

Fundamentals of Database Systems

[SQL – III]

Malay Bhattacharyya

Assistant Professor

Machine Intelligence Unit
Indian Statistical Institute, Kolkata

August, 2019

1 Dealing with Nullity and Duplicity

2 View

3 Problems

Handling empty tuples

SQL can test whether a subquery has any tuples in its result. The `exists` construct returns the value `true` if the argument subquery is nonempty.

```
select ...
from ...
where exists (
  select ...
  from ...
  where ...);
```

Note: We can test for the nonexistence of tuples in a subquery by using the `not exists` construct.

Handling empty tuples

Suppose the customer details of two banks are provided in the form of two separate tables (say bank1 and bank2). The names of all the customers who have an account in both the banks can be retrieved with the following SQL query.

```
select customer_name
from bank1
where exists (
select *
from bank2
where bank1.customer_name = bank2.customer_name);
```

Note: Existence of tuples in the relation returned by the subquery is verified tupewise from the main query (not as a whole).

Handling duplicate tuples

SQL can test whether a subquery has any duplicate tuples in its result. The `unique` construct returns the value `true` if the argument subquery contains no duplicate tuples.

```
select ...
from ...
where unique (
select ...
from ...
where ...);
```

Note: We can test for the existence of duplicate tuples in a subquery by using the `not unique` construct.

Creating views

We can define a view in SQL by using the `create view` construct. View names may appear in any place that a relation name may appear. To define a view, we must give the view a name and must state the query that computes the view.

The form of the `create view` command is as follows:

```
create view <view_name> as <query_expression>;
```

where `query_expression` is any arbitrary query expression which is legal.

Updating views

SQL allows a view name to appear wherever a relation may appear. Hence, we can use the following constructs.

```
select * from <view_name> where ...;
```

or

```
insert into <view_name> values (...);
```

etc.

Problems

- 1 Consider the following schema of a bank:
 - Branch = $\langle \underline{\text{branch_id}} : \text{integer}, \text{branch_name} : \text{string}, \text{branch_city} : \text{string}, \text{assets} : \text{real} \rangle$
 - Customer = $\langle \underline{\text{customer_id}} : \text{integer}, \text{customer_name} : \text{string}, \text{customer_street} : \text{string}, \text{customer_city} : \text{string} \rangle$
 - Account = $\langle \underline{\text{account_number}} : \text{integer}, \text{customer_id} : \text{integer}, \text{branch_name} : \text{string}, \text{balance} : \text{real} \rangle$
 - Depositor = $\langle \underline{\text{customer_id}} : \text{integer}, \text{customer_name} : \text{string}, \text{account_number} : \text{integer} \rangle$
 - Borrower = $\langle \text{customer_name} : \text{string}, \text{loan_number} : \text{integer} \rangle$

Write the following queries in SQL.

- (i) Find all customers who have both an account and a loan at the bank.
- (ii) Find all customers who have an account at all the branches located in Kolkata.

Problems

- 2 Consider the following schema of an online code repository system like GitHub:
- Contributor = $\langle \underline{\text{contributor_id}} : \text{integer}, \text{contributor_name} : \text{string} \rangle$
 - Code-Group = $\langle \text{contributor_id} : \text{integer}, \text{code_group} : \text{string}, \text{count_submissions} : \text{integer} \rangle$

Write the following queries in SQL.

- Find all the contributors who have made no submissions within the Java code group.
- Find all the contributors who have made at most one submission within the R code group.
- Find all the contributors who have made at least three submissions within the Scala code group.