

# 1 SQL

1. (36 points)

A company maintains a database about their employees and projects with the following schema.

```
Employee(eid, name, salary)
```

```
Project(pid, title, budget)
```

```
WorksOn(eid, pid, year)
```

`WorksOn` records which employee worked on which projects. `salary` and `budget` represent yearly salary and budget respectively. An employee may work on multiple project during the same year, and may also work on the same project during multiple years. All keys are underlined, and `WorksOn.eid`, `WorksOn.pid` are foreign keys to `Employee` and `Project` respectively.

- (a) (10 points) The yearly salary expenses of a project is the sum of all salaries of the employees who worked on that project during that year. Write a SQL query to find all the projects whose yearly salary expenses exceeded its budget. Return only the project title, and sort the projects in ascending order alphabetically by their title.

```
Employee(eid, name, salary)
Project(pid, title, budget)
WorksOn(eid, pid, year)
```

- (b) (10 points) We say that an employee worked intermittently on a project  $p$  if she worked on  $p$  during one year, then did not work on  $p$  during a later year, then worked again on  $p$  during an even later year. For example if Alice worked on the project during 2012, 2013, and 2017 then we say that she worked intermittently on  $p$ ; if Bob worked on that project during 2013, 2014, 2015 and no other years, then we say he worked continuously. Write a SQL query to retrieve all employees that worked intermittently on some project. Return the employee name, and the project title.

```
Employee(eid, name, salary)
Project(pid, title, budget)
WorksOn(eid, pid, year)
```

(c) For each question below indicate whether the two SQL queries are equivalent. Assume that the database does not contain any NULL values.

i. (2 points) Are Q1 and Q2 equivalent?

```
Q1: select W.year, W.pid, count(*)
     from WorksOn W
     group by W.year, W.pid;
```

```
Q2: select distinct W.year, W.pid,
     (select count(*)
      from WorksOn W2
      where W.year=W2.year and W.pid = W2.pid)
     from WorksOn W;
```

i. \_\_\_\_\_

Yes/No:

ii. (2 points) Are Q3 and Q4 equivalent?

```
Q3: select W.year, W.pid, count(*)
     from WorksOn W, Employee E
     where W.eid = E.eid and E.salary > 100000
     group by W.year, W.pid;
```

```
Q4: select W.year, W.pid,
     (select count(*)
      from Employee E
      where W.eid = E.eid and E.salary > 100000)
     from WorksOn W;
```

ii. \_\_\_\_\_

Yes/No:

iii. (2 points) Are Q5 and Q6 equivalent?

```
Q5: select distinct E.name
      from Employee E, WorksOn W, WorksOn W2
      where E.eid = W.eid and E.eid = W2.eid
            and W.pid = W2.pid
            and W.year > 2010 and W2.year < 2015;
```

```
Q6: select distinct E.name
      from Employee E, WorksOn W
      where E.eid = W.eid
            and W.year > 2010;
```

iii. \_\_\_\_\_

Yes/No:

iv. (2 points) Are Q7 and Q8 equivalent?

```
Q7: select distinct E.name
      from Employee E, WorksOn W, WorksOn W2
      where E.eid = W.eid and E.eid = W2.eid
            and W.pid = W2.pid
            and W.year < 2010 and W2.year < 2015;
```

```
Q8: select distinct E.name
      from Employee E, WorksOn W
      where E.eid = W.eid
            and W.year < 2010;
```

iv. \_\_\_\_\_

Yes/No:

# 1 SQL and Indexing

1. (50 points)

Consider the following database about picture tagging Website:

Member(mid, name, age)

Picture(pid, year)

Tagged(mid, pid)

Relation **Member** contains a tuple for each registered member of the Website; **Picture** contains the metadata of the pictures (its key **pid**, and the year when the picture was taken), and **Tagged** says which user was tagged in what picture. Primary keys are underlined. Attributes values may be null.

- (a) (10 points) Write a SQL query that retrieves the names of all members that were tagged both in the year 2011 and in 2014. Return them sorted in lexicographic order.

```
Member(mid, name, age)
Picture(pid, year)
Tagged(mid, pid)
```

- (b) (10 points) For each query below, indicate if the query correctly returns the name of all users who were never tagged in 2015. Anywhere between zero and all queries compute the correct answer. Each query is syntactically correct.

```
Q1 = select distinct x.name
      from Member x, Tagged y
      where x.mid = y.mid
            and not exists
                  (select *
                   from Picture z
                   where y.pid = z.pid
                        and z.year = 2015);
```

```
Q2 = select distinct x.name
      from Member x
      where not exists
            (select *
             from Tagged y, Picture z
             where x.mid = y.mid
                  and y.pid = z.pid
                  and z.year = 2015);
```

```
Q3 = select distinct x.name
      from Member x
      where not exists
            (select *
             from Tagged y
             where x.mid = y.mid
                  and not exists
                        (select *
                         from Picture z
                         where y.pid = z.pid
                              and z.year = 2015));
```

```
Q4 = select distinct x.name
      from Member x, Tagged y, Picture z
      where x.mid = y.mid and y.pid = z.pid and z.year = 2015
      group by x.name
      having count(z.pid) = 0;
```

(b) \_\_\_\_\_

Answer with any subset of Q1, Q2, Q3, Q4:

Member(mid, name, age)  
Picture(pid, year)  
Tagged(mid, pid)

- (d) (10 points) Write a SQL query that returns the mid's and names of all members that were tagged only in pictures were Alice was also tagged. Hint: you may want to first answer question 2.b.

## 2 Relational Algebra

2. (29 points)

Consider the same relational schema as before, including the key/foreign key constraints:

Employee(eid, name, salary)


Project(pid, title, budget)

WorksOn(eid, pid, year)

- (a) (5 points) Write a Relational Algebra expression in the form of a logical query plan (i.e., draw a tree) that is equivalent to the SQL query below. Your query plan does not have to be necessarily “optimal”: however, points will be taken off for overly complex solutions.

```
select x.eid, x.name
from Employee x
where not exists (select *
                  from WorksOn y, Project z
                  where x.eid = y.eid and y.pid = z.pid
                  and z.budget > 1000000);
```

Hint: to avoid renaming, use aliases in the query plan, like this





- (b) Assume that the database instance does not contain NULL's, but otherwise can be any valid instance (satisfying all key/foreign key constraints). Answer the questions below, when the relational algebra expressions have bag semantics.

- i. (2 points) Are these two expressions equivalent?

$$\Pi_{\text{year}}(\sigma_{\text{salary}>40000}(\text{Employee} \bowtie_{\text{eid}=\text{eid}} \text{WorksOn})) = \Pi_{\text{year}}(\text{WorksOn})$$

i. \_\_\_\_\_

Yes/No:

- ii. (2 points) Are these two expressions equivalent?

$$\Pi_{\text{year}}(\sigma_{\text{year}>2015}(\text{Employee} \bowtie_{\text{eid}=\text{eid}} \text{WorksOn})) = \Pi_{\text{year}}(\sigma_{\text{year}>2015}(\text{WorksOn}))$$

ii. \_\_\_\_\_

Yes/No:

- iii. (2 points) Are these two expressions equivalent?

$$\Pi_{\text{name}}(\sigma_{\text{salary}>40000}(\text{Employee} \bowtie_{\text{eid}=\text{eid}} \text{WorksOn})) = \Pi_{\text{name}}(\sigma_{\text{salary}>40000}(\text{Employee}))$$

iii. \_\_\_\_\_

Yes/No:

iv. (2 points) Are these two expressions equivalent?

$$\Pi_{\text{name}}(\sigma_{\text{year}>2015}(\text{Employee} \bowtie_{\text{eid}=\text{eid}} \text{WorksOn})) = \Pi_{\text{name}}(\text{Employee})$$

Yes/No:

iv. \_\_\_\_\_

v. (2 points) The notation  $|S|$  means the cardinality of a set  $S$  (number of tuples in  $S$ ). Does the following always hold?

$$|\text{Employee} \bowtie_{\text{eid}=\text{eid}} \sigma_{\text{year}<2015}(\text{WorksOn})| \leq |\text{Employee}|$$

Yes/No:

v. \_\_\_\_\_

vi. (2 points) Does the following always hold?

$$|\sigma_{\text{salary}<5000}(\text{Employee}) \bowtie_{\text{eid}=\text{eid}} \text{WorksOn}| \leq |\text{WorksOn}|$$

Yes/No:

vi. \_\_\_\_\_

CSE 344 Midterm Exam 2016

Product (pid, name, cid) -- cid is foreign key to Company.cid  
Company (cid, cname, city)  
Purchase(pid, custId, quantity, price)  
-- pid is foreign key to Product.pid, custId is foreign key to Customer.custId  
Customer(custId, name, city)

- a) Write a SQL query that returns the name of companies, along with the number of products sold, for companies that have sold at least 2 different types of products anywhere. (13 points)

- b) Write a relational algebra query that returns the distinct names of all customers from Seattle who purchased any one type of product with quantity > 10. Write your query as a tree or a single relational algebra expression. (13 points)

Schema repeated here for your reference:

Product (pid, name, cid) -- cid is foreign key to Company.cid  
Company (cid, cname, city)  
Purchase(pid, custId, quantity, price)  
-- pid is foreign key to Product.pid, custId is foreign key to Customer.custId

### 3 Relational Data Model and Query Evaluation

3. (30 points)

Answer each question below.

- (a) (2 points) In the relational data model an attribute of a relation cannot be another relation.

(a) \_\_\_\_\_

True or false?

- (b) (2 points) Consider the relation `Product`(productName, manufacturer, price, weight) where the key is (productName, manufacturer). Can there be two different products with the same value of `manufacturer`?

(b) \_\_\_\_\_

Yes or no?

- (c) (2 points) If an attribute in a table is a foreign key, then the table cannot contain two tuples with the same value of that attribute.

(c) \_\_\_\_\_

True or false?

- (d) (2 points) In this and the following question we will assume that the relational algebra has set semantics. If the relation  $R(A, B)$  has  $m$  tuples, and the relation  $S(B, C)$  has  $n$  tuples, what is the largest possible size of the output of the natural join  $R \bowtie S$ ? Choose one of the following:

- (a)  $m$
- (b)  $n$
- (c)  $mn$
- (d)  $m + n$
- (e)  $\max(m, n)$
- (f)  $(m + n)/2$
- (g) None of the above.

(d) \_\_\_\_\_

Answer a-g:

- (e) (2 points) If the relation  $R(A, B)$  has  $m$  tuples, and the relation  $S(B, C)$  has  $n$  tuples, what is the largest possible size of the output of  $\Pi_{AB}(R \bowtie S)$ ? Here  $\bowtie$  represents the natural join. Choose from the same options as in the previous question.

(e) \_\_\_\_\_

Answer a-g:

- (f) Consider two relations  $R(A, B), S(C, D)$ , where all attributes are integers and cannot be NULL. For each identity below, indicate whether it is true or false.

i. (2 points)  $R \bowtie_{B=C} S = S \bowtie_{B=C} R$

i. \_\_\_\_\_

True or false?

ii. (2 points)  $(\sigma_{B < 1000}(R)) \bowtie_{B \neq C} S = (\sigma_{B < 1000}(R)) \bowtie_{B \neq C} (\sigma_{C \geq 1000}(S))$

ii. \_\_\_\_\_

True or false?

iii. (2 points)  $R \times S - R \bowtie_{B=C} S = R \bowtie_{B \neq C} S$

iii. \_\_\_\_\_

True or false?

iv. (2 points)  $R - \Pi_{AB}(R \bowtie_{B=C} S) = \Pi_{AB}(R \bowtie_{B \neq C} S)$

iv. \_\_\_\_\_

True or false?

v. (2 points)  $\gamma_{A, \text{count}(*), \rightarrow F}(R \bowtie_{B=C} S) = \gamma_{A, \text{sum}(E), \rightarrow F}(R \bowtie_{B=C} (\gamma_{C, \text{count}(*), \rightarrow E}(S)))$

v. \_\_\_\_\_

True or false?

- (g) (2 points) If a relation is stored in 1000 blocks on disc, then it cannot have more than 1000 records.

(g) \_\_\_\_\_

True or false?

- (h) (2 points) Alice is a database administrator. She wants to add indexes to a table that is accessed frequently by her applications. She knows that a clustered index delivers better performance than an unclustered index, but most indexes she creates are unclustered. Why? Choose one of the following:

- (a) Because clustered indexes take more space than unclustered indexes and Alice's database servers has little space.
- (b) Because a relation can have only one clustered index, but many unclustered indexes.
- (c) Because clustered indexes will slow down updates and Alice's application has many updates.
- (d) None of the above.

(h) \_\_\_\_\_

Choose one of a-d:

- (i) (2 points) There is an important distinction between logical operators and physical operators. Indicate which of the operators below are physical operators:

- (a) Theta join
- (b) Hash join
- (c) Natural join
- (d) Eq join
- (e) Index-based join
- (f) Merge join

(i) \_\_\_\_\_

Answer a-f (choose all that apply):

- (j) (2 points) Did the speed of sequential reads from disk increase significantly over the years?

(j) \_\_\_\_\_

Yes or no?

- (k) (2 points) Did the speed of random reads from disk increase significantly over the years?

(k) \_\_\_\_\_

Yes or no?

## 4 Semistructured Data

4. (20 points)

(a) Answer the following questions. In your answer you may omit apostrophes, i.e. you may write `{A:a1}` instead of `{'A':'a1'}`.

i. (2 points) What does the following SQL++ query return?

```
with t as
  [{'A':'a1', 'F':[{ 'B':'1'}, { 'B':'2'}], 'G':[{ 'C':'1'}, { 'C':'2'}]},
   {'A':'a2', 'F':[{ 'B':'3'}, { 'B':'4'}, { 'B':'5'}], 'G':[ ]},
   {'A':'a3', 'F':[{ 'B':'2'}], 'G':[{ 'C':'1'}, { 'C':'3'}]}]
Select x.A
From t x, x.G y
where y.C='1';
```

ii. (2 points) What does the following SQL++ query return?

```
with t as
  [{'A':'a1', 'F':[{ 'B':'1'}, { 'B':'2'}], 'G':[{ 'C':'1'}, { 'C':'2'}]},
   {'A':'a2', 'F':[{ 'B':'3'}, { 'B':'4'}, { 'B':'5'}], 'G':[ ]},
   {'A':'a3', 'F':[{ 'B':'2'}], 'G':[{ 'C':'1'}, { 'C':'3'}]}]
Select x.A, y.B
From t x, x.F y;
```

iii. (2 points) What does the following SQL++ query return?

```
with t as
  [{'A':'a1', 'F':[{'B':'1'}, {'B':'2'}], 'G':[{'C':'1'},{'C':'2'}]},
   {'A':'a2', 'F':[{'B':'3'}, {'B':'4'},{'B':'5'}], 'G':[ ]},
   {'A':'a3', 'F':[{'B':'2'}], 'G':[{'C':'1'},{'C':'3'}]}]
Select x.A, y.B
From t x, x.F y, x.G z
where y.B=z.C ;
```

iv. (2 points) What does the following SQL++ query return?

```
with t as
  [{'A':'a1', 'F':[{'B':'1'}, {'B':'2'}], 'G':[{'C':'1'},{'C':'2'}]},
   {'A':'a2', 'F':[{'B':'3'}, {'B':'4'},{'B':'5'}], 'G':[ ]},
   {'A':'a3', 'F':[{'B':'2'}], 'G':[{'C':'1'},{'C':'3'}]}]
Select x1.A as A1, x2.A as A2, y.B
From t x1, t x2, x1.F y, x2.G z
where y.B=z.C ;
```



## Problem 1: Warm up (10 points total)

Select either True or False for each of the following questions. For each question you get 1 point for answering it correctly, -0.5 point for an incorrect answer, and 0 point for no answer. The minimum you will get for this entire problem is 0.

- a) The arity of the relation  $R(A \text{ int}, B \text{ int})$  with tuples  $(4,2), (2,2), (4,2), (3,3), (4,2)$  is 3.
- True                  False
- b) Data is encoded in JSON documents using key-value pairs.
- True                  False
- c) A relation can have a clustered index on a set of attributes.
- True                  False
- d) A datalog rule is safe if and only if every variable appears in a relational atom.
- True                  False
- e) A SQL query with `GROUP BY` can select all attributes that are grouped on, joined upon, or are aggregates.
- True                  False
- f) Every natural join can be rewritten using cross product and selection.
- True                  False
- g) A foreign key does not have to reference a primary key.
- True                  False
- h) An inner join between relations  $R$  and  $S$  always includes all tuples from  $R$  in the result.
- True                  False
- i) Queries with universal quantifiers cannot be unnested.
- True                  False
- j) A subquery in the `SELECT` clause in SQL can refer to tuple variables defined in the `WHERE` clause.
- True                  False

## Problem 2: Writing Queries (52 points total)

Write the following queries using the schema below: