

Name

CWID

Final Exam

December 10th, 2014

8:00-10:00

CS425 - Database Organization

Please leave this empty!

1.1 1.2 1.3 1.4 1.5 1.6

Sum

Instructions

- Try to answer all the questions using what you have learned in class. Keep hard questions until the end.
- **When writing a query, write the query in a way that it would work over all possible database instances and not just for the given example instance!**
- The exam is closed book and closed notes!

Consider the following database schema and example instance about music albums:

song

title	length	writtenBy	writtenYear
De Charybde en Scylla	5:35	Louis Sclavis	2008
Un Vent Noir	3:40	Louis Sclavis	2008
Standing Ovation	4:26	Louis Sclavis	1995
moto [ein weltbewusstsein.]	6:00	Michael Wertmueller	2000

songInAlbum

song	album	trackNumber
De Charybde en Scylla	Lost on the Way	1
Un Vent Noir	Lost on the Way	9
Standing Ovation	Carnet De Routes	1
moto [ein weltbewusstsein.]	Werthmueller	4

album

aTitle	releaseYear	producedBy	playedBy
Lost on the Way	2008	ECM	Louis Sclavis
Carnet de Routes	1995	Label Blue	Louis Sclavis
Werthmueller	2005	GROB	Steamboat Switzerland

artist

name	nationality
Louis Sclavis	french
Steamboat Switzerland	swiss
Michael Wertmueller	swiss

label

name	revenue
Sony	50,000,000
ECM	3,000,000
Label Bleu	150,000
GROB	10,000

Hints:

- Attributes with black background form the primary key of a relation (.e.g, *title* for relation *song*)
- The attribute *writtenBy* of relation *song* is a foreign key to relation *artist*. Note that the artist writing a song is not necessarily the artist playing the song.
- The attribute *album* of relation *songInAlbum* is a foreign key to relation *album*.
- The attribute *song* of relation *songInAlbum* is a foreign key to relation *song*.

- The attribute *producedBy* of relation *album* is a foreign key to relation *label*.
- The attribute *playedBy* of relation *album* is a foreign key to relation *artist*.
- All foreign keys have been created with the **CASCADE** option.

Part 1.1 Relational Algebra (Total: 20 Points)

The queries in this part should be written using the set semantics version of relational algebra.

Question 1.1.1 (4 Points)

Write a **relational algebra** expression that returns the titles and length of all songs written by french artists (*nationality*) in 2008.

Question 1.1.2 (4 Points)

Write a **relational algebra** expression that returns the names of labels that produced at least one album in 1993.

Question 1.1.3 (5 Points)

Write a **relational algebra** expression that returns the number of albums produced by the ECM label per year.

Question 1.1.4 (7 Points)

Write a **relational algebra** expression that returns labels that have not released (*producedBy*) any albums by swiss artists.

Part 1.2 SQL - DDL (Total: 7 Points)

Question 1.2.1 (7 Points)

Write an **SQL statement** that creates a new relation *reviews* that records information about reviews of albums written by critics. For each review we want to store the magazine the review was published in as well as the release year and issue. Furthermore, we store the name of the critic who has written the review, the album title, the rating (one out of “positive”, “negative”, “neutral”), and the actual text of the article. A review can be uniquely identified by the magazine and album. We require that every review has a critic, rating, and text.

Part 1.3 SQL - Queries (Total: 20 Points)

Question 1.3.1 (3 Points)

Write an **SQL query** that returns the total length for each album, i.e., the total length of all tracks of each album. Return two attributes: album title and the total length.

Question 1.3.2 (5 Points)

Write an **SQL query** that returns the names of all artists that have only played songs (have songs included in albums they have played on) that they have written themselves.

Question 1.3.3 (6 Points)

Write an **SQL query** that returns artists that have worked for all labels, i.e., for which there exists for every label at least one album published by this label and played by this artist.


Question 1.3.4 (6 Points)

Write an **SQL query** that returns the year in which the maximum amount of music was produced. That is, the year with the maximal total length of all songs written in this year.

Part 1.4 Translation of ER into Relational Model (Total: 22 Points)

Question 1.4.1 (22 Points)

Take the following ER-model and translate it into a relational schema using the rules presented in class. Present the relational schema using the notation from the slides. For example, a relation R with attributes a_1 and a_2 where a_2 is the primary key is written as $R(a_1, \underline{a_2})$. You do not need to specify foreign key constraints. You do not have to show intermediate results.



Part 1.5 Normalization and Functional Dependencies (Total: 22 Points)

Consider the following relation $R(A, B, C, D, E)$ and functional dependencies F that hold over this relation.

$$F = A, B \rightarrow C$$

$$A \rightarrow E$$

$$C \rightarrow D, E$$

$$D \rightarrow B, E$$

Question 1.5.1 (4 Points)

Determine all candidate keys of R .

Question 1.5.2 (8 Points)

Compute the canonical cover of F . Show each step of the generation according to the algorithm shown in class.

Question 1.5.3 (6 Points)

Apply the 3NF decomposition algorithm to relation R .

Question 1.5.4 (4 Points)

Which of the statements below are true and which ones are false? Put a check mark into the box to indicate that a statement is correct or an X to indicate that a statement is wrong. You also have the option to leave the check box empty if you are unsure. Incorrect answers will result in negative points (you will get at least 0 points for this question).

- In 2NF, for every non-trivial dependency $\alpha \rightarrow \beta$, α is a superkey
- Every superkey is a candidate key
- Every candidate key is a superkey
- A relation can have at most one superkey
- Let α and β be candidate keys. A functional dependency $\alpha \rightarrow \beta$ does not violate 3NF.
- Testing whether a relation is in 3NF is harder than decomposing a relation into 3NF. This is because testing requires the generation of all candidate keys for a relation.
- Every relation in BCNF is also in 2NF.
- The difference between BCNF and 3NF is that 3NF allows dependencies between parts of candidate keys.

Part 1.6 Concurrency Control (Total: 9 Points)

Question 1.6.1 (9 Points)

For each of the following schedules determine which properties this schedule has. E.g., a schedule may be *recoverable* and *cascade-less (strict)*. Consider the following notation for operations of transactions:

$w_1(A)$ transaction 1 wrote item A
 $r_1(A)$ transaction 1 read item A
 c_1 transaction 1 commits
 a_1 transaction 1 aborts

$$S_1 = r_1(C), r_1(B), r_2(A), w_2(A), w_1(B), r_2(C), w_2(C), c_1, w_2(B), c_2$$

$$S_2 = r_2(C), w_2(A), w_1(A), w_2(B), c_2, r_3(B), c_1, c_3$$

$$S_3 = r_2(A), r_1(B), w_2(A), r_2(B), r_3(A), w_1(B), c_1, w_3(A), c_3, w_2(B), c_2$$

- S_1 is recoverable
 - S_1 is cascade-less
 - S_1 is conflict-serializable
 - S_1 is 2PL
 - S_1 is S2PL
 - S_1 is SS2PL
-

- S_2 is recoverable
 - S_2 is cascade-less
 - S_2 is conflict-serializable
 - S_2 is 2PL
 - S_2 is S2PL
 - S_2 is SS2PL
-

- S_3 is recoverable
- S_3 is cascade-less
- S_3 is conflict-serializable
- S_3 is 2PL
- S_3 is S2PL
- S_3 is SS2PL

