

Exercise sheet for Session 1

Uncertain data management

Antoine Amarilli

November 30th, 2015

1 Exercise 1.

Consider a database schema that consists of the following tables:

- Class(**class**, **date**, **teacher**, **room**), indicating the planned classes
- Sick(**teacher**, **date**), indicating the dates at which a teacher is sick
- Unavail(**teacher**, **date**), indicating when a teacher is more generally unavailable
- Closed(**date**), indicating dates at which the entire school is closed
- Canceled(**class**, **date**), the occurrences of classes that have to be canceled.

Remember that a *tuple-generating dependency* is a rule of the form:

$$\forall \mathbf{x} \phi(\mathbf{x}) \Rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y})$$

where ϕ and ψ are conjunctions of atoms. An *inclusion dependency* is a tuple-generating dependency where ϕ and ψ consist of a single atom without repeated variables.

Question 1. Write tuple-generating dependencies to express the following:

- When the school is closed, all planned classes on that day are canceled
- When a teacher is sick on a day, then they are unavailable on that day
- When a teacher is unavailable on a day, all classes that they planned to give on that day are canceled

Question 2. Which of these tuple-generating dependencies are inclusion dependencies?

Question 3. Write a conjunctive query Q that asks which classes are canceled on November 30th. (Do not assume that the Canceled table only contains classes; only return answers that occur in the Class table.) Write it in the relational calculus, and in the relational algebra.

Question 4. Consider the database instance that contains the following facts:

- John is sick on November 30th
- The class with **id** “UDM” is taught by Antoine on November 30th in room C017
- The class with **id** “FOO” is taught by John on November 30th in room C42
- The class with **id** “UDM” is taught by Antoine on December 7th in room C47
- The school is closed on December 7th

Construct the chase of this instance by the dependencies of Question 1.

Question 5. Evaluate Q on the chase. What can we deduce from this?

Question 6. The chase in Question 4 was finite. Would the chase be finite for any database instance? Why, or why not?

Question 7. Rewrite the query Q (in the relational calculus) to a union of conjunctive queries Q' such that, for any instance, Q' holds on the instance iff Q is entailed by the instance and the dependencies of Question 1.

2 Exercise 2.

Consider a database schema that consists of the following tables:

- Jedi(**jedi**), indicating the list of known Jedis
- Teach(**master**, **padawan**), indicating which Jedi trained which Jedi
- Light(**jedi**), indicating which Jedis are on the light side of the force
- Dark(**jedi**), indicating which Jedis are on the dark side of the force

Question 1. Write tuple-generating dependencies Σ that express the following:

- Anybody on the light side of the force is a Jedi
- Likewise for anybody on the dark side of the force
- If a master teaches a padawan, then both are Jedis
- Every Jedi was taught by some master
- Whenever some padawan is on the light side of the force and was taught by a master, then the master is on the dark side of the force.
- Conversely, when a padawan is on the dark side of the force, any master is on the light side of the force.

Question 2. Which one of these dependencies are inclusion dependencies? Which ones can be rewritten to be inclusion dependencies?

Question 3. Consider the instance I where the Jedis are Obi-wan (light side) and Anakin (dark side), and the first taught the second. Is the chase of this instance by Σ finite? Why?

Is the chase accurate with respect to the Star Wars movies?

Question 4. Is there an instance I whose chase by Σ is finite?

Question 5. Write a conjunctive query Q that asks whether a dark Jedi trained a dark Jedi. Write it both in the relational algebra and in the relational calculus.

Question 6. Is Q entailed by the instance I and tuple-generating dependencies Σ ? Why (not)?

Question 7. Is there an instance I' where Q does not hold, but such that I' entails Q under Σ ?

Question 8. Write a formula in first-order logic that asserts that a Jedi either follows the light side or the dark side, but not both. Can this be expressed as a tuple-generating dependency?