

Exercise sheet for Session 3

Uncertain data management

Antoine Amarilli

December 5th, 2016

1 Exercise 1: Codd tables

Consider the following instance I :

Events				Class		Exam	
class	session	teacher	room	class	session	class	session
UDM	1	NULL	NULL	UDM	1	UDM	9
FOO	NULL	NULL	C42	FOO	NULL	FOO	NULL
UDM	9	NULL	NULL				
FOO	NULL	NULL	C43				

Question 1. Rephrase in plain English what we know about the UDM class, and what we know about the FOO class.

Answer. The UDM class has a normal session (“1”) and an exam session (“9”) but we do not know where it takes place or who teaches it.

The FOO class has one normal session and one exam session, and besides it has one session in C42 and one session in C43 (but we do not know who teaches them).

Question 2. Write a Boolean conjunctive query that asks whether some class has a scheduled session in the same room as the exam for that class. Write it in the relational algebra and in the relational calculus.

Answer. In the relational calculus:

$$Q : \exists c s t r s' t' Events(c, s, t, r) \wedge Events(c, s', t', r) \wedge Class(c, s) \wedge Exam(c, s')$$

In the relational algebra:

$$\Pi_{\emptyset} (Events \bowtie \rho_{\text{session} \rightarrow \text{s1}, \text{teacher} \rightarrow \text{t1}}(Events) \bowtie Class \bowtie \rho_{\text{session} \rightarrow \text{s1}}(Exam))$$

Another possibility in the relational algebra:

$$\Pi_{\emptyset} ((\Pi_{\text{class}, \text{session}} (Event \bowtie Class)) \bowtie (\Pi_{\text{class}, \text{session}} (Event \bowtie Exam)))$$

Question 3. Is this query *possible* on the instance? If yes, what is a witnessing possible world?

Answer. The query is possible and a witnessing possible world is obtained by mapping all NULLs to “a”.

Question 4. Is this query *certain* on the instance? If no, what is a counterexample possible world?

Answer. The query is not certain and a counterexample possible world is obtained by mapping all NULLs to different values.

Question 5. Write a tuple-generating dependency that says that whenever a class has a scheduled session then it has a scheduled exam with the same teacher (but possibly in a different room).

Answer.

$$\phi : \forall c s t r \text{ Events}(c, s, t, r) \wedge \text{Class}(c, s) \Rightarrow \exists s' r' \text{ Events}(c, s', t, r') \wedge \text{Exam}(c, s')$$

Question 6. Is there a possible world that satisfies this constraint? Do all possible worlds satisfy this constraint?

Answer. Use the same worlds as for Questions 3 and 4.

Question 7. Replace three pairs of NULLs by named NULLs to obtain a v-table where the constraint is always respected.

Answer. The following is a solution:

<i>Events</i>				<i>Class</i>		<i>Exam</i>	
<i>class</i>	<i>session</i>	<i>teacher</i>	<i>room</i>	<i>class</i>	<i>session</i>	<i>class</i>	<i>session</i>
<i>UDM</i>	<i>1</i>	<i>NULL₁</i>	<i>NULL</i>	<i>UDM</i>	<i>1</i>	<i>UDM</i>	<i>9</i>
<i>FOO</i>	<i>NULL</i>	<i>NULL₄</i>	<i>C42</i>	<i>FOO</i>	<i>NULL</i>	<i>FOO</i>	<i>NULL₃</i>
<i>UDM</i>	<i>9</i>	<i>NULL₁</i>	<i>NULL</i>				
<i>FOO</i>	<i>NULL₃</i>	<i>NULL₄</i>	<i>C43</i>				

2 Exercise 2: v-tables and c-tables

Consider the following instance of a v-table:

Events			
class	session	teacher	room
UDM	1	NULL ₁	NULL ₂
NULL ₃	2	NULL ₁	NULL ₄
FOO	1	John	NULL ₂

Question 1. What is this relation saying in plain English?

Answer. There is a class UDM that has a first session, and a class FOO that has a first session in the same room. Further some class has a second session and its teacher is the same as that of the first session of the UDM class.

Question 2. Write a query Q_1 in the relational algebra that returns the triples of a class *class* and two sessions $s_1 < s_2$ such that s_1 and s_2 are sessions of *class* that have the same teacher.

Answer.

$$\sigma_{s_1 < s_2} (\rho_{\text{session} \mapsto s_2} (\Pi_{\text{class, session, s1}} (\rho_{\text{room} \mapsto r_1, \text{session} \mapsto s_1} (\text{Events}) \bowtie \text{Events})))$$

Question 3. Evaluate Q_1 on the instance to obtain a c-table R_1 .

Answer. The result is:

<i>class</i>	<i>s1</i>	<i>s2</i>	
UDM	1	2	NULL ₃ = "UDM"
FOO	1	2	NULL ₃ = "FOO" \wedge NULL ₁ = "John"

Question 4. Write an analogous query Q_2 that returns the triples of a class and two sessions $s_1 < s_2$ of the class that take place in the same room (but may have different teachers).

Answer. Replace **room** \mapsto **r1** by **teacher** \mapsto **t1** in Q_1 .

Question 5. Evaluate Q_2 on the instance to obtain a c-table R_2 .

Answer. The result is:

<i>class</i>	<i>s1</i>	<i>s2</i>	
UDM	1	2	NULL ₃ = "UDM" \wedge NULL ₃ = NULL ₄
FOO	1	2	NULL ₃ = "FOO" \wedge NULL ₃ = NULL ₄

Question 6. Compute a c-table representation of the union R of R_1 and R_2 .

Answer. The result (after Boolean simplifications) is:

<i>class</i>	<i>s1</i>	<i>s2</i>	
<i>UDM</i>	<i>1</i>	<i>2</i>	$\text{NULL}_3 = \text{"UDM"}$
<i>FOO</i>	<i>1</i>	<i>2</i>	$\text{NULL}_3 = \text{"FOO"} \wedge (\text{NULL}_1 = \text{"John"} \vee \text{NULL}_3 = \text{NULL}_4)$

Question 7. How many rows may the possible worlds of R have?

Answer. The possible worlds may be empty (if NULL_3 is neither "FOO" nor "UDM"), or may contain a single tuple (for instance if NULL_3 is "UDM"). The possible worlds cannot contain two different tuples, because we cannot have both $\text{NULL}_3 = \text{"UDM"}$ and $\text{NULL}_3 = \text{"FOO"}$.

3 Exercise 3: Boolean c-tables

Consider the following instance:

Classes			
session	date	prof	room
2	Nov 28	Antoine	C017
3	Dec 5	Antoine	C47
4	Dec 12	Silviu	C47
5	Jan 9	Silviu	C47
6	Jan 16	Silviu	C47

Consider the following uncertain Boolean events:

- x_1 : Room C47 collapses. All UDM classes in room C47 must be canceled.
- x_2 : D&K students accept to return from vacation. If this does *not* happen, all UDM classes in January are cancelled.
- x_3 : Silviu wins the lottery and escapes to the Bahamas. All of Silviu's classes must be canceled.

Question 1. Annotate the rows of the instance to make a Boolean c-table that describes the correct outcome depending on the value of the events.

Answer.

Classes				
<i>session</i>	<i>date</i>	<i>prof</i>	<i>room</i>	
<i>2</i>	<i>Nov 28</i>	<i>Antoine</i>	<i>C017</i>	
<i>3</i>	<i>Dec 5</i>	<i>Antoine</i>	<i>C47</i>	$\neg x_1$
<i>4</i>	<i>Dec 12</i>	<i>Silviu</i>	<i>C47</i>	$\neg x_1 \wedge \neg x_3$
<i>5</i>	<i>Jan 9</i>	<i>Silviu</i>	<i>C47</i>	$\neg x_1 \wedge x_2 \wedge \neg x_3$
<i>6</i>	<i>Jan 16</i>	<i>Silviu</i>	<i>C47</i>	$\neg x_1 \wedge x_2 \wedge \neg x_3$

Question 2. How many possible worlds does the table have?

Answer. There are four possible worlds.

Optional justification: Indeed, we observe that sessions 5 and 6 have the same condition, tuple 2 is always present, and the following implications hold:

- the condition of tuple 5 implies that of tuple 4
- the condition of tuple 4 implies that of tuple 3
- the condition of tuple 3 implies that of tuple 2

Hence, for any possible world, its set of tuples is a prefix of the initial set, so it is uniquely defined by the last tuple of the world (which may be either 2, 3, 4, or 6).

Now, we observe that these four possible worlds can actually be obtained; for instance, we can use the following valuations:

- $\{x_1 \mapsto 1, x_2 \mapsto 0, x_3 \mapsto 0\}$;
- $\{x_1 \mapsto 0, x_2 \mapsto 0, x_3 \mapsto 1\}$;
- $\{x_1 \mapsto 0, x_2 \mapsto 0, x_3 \mapsto 0\}$;
- $\{x_1 \mapsto 0, x_2 \mapsto 1, x_3 \mapsto 0\}$;

Question 3. Using only two Boolean variables x and y , create a different Boolean c-table on the same rows that describes the same set of possible outcomes.

Answer. We number each possible world in binary. For brevity, we only write the **session** column when writing out possible worlds:

00	01	10	11
session ...	session ...	session ...	session ...
2 ...	2 ... 3 ...	2 ... 3 ... 4 ...	2 ... 3 ... 4 ... 5 ... 6 ...

We now write, for each tuple (again denoted by its **session** value), the possible worlds where it occurs:

2	00	01	10	11
3		01	10	11
4			10	11
5				11
6				11

We now introduce our two variables, x and y , which we will use to code the choice of possible worlds. Intuitively, x codes the choice of the most significant bit of the possible world number (written in binary), and y codes the choice of the least significant bit. This allows us to write the condition of each tuple (again denoted by its **session** value):

2	$\neg x \wedge \neg y$	\vee	$\neg x \wedge y$	\vee	$x \wedge \neg y$	\vee	$x \wedge y$
3	\perp	\vee	$\neg x \wedge y$	\vee	$x \wedge \neg y$	\vee	$x \wedge y$
4	\perp	\vee	\perp	\vee	$x \wedge \neg y$	\vee	$x \wedge y$
5	\perp	\vee	\perp	\vee	\perp	\vee	$x \wedge y$
6	\perp	\vee	\perp	\vee	\perp	\vee	$x \wedge y$

We can simplify these conditions, which yields:

2	\top
3	$x \vee y$
4	x
5	$x \wedge y$
6	$x \wedge y$

We annotate the initial table with these conditions, yielding our final answer (the \top , which indicates the true annotation, can be omitted):

<i>Classes</i>				
<i>session</i>	<i>date</i>	<i>prof</i>	<i>room</i>	
2	Nov 28	Antoine	C017	
3	Dec 5	Antoine	C47	$x \vee y$
4	Dec 12	Silviu	C47	x
5	Jan 9	Silviu	C47	$x \wedge y$
6	Jan 16	Silviu	C47	$x \wedge y$