

University of Illinois at Urbana-Champaign
Midterm Examination
CS411 Database Management Systems

Time Limit: 75 minutes
Exam Date: Feb. 26, 2008

- Closed notes; closed book; no sheet of formulas permitted.
- Please write your answers directly on the exam sheet. The space we left for your answers is often more than what you actually need. Please use the back side of the exam as scratch paper.
- In case you find a question ambiguous, please write down your assumption and answer the question accordingly.
- *Answers to queries should not contain duplicates, for any query on this exam. Do not forget to remove duplicates!*
- You may use temporary relations, if you like, for any of the queries below. If you use the same temporary relation for a second exam question, you must redefine the relation in your answer to the second question. In other words, your answer to each question should be self-contained.

Your Name:

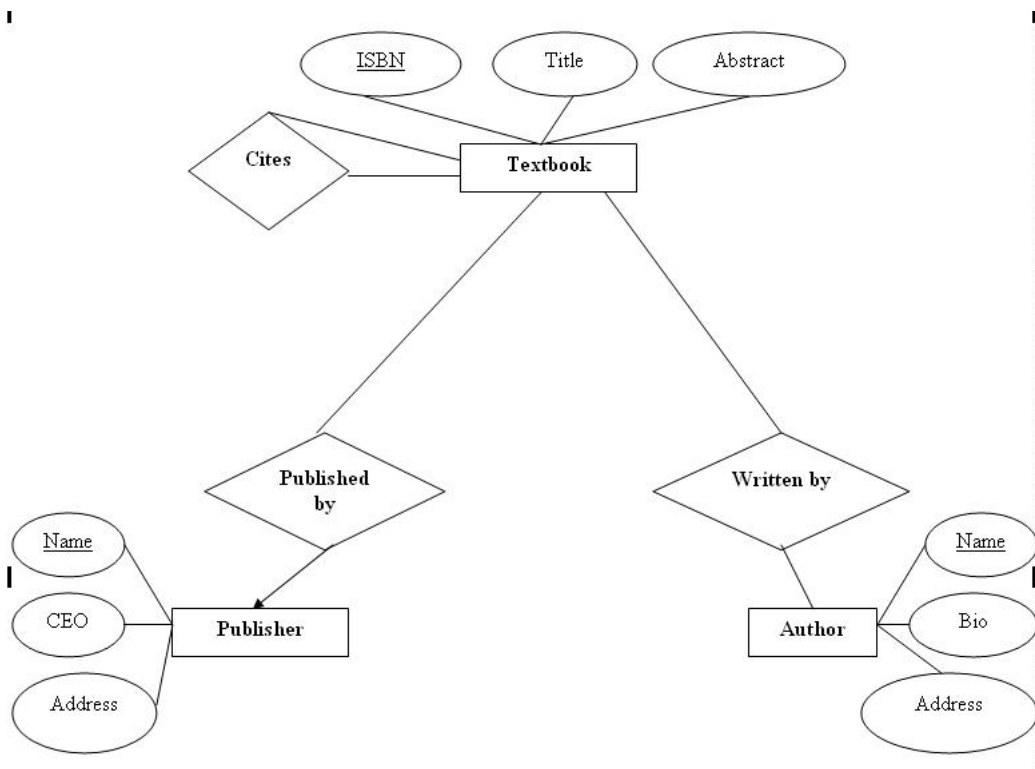
Your NetID:

Good Luck!

1. [20 points] Textbook bibliographical databases consist of information about publishers, books, authors, and citations, telling

1. which textbooks are published by which publisher;
2. which textbooks are written by which authors;
3. which textbooks cite which other text books.

Devise an entity-relationship diagram for this database. List the keys. Tell whether each relationship is many-many, many-one, one-many, or one-one. To keep your answer small, you only need to include one attribute per entity set in addition to its key attribute(s). You can assume textbooks come in one volume only, so you do not have to model volume information.



2. [35 points]. A small group has instituted a betting pool for University of Illinois football games. Bettors are given even odds (i.e., if you bet on Illinois to win, and Illinois wins, then you win the amount that you bet. If Illinois loses, you have to pay the amount you bet). There are two relations in the database for this activity: *Out(game, outcome)*, which lists, for each game played so far, whether Illinois won the game or not, and *Bets(who, amt, outcome, game)*, which keeps track of who bet how much on which game for which team to win. For the *Out* relation, the *game* attribute is the only key. For the *Bets* relation, the attributes *who, game, and outcome* together form the only key. This means that a bettor can make only one bet on each game. Example instances of these two relations are given here:

OUT Relation

Game	Outcome
Michigan	W
Yale	L
Georgia	W
USC	L
Notre Dame	W
Kansas	L

Note that in the *Out* relation, 'W' means that Illinois won, and 'L' means that Illinois lost.

BETS Relation

Who	Amt	Outcome	Game
Liu	100	W	Michigan
Liu	20	W	Yale
Liu	30	W	Georgia
Liu	40	W	USC
Winslett	20	L	Michigan
Winslett	20	L	Georgia
Winslett	150	W	Notre Dame
Jones	250	W	USC
Harandi	110	L	USC
Reed	5	W	USC
Reed	5	L	USC
Reed	5	W	Tennessee
Reed	5	L	Tennessee
Campbell	100	L	Georgia
Campbell	100	W	Michigan
Campbell	100	W	Yale
Ng	100	W	Georgia
Ng	5	L	LSU
Kale	5	L	Auburn

Note that in the *Bets* relation, 'W' means that the bettor is betting that Illinois will win, and 'L' means that the bettor is betting that Illinois will lose. Also note that not all the games have been played yet: the LSU, Arizona, Tennessee, and Auburn games have not yet taken place.

Answer the following queries using relational algebra. Your queries should work for any instance of the database, not just this one.

a. [5 points] List all the games where someone bets that Illinois will win. For the instance given above, your answer should be:

Georgia
Michigan
Notre Dame
Tennessee
USC
Yale

Solution:

$\pi_{Game} \sigma_{outcome='W'} BETS$

b. (5 points) List the completed games that no one bet on. For the instance given above, your answer should be:

Kansas

Solution:

$\pi_{Game} OUT - \pi_{Game} BETS$

c. [10 points] List the people who have made two or more bets on Illinois to lose. For the instance given above, your answer should be:

Winslett
Reed

Solution:

$\pi_{BETS_1.Who}(BETS_1 \bowtie_{(BETS_1.Who=BETS_2.Who)and(BETS_1.Game!=BETS_2.Game)and(BETS_1.Outcome='L')and(BETS_2.Outcome='L')} BETS_2)$

$BETS_1$ and $BETS_2$ are both of type $BETS$.

d. [15 points] List the games that all the bettors agree on (i.e., games where all the bets are for Illinois to lose or all the bets are for Illinois to win). For the instance given above, your answer should be:

Yale
Notre Dame
Auburn
LSU

Solution:

$(\pi_{Game} BETS_1) - (\pi_{BETS_2.Game}(BETS_2 \bowtie_{(BETS_2.Outcome!=BETS_3.Outcome)and(BETS_2.Game=BETS_3.Game)} BETS_3))$

$BETS_1$, $BETS_2$ and $BETS_3$ are of type $BETS$.

3. [45 points]. Using the schema of problem 2, write the following queries in SQL. Your queries must return the correct answer for any instance of the database, not just the one given above.

a [10 points]. Who has bet on every game? For the instance given above, your answer must be empty relation.

Solution:

```

Create View AllGames As
(Select Game
from OUT)
UNION
(Select Game
from BETS)

Select Who
from BETS
Group By Who
Having count (distinct Game) =
(Select count (distinct Game)
form AllGames)

```

b [15 points]. For each game, list the number of people betting on Illinois to win and the number betting on Illinois to lose. For the instance given above, your answer should be:

Michigan	W	2
Michigan	L	1
Yale	W	2
Georgia	W	2
Georgia	L	2
USC	W	2
USC	L	3
Notre Dame	L	1
LSU	L	1
Tennessee	W	1
Tennessee	L	1
Auburn	L	1

Solution:

```

Select Game, Outcome, Count (Who)
from Bets
Group By Game, Outcome

```

c [20 points]. What game or games so far won the most money for people betting on Illinois to win? For the instance given above, your answer should be:

: Solution:

```
Create View Success-Win As
Select Bets.Game, Sum(Bets.Amt) As SumAmt
From Bets, Out
Where Out.Game = Bets.Game and Bets.Outcome='W' and Out.Outcome='W'
Group By Game
```

```
Select distinct Game
From Success-Win
Where SumAmt >= ALL
(Select SumAmt
From Success-Win)
```