Department of Informatics
Prof. Dr. Michael Böhlen
Binzmühlestrasse 14
8050 Zurich
Phone: +41 44 635 4333
Email: boehlen@ifi.uzh.ch

University of Zurich[UZH]

Database Systems — Exam
Spring 2016 — 21.06.2016

Name: _____ Matriculation number: _____

Advice

- Please check the completeness of your exam (14 numbered pages).
- Put your name and student id on the top of each page you hand in.
- Do not use pencil.
- Do not use a red pen.
- Stick to the notations and solutions used in the lecture.
- If you make any assumption for a solution, declare it clearly.
- Exercises with more than one solution will not be considered.
- Make sure to hand in all sheets at the end of your exam.
- You are allowed to use one A4 sheet with your personal notes and a pocket calculator.
- Time for the exam: 90 minutes.

Signature:

Correction slot — Please do not fill out the part below

Exercise	1	2	3	4	Total
Maximum Points	24	22	12	12	70
Points Achieved					

Exercise 1 — 24 Points

1.1 A database stores information about flights and passengers. The value of attribute Gender is either 'M' or 'F' and cannot be NULL. The primary key of each relation is underlined. A sample database instance is illustrated below:

flight

FlightNumber	**Departure**	**Destination**
'WK432'	'Lugano'	'Zurich'
'KL516'	'Basel'	'London'

trip

FlightNumber	**Date**	**Time**	**Aircraft**
'WK432'	'2015-10-19'	'16:40'	'A380'
'WK432'	'2015-10-26'	'16:40'	'A318'
'KL516'	'2015-05-04'	'18:28'	'737'

passenger

FlightNumber	**Date**	**Passport**	**Seat**
'WK432'	'2015-10-19'	'1427690'	'25B'
'WK432'	'2015-10-19'	'8365002'	'10C'
'KL516'	'2015-05-04'	'1427690'	'10A'

person

Passport	**LastName**	**BirthDate**	**Gender**
'1427690'	'Cerf'	'1943-06-23'	'M'
'4067231'	'Kahn'	'1938-12-23'	'M'
'8365002'	'Liskov'	'1939-11-07'	'F'

Tasks:

(a) Write a **relational algebra expression** that returns the flight number and the date of each trip with at most 10 passengers where Basel is either the departure or destination city. All relational algebra operations **except theta join and Cartesian product** are allowed (5 points).

(b) Write a **domain relational calculus expression** that returns the passport and the last name of every passenger who has landed in Zurich more than once (6 points).

(c) Write an **SQL query** that returns a list with unique flight numbers that includes all flights with at least one trip and one passenger, such that **all** trips of a flight are taken by both male and female passengers (7 points).

1.2 Consider relations p and q with schemas $P(A, B, C)$ and $Q(A, C, D, E)$, respectively. Assume the domain of all attributes are the integers. Find all mistakes in the following relational algebra expressions and suggest a possible way how the mistakes can be fixed.

For each incorrect answer 0.5 points are subtracted. A total negative amount of points yields zero points for this task. (4 points).

(a) $\pi_{A,B}(p \times q) - \sigma_{A=5,B=10}(\pi_{A,B}(p))$

(b) $(p \bowtie_{A>5} \pi_{E,D}(q)) - p$

(c) $\pi_B(p) \bowtie_{F>D} q$

1.3 Consider two relations r and s with the same relation schema. Let X be a strict subset of the attributes of the relation schema. Prove that the following equality does not hold (2 points):

$$\pi_X(r - s) = \pi_X(r) - \pi_X(s)$$

Name: Matriculation number:

Exercise 2 22 Points

2.1 Given is the following set F of functional dependencies for relation schema $R(A, B, C, D, E, F)$:

$$
\begin{aligned}
F \;=\; \{ & \\
& A \;\to\; D\,E, \\
& B \;\to\; F, \\
& C\,E \;\to\; D, \\
& E \;\to\; A, \\
& F \;\to\; B\,D \\
\} &
\end{aligned}
$$

Determine for each functional dependency f in the table below if it can be derived from F, i.e., $F \vDash f$.

For each correct answer 0.5 points are added; for each incorrect answer 0.5 points are subtracted; for each unanswered question 0 (zero) points are added. A total negative amount of points yields 0 (zero) points for this task. (3 points)

Functional Dependency f	derivable	not derivable
$C\,E\,F \to A\,B\,C\,D\,E\,F$	☐	☐
$A\,F \to A\,B\,C\,D\,E\,F$	☐	☐
$B\,E \to D$	☐	☐
$A\,C\,F \to B\,C\,D$	☐	☐
$A\,B\,C \to D\,E\,F$	☐	☐
$D \to B$	☐	☐

2.2 Given is relation schema $R(A, B, C, D)$ with the following set F of functional dependencies. All attributes are atomic.

$$
\begin{aligned}
F \;=\; \{ & \\
& A \;\; \rightarrow \;\; C, \\
& C \;\; \rightarrow \;\; A\ B, \\
& A\ D \;\; \rightarrow \;\; B, \\
\} &
\end{aligned}
$$

(a) Determine the set S of all candidate keys for relation schema R. Justify for each element in S why it is a candidate key for R. (3 points)

(b) Determine the highest normal form (1NF, 2NF, 3NF, BCNF) that relation schema R is in. Explain your answer. (2 points)

(c) Relation schema R is transformed into two relation schemas: $R_1(A, B, C)$ and $R_2(A, C, D)$.

i. Determine the sets S_1 and S_2 of all candidate keys for relation schemas R_1 and R_2, respectively. Specify the minimal covers F_1 and F_2 for relation schemas R_1 and R_2, respectively. (2 points)

ii. Is the transformation *dependency preserving*? Justify your answer. (3 points)

2.3 Consider the relation schema $R(A, B, C, D)$ with the following functional dependencies. All attributes are atomic. R is in 2NF.

$$
\begin{aligned}
F \;=\; \{ & \\
& A \;\; \rightarrow \;\; C, \\
& B\,D \;\; \rightarrow \;\; A, \\
\} &
\end{aligned}
$$

(a) Extend F to a new set F_2 of functional dependencies, so that R is **at most in 1NF**. F_2 may contain exactly one additional functional dependency compared to F. None of the functional dependencies of F may be removed. Explain why relation schema R with your set F_2 is at most in 1NF. (3 points)

(b) Extend F to a new set F_3 of functional dependencies, so that R is in **3NF**. F_3 may contain exactly one additional functional dependency compared to F. None of the functional dependencies of F may be removed. Explain why relation schema R with your set F_3 is in 3NF. (3 points)

2.4 Assume $B \twoheadrightarrow A$ is a multivalued dependency on relation schema $R(A, B, C)$ of relation r. Based on the three tuples in relation r, add all other tuples that **must belong** to the relation. (3 points)

r

A	B	C
"a"	1	7
"c"	1	8
"e"	1	9

Exercise 3 — 12 Points

3.1 Consider the following DDL statements to create four tables. Construct an ER-diagram for this database. If relationship cardinalities are unclear, pick the ones that do not violate the DDL statements. Each table must map to one entity (either weak or strong) in the ER-diagram (10 points).

```
CREATE TABLE a(
  A INTEGER,
  B INTEGER,
  C INTEGER,
  PRIMARY KEY (A)
);

CREATE TABLE c(
  F INTEGER,
  G INTEGER,
  H INTEGER,
  A INTEGER,
  PRIMARY KEY (F,G),
  FOREIGN KEY (A)
    REFERENCES a(A)
);
```

```
CREATE TABLE b(
  D INTEGER,
  E INTEGER,
  A INTEGER,
  PRIMARY KEY (A,D),
  FOREIGN KEY (A)
    REFERENCES a(A)
);

CREATE TABLE d(
  I INTEGER,
  J INTEGER,
  F INTEGER,
  G INTEGER,
  PRIMARY KEY (I),
  FOREIGN KEY (F,G)
    REFERENCES c(F,G)
);
```

3.2 Consider the following DDL statements and the corresponding ER-diagram. How does the ER-diagram change if we add a **NOT NULL** constraint on attribute K in table y? You may draw your changes into the existing ER-diagram (2 points).

```
CREATE TABLE x(
  K INTEGER PRIMARY KEY
);

CREATE TABLE y(
  L INTEGER PRIMARY KEY,
  K INTEGER,
  FOREIGN KEY (K)
    REFERENCES x(K)
);
```

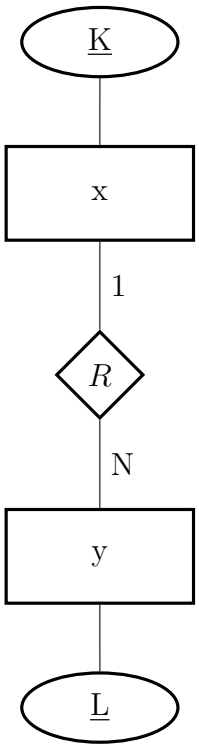

Exercise 4 — 12 Points

The following notation is used in this exercise:

Operation example	Explanation
$r_1(A)$	transaction T_1 reads data item A
$w_1(A)$	transaction T_1 writes data item A
$sl_1(A)$	transaction T_1 takes a shared lock on data item A
$xl_1(A)$	transaction T_1 takes an exclusive lock on data item A
$u_1(A)$	transaction T_1 unlocks data item A
c_1	transaction T_1 commits

4.1 Draw the precedence graph of the following schedule (4 points):

$S = \langle w_4(B), r_1(A), c_4, r_2(B), w_2(A), r_3(A), r_1(B), w_3(A), w_1(B), r_3(B), c_1, c_2, c_3 \rangle$

4.2 Consider schedule S from exercise 4.1. Check for each of the properties in the table below if S fulfills it or not.

For each property points are awarded as follows, correct answer: $+0.5$, incorrect answer: -0.5, no answer: 0. A negative amount of points yields zero points for this task (2 points).

	Fulfilled	Not fulfilled
Serializable		
Cascadeless		
Recoverable		
Serial		

4.3 Draw the wait-for graph of the following schedule. Is the schedule in a deadlock state? (3 points):

$S = \langle sl_1(A), r_1(A), xl_3(B), sl_3(A), xl_2(C), r_3(A), r_2(C), xl_2(A), xl_1(B) \rangle$

4.4 Consider the following two schedules.

$S_1 = \langle sl_2(B), sl_1(B), r_1(B), xl_2(C), u_1(B), w_2(C), u_2(C), c_1, r_2(B), u_2(B), c_2 \rangle$

$S_2 = \langle sl_1(A), xl_1(B), xl_2(C), r_1(B), sl_2(A), u_2(A), u_1(A), w_2(C), u_1(B), c_1, u_2(C), c_2 \rangle$

(a) Does schedule S_1 adhere to the 2PL protocol? If yes, show the lock point for each transaction. Otherwise explain why 2PL is violated (1 point).

(b) Does schedule S_2 adhere to the S2PL protocol? If yes, show that the definition of S2PL is fulfilled. Otherwise, explain why S2PL is violated (1 point).

(c) Does schedule S_2 adhere to the SS2PL protocol? If yes, show that the definition of SS2PL is fulfilled. Otherwise, explain why SS2PL is violated (1 point).