
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
Spring 2015

Exam
16.06.2015

Name: _____ Matriculation number: _____

Advice

- Please check the completeness of your exam (18 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	5	Total
Maximum Points	24	20	17	13	12	86
Points Achieved						

A database stores information about PCs. Each PC has a given brand and model, and it is sold at a given price. Each PC is sold with a specific configuration that includes its components, i.e., the hardware parts of a PC. The primary key of each relation is underlined. A sample database instance is illustrated below:

hardware

<u>ID</u>	Name	Brand	Type
1	Disk	Sony	Flash
2	Disk	Apple	Solid
3	Touchpad	Logitech	Touch
4	Ram	Samsung	DCR

component

<u>Brand</u>	<u>Model</u>	<u>HW_ID</u>
Lenovo	X1	3
Lenovo	X1	2
Lenovo	X250	1
Apple	MacBook Pro	4

pc

<u>Brand</u>	<u>Model</u>	Price
Lenovo	X1	1200
Lenovo	X250	1400
Apple	MacBook Pro	1030

Tasks:

- 1.1 Write a **relational algebra expression** that returns name, brand and type of the hardware parts of PCs that cost more than the average price of Apple PCs. (8 points)

Name:

Matriculation number:

1.2 Write an **SQL query** that returns the brand and model of the PCs with a price smaller than 1500 and having more than 5 components. No subqueries are allowed. (8 points).

1.3 Give a **domain relational calculus expression** that returns name, brand and type of the hardware included in the most expensive PC. (8 points).

2.1 Assume relations A and B, where relation A stores 100 tuples, and relation B stores 100'000'000 tuples. The average tuple size (of A as well as B) is 20 bytes. The block size is 1000 bytes. Each tuple of A has on average 2 join matches in B. A B-Tree of three levels is available on the join attribute of each relation. The buffer size is 3 blocks.

Calculate the cost in terms of the number of block I/Os for the following query evaluation strategies. If there are multiple possibilities to choose outer and inner relation, use the one with the lower execution time.

a) Nested Loop Join (2 points)

b)Block Nested Loop Join (2 points)

c) Indexed Nested Loop Join (3 points)

d) SortMerge (3 points)

e) HashJoin (2 points)

Name:

Matriculation number:

- 2.2 Given a hash function such that $h(1) = 0001$, $h(2) = 0010$, $h(3) = 0011$, $h(15) = 1111$. Show the extendable hash structure with bucket-size = 2, after inserting 1, 2, 3 and 15 (12 points).

Use the following notation for this exercise:

Operation example	Explanation
$r_1(A)$	transaction T_1 reads data item A
$w_2(B)$	transaction T_2 writes data item B
$sl_1(A)$	transaction T_1 takes a shared lock on data item A
$xl_2(B)$	transaction T_2 takes an exclusive lock on data item B
$u_1(C)$	transaction T_1 unlocks data item C
c_1	transaction T_1 commits
a_2	transaction T_2 aborts

3.1 Consider the schedules:

$$S_1 = \langle r_1(A), r_1(C), A := A + 10, C := C + 21, r_2(B), B := B * 2, w_2(B), w_1(A), r_2(A), A := A/2, w_2(A), w_1(C) \rangle$$

$$S_2 = \langle r_2(B), r_1(C), r_1(A), B := B * 2, A := A + 10, C := C + 21, w_1(A), w_1(C), w_2(B), r_2(A), A := A/2, w_2(A) \rangle$$

$$S_3 = \langle r_1(A), A := A + 10, r_1(C), C := C + 21, r_2(A), A := A/2, w_2(A), r_2(B), B := B * 2, w_2(B), w_1(A), w_1(C) \rangle$$

- (a) Use the definition of conflict equivalence to find all pairs of conflict equivalent schedules. (3 points)

Name:

Matriculation number:

- (b) Decide if the given schedules are serializable. Explain your answer. (3 points)

3.2 Consider the schedules:

$$S_1 = \langle r_1(A), r_2(A), r_1(B), r_1(C), w_2(A), c_2, w_1(B)r_3(A), w_3(A), c_3, w_1(A), c_1 \rangle$$

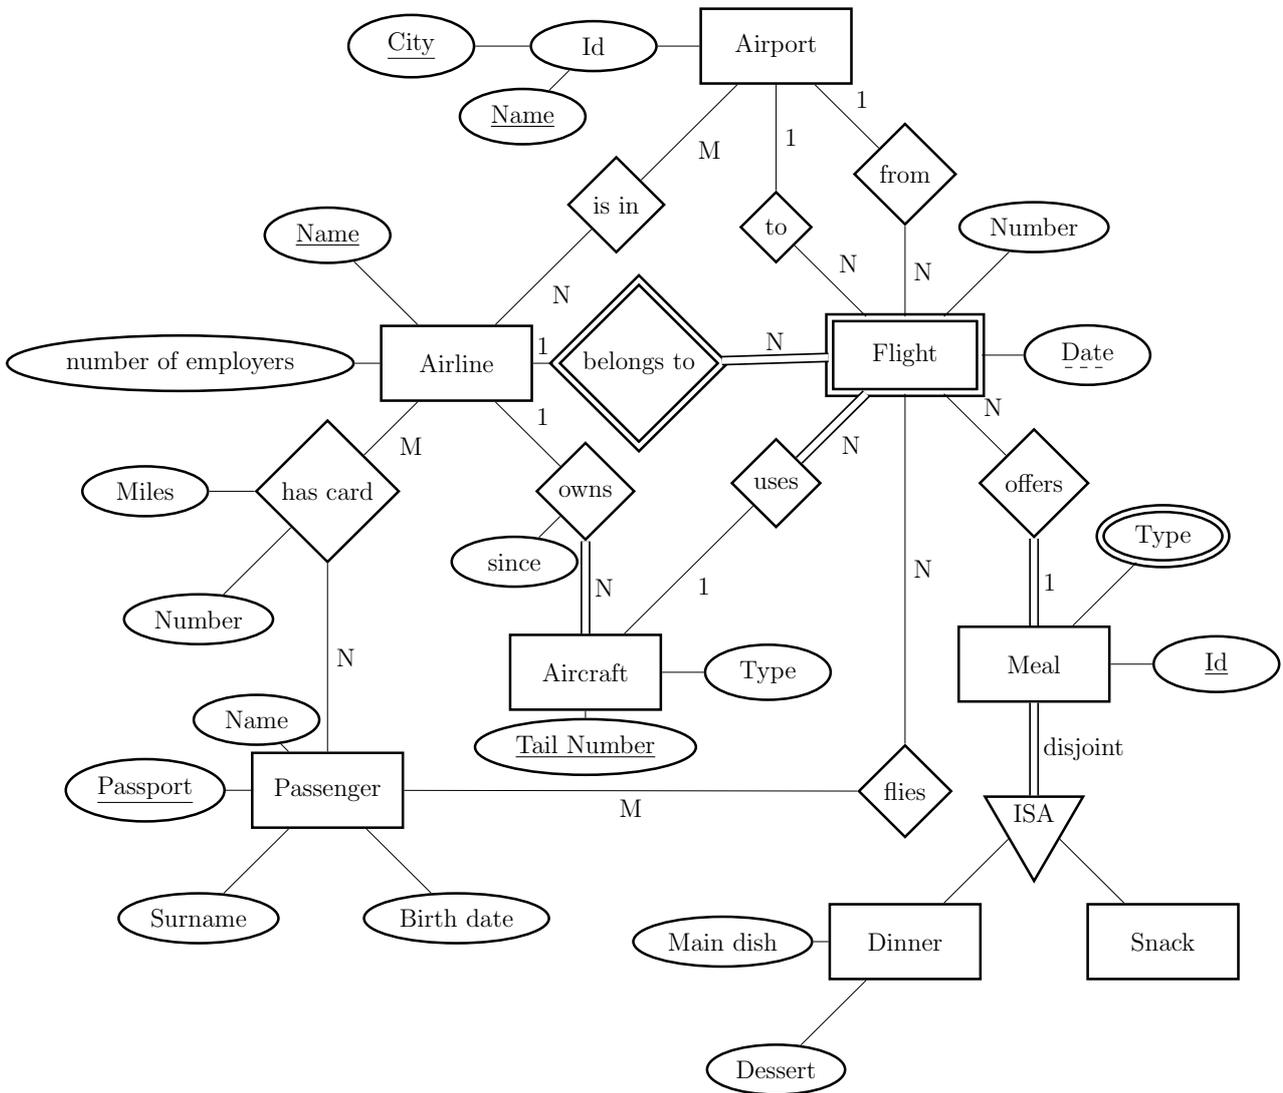
$$S_2 = \langle r_1(A), r_2(B), w_1(A), w_2(B), r_3(A), r_4(B)w_4(B), w_3(A) \rangle$$

(a) For each schedule show the corresponding conflict graph. (2 points)

(b) Is S_1 a recoverable schedule? Is S_1 a cascadeless schedule? Explain why. (2 points)

(c) Insert commits into S_2 to make it recoverable, but not cascadeless. (1 point)

4.1 The following ER diagram describes a flight management system. Map the diagram, including primary and foreign keys, to a corresponding relational database schema. Avoid NULL values and redundancy as much as possible (9 points).



Name:

Matriculation number:

Name:

Matriculation number:

5.1 Consider the database instance *artist*:

artist

ArtistId	Name	Play	Role	Type
'B1021'	'Ana Soboleva'	'Swan Lake'	'Odile', 'Odette'	'ballet'
'B1048'	'Vik Lebedev'	'Swan Lake'	'Siegfried'	'ballet'
'OH349'	'Spenser Lang'	'Robin Hood'	'Daniel'	'ballet'

and the set of functional dependencies:

Play \rightarrow Type

ArtistId \rightarrow Name

Role, ArtistId \rightarrow Play

Type, ArtistId \rightarrow Name

For each database instance, starting with the initial relation shown above, iterate the following steps until the database schema is in third normal form: (6 points)

- List all candidate keys of each relation schema
- Determine the current normal form of each relation schema
- Identify all functional dependencies that violate the next higher normal form
- Make a normalization step
- List all functional dependencies of the resulting relation schemas
- List all candidate keys of the resulting relation schemas

Note: While applying the algorithm, use the violating functional dependencies in the order listed above.

Name:

Matriculation number:

5.2 Which of the following properties are true for your decomposition? (3 points)

- (a) Losslessness
- (b) Dependency preservation

5.3 Assume a relation schema $R(A, B, C, D)$ and the functional dependencies:

$A \rightarrow C$

$B \rightarrow C$

Compute F^+ . Explain your solution. Include only the nontrivial functional dependencies in your answer. (3 points)

Name:

Matriculation number:
