

Final exam practice questions

CSCC43 – Introduction to Databases

These are actual questions from last year's midterm and final exam. They should give a good indication of the sorts of questions I will ask this year.

Solutions are inside boxes to distinguish them from the questions

NOTE: There are not very many questions about functional dependencies or normal forms because this year I completely changed how I teach those, and the old questions would be useless to you (they asked busywork questions about situations that only arise if you fail to start from a proper minimal basis). The practice FD problems posted online this semester are both numerous and very close to what you can expect to see on the final.

1. Intersection vs. join vs. Cartesian product

Consider relations R and S shown to the right. Compute the cardinality and arity of each of the following:

- a) $R \times S$
- b) $R \cap S$
- c) $R \bowtie_{\rho_{C=X}}(S)$
- d) $R \bowtie_{\rho_{B=X, C=Y}}(S)$

R			S		
A	B	C	A	B	C
1	0	1	2	0	0
1	0	0	2	0	1
1	1	1	0	2	1
1	2	0	0	2	0
0	2	2	2	2	0
0	0	1	0	2	0
2	0	0	2	2	0
0	2	0	2	1	2

2. Joins and Aggregation

Consider the following relations:

F-amousPerson(personID, name, ...)

R-eporter(reporterID, name, ...)

C-aughtOnCamera(personId, reporterId, timestamp, ...)

Give SQL queries that give the following information. Timestamps can be converted to Julian days with the function 'julianday()' and the function 'now()' returns the current timestamp. For brevity you may use F/R/C as table names and pid/rid for person and reporter ids.

- a) Name each famous person who was photographed at least five times in the last 90 days. Ignore people photographed by only one reporter. Your query should also compute the number of qualifying photos taken and how many reporters were involved.

- b) Name each reporter and the last person (if any) they photographed on 1 Aug 2011.

3. Relational algebra expression trees

- a) Convert the following SQL query into a relational algebra expression tree that places σ and π operators so as to minimize the amount of data the system must process.

```
select C.name, sum(L.extPrice)/1e6 sales
from LineItem L, Orders O, Customer C, Nation N
where L.oid=O.oid and O.cid=C.cid and C.nid=N.nid
      and N.name = 'Canada' and O.orderdate > '2010-12-31'
      and julianday(L.shipdate) - julianday(O.orderdate) > 90
group by C.name having average(L.qty) < 10
order by sales desc
```

4. Entity-relationship data model

Create an E/R diagram which captures the following information:

- One Superhero protects each city
- Every supervillain operates out of a secret lair at some location (not always a city)
- Some cities host multiple supervillains, some host none.
- Every superhero has a supervillain nemesis.
- A superhero and supervillain always fight in some city.
- Each superhero can team up with other superheroes; supervillains form conspiracies.
- Teams of superheroes can battle conspiracies at any location.
- Battles in cities, and nemeses that live in the same city, are especially interesting.

Hint: this is pretty straightforward; just take it one step at a time.

5. Functional dependencies and normal forms

Consider the relation and functional dependencies (FDs) given below:

`R(address, author, genre, publisher, title)`

`author title -> publisher`
`genre title -> author`
`publisher -> address genre`

- a) List all candidate keys for R.

- b) Identify any FD(s) that prevent R from satisfying 3NF.

6. Outer joins

Consider the following three relations:

R		S			T	
a	b	b	c	d	d	e
0	5	3	5	4	3	2
1	2	1	2	3	4	3
3	4	0	5	0	5	0
0	2	2	1	1	2	1
1	3					

Populate the leftmost column of the table below by computing $(R \bowtie S) \bowtie T$ – the full outer join of all three relations – using a “-” to represent NULL values where appropriate. Then, for each remaining column, put a check mark by any rows that also appear in the result of the specified operation:

$(R \bowtie S) \bowtie T$	$R \bowtie (S \bowtie T)$	$(T \bowtie_L S) \bowtie R$	$(T \bowtie_L S) \bowtie_L R$	$(R \bowtie_L S) \bowtie_L T$	$(R \bowtie S) \bowtie_L T$

7. Sudoku revisited

In the Sudoku puzzle to the right, the gray-filled cells form a “locked” 3-6 pair. The circled cells also form a locked pair (1-6), and empty boxes mark three more. The locked numbers cannot appear in other cells of a locked pair’s group (row/col/block). Example: the gray 3-6 pair “shadows” the nearby circled cell (id:70), so cell 70 cannot legally take the value 6. This is useful because deleting candidate (70,6) solves the cell: only candidate (70,1) remains.

3 1 5 · · 7 6 · 9	7 8 · 3 9 6 5 · ·	4 9 · 1 5 2 8 3 ·
· 5 6 · · 8 · · 3	8 · · · · 5 · · 7	9 7 · 2 · · 5 8 ·
5 · 2 · · 4 · · 1	1 7 8 2 5 · 6 4 ·	· 4 9 7 · 8 · 2 5

We developed a query in class to find shadowed candidates like (70,6) that should be deleted. Now do the reverse: write a query that returns the ids of any locked pair(s) – **and only those pairs** – which require us to delete candidate (70,6). In the example above, only the gray-filled pair should be returned, even though other locked pairs shadow other candidates as well. Assume that CN(id,n) contains all remaining candidates and C(id, x, y, b) maps cell ids to locations on the board. Your solution may create views if that helps.

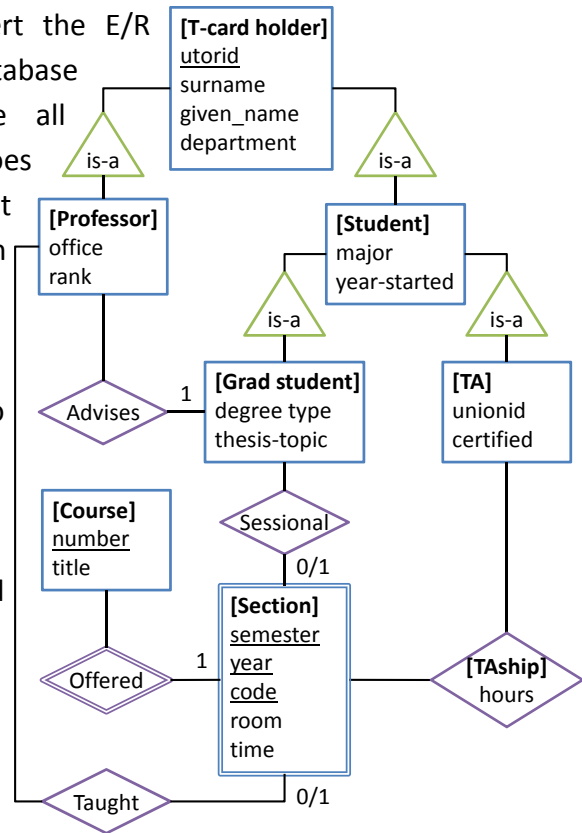
8. E/R diagrams, relational schemata, and constraints

- a) Using the space on the next pages, convert the E/R diagram on the right into an efficient SQL database schema. Your schema should incorporate all constraints specified in the diagram, and does not need to incorporate any others. Correct schemata defining fewer than ten tables earn bonus marks (eight tables is possible).

Answers should be formatted nicely, wrapping and indenting if they're longer than one line. To save space, you may shorten long table and attribute names and abbreviate SQL syntax. The intent must be clear, however. As an example, the tables below would be considered equivalent for marking purposes:

```
CREATE TABLE movie_star(
    star_name PRIMARY KEY,
    studio_id REFERENCES studio(id)
); -- too long for one line
```

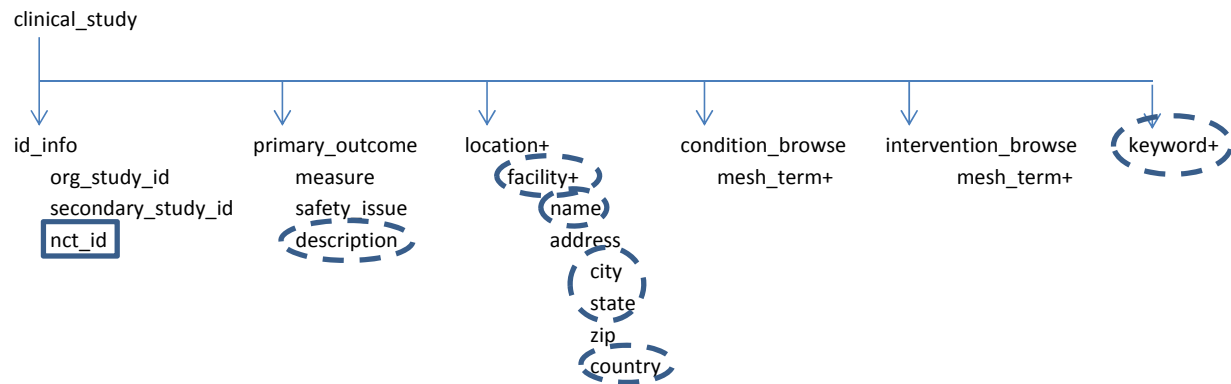
```
TABLE star(name PK, sid REF studio(id)); -- short enough for one line
```



- b) Assume that CHECK constraints do **not** allow nested queries. Extend your schema from part (a) to enforce the following additional constraints:
- Course titles are unique, as are TA unionids.
 - Semesters are “F” or “W” and TAship hours are between 10 and 70
 - Only graduate TAs can hold more than three TAships (lifetime total).
 - A professor and a sessional cannot both teach the same section.
- c) Extend your schema from part (a) with views that make the following information easily accessible to a web application that uses the database:
- A listing of graduate TAships (TA name, course code, semester, year, etc.).
 - A full description (location, instructor, course title) of all sections offered.
 - Instructors (whether professor and sessional) and semesters when they have taught.
- d) You may assume that all primary and foreign key fields already have indexes. Create additional indexes as needed to accelerate searches for the following:
- people by surname
 - sections by year and semester

9. Database operations in XQuery

Suppose you are given a variable, `$input`, which contains a sequence of `clinical_study` elements structured according to the diagram below:



- Write an XQuery expression that returns any study that mentions “cancer” as part of a keyword, condition, or outcome (assume string tests are case insensitive).
- Write an XQuery expression that converts `$input` into a sequence of `study` elements containing only the elements circled on the diagram. The `nct_id` should become an attribute of `study`. If the parent of an element is not selected, make that element a child of the lowest ancestor that is selected.

- c) Give an Xquery expression that identifies the `city` involved in the most clinical studies? Give the city's name and study count, breaking ties by returning all such cities. NOTE: count studies, not locations.

- d) List pairs of cities which are joint locations of any study. Keep in mind that many studies have 3+ locations. No pair of cities should appear more than once in the output.