB]\}$, and $\{R[AB], R[AC], R[AD], R[DB]\}$, answer the following questions.

- State whether or not the decomposition is lossless.
- If the decomposition is lossless, state the rule which mandates it to be so.
- If the decomposition is not lossless, give a specific example of an instance r of $R[ABCD]$ whose projections, as defined by the decomposition, do not join together to obtain r . For example, if $\{R[AB], R[BC], R[CD]\}$ is not lossless, then provide an example r for which $\pi_{AB}(r) \bowtie \pi_{BC}(r) \bowtie \pi_{CD}(r)$ is not the same as r .
- State whether or not the decomposition is dependency preserving. If it is not dependency preserving, identify a dependency which is not recapturable from the decomposition.
- For each of 2NF, 3NF, and BCNF, state whether or not the decomposition is in the given normal form.

Notes:

- As stipulated in the course syllabus, this exercise may be done either individually, in a group of two, or in a group of three.
- Remember that there are point penalties for late submission. See the course syllabus.
- It is strongly recommended that you use a text-processing tool to display your results. If you write them by hand, they must be very neat.
- It is not allowed to copy the work of others. The submission must be the original work of the individual(s) in the working group.
- The grader reserves the right to interview members of the working group about the solution.
- So that solutions may be discussed in a class meeting, students and/or groups that are very late in preparing a solution may be required to solve an alternate problem to receive credit for this exercise.
- If you have solved some of these problems for a previous offering of the course, you may re-use your old solutions, subject to the following conditions: (a) You may not work with any partners, except possibly those with whom you worked to prepare the solution in the previous course. (b) You must explicitly note any partners from the previous course with whom you submitted a joint solution for that course. Note that grading criteria may not be identical between years, so that a solution which was found to be satisfactory last year may not be evaluated similarly this year. Note: You must solve the problems given above. Credit will not be given for solutions to different problems from previous years under any circumstances.
- Once a solution to this exercise is submitted, a carryover of Exercise 3 from last year is not permitted, either for points or for a satisfactory rating.

5DV021, Obligatory Exercise 4, page 3

- You are welcome to test your solutions using LDBN, but remember that you will not be able to use LDBN on the examination, and there will be similar questions, so it is important that you understand the underlying notions and algorithms.