

CARNEGIE MELLON UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE
15-415/615 - DATABASE APPLICATIONS
C. FALOUTSOS & A. PAVLO, FALL 2016

Homework 1 (by Prashanth Menon)

Due: hard copy, in class at 3:00pm, on Monday, Sep. 19

VERY IMPORTANT: Deposit **hard copy** of your answers, in class. For ease of grading, please

1. **Separate** your answers, on different page(s) for each question (staple additional pages, if needed).
2. **Type** the full info on **each** page: your **name**, **Andrew ID**, **course#**, **Homework#**, **Question#** on each of the 5 pages.

Reminders:

- *Plagiarism:* Homework is to be completed *individually*.
- *Typeset* all of your answers whenever possible. Illegible handwriting may get zero points, at the discretion of the graders.
- *Late homeworks:* in that case, please email it
 - to all TAs
 - with the subject line exactly 15-415 Homework Submission (HW 1)
 - and the count of slip-days you are using.

For your information:

- Graded out of **100** points; **5** questions total
- Rough time estimate: *approx. 6 hours* - 1 to 2 hours per question

Revision : 2016/09/09 14:57

Question	Points	Score
Entity-Relationship Diagram	25	
SQL Tables from the ER Model	15	
Relational Algebra	30	
Relational Tuple Calculus (RTC)	10	
Relational Domain Calculus (RDC)	20	
Total:	100	

Question 1: Entity-Relationship Diagram [25 points]

On separate page, with '[course-id] [hw#] [question#] [andrew-id] [your-name]'

Consider a database to store information for a social networking website. The database has the following properties:

- Every user has a unique user ID (integer) along with a full name, age and phone number.
- Every group has a unique group ID (integer) and a name. Every group must have at least one user that serves as moderator of the group.
- A user may be a member of zero or more groups; groups may contain zero or more members (and one or more moderators).
- Users are allowed to create zero or more albums. An album has a unique album ID (integer), a creation date, and a name. An album is owned by exactly one user: the user that created it.
- An album can contain zero or more media files. For every media file, we record its unique URL, the date the file was added to the album, and a caption (if one exists).
- Users can zero or more photos to albums. Photos are a type of media file, but we also track the encoding (e.g., JPEG, PNG, etc.) and the size of the photo (in bytes).
- Users may add zero or more videos to albums. Videos are a type of media file, and we track the codec used to encode the video (e.g., MPEG-4), the length of the video (in seconds), and the video's bitrate.
- A media file may belong to at most one album.

Given this description of the database and its constraints, we have created a mostly correct Entity-Relationship Diagram, shown in Figure 1. This diagram is at this [hyperlink](#) - feel free to use it as a starting point.

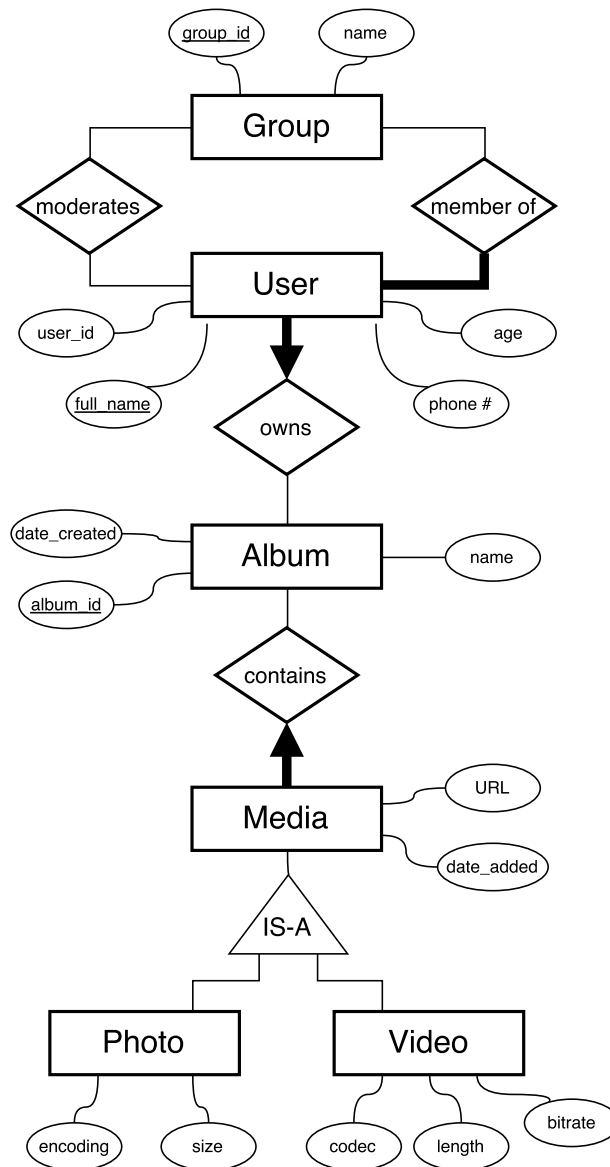


Figure 1: Almost correct ER diagram

- (a) **[10 points]** Find and correct any mistakes in the given ER diagram. For this sub-question, list all the modifications (in the next one, we will ask you to list your additions/deletions). Specifically, here, number and list the modifications, like, e.g.
1. delete: arrow, from x to y
 2. change to bold line: thin line, from z to w
 3. change to bold box: entity e
 4. underline: attribute a of entity e
- (b) **[5 points]** There may also be some missing element(s). If none, say 'none' - otherwise, add them to the picture, **and** list them, numbered. E.g.
1. add: attribute a , to entity e

2. add: bold line, arrow, from c to d .
 3. add: weak entity, f , with attributes . . .
- (c) [**10 points**] List and number all the bold lines and all the arrows that are in the final, corrected version of the diagram. E.g.
1. **bold**, line, from **Department** to **Employee**
 2. thin, **arrow**, from x to y

Clarifications/Hints:

- List your assumptions, if any. We will accept all reasonable assumptions.
- OPTIONALLY, to help graders save time, you may give your corrected ER diagram.

Question 2: SQL Tables from the ER Model..... [15 points]

On separate page, with '[course-id] [hw#] [question#] [andrew-id] [your-name]'

Consider a database for an apartment-tracking application. It records information about apartment buildings, units and tenants. The constraints are exactly as shown in Figure 2. Apartment buildings, units and tenants all have unique identifiers as shown in the figure.

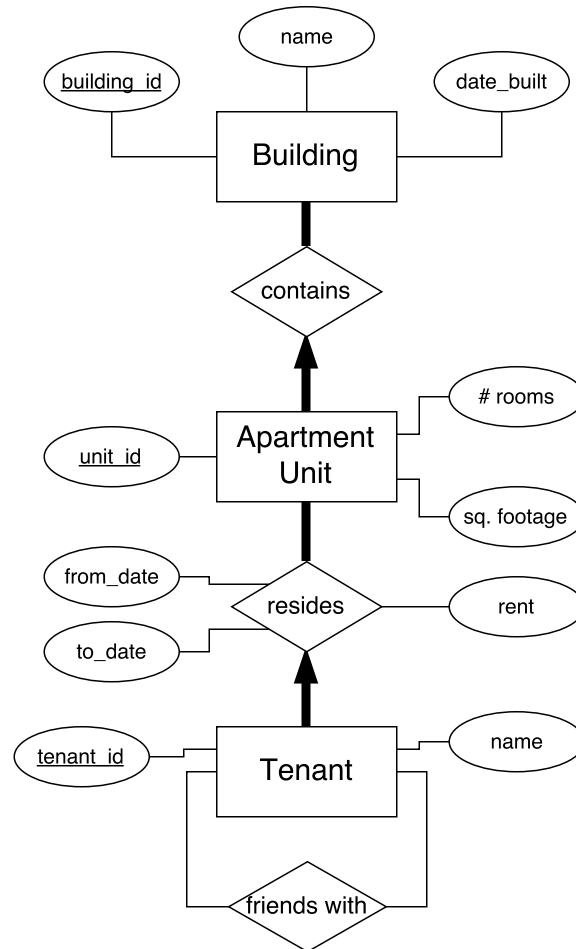


Figure 2: ER diagram for apartment tracking: turn to SQL tables

- (a) [15 points] Give the DDL statements, that correspond to the above ER diagram.
- Use proper data types (we'll accept all reasonable choices).
 - Avoid syntax errors (we'll forgive missing semicolons).
 - Specify your decisions with respect to **CASCADE** deletions.
 - **Without** using **CHECK** statements, enforce as many as possible of the implied integrity constraints as you can.

Question 3: Relational Algebra [30 points]

On separate page, with '[course-id] [hw#] [question#] [andrew-id] [your-name]'

Consider the relations of a database for the 2016 Olympics as shown in Table 1. These relations record the athletes, events, and outcomes/results of the 2016 Olympic games. FYI, the *optional*, zero-point questions, will be used for partial credit, if needed

athlete_id	country	name	age
A1	U.S.A.	Michael Phelps	31
A2	U.S.A.	Justin Gatlin	34
A3	U.S.A.	Ryan Lochte	32
A4	Canada	Andre De Grasse	21
A5	Jamaica	Usain Bolt	30
A6	France	Christophe Lemaitre	26
A7	Japan	Masato Sakai	24
A8	Japan	Naito Ehara	60
A9	GBR	Duncan Scott	35
A10	GBR	James Guy	32

(a) Athletes Table

event_id	name
E1	100m Sprint
E2	200m Sprint
E3	200m Butterfly
E4	4x200 Freestyle Relay

(b) Events Table

event_id	athlete_id	result
E1	A5	Gold
E1	A2	Silver
E1	A4	Bronze
E2	A5	Gold
E2	A4	Silver
E3	A1	Gold
E3	A7	Silver
E3	A9	Bronze
E4	A1	Gold
E4	A3	Gold
E4	A7	Silver
E4	A8	Silver
E4	A9	Bronze
E4	A10	Bronze

(c) Event_Results Table: the outcome of every event

Table 1: Relations of the Olympic database.

We have the following tables:

- **Athletes:** For every athlete, we record a unique `athlete_id`, the country they represent, their name, and their age.
- **Events:** This table lists all the events that are part of the 2016 Olympic games. Every event has a unique integer `event_id` and a name.
- **Event_Results:** Lists the outcomes of all events. Every outcome records the `event_id` of the event, the `athlete_id` of the athlete that won a medal in the event, (or was part of the team that won), and the standing of the athlete (i.e., gold, silver or bronze).

Given this database instance, answer the following questions:

- (a) **[2 points]** Which of the following is the meaning of the expression:

$$\sigma_{\text{age} \geq 25}(\text{Athletes})$$

1. It lists the `athlete_id` and `name` of all athletes that are at least 25 years old.
 2. It lists all attributes of all athletes whose `age` is greater than or equal to 25.
 3. It lists all athlete `ages` that are greater than or equal to 25 .
 4. None of the above. The real answer is
- (b) **[5 points]** We want to list the `names` of the athletes that have won at-least one gold medal (eliminating duplicate names) . Which, if any, of the following expressions achieve that ? Mark all valid expressions.

1. $\pi_{\text{name}}(\sigma_{\text{result}='Gold'}(\text{Athletes} \bowtie \text{Event_Results}))$
2. $\pi_{\text{name}}(\text{Athletes} \bowtie \sigma_{\text{result}='Gold'}(\text{Event_Results}))$
3. $\pi_{\text{name}}(\sigma_{\text{result}='Gold'}(\text{Athletes} \bowtie \pi_{\text{athlete_id,result}}(\text{Event_Results})))$
4. $\pi_{\text{name}}(\text{Athletes}) \bowtie \sigma_{\text{result}='Gold'}(\text{Event_Results})$
5. $\pi_{\text{name}}(\text{Athletes}) - \pi_{\text{name}}(\text{Athletes} \bowtie \sigma_{\text{result} \neq 'Gold'}(\text{Event_Results}))$

- (c) For the following expression:

$$\sigma_{\text{age} < 25}(\text{Athletes} \bowtie \text{Event_Results})$$

- i. **[0 points]** *Optional:* describe in English what the expression does
 - ii. **[1 point]** How many, and which are the columns (= attributes) in the answer?
 - iii. **[3 points]** How many tuples are in the answer?
 - iv. **[3 points]** List all the tuples in the answer, as a table.
- (d) For the following expression:

$$\pi_{\text{athlete_id,event_id}}(\text{Event_Results}) \div \pi_{\text{event_id}}(\sigma_{\text{athlete_id}='A5'}(\text{Event_Results}))$$

- i. **[0 points]** *Optional:* describe in English what the expression does
- ii. **[2 points]** How many, and which are the columns (= attributes) in the answer?

- iii. [3 points] How many tuples are in the answer?
 - iv. [3 points] List all the tuples in the answer, as a table.
- (e) For the following expression:

$$\pi_{A.athlete_id}(\rho_A(\text{Athletes})) - \pi_{ER1.athlete_id}(\rho_{ER1}(\text{Event_Results}) \bowtie_{ER1.athlete_id=ER2.athlete_id \wedge ER1.result \neq ER2.result} \rho_{ER2}(\text{Event_Results}))$$

- i. [0 points] *Optional:* describe in English what the expression does
- ii. [2 points] How many, and which are the columns (= attributes) in the answer?
- iii. [3 points] How many tuples are in the answer?
- iv. [3 points] List all the tuples in the answer, as a table.

Question 4: Relational Tuple Calculus (RTC) [10 points]

On separate page, with '[course-id] [hw#] [question#] [andrew-id] [your-name]'

We will again use the Olympic database from the previous question (see Table 1).

(a) For the following RTC expression

$$\{t \mid \exists er \in \text{Event_Results} (er.athlete_id = \text{"A4"} \wedge t.event_id = er.event_id)\}$$

- i. [0 points] *Optional:* describe in English what the expression does
- ii. [1 point] How many, and which are the columns (= attributes) in the answer?
- iii. [2 points] How many tuples are in the answer?
- iv. [2 points] List all the tuples in the answer, as a table.

(b) For the following RTC expression

$$\{t \mid \exists er1, er2 \in \text{Event_Results} \\ (er1.event_id \neq er2.event_id \\ \wedge er1.athlete_id = er2.athlete_id \\ \wedge t.athlete_id = er1.athlete_id)\}$$

- i. [0 points] *Optional:* describe in English what the expression does
- ii. [1 point] How many, and which are the columns (= attributes) in the answer?
- iii. [2 points] How many tuples are in the answer?
- iv. [2 points] List all the tuples in the answer, as a table.

Question 5: Relational Domain Calculus (RDC) [20 points]

On separate page, with '[course-id] [hw#] [question#] [andrew-id] [your-name]'

For this question, we will reuse the 2016 Olympic database from Question 3 (see Table 1).

(a) For the following RDC expression

$$\{\langle n \rangle \mid \exists ai, \exists c, \exists n, \exists a (\langle ai, c, n, a \rangle \in \mathbf{Athletes} \wedge a \geq 35)\}$$

- i. [0 points] *Optional:* describe in English what the expression does
- ii. [1 point] How many, and which are the columns (= attributes) in the answer?
- iii. [2 points] How many tuples are in the answer?
- iv. [2 points] List all the tuples in the answer, as a table.

(b) For the following RDC expression:

$$\{\langle en \rangle \mid \exists ai, \exists an, \exists aa, \exists ei, \exists en (\\ \langle ai, \text{"Canada"}, an, aa \rangle \in \mathbf{Athletes} \\ \wedge \langle ei, en \rangle \in \mathbf{Events} \\ \wedge \langle ei, ai, \text{"Silver"} \rangle \in \mathbf{Event_Results})\}$$

- i. [0 points] *Optional:* describe in English what the expression does
- ii. [1 point] How many, and which are the columns (= attributes) in the answer?
- iii. [2 points] How many tuples are in the answer?
- iv. [2 points] List all the tuples in the answer, as a table.

(c) For the following RDC expression:

$$\left\{ \langle c, an1, aa1 \rangle \mid \exists ai1 \left(\langle ai1, c, an1, aa1 \rangle \in \mathbf{Athletes} \right. \right. \\ \left. \left. \wedge \forall ai2 \left(\exists aa2, \exists an2 ((\langle ai2, c, an2, aa2 \rangle \in \mathbf{Athletes}) \Rightarrow (aa1 \geq aa2)) \right) \right) \right\}$$

- i. [0 points] *Optional:* describe in English what the expression does
- ii. [1 point] How many, and which are the columns (= attributes) in the answer?
- iii. [4 points] How many tuples are in the answer?
- iv. [5 points] List all the tuples in the answer, as a table.