

Name

CWID

Midterm Exam

March 19th, 2018

1:50-3:05

CS520 - Data Integration, Warehousing, and Provenance

Please leave this empty!

1.1

1.2

1.3

1.4

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:

user

nickname	name	postVisibility	country
BobAwesome	Bob	FOF	USA
Ali12	Alice	friends	France
Peter	Peter	friends	India
Pokegert	Gert	public	China

friends

person	friend
BobAwesome	Ali12
Ali12	BobAwesome
BobAwesome	Peter
Peter	Pokegert

posts

pid	user	text	time
1	BobAwesome	Hello just brought ...	2018-01-10
2	BobAwesome	meet @Ali12 at ...	2018-01-11
3	Peter	... is great, would recommend.	2018-01-15

Hints:

- Attributes with black background form the primary key of a relation (e.g., *nickname* for relation *user*)
- The attributes *person* and *friend* of relation *friends* are a foreign key to relation *user*.
- The attribute *user* of relation *posts* is a foreign key to relation *user*.

Part 1.1 Datalog (Total: 25 Points)

Recall that Datalog applies set semantics.

Question 1.1.1 (4 Points)

Write a **Datalog program** that returns the *nicknames* of users which have posted at least one post.

Question 1.1.2 (6 Points)

Write a **Datalog program** that returns the nicknames of users that are friends of a user with nickname BobAwesome or are friends of friends of BobAwesome. For instance, in the example database Pokegert is a friend of one of BobAwesome's friends.

Question 1.1.3 (6 Points)

Write a **Datalog program** that returns the names of users that have not written any posts. Note that we are asking for names and not nicknames.

Question 1.1.4 (9 Points)

Write a **Datalog program** that returns all pairs of users (X, Y) such that user X can see the posts of user Y . Represent users using their nicknames. Whether the posts of a user are visible to another user is determined based on the `postVisibility` attribute for the user. The possible values are

- “friends”: only friends of the user can see the users posts
- “FOF”: in addition to friends, also friends of friends of the user can see the users posts
- “public”: everybody can see the users posts

For example, one result returned for the example database would be `(Pokegert,BobAwesome)`, because user `BobAwesome` has set his post visibility to `FOF` and `Pokegert` is a friend of `Peter` who is a friend of `BobAwesome`.

Part 1.2 Constraints (Total: 30 Points)

Question 1.2.1 Expressing Constraints in First-Order Logic (15 Points)

Recall the representation of constraints as universally quantified formulas in first-order logic introduced in class. Write down the logical formalisms encoding the following constraints over the example schema:

- The foreign key from attribute `person` of relation `friends` to relation `user`.
- The foreign key from attribute `friend` of relation `friends` to relation `user`.
- The primary key of relation `user`
- The following functional dependency for relation `posts`: $user, time \rightarrow text$

Question 1.2.2 Creating Denial Constraints (15 Points)

Create denial constraints over the example schema based on the following descriptions. **Some of the constraints may require you to write more than one denial constraint!**

- Users from France cannot have a `postVisibility` of “*public*”.
- Users from USA cannot be friends with users from France and vice versa.
- Implement the functional dependency `user, time → text` over relation `posts`

Part 1.3 Query Containment And Equivalence (Total: 27 Points)

Question 1.3.1 (27 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.

$Q_1(X) :- R(X, Z), R(X, X), R(Z, A).$

$Q_2(X) :- R(X, Y), R(Y, Z).$

$Q_3(X) :- R(X, Z), R(Z, A), R(A, Y).$

Part 1.4 Virtual Data Integration (Total: 18 Points)

Question 1.4.1 (18 Points)

Given are the views and schema shown below. Find a maximally contained rewriting of the following query using these views. The query returns titles of books, their authors, price, and year where the author is from IIT. **You are allowed to use any of the algorithms discussed in class such as the Bucket algorithm. However, only the result counts, you are not required to use any of these algorithms.**

```
Q(Title, Author, Price, Genre) :- Books(Title, X, Y, Genre, Z, Price),
    Author(Title, Author),
    Person(Author, A, 'IIT').
```

Views

```
V1(Title, Pub, Genre, Price, Year) :- Books(Title, Pub, X, Genre, Year, Price),
    Genre = 'CS'.
```

```
V2(Title, Pub, Price, Ed) :- Books(Title, Pub, Ed, X, Y, Price).
```

```
V3(Title, Author, Affl) :- Author(Title, Author),
    Person(Author, 'Dr.', Affl).
```

```
V4(Author) :- Author(X, Author).
```

```
V5(Title, Author) :- Books(Title, X, Y, Z, Year, A),
    Author(Title, Author),
    Person(Author, B, 'IIT'),
    Year = 2017.
```

Schema

```
Books(title, publisher, edition, genre, year, price)
Author(bookTitle, authorName)
Person(name, title, affiliation)
Publisher(name, country)
```


