

CS 450 - Database Concepts

Spring 2012

Instructor: Dr. Jessica Lin

Homework 2 – Due Feb 29 at 4:35pm, both in class and electronically

Problem 1.

Consider the following schema for a discussion forum like Piazza. Note this is a modified version from the one given in HW1. In HW1, the courses do not have unique identifiers. For simplicity, in this assignment each course is given a unique identifier. Also, login names are used instead of email addresses.

For each relation, the attribute(s) of the primary key is(are) underlined. The foreign key constraints are also specified.

Users(Login, First_name, Last_name, Age, Member_since)

Schools(SID, Sname, City, State, Zip)

Courses(CID, SID, Code, Number, Sec, Semester, Title, Status)

Foreign Key: SID References Schools(SID)

Enroll_in(Login, CID, Role, Notification_Preference)

Foreign Key: Login References Users(Login)

Foreign Key: CID References Courses(CID)

Posts(PID, Posted_By, CID, Content, Post_Time, Original_PID)

Foreign Key: Posted_By References Users(Login)

Foreign Key: CID References Courses(CID)

Foreign Key: Original_PID References Posts(PID)



(Note the self-referencing foreign key in the last relation.)

Write the following queries in Relational Algebra. Note “dsmith” is a unique user ID.

1. Print the PIDs of all posts for this class (CS450, Section 001, Spring12, GMU).
2. Print the full names (First name, Last name) of all users who are enrolled as an Instructor in at least one course.
3. Print the full names of all users who are enrolled as both an Instructor and as a Student (in different classes, of course).
4. Find the full name of the youngest user registered on this site, and the CIDs all active courses that he/she is enrolled in.

5. Print the CIDs of courses that have members but no posts.
6. A new post has just been made for this class. In order to determine whom to send the notification to, list the names of members for this class who set their notification preference to “Real-time”.

For the following questions, you can consider the example posts hierarchy and the corresponding relation instance provided below. **However, your query should be general and not be based on this (or any) relation instance.**



P07 (chris)

| PID | Posted_by | CID | Content | Post_Time | Original_PID |
|-----|-----------|------|---------|-----------|--------------|
| P01 | dsmith | C001 | XXX | 2/17/12 | Null |
| P02 | dsmith | C001 | XXX | 2/18/12 | Null |
| P03 | erin | C001 | XXX | 2/18/12 | P01 |
| P04 | jane | C001 | XXX | 2/19/12 | P03 |
| P05 | chris | C001 | XXX | 2/19/12 | P01 |
| P06 | erin | C001 | XXX | 2/19/12 | P02 |
| P07 | chris | C001 | XXX | 2/19/12 | Null |
| P08 | dsmith | C001 | XXX | 2/19/12 | P05 |
| P09 | chris | C001 | XXX | 2/20/12 | P08 |

7. For posts made by *dsmith*, print the full names of users that posted follow-ups on those posts. In this example, *dsmith* has 3 posts: P01, P02, and P08. P01 has two direct follow-ups (P03 and P05); P02 has one direct follow-up (P06); and P08 has one direct follow-up (P09). Your query should print the full names of users who posted P03, P05, P06, and P09, respectively.
8. Print the PIDs of all original posts (i.e. not a follow-up) made for this class. In this example, the original posts are P01, P02, and P07.
9. Print the PIDs of all *original* posts made for this class that have been followed up on (i.e. have at least one reply). In this example, the query should print P01 and P02.
10. Print the full name(s) of user(s) in this class, if any, who replied to every single *original* post made for this class that been followed up on. In this example, your query should return the full name of “erin” (who followed up on both P01 and P02).