Final exam in 1DL301 Database Design I

Department of Information Technology, Uppsala University January 4, 2018, 08.00 – 13.00

This is a multiple-choice exam. Each question might have several correct choices, but you **must select only one choice**. If you select more than one, your answer will be marked as incorrect (even if you have selected only correct choices).

The questions are divided into four sections:

Section	Questions
Data modeling	10
SQL	10
FDs, CKs, NFs and normalization	10
Other	6

To pass the exam you need:

- to answer correctly on at least 50% of questions in each section, and
- at least 22 correct answers in total.

Failing to meet any of these criteria means failing the exam (i. e. grade U).

If you fulfill both criteria, your grade will be determined by the number of correct answers:

Correct answers	Grade
22 - 26	3
27 – 31	4
32 - 36	5

Your answers must be given on the answer sheet which will be handed in. Don't forget to fill out your exam code. To mark your answer **fill the answer box entirely using a dark colored pen** (black or blue). The optical character recognition system will not recognize ticks, crosses, circles or any other additional notes. Remember, do not to fill more than one answer box per question.

If you make a mistake on the answer sheet, request a new one and make sure you hand in the correct sheet (if you hand in several answer sheets, your exam will not be graded).

You can keep the question sheets. We recommend that you first note your answers on these and fill out the answer sheet just before handing it in.

Allowed aids One A4 sheet with handwritten notes (both sides can be used) which must be handed in with your exam (remember to fill out your exam code). An English explanatory dictionary and/or a translation dictionary between English and your mother tongue.

1 Information about you

Question i (Not assessed.) What is your study program?

- A IT
- B STS
- C X
- D Other CS/math/physics
- E Other non-CS/math/physics

Question ii (Not assessed.) In which period were you registered in the course for the first time?

- A Fall 2017, period 2 (this period)
- B Fall 2017, period 1 (previous period)
- C Earlier period

Question iii (Not assessed.) How many lectures have you attended?

- A None or very few
- B Around 25%
- C Around 50%
- D Around 75%
- E Almost all or all

2 Data modeling

Question 1 What ER element is depicted in the following figure?



A Weak attribute

- B Composite attribute
- C Key attribute
- D Derived attribute
- E Composite multi-valued attribute





A No

Page 3

Question 3 The Internet Movie Database (IMDb) is an online database of information related to movies. Registered users can rate movies on a scale of 1 to 10 stars (each user can rate each movie once). Based on the ratings IMDb shows the average rating as well as average ratings for different combinations of gender and age (e. g. the average rating given by males between 18 and 29 years old) for each movie. Which of the following ER diagrams shows a correct way how to store the ratings? (Only the relevant portion of each diagram is shown.)



Question 4 In a medical experiment we give a patient an experimental drug and measure his or her blood pressure 1, 6, 12 and 24 hours after taking the drug. These measurements must be stored in the database and it must be possible to say which measurement was first, which second etc. Can the following model be used? (Only the relevant portion of the diagram is shown.)



Question 5 Which of the following statements about the depicted ER diagram is correct?



- A Each E_1 entity must be related to one or more E_2 entities. An E_2 entity can be related to zero or one E_1 entity.
- B The ER model is not semantically correct.
- C Each E_1 entity must be related to exactly one E_2 entity but it can be related to it many times. An E_2 entity can be related to zero or one E_1 entity.
- D An E_1 entity can be related to zero, one or more E_2 entities. Each E_2 entity must be related to exactly one E_1 entity.
- E An E_1 entity can be related to zero or one E_2 entity. Each E_2 entity must be related to one or more E_1 entities.



Question 6 The ER diagram from the previous question is equivalent to which of the following diagrams?

Question 7 Based on the depicted ER diagram and the semantics of ER models, which of the following statements is correct? (Hint: only one choice is correct.)



A Same person cannot be the mayor of two cities.

B No two countries have exactly the same area and population.

C Each city is exactly in one country.

D The population of a country is always larger than 0.

E The name of a city must be different from the name of a country.

Question 8 Convert the ER model from the previous question to a relational model.

- A Country(<u>Name</u>, Area, Population) City(<u>CountryName</u>, Name, Mayor) with CountryName $^{FK} \rightarrow$ Country(Name)
- B Country(<u>Name</u>, Area, Population) City(<u>Name</u>, CountryName, Mayor) with CountryName $^{FK} \rightarrow$ Country(Name)
- C The ER model cannot be converted to a relational model because both City and Country have an attribute with the same name.
- D Country(<u>Name</u>, Area, Population) City(<u>Name</u>, Mayor)
- $[\underline{E}]$ Country(<u>Name</u>, CityName, Area, Population) with CityName $^{FK} \rightarrow$ City(Name) City(<u>Name</u>, Mayor)

Page 4

Question 9 Consider the following ER model for a database of soccer games in the next season of Allsvenskan (a Swedish professional soccer league). Convert the model to a relational model.



- $\begin{array}{||c||}\hline A & Team(\underline{Id, City}, Coach) \\ & Game(\underline{HomeId, AwayId}, Date, Score) \ with \ HomeId \ ^{FK} \rightarrow Team(Id, City), \ AwayId \ ^{FK} \rightarrow Team(Id, City) \\ \end{array}$
- $\begin{array}{|c|} \hline C & Team(Id, City, Coach) \\ & Game(\underline{HomeId}, AwayId, Date, Score) with HomeId {}^{FK} \rightarrow Team(Id), AwayId {}^{FK} \rightarrow Team(Id) \\ \end{array}$
- D Team(Id, City, Coach) Game(Date, Score)
- E Team(Id, City, Coach, AwayId) with AwayId $^{FK} \rightarrow$ Team(Id) Game(<u>HomeId</u>, Date, Score) with HomeId $^{FK} \rightarrow$ Team(Id)
- **Question 10** Which of the following relational models can be used for the ER model depicted below:



- A Car(<u>VIN</u>, Color, MaxSpeed) Truck(<u>VIN</u>, Color, Tonnage)
- B Vehicle(<u>VIN</u>, Color) Car(<u>VIN</u>, MaxSpeed) with VIN $^{FK} \rightarrow$ Vehicle(VIN) Truck(<u>VIN</u>, Tonnage) with VIN $^{FK} \rightarrow$ Vehicle(VIN)
- C Vehicle(<u>VIN</u>, Color) Car(MaxSpeed) Truck(Tonnage)
- D None of the other answers
- E Vehicle(<u>VIN</u>, Color, MaxSpeed, Tonnage)

3 SQL

For questions 11 – 15 consider the following database consisting of two tables T1 and T2:

T1	
1	C2
	1
	2
3	2
7	1

Question 11 How many rows are in the result of the following SQL query?

SELECT * FROM T1, T2

A The SQL is invalid
B None
C 8
D 16
E 4

Hint for the following 4 questions: First, write down the result of the following queries:

SELECT T1.C1, T1.C2, T2.C1, T2.C2, T1.C2*T2.C2 FROM T1 JOIN T2 ON T1.C1 = T2.C1
SELECT T1.C1, T1.C2, T2.C1, T2.C2 FROM T1 LEFT JOIN T2 ON T1.C1 = T2.C1

Question 12 What is the result of the following SQL query? (Not showing the header of the result.)

SELECT DISTINCT T1.C1, T1.C2, T2.C1, T2.C2 FROM T1 JOIN T2 ON T1.C1 = T2.C1 WHERE T1.C2 IN (2, 4)



C Empty table



Question 13 What is the result of the following SQL query? (Not showing the header of the result.)

```
SELECT T1.C1, SUM(T1.C2*T2.C2)
FROM T1, T2
WHERE T1.C1 = T2.C1
GROUP BY T1.C1
```



Question 14 What is the result of the following SQL query? (Not showing the header of the result.)

```
SELECT T1.C1, SUM(T1.C2*T2.C2)

FROM T1, T2

WHERE T1.C1 = T2.C1 AND SUM(T1.C2*T2.C2) > 10

GROUP BY T1.C1

\overrightarrow{A} \overrightarrow{A} 22

\overrightarrow{D} \overrightarrow{A} 12

\overrightarrow{E} Empty table

\overrightarrow{C} \overrightarrow{A} 18
```

Question 15 What is the result of the following SQL query? (Not showing the header of the result.)

SELECT T1.C2, T2.C2 FROM T1 LEFT JOIN T2 ON T1.C1 = T2.C1 WHERE T1.C2 < 2 OR T1.C2 > 4





For questions 16 – 20 consider a set of data organized in a hierarchical tree structure. A tree can be defined as a collection of nodes, where each node consists of data and the reference to its *parent* node. A node can be a parent for many nodes, these nodes are called *children* of that node. Each non-empty tree has one node without any parent, such a node is called the *root*. Nodes without any children are called *leaves*. (There are also some additional constraints but these are not relevant for these questions.)

To represent such a hierarchical tree structure, we will use the following relation:

Node(<u>id</u>, data, parent_id)

Question 16 Assume that data is a short string. What is the correct SQL to create the Node table?

- A CREATE TABLE Node(id int NOT NULL PRIMARY KEY, data varchar(255), parent_id int, FOREIGN KEY(id) REFERENCES Node(parent_id))
- B CREATE TABLE Node(id int NULL PRIMARY KEY, data varchar(255), parent_id int NOT NULL, FOREIGN KEY(parent_id) REFERENCES Node(id))
- C CREATE TABLE Node(id int NOT NULL PRIMARY KEY, data varchar(255), parent_id int NOT NULL, FOREIGN KEY(id) REFERENCES SELF)
- D CREATE TABLE Node(id int NOT NULL PRIMARY KEY, data varchar(255), parent_id int REFERENCES SELF)
- E CREATE TABLE Node(id int NOT NULL PRIMARY KEY, data varchar(255), parent_id int, FOREIGN KEY(parent_id) REFERENCES Node(id))

Question 17 Finish the SQL query to select the ID and the data of the root:

```
SELECT ①
FROM Node
WHERE ②
A ① id, data ② parent_id IS NULL
B None of the other answers
C ① id, data ② parent_id = NULL
D ① id, data ② NOT EXISTS parent_id
E ① parent_id, data ② id != NULL
```

Question 18 Finish the SQL query to select the ID and the data of the leaves:

```
SELECT id, data FROM Node n where \oplus
```

A ① EXISTS (SELECT * FROM Node m WHERE m.parent_id!=n.id)
 B ① parent_id IN (SELECT id FROM Node WHERE id IS NOT NULL)
 C ① NOT EXISTS (SELECT * FROM Node m WHERE m.parent_id=n.id)
 D ① parent_id IS NOT NULL
 E ① id NOT IN (SELECT parent_id FROM Node WHERE parent_id IS NOT NULL)

Question 19 Finish the SQL query to select the ID and the data of all grandchildren (i. e. children of children) of the node with the ID 123:

```
SELECT ①
FROM node n, node m
WHERE ②
A ① n.id, n.data
B ① n.id, m.data
C ① m.id, m.data
D ① n.id, n.data
E None of the other answers
a m.parent_id=123 AND n.parent_id=m.id
C ① m.id, n.data
C ① m.id,
```

Question 20 Finish the query to select the ID and the children count for each node:

```
SELECT parent_id AS node, COUNT(id) AS count
FROM node
WHERE ①
A False
B parent_id IS NOT NULL
C None of the other answers
D id != NULL
E True
```

4 FDs, CKs, NFs and normalization

For questions 21 - 24, consider the following relation in 1NF:

Employee(ID, SSN, FirstName, LastName, DepartmentID, DepartmentName)
with the following non-trivial FFDs:

- $\{ID\} \rightarrow \{SSN, FirstName, LastName, DepartmentId, DepartmentName\}$
- $\{SSN\} \rightarrow \{ID, FirstName, LastName, DepartmentId, DepartmentName\}$
- {DepartmentId} \rightarrow {DepartmentName}
- {DepartmentName} \rightarrow {DepartmentId}

Question 21 Which of the following options lists all candidate keys?

- A {ID}, {SSN}, {DepartmentId}, {DepartmentName}
- B {DepartmentId}, {DepartmentName}
- \mathbb{C} {ID}, {SSN}
- D {ID}
- $E \{SSN\}$

Question 22 Which of the following options lists all prime attributes?

- A All attributes are prime
- B ID, SSN, DepartmentId, DepartmentName

C ID

- D There are no prime attributes
- E ID, SSN

Question 23 In which NF is the Employee table?

- A In 1NF but not in 2NF
- B In 2NF but not in 3NF
- C In 3NF but not in 2NF
- D In 3NF but not in BCNF

E In BCNF

Question 24 Which of the following normalized databases contains all the information contained in the original table, with all relations in BCNF (not showing the foreign key constraints)?

- A Employee(<u>ID</u>, SSN, FirstName, LastName) Department(<u>DepartmentID</u>, DepartmentName)
- B Employee(ID, SSN, FirstName, LastName, DepartmentID, DepartmentName)
- C Employee(ID, SSN, FirstName, LastName, DepartmentName) Department(DepartmentName, DepartmentID)
- D Employee(ID, SSN, FirstName, LastName) EmployeeBySSN(<u>SSN</u>, ID, FirstName, LastName) Department(<u>DepartmentID</u>, DepartmentName) DepartmentByName(<u>DepartmentName</u>, DepartmentID)
- E Employee(<u>ID</u>, SSN, FirstName, LastName, DepartmentID) Department(<u>DepartmentID</u>, DepartmentName)

For questions 25 – 28, consider a relation in 1NF R(A, B, C, D, E) with the following dependencies:

- $\{A, B\} \rightarrow \{C\}$
- $\{B\} \rightarrow \{E\}$
- $\{E\} \rightarrow \{D\}$

Question 25 What is $\{B, C\}^+$ (i. e. the set of all attributes which can be determined by $\{B, C\}$)?

- $\begin{array}{l} \hline \mathbf{A} & \{B, C, D, E\} \\ \hline \mathbf{B} & \text{Empty set} \\ \hline \mathbf{C} & \text{None of the other answers} \\ \hline \mathbf{D} & \{B, C, E\} \\ \hline \mathbf{E} & \{A, B, C, D, E\} \end{array}$
- $\mathbf{F} \{B, C\}$

Question 26 Which of the following options lists all candidate keys?

- $\boxed{\mathbf{A}} \{A, B, E\}$
- **B** $\{A, B\}$ and $\{E\}$
- $\boxed{\mathsf{C}} \{A, B\}$
- **D** All sets including both A and B
- E There are no candidate keys
- $\mathbf{F} \{A, B, C, D, E\}$

Question 27 In which NF is *R*?

- A In BCNF but not in 3NF
- B In 2NF but not in 3NF
- C In 3NF but not in BCNF
- D In 1NF but not in 2NF
- E In BCNF
- F In 3NF but not in 2NF

Question 28 Which of the following normalized databases contains all the information contained in the original table, with all relations in BCNF (not showing the primary and foreign key constraints)?

- $\mathbb{E} R(A, B, C, D, E)$

Question 29 The following table shows the current state of a relation. Is it true that $\{C_2\} \rightarrow \{C_3\}$?

C_1	C_2	C_3	C_4
1	1	XX	157
2	1	XX	178
3	1	XX	192
4	2	XY	183
5	2	XY	166

A No, values of C_3 are not unique.

B Yes, since if C_2 is 1, then C_3 is XX, and if C_2 is 2, then C_3 is XY.

C A dependency cannot be confirmed by just checking one state of the relation.

Question 30 Consider a relation R(A, B, C, D) in 1NF, with two candidate keys $\{A\}$ and $\{B\}$. Which of the following statements is correct? (Assume that no attribute can be determined by the empty set.)

A R is in BCNF

B R is at least in 3NF

- C R is not in 2NF
- D R is in 3NF
- E R is at least in 2NF

5 Other

Question 31 In the relational model, if *K* is a candidate key of relation *R* and *X* is one of its attributes (i. e. $X \in K$), then:

A None of the other answers

B $K \setminus \{X\}$ is a candidate key (\ indicates the set difference)

 $[\mathbf{C}] K \cap \{X\}$ is a superkey

D $K \cap \{X\}$ is a candidate key

 $\begin{bmatrix} E \end{bmatrix}$ K is the primary key

Question 32 If P_1 is NULL and P_2 is TRUE, which of the following statements is correct?

A P_1 AND P_2 is NULL, P_1 OR P_2 is NULL

B P_1 AND P_2 is TRUE, P_1 OR P_2 is TRUE

C P_1 AND P_2 is NULL, P_1 OR P_2 is TRUE

D P_1 AND P_2 is TRUE, P_1 OR P_2 is NULL

E P_1 AND P_2 is FALSE, P_1 OR P_2 is TRUE

Question 33 When transferring money from one bank account to another we need to run two SQL statements, one that reduces the balance of the first account and one that increases the balance of the second account. Which of the ACID properties guarantees that either both SQL statements get executed (and the money is transferred) or none does (not transferring the money at all)?

A Atomicity

- **B** Isolation
- C Durability
- D Consistency
- E Integrity

Question 34 Consider a relation, $R(A_1, A_2, ..., A_n)$, and a subset of its attributes, $K \subseteq \{A_1, A_2, ..., A_n\}$. If for any legal instance r and any two distinct tuples $t_1, t_2 \in r, t_1 \neq t_2$ it holds that $t_1[K] \neq t_2[K]$ (where t[K] denotes the projection of t onto the attributes in K), then K is:

- A Superkey
- B Primary key
- C Foreign key
- D Determinant
- E Candidate key

Question 35 Bob creates the table X. Then, the following sequence of statements is executed, in the given order (the name of the user executing the statement is indicated at the beginning of each statement):

Bob: GRANT SELECT ON X TO Jim WITH GRANT OPTION Bob: GRANT SELECT, UPDATE ON X TO Ann WITH GRANT OPTION Jim: GRANT SELECT ON X TO Tim Ann: GRANT SELECT ON X TO Tim Jim: REVOKE SELECT ON X FROM Tim

Which privileges (concerning the table X) does Tim have?

A SELECT, UPDATE, both with the grant option

B SELECT, UPDATE, both without the grant option

C SELECT with the grant option

- D None
- E SELECT without the grant option

Question 36 Consider a relation $R(\underline{A}, B, C, D)$ containing 10^7 records. A is the primary key, and B contains 10^5 distinct values. The following prepared statement is executed very frequently:

SELECT C FROM R WHERE B=?

Which index would help to avoid the full table scan?

A One index on B

B One index on (B, C)

C One index on C

D One index on (C, B)

E One index on A and one on C



Final exam in 1DL301 Database Design I – Answer sheet

Department of Information Technology, Uppsala University January 4, 2018, 08.00 – 13.00

Your exam code	
----------------	--

|--|

Encode the number in your exam code in the table on the right side. For example, if your exam code is AB2097, fill in the box no. 2 in the first column, the box no. 0 in the second, the box no. 9 in the third, and the box no. 7 in the last column. Fill in the boxes entirely!

Use a **dark colored pen** (blue or black). To mark your answer **fill in the box entirely** (■)! The OCR software will not recognize ticks, crosses, circles, etc. **Do not make any additional notes** on this sheet!

If you make an error, ask for a new answer sheet!

Remember, choose only one answer for each question!



	. —		~ □			
18	A:	B:	C:	D:	E:	
19	A:	B:	C:	D:	E:	
20	A:	B:	C:	D:	E:	
21	A:	B:	C:	D:	E:	
22	A:	B:	C:	D:	E:	
23	A:	B:	C:	D:	E:	
24	A:	B:	C:	D:	E:	
25	A:	B:	C:	D:	E:	F:
26	A:	B:	C:	D:	E:	F:
27	A:	B:	C:	D:	E:	F:
28	A:	B:	C:	D:	E:	
29	A:	B:	C:			
30	A:	B:	C:	D:	E:	
31	A:	B:	C:	D:	E:	
32	A:	B:	C:	D:	E:	
33	A:	B:	C:	D:	E:	
34	A:	B:	C:	D:	E:	
35	A:	B:	C:	D:	E:	
36	A:	B:	C:	D:	E:	

