

CSE 344 Final Examination

Monday, December 11, 2017, 2:30-4:20

Name: _____

Question	Points	Score
1	30	
2	20	
3	30	
4	40	
5	40	
6	40	
Total:	200	

- This exam is CLOSED book and CLOSED devices.
- You are allowed TWO letter-size sheets with notes (both sides).
- You have 110 minutes;
- Answer the easy questions before you spend too much time on the more difficult ones.
- Good luck!

1 Relational Data Model

1. (30 points)

A university maintains a database of students, the courses they took, and their summer internships, with the following schema:

```
Student(sid, name, zip)
Takes(sid, cid, year)
Course(cid, title, dept)
Intern(sid, companyName)
```

The keys are underlined; `Takes.sid` and `Intern.sid` are foreign keys to `Student`, while `Takes.cid` is a foreign key to `Course`.

- (a) (10 points) If two students took the same course in the same year then we call them *pals*. Write a SQL query that computes, for each student, the number of his/her pals. Your query should return all students indicating their student ID and name, and the number of their pals, ordered lexicographically by their names. (Hint: students who never took any course have zero pals; you need to return these too. Students who took at least one course have at least one pal, naturally (why?); you should not exclude himself/herself when you count their number of pals.)

```
Student(sid, name, zip)
Takes(sid, cid, year)
Course(cid, title, dept)
Intern(sid, companyName)
```

(b) Consider the following SQL query:

```
select *
from Student s
where not exists
  (select *
   from Takes t, Course c
   where s.sid = t.sid
        and t.cid = c.cid
        and c.dept = 'CSE');
```

i. (5 points) Write this query in relational algebra.

ii. (5 points) Write this query in datalog.

```
Student(sid, name, zip)
Takes(sid, cid, year)
Course(cid, title, dept)
Intern(sid, companyName)
```

- (c) (10 points) We say that a student X *has a connection* to Y if either X, Y are pals, or they have a common connection. Alice wants to find a connection who was an intern at Microsoft. Write a datalog program to find all Alice's connections who were Microsoft interns. In other words, if Alice is a pal with someone, who is a pal with someone, who is a pal . . . , who is a pal with someone who interned at Microsoft, then you return the Microsoft intern; you should return the sid and name. Recall that *pal* means that the two students took a common course in a common year.

2 NoSQL, JSon, SQL++

2. (20 points)

(a) A NoSQL database stores one billion (10^9) key-value pairs. The system runs on 100 servers, and replicates every key-value pair on three randomly chosen servers, to increase reliability. Replicas may be out of sync, but they will eventually be reconciled by the system, i.e. if a value is updated on one server, eventually the other two replicas of that value will also be updated. Answer the following questions. In all cases, you may round your answer to the closest integer.

i. (2 points) How many key-value pairs are stored on one server?

i. _____

Number of key-value pairs:

ii. (2 points) A transaction needs to access 5 different key-value pairs, and does not necessarily need to get the most up to date values. From how many servers does it need to read (in expectation)?

ii. _____

Number of servers:

iii. (2 points) A transaction needs to access 5 different key-value pairs, and must get the most up to date values. From how many servers does it need to read (in expectation)?

iii. _____

Number of servers:

iv. (4 points) One day three servers fail catastrophically, simultaneously. In expectation, how many key value pairs are lost after this failure? Your answer may be approximate. Note that the NoSQL database can recover a key-value pair if at least one replica survives; you need to compute the (approximate) number of key-value pairs that cannot be recovered after three servers fail.

iv. _____

Expected number:

(b) For each statement below indicate if it is true or false.

- i. (2 points) The main reason why NoSQL database were introduced was because relational databases did not scale up to very large number of servers.

i. _____

True or false?

- ii. (2 points) In JSon one can represent nested collections.

ii. _____

True or false?

- iii. (2 points) The JSon data model is in First Normal Form.

iii. _____

True or false?

- iv. (2 points) It is easy to represent a many-to-one relationship in JSon, but it is difficult to represent a many-to-many relationship.

iv. _____

True or false?

- v. (2 points) SQL++ can express recursive queries, such as transitive closure.

v. _____

True or false?

3 Query Evaluation

3. (30 points)

Consider the relations below:

```
Student(sid, name, zip)
Takes(sid, cid, year)
Course(cid, title, dept)
```

(a) We run the following SQL command:

```
CREATE INDEX newI on Student(zip,name);
```

For each of the following SQL queries indicate whether the index `newI` will speed it up, or will not affect its performance, or will slow it down. Assume that the optimizer is super-smart, and will always chose the absolute best possible plan.

i. (2 points)

```
select * from Student where name = 'Alice';
```

i. _____

Faster? The same? Or slower?

ii. (2 points)

```
select * from Student where zip = 98195;
```

ii. _____

Faster? The same? Or slower?

iii. (2 points)

```
select Student.name, Student.zip
from Student, Takes, Course
where Student.sid=Takes.sid
and Takes.cid = Course.cid
and Course.title = 'Databases';
```

iii. _____

Faster? The same? Or slower?

iv. (2 points)

```
insert into Student values(1234,'Alice',98195);
```

iv. _____

Faster? The same? Or slower?

v. (2 points)

```
delete from Student where name = 'Alice';
```

v. _____

Faster? The same? Or slower?

Student(sid, name, zip)
Takes(sid, cid, year)
Course(cid, title, dept)

(b) (5 points) Consider the following statistics on the relations and attributes:

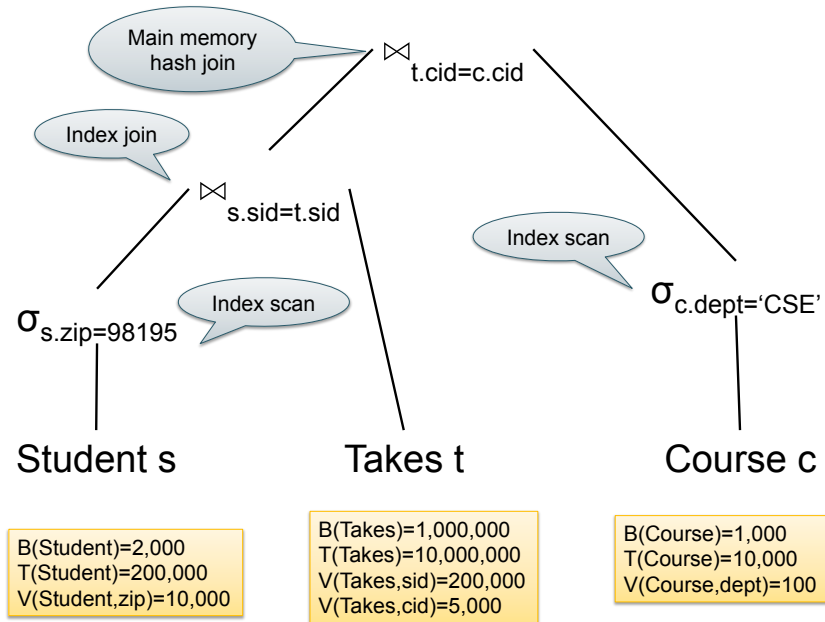
$B(\text{Student}) = 2,000$	$B(\text{Takes}) = 1,000,000$	$B(\text{Course}) = 1,000$
$T(\text{Student}) = 200,000$	$T(\text{Takes}) = 10,000,000$	$T(\text{Course}) = 10,000$
$V(\text{Student}, \text{zip}) = 10,000$	$V(\text{Takes}, \text{sid}) = 200,000$	$V(\text{Course}, \text{dept}) = 100$
	$V(\text{Takes}, \text{cid}) = 5,000$	

Estimate the size of the answer of the following SQL query:

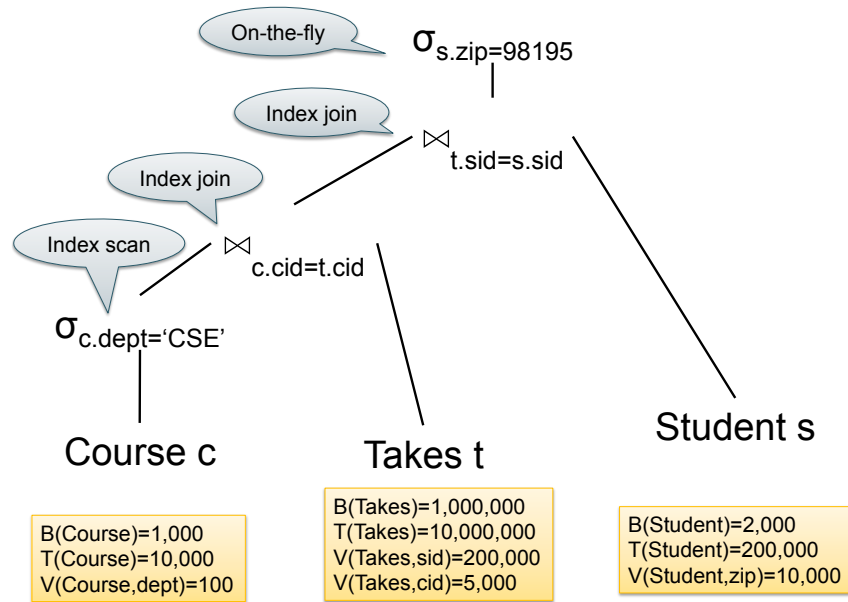
```
select *  
from Student s, Takes t, Course c  
where s.sid = t.sid and t.cid = c.cid  
and s.zip = 98195 and c.dept = 'CSE';
```


- (c) For each of the physical plans below, estimate its cost. Assume the size of the main memory is $M = 2000$ pages and that all indexes are unclustered, and are stored in main memory (hence accessing them requires zero disk I/O's).

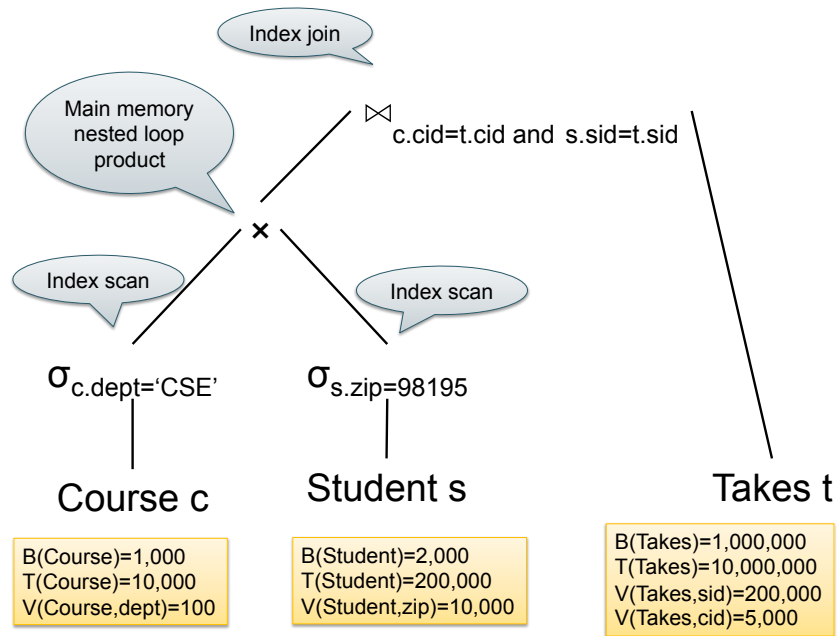
i. (5 points)



ii. (5 points)



iii. (5 points)



Example

- Counting the number of occurrences of each word in a large collection of documents
- Each Document
 - The **key** = document id (**did**)
 - The **value** = set of words (**word**)

```
map(String key, String value):  
  // key: document name  
  // value: document contents  
  for each word w in value:  
    EmitIntermediate(w, "1");
```

```
reduce(String key, Iterator values):  
  // key: a word  
  // values: a list of counts  
  int result = 0;  
  for each v in values:  
    result += ParseInt(v);  
  Emit(AsString(result));
```

4 Parallel Databases

4. (40 points)

(a) (10 points) Write a Map/Reduce program to compute the following SQL query:

```
select p.category, count(*)  
from Product p  
where p.price > 100  
group by p.category  
having sum(p.quantity) > 200
```

You need to write two functions: **Map** and **Reduce**. Use pseudo-code to write both functions, as we did in class (to refresh your memory, the relevant slide from the lectures is shown at the top of the page). You may assume that the **Map** function has a single input, namely the **Product** record.

- (b) Consider the following SQL query, computing for each person p the number of its followers under 20 years old

```
select p.pid, p.name, count(*)
from Person p, Follower f
where p.pid = f.follows and f.age <= 20
group by p.pid, p.name
```

We have $T(\text{Person}) = 10^7$, $T(\text{Follower}) = 10^9$, $\min(\text{age}) = 1$, $\max(\text{age}) = 99$, and $V(\text{Follower}, \text{follows}) = 10^7$. Assume the logical query plan is:

$$\gamma_{\text{pid,name,count}(\ast)}(\text{Person} \bowtie \sigma_{\text{age}<20}(\text{Follower}))$$

The system computes the query on 200 servers as follows.

Step 1 Initially both relations are distributed randomly and uniformly on all servers.

The selection is computed on the fly

Communication Next, there is a communication step where **Person** is hash-partitioned on **pid**, and **Follower** is hash partitioned on **follows**.

Step 2 Each server computes the join and group by locally.

Assume the values of **age** are uniformly distributed among **Follower** and that the hash function is uniform. For anything else, assume the worst case distribution.

- i. (2 points) How many **Person** records does each server hold during each step of the computation? Report the maximum number of records during steps one and two, i.e. do not add up the number of records held during steps one and two.

i. _____

Number of records:

- ii. (2 points) How many **Follower** records does each server hold during each step of the computation? Like in the previous question, report the maximum number, not their sum.

ii. _____

Number of records:

(c) Answer the following questions:

- i. (2 points) If you can compute a query in parallel on 1000 servers in 15 minutes, then you can always compute it in parallel on 500 servers in at most 30 minutes.

i. _____

True or false?

- ii. (2 points) If you can compute a query in parallel on 500 servers in 30 minutes, then you can always compute it in parallel on 1000 servers in at most 15 minutes.

ii. _____

True or false?

- iii. (2 points) If one server of a parallel database system fails, then the entire job that the system was computing is restarted; if one server of a Map/Reduce job fails then the M/R system can recover and continue the job. Why the difference? Choose one answer below:

1. Parallel databases can answer complex SQL queries, while an M/R job is restricted to two functions (map and reduce).
2. Parallel databases were developed in the 80's when fault tolerance techniques were not yet developed, while M/R was developed in the early 2000's, when the technology exists
3. Parallel database systems run on a small number of servers (10-20) while M/R jobs can run on a very large number of servers (10000) making failure much more likely.
4. Parallel database systems must maintain statistics on the database to be used by the query optimizer, while M/R systems do not maintain statistics on the data.

iii. _____

Choose one answer

(d) Answer the following questions about Map/Reduce:

- i. (2 points) The number of map tasks is usually equal to the number of blocks of the input file.

i. _____

True or false?

- ii. (2 points) The number of reduce tasks is usually equal to the number of blocks of the output file.

ii. _____

True or false?

- iii. (2 points) The first reduce task can start as early as the first map task finishes.

iii. _____

True or false?

- iv. (2 points) The number of reduce tasks must be less than or equal to the number of servers.

iv. _____

True or false?

- v. (2 points) The combine function is executed by the map task.

v. _____

True or false?

- vi. (2 points) If a server running a reduce task fails during the computation, then the master will start that reduce task on another server (or the same server, if it recovers).

vi. _____

True or false?

- vii. (2 points) In a Map/Reduce job, *shuffle* refers to the process of distributing the input data to the map tasks.

vii. _____

True or false?

(e) Answer the following questions about Spark:

- i. (2 points) A *Resilient Distributed Dataset* (RDD) is another term for a distributed file on disks.

i. _____

True or false?

- ii. (2 points) Consider the function $\text{map}(f) : \text{RDD} \langle T \rangle \rightarrow \text{RDD} \langle U \rangle$, where $f : T \rightarrow U$. The output RDD can never be strictly smaller than the input RDD. (In other words, the number of U elements in the output can never be strictly smaller than the number of T elements in the input.)

ii. _____

True or false?

- iii. (2 points) Consider the function $\text{flatMap}(f) : \text{RDD} \langle T \rangle \rightarrow \text{RDD} \langle U \rangle$, where $f : T \rightarrow \text{Seq}(U)$. The output RDD can never be strictly smaller than the input RDD. (Recall: $\text{Seq}(T)$ means *sequence of T* , or *array of T* .)

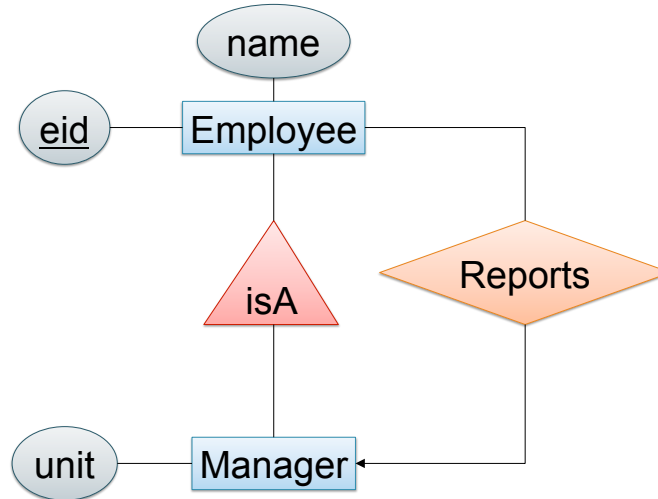
iii. _____

True or false?

5 Conceptual Design

5. (40 points)

(a) (5 points) Consider the following E/R diagram:



Write the CREATE TABLE statements to implement the E/R diagram in SQL. Your statements should contain all key and foreign key declarations. Assume that `eid` and `unit` are integers, and `name` is a text.

- (b) (5 points) The relation $R(A, B, C, D, E, F)$ satisfies the following functional dependencies:

$$AB \rightarrow C$$

$$CD \rightarrow E$$

Find a BCNF decomposition for R . In each of the new relation indicate the key (by underlining the attribute, or set of attributes that form a key).

(c) For each statement below indicate if whether it is true or false.

- i. (3 points) The reason why we decompose a relation in Boyce Codd Normal Form is to save disk space when the relation becomes too large.

i. _____

True or false?

- ii. (3 points) Suppose that the BCNF decomposition of R consists of three relations S, T, K : then S, T, K are a lossless decomposition of R .

ii. _____

True or false?

- iii. (3 points) If $X \subseteq Y$ then $X^+ \subseteq Y^+$.

iii. _____

True or false?

- iv. (3 points) $(X^+)^+ = X$.

iv. _____

True or false?

- v. (3 points) If $(X \cup Y)^+ = X^+ \cup Y^+$

v. _____

True or false?

- (d) Answer the questions below:
- i. (3 points) Consider the following three relations:

$$R(\underline{A}, B), S(B, C), T(C, A)$$

where $R.A$ is a key. Suppose they have the same cardinality N :

$$|R| = |S| = |T| = N$$

What is the largest possible size of the natural join $R \bowtie S \bowtie T$? Your answer should be a mathematical formula that depends on N , e.g. $N^2 \log(N)$ or $N + 5$ (not a real answer).

i. _____

Maximum size of $R \bowtie S \bowtie T$ is:

- ii. (3 points) Consider two relations $R(A, B, C)$ and $S(D, E, F)$ satisfying the following functional dependencies:

$$R : A \rightarrow B$$

$$S : DE \rightarrow F$$

Let $T(A, B, C, D, E, F)$ be the following table:

```
create table T as
(select R.A, R.B, R.C, S.D, S.E, S.F
 from R, S
 where R.C = S.D and S.E = 5)
```

Find the key or keys of T . If T has more than one key, list all keys.

- iii. (3 points) Consider a relation with the schema $R(A, B, C)$. Suppose the relational instance satisfies the functional dependency

$$A \rightarrow B$$

and does not satisfy any other functional dependencies. We know the following statistics on R :

$$|R| = 512$$

$$V(R, B) = 510$$

Compute $V(R, A)$; your answer should consist of a number.

$V(R, A) =$ iii. _____

- iv. (3 points) True or false? If R is a relation that satisfies a set of functional dependencies, then we can declare all of them in the CREATE TABLE R statement by using PRIMARY KEY and UNIQUE declarations.

True or false? iv. _____

- v. (3 points) True or false? If R is a relation that satisfies a set of functional dependencies and is in BCNF, then we can declare all of them in the CREATE TABLE R statement by using PRIMARY KEY and UNIQUE declarations.

True or false? v. _____

6 Transactions

6. (40 points)

(a) For each schedule below indicate whether it is conflict serializable and, if it is, indicate the equivalent serial schedule.

i. (5 points)

$$R_1(A), R_4(D), W_3(C), R_2(B), W_2(A), R_1(C), W_3(B)$$

ii. (5 points)

$$R_1(A), W_2(B), R_3(C), W_3(A), R_3(D), R_4(B), W_4(D), W_2(C)$$

(b) An application uses a database of customers and their accounts:

```
Customer(cid, name, acc)
Account(acc, balance)
```

The application runs concurrently many instances of the following transaction (shown here in pseudocode):

```
function updateAccount(c, delta)
  BEGIN TRANSACTION
  myCustomer = "SELECT * FROM Customer WHERE cid = [c]";
  myAccount  = "SELECT * FROM Account WHERE acc = [myCustomer.acc]";
  if myAccount.balance + delta < 0 then ROLLBACK
  "UPDATE Account
   SET balance a = [a.balance + delta]
   WHERE acc = [myCustomer.acc]"
  COMMIT
```

Thus, the transaction first looks up the customer with $cid = c$, call him/her `myCustomer`, then looks up the account `myCustomer.acc`, and adds `delta` to the account balance, checking first that it doesn't become negative.

The concurrency control manager of the database uses a locking based concurrency control mechanism. It uses the strict 2PL mechanism for write locks, the mechanism decided by the isolation level for read locks, and table-level locking for phantoms. The only transactions on the system are instances of the `updateAccount` shown above.

Answer the questions on the next page.

In each case below you are asked to indicate (a) whether the scheduled is guaranteed to be serializable, and (b) whether a deadlock may occur.

- i. (3 points) Suppose we run all transactions with:
SET ISOLATION LEVEL READ UNCOMMITTED

It is always serializable? i. _____

Can deadlock occur? i. _____

- ii. (3 points) Suppose we run all transactions with:
SET ISOLATION LEVEL READ COMMITTED

It is always serializable? ii. _____

Can deadlock occur? ii. _____

- iii. (3 points) Suppose we run all transactions with:
SET ISOLATION LEVEL REPEATABLE READ

It is always serializable? iii. _____

Can deadlock occur? iii. _____

- iv. (3 points) Suppose we run all transactions with:
SET ISOLATION LEVEL SERIALIZABLE

It is always serializable? iv. _____

Can deadlock occur? iv. _____

- v. (3 points) Assume that the concurrency control mechanism uses strict 2PL with exclusive locks only; in other words, it no longer uses separate read and write locks.

It is always serializable? v. _____

Can deadlock occur? v. _____

(c) For each of the following statements indicate whether it is true or false:

- i. (3 points) In a static database, every serializable schedule is conflict serializable.

True or false?

i. _____

- ii. (3 points) In a dynamic database, every serializable schedule is conflict serializable.

True or false?

ii. _____

- iii. (3 points) In a static database, every conflict serializable schedule is serializable.

True or false?

iii. _____

- iv. (3 points) In a dynamic database, every conflict serializable schedule is serializable.

True or false?

iv. _____

- v. (3 points) In Sqlite all schedules are guaranteed serializable.

True or false?

v. _____