

Name

CWID

# Final Exam

May 4th, 2014

8:00-10:00

## CS520 - Database Organization Results

---

*Please leave this empty!*

1.1  1.2  1.3  1.4  1.5

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! No calculator, smartphones, or similar allowed!

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 Datalog (Total: 24 Points)

Recall that datalog applies set semantics.

### Question 1.1.1 (3 Points)

Write a **datalog program** that returns the titles and artists (attribute *playedBy*) of all albums in the database.

### Solution

$$Q(X, Y) : \text{--album}(X, A, B, Y).$$

### Question 1.1.2 (3 Points)

Write a **datalog program** that returns the titles of albums played by swiss artists (*nationality* attribute)

### Solution

$$Q(X) : \text{--album}(X, Y, Z, A), \text{artist}(A, \text{swiss}).$$

### Question 1.1.3 (4 Points)

Write a **datalog program** that returns the length of all songs written in 2008 (attribute *writtenYear*) or written by Louis Sclavis (attribute *writtenBy*).

### Solution

$$Q(Y) : \text{--song}(X, Y, Z, 2008).$$
$$Q(Y) : \text{--song}(X, Y, \text{LouisSclavis}, Z).$$

### Question 1.1.4 (4 Points)

Write a **datalog program** that returns names of artists that have released at least one album produced by label ECM (attribute *producedBy*) and at least one album released by label GROB.

### Solution

$$Q(X) : \text{--album}(A, B, \text{ECM}, X), \text{album}(A', B', \text{GROB}, X).$$

### Question 1.1.5 (5 Points)

Write a **datalog program** that returns songs (their title) played by (songs is on an album played by) french artists or written by french artists (attribute *writtenBy*).

### Solution

$$Q(X_1) : \text{--song}(X_1, X_2, X_3, X_4), \text{songInAlbum}(X_1, Z_1, Y_1), \text{album}(Z_1, Z_2, Z_3, A_1), \text{artist}(A_1, \text{french}).$$
$$Q(X_1) : \text{--song}(X_1, X_2, A_1, X_4), \text{artist}(A_1, \text{french}).$$

### Question 1.1.6 (5 Points)

Consider two artists to be connected if they have released albums with at least one common label. For instance, if Peter has released an album with ECM and Bob has also released an album with ECM then they are considered connected. Write a **datalog program** that computes all pairs of artists that are directly or indirectly connected. For example, if Bob is connected to Alice and Alice is connected to Peter, then Bob is also connected to Peter.

### Solution

$$\text{artistLabel}(X, Y) : \text{--artist}(X, Z), \text{album}(A, B, Y, X).$$
$$\text{direct}(X, Y) : \text{--artistLabel}(X, Z), \text{artistLabel}(Y, Z).$$
$$\text{connected}(X, Y) : \text{--direct}(X, Y).$$
$$\text{connected}(X, Y) : \text{--connected}(X, Z), \text{direct}(Z, Y).$$

## Part 1.2 Constraints (Total: 20 Points)

### Question 1.2.1 Expressing Constraints in First-Order Logic (10 Points)

Recall the logical representation of constraints introduced in class. Write down the logical definition for the following constraints for the example schema:

- The primary key of relation *song*
- The foreign key from attribute *playedBy* of relation *album* to relation *artist*.
- The foreign key from attribute *album* of relation *songInAlbum* to relation *album*.
- The foreign key from attribute *song* of relation *songInAlbum* to relation *song*.
- The following functional dependency for relation *album*:  $releaseYear, playedBy \rightarrow producedBy$

### Solution

$$\begin{aligned}PK(song) &: \forall t, l, w, y, l', w', y' : song(t, l, w, y), song(t, l', w', y') \rightarrow l = l' \wedge w = w' \wedge y = y' \\FK(playedBy) &: \forall t, r, p, b : album(t, r, p, b) \rightarrow \exists n : artist(b, n) \\FK(album) &: \forall s, a, t : songInAlbum(s, a, t) \rightarrow r, p, b : album(a, r, p, b) \\FK(song) &: \forall s, a, t : songInAlbum(s, a, t) \rightarrow l, w, y : song(s, l, w, y) \\FD &: \forall t, r, p, b, t', p' : album(t, r, p, b) \wedge album(t', r, p', b) \rightarrow p = p'\end{aligned}$$

## Question 1.2.2 Creating Denial Constraints (10 Points)

Create denial constraints over the example schema based on the following descriptions.

- Albums released before 2000 cannot have more than 20 tracks. You can assume that tracks are numbered consecutively starting from 1. For instance, the track numbers for an album with three tracks would be numbered 1,2, and 3.
- No label has a revenue higher than 50,000,000.
- All songs are between 2:00 and 20:00 minutes long.
- The first track of an album is never longer than 5:00 minutes.

## Solution

$$d_1 : \forall x, y, z, a, b, c : \neg(\text{songInAlbum}(x, y, z) \wedge \text{album}(y, a, b, c) \wedge a < 2000 \wedge z > 20)$$

$$d_2 : \forall x, y : \neg(\text{label}(x, y) \wedge y > 50,000,000)$$

$$d_3' : \forall a, b, c, d : \neg(\text{song}(a, b, c, d) \wedge b < 2:00)$$

$$d_3'' : \forall a, b, c, d : \neg(\text{song}(a, b, c, d) \wedge b > 20:00)$$

$$d_4 : \forall x, y, z, a, b, c : \neg(\text{songInAlbum}(x, y, z) \wedge \text{song}(x, a, b, c) \wedge z = 1 \wedge a > 5:00)$$



## Part 1.3 Query Containment And Equivalence (Total: 18 Points)

### Question 1.3.1 (18 Points)

Consider the 3 queries shown below. Check all possible containment relationships. If there exists a containment mapping from  $Q_i$  to  $Q_j$  then write down the mapping.

$$Q_1(X, Y) : \neg R(X, Y), R(Z, Y).$$

$$Q_2(X, Y) : \neg R(X, Y), R(Y, X).$$

$$Q_3(X, Y) : \neg R(X, X), R(Y, Y).$$

### Solution

$Q_1 \rightarrow Q_2$ :

$$X \rightarrow X$$

$$Y \rightarrow Y$$

$$Z \rightarrow X$$

$Q_1 \rightarrow Q_3$ :

**no containment mapping exists**

$Q_2 \rightarrow Q_1$ :

**no containment mapping exists**

$Q_2 \rightarrow Q_3$ :

**no containment mapping exists**

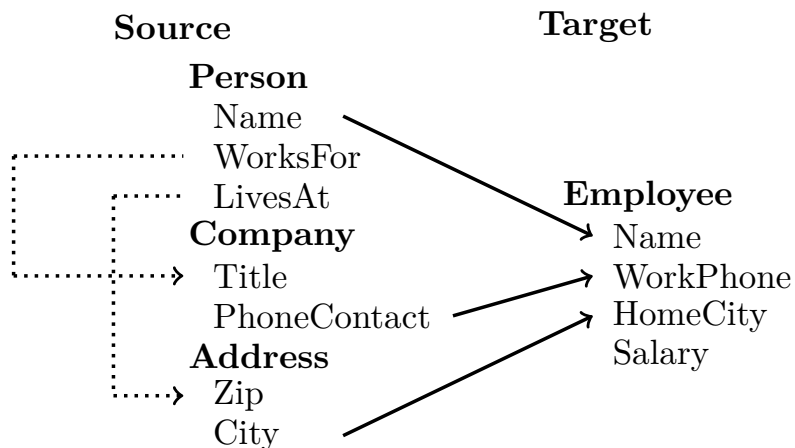
$Q_3 \rightarrow Q_1$ :

**no containment mapping exists**

$Q_3 \rightarrow Q_2$ :

**no containment mapping exists**

## Part 1.4 Schema Mappings (Total: 16 Points)



### Question 1.4.1 (16 Points)

Consider the source and target schemas shown above. Arrows from source attributes to target attributes represent schema matches. The dotted arrows represent foreign key constraints. Use the Clio algorithm presented in the data exchange part of class to 1) determine all source and target associations and 2) construct schema mappings using these associations and the schema matches (attribute correspondences). Do not generate mappings that are subsumed by other mappings.

#### Solution

Since there are no foreign keys all source relations are single relation. Source associations are:

$$\begin{aligned}
 & Person(X_1, X_2, X_3), Company(X_2, X_4), Address(X_3, X_5) \\
 & Company(X_1, X_2) \\
 & Address(X_1, X_2)
 \end{aligned}$$

Target associations are:

$$Employee(Y_1, Y_2, Y_3, Y_4)$$

Based on this we get 3 possible combinations of source with target associations. The resulting mappings are:

$$\begin{aligned}
 m_1 & : \forall x, y, z, a, b : Person(x, y, z) \wedge Company(y, a) \wedge Address(z, b) \rightarrow \exists c : Employee(x, a, b, c) \\
 m_2 & : \forall x, y : Company(x, y) \rightarrow \exists a, b, c : Employee(a, y, b, c) \\
 m_3 & : \forall x, y : Address(x, y) \rightarrow \exists a, b, c : Employee(a, b, y, c)
 \end{aligned}$$

For this question there are no subsumed mappings.



## Part 1.5 Virtual Data Integration (Total: 22 Points)

### Question 1.5.1 (24 Points)

Use the mappings from the previous part (recall the st-tgds are equivalent to GLAV mappings under the open world assumption). As a first step write down the GLAV expressions equivalent to these st-tgs. Then find a maximally contained rewriting for the query shown below using the bucket algorithm.

$$Q(W) : \neg Employee(U, V, W, X)$$

### Solution

The GLAV versions of the mappings are

$$\begin{aligned}
Q(X, A, B)_1^G &: -Employee(X, A, B, C) \supseteq Q_1^L(X, A, B) : -Person(X, Y, Z), Company(Y, A), Address(Z, B) \\
Q(Y)_2^G &: -Employee(A, Y, B, C) \supseteq Q_2^L(Y) : -Company(X, Y) \\
Q(Y)_3^G &: -Employee(A, B, Y, C) \supseteq Q_3^L(Y) : -Address(X, Y)
\end{aligned}$$

## Rewriting $Q$

This query returns home cities. The query has one atom, so we create only one bucket containing all views that contain relation **Employee**.

All three views contain this relation, but only  $Q_1^G$  and  $Q_3^G$  have the variable at the 3rd position of relation **Employee** in their head. Thus, we generate two potential rewritings:

$$\begin{aligned}
Q'(W) &= Q_1^G(X', A', W) \\
Q''(W) &= Q_3^G(W)
\end{aligned}$$

Both rewritings fulfill the containment check after unfolding:

$$\begin{aligned}
Q'(W) &= Employee(X', A', W, C') \\
Q''(W) &= Employee(A', B', W, C')
\end{aligned}$$

with containment mappings

$$\begin{aligned}
U &\rightarrow X' \\
V &\rightarrow A' \\
W &\rightarrow W \\
X &\rightarrow C'
\end{aligned}$$

and

$$\begin{aligned}
U &\rightarrow A' \\
V &\rightarrow B' \\
W &\rightarrow W \\
X &\rightarrow C'
\end{aligned}$$

Thus, both are included in a maximally contained rewriting. After replacing the global with the local views and unfolding them we get:

$$\begin{aligned}
Q(W) &: -Person(X', Y, Z), Company(Y, A'), Address(Z, W) \\
Q(W) &: -Address(X, W)
\end{aligned}$$



