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 2014 — 17.06.2014

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	16	24	20	15	75
Points Achieved					

Exercise 1 — 16 Points

1.1 Map the following ER diagram, including keys, to a corresponding relational database schema. Avoid NULL values and redundancy as much as possible (7 points).

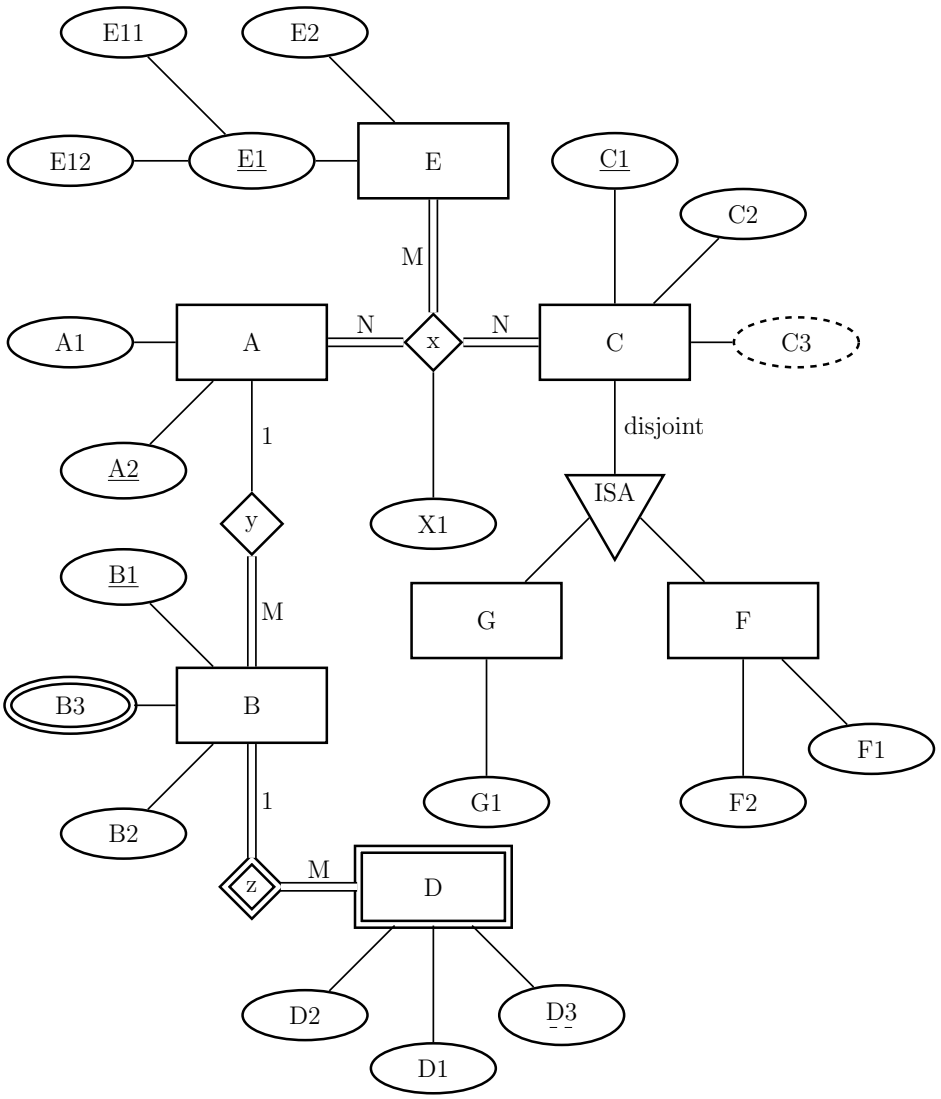

1.2 Use SQL to create the relation schemas that map entity and relationship types B and X. The definitions shall include primary and foreign keys. Assume that all attributes are Integer and all referenced relations have been defined accordingly. (2 points)

1.3 Assume the following ER diagram:

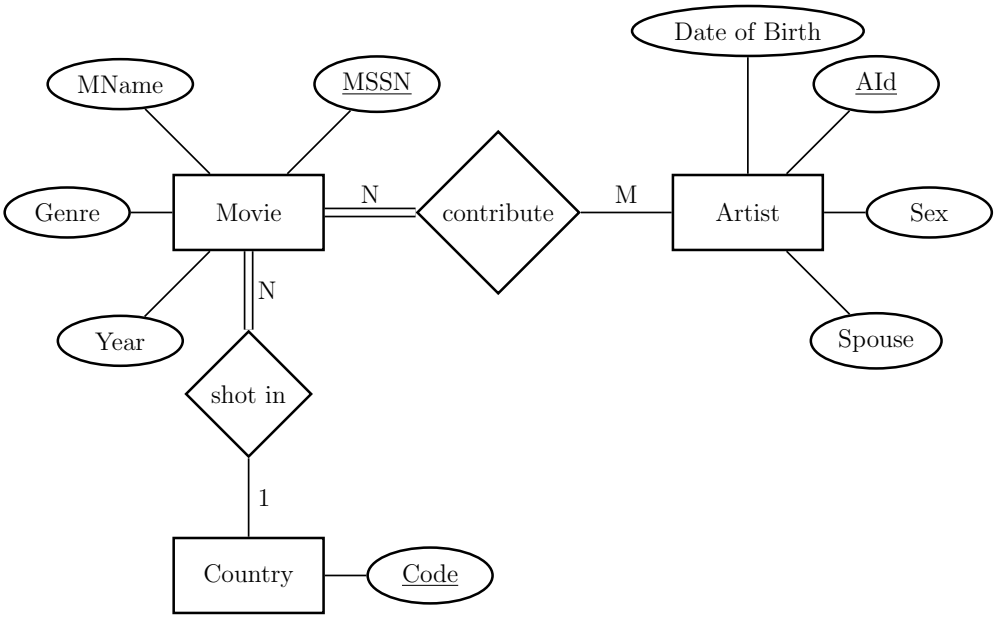

Tasks:

Modify the above ER diagram to satisfy the following new or changed requirements (7 points):

- A movie may have many genres.
- A spouse must be an artist.
- A movie may be shoot in many countries.
- An artist must be an actor or a director or both of them in a movie.
- For each actor the character name (*Character*) that he/she plays in a movie (for example: the character Tony Stark by actor Robert Downey Jr. in the movie Iron Man) should be modeled.
- For each actor the set of awards (*Awards*) shall be stored.
- A director may direct movies.

Exercise 2 — 24 Points

A database stores the nutritive content of food. Here follows an extract.

Nutrient

ID	NAME	TYPE
1	Vit. C	Vitamine
2	Omega 3	Fat

Food

ID	NAME	TYPE
1	Tofu	Vegi-Cheese
2	Strawberry	Fruit
3	Focaccia	Bread

Content

NID	FID	VALUE
1	2	58.8
1	1	0.1
1	2	0.4

Tasks:

2.1 Write a **relational algebra expression** that returns, for each food type, the average content of the nutrient "Vit. A". (8 points)

2.2 Write an **SQL query** that returns the name of the foods having no value of "Vit. A" or a value of "Vit. A" smaller than 0.01 (8 points).

2.3 Give a **domain relational calculus expression** that, for each food type, returns the name of the food with the highest "Vit. A" value (8 points).

Exercise 3 — 20 Points

Use the following notation for the following exercises:

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 shared locks data item A
$xl_2(B)$	transaction T_2 exclusive locks 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 Construct schedules for the following tasks:

a) A schedule for two transactions T_1 and T_2 that leads to a deadlock. Make sure to add the necessary lock operations. (3 points)

b) A schedule for two transactions T_1 and T_2 that is not recoverable. Make sure to add the necessary commit and abort operations. (3 points)

3.2 Consider the following schedule S_1:

$$S_1 = \langle r_4(A), r_1(A), r_2(A), r_3(A), w_1(A), w_2(B), \\ r_3(B), w_3(B), r_2(C), w_2(C), r_4(C), w_4(C) \rangle$$

Tasks:

a) Show the precedence graph (conflict graph) of schedule S_1. (2 points)

b) State if S_1 is conflict serializable. If yes give **all** serializability orders; If not give the first conflict. (2 points)

3.3 Consider the following schedule S_2:

$$S_2 = \langle r_1(A), r_1(B), r_2(A), r_3(A), w_3(A), r_1(C), \\ w_1(C), c_1, w_2(B), c_2, r_3(E), c_3 \rangle$$

Tasks:

a) State if S_2 is a valid strict two-phase locking (S2PL) schedule. If yes add **all** required lock and unlock instruction (sl, xl and u). If no explain why. (6 points)

b) Assume $r_2(A)$ is moved immediately after $w_3(A)$ in S_2. Is it a valid S2PL schedule? If yes add **all** required lock and unlock instruction (sl, xl and u). If no explain why. (4 points)

Exercise 4 — 15 Points

4.1 Consider a relation R (A, B, C, D, E, F) for which the functional dependencies **ACD → A, C → A, C → E, CD → F, F → B** hold.

4.1.1 Prove that DC is a candidate key of R. Write the necessary conditions and elaborate them. (5 points)

4.1.2 Assume R is in 1NF. For each normal form that R does not satisfy, identify all functional dependencies that violate it. (3 points)

4.1.3 Compute the Boyce-Codd Normal Form (BCNF) decomposition of R. Indicate each step of the algorithm, by showing the schema and the violation of BCNF on which you focus during that step. Also, indicate clearly your final result: the relations and their attributes. (3 points)

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

4.2 In the following figure, you see an instance of relation R (B,C,D,E).

	B	**C**	**D**	**E**
r_1	1	4	1	2
r_2	5	4	2	4
r_3	3	8	1	2
r_4	3		1	2
r_5	1	8	1	2
r_6	5	6	2	
r_7	8	4	2	4

4.2.1 Fill in the missing values for attributes C and E in the above instance of relation R so that the multivalued dependency $\mathbf{DE} \twoheadrightarrow \mathbf{C}$ holds. (4 points)