SQL – Subqueries

CS 4750 Database Systems

[A. Silberschatz, H. F. Korth, S. Sudarshan, Database System Concepts, Ch.5.3]

© Praphamontripong

Subqueries: Core Idea



The smaller the problem, the simpler to solve, the easier to debug

Subqueries

- **Subquery** = a query that is part of another query
- A subquery can have subqueries

Usage:

- Return a single constant that can be used to compute an associated value in a SELECT clause
- Return a single constant that can be compared to another value in a WHERE clause
- Return a relation that can be compared or evaluated in a WHERE clause
- Return a relation that can be used as input for another query, in a FORM clause

Equivalent Query Example

Find the average salary for each job

practice_	_emp		
empno	ename	job	sal
7369	Smith	Clerk	1200
7499	Allen	Salesman	2000
7521	Ward	Salesman	1650
7566	Jones	Manager	3375
7654	Martin	Salesman	1650
7698	Blake	Manager	3250
7782	Clark	Manager	2850
7788	Scott	Analyst	3500
7839	King	President	6500
7844	Turner	Salesman	1900
7876	Adams	Clerk	1500
7900	James	Clerk	1350
7902	Ford	Analyst	3500
7934	Miller	Clerk	1700

SELECT job, AVG(sal)
FROM practice_emp
GROUP BY job



jobAvgSalAnalyst3500.0000Clerk1437.5000Manager3158.3333President6500.0000Salesman1800.0000

Idea (self join):

- 1. Self-join practice_emp
- 2. Use one copy to aggregate, group by job
- 3. Use one copy to keep the original job

SELECT E1.job, AVG(E2.sal) AS AvgSal
FROM practice_emp E1, practice_emp E2
WHERE E1.job = E2.job
GROUP BY E1.job

(Equivalent) Subquery (SELECT)

Find the average salary for each job

practice_emp

empno	ename	JOD	sal
7369	Smith	Clerk	1200
7499	Allen	Salesman	2000
7521	Ward	Salesman	1650
7566	Jones	Manager	3375
7654	Martin	Salesman	1650
7698	Blake	Manager	3250
7782	Clark	Manager	2850
7788	Scott	Analyst	3500
7839	King	President	6500
7844	Turner	Salesman	1900
7876	Adams	Clerk	1500
7900	James	Clerk	1350
7902	Ford	Analyst	3500
7934	Miller	Clerk	1700

of the outer query)					
SELECT E1.job,					
	(SELECT AVG(E2	.sal)			
	FROM practice	_emp AS E2	2		
	WHERE E1.job	= E2.job)	AS AvgSal		
FROM	practice_emp E	1			
GROUF	• BY E1.job	job	AvgSal		
GROUF	• BY E1.job	job Analyst	AvgSal 3500.0000		
GROUF	• BY E1.job	job Analyst Clerk	AvgSal 3500.0000 1437.5000		
GROUF	• BY E1.job	job Analyst Clerk Manager	AvgSal 3500.0000 1437.5000 3158.3333		
GROUF	P BY E1.job	job Analyst Clerk Manager President	AvgSal 3500.0000 1437.5000 3158.3333 6500.0000		
GROUF	P BY E1.job	job Analyst Clerk Manager President Salesman	AvgSal 3500.0000 1437.5000 3158.3333 6500.0000 1800.0000		

"**Correlated**" query Recomputed for each tuple (can't be run independently

- 1. Group by job
- 2. For each tuple, compute aggregate

A subquery in **SELECT** returns a **single value** – used to compute an associated value

(Equivalent) Subquery (FROM)

Find the average salary for each job

"Uncorrelated" query Independent of outer query

practice_emp

empno	ename	job	sal
7369	Smith	Clerk	1200
7499	Allen	Salesman	2000
7521	Ward	Salesman	1650
7566	Jones	Manager	3375
7654	Martin	Salesman	1650
7698	Blake	Manager	3250
7782	Clark	Manager	2850
7788	Scott	Analyst	3500
7839	King	President	6500
7844	Turner	Salesman	1900
7876	Adams	Clerk	1500
7900	James	Clerk	1350
7902	Ford	Analyst	3500
7934	Miller	Clerk	1700

SELE FROM	CT E1.job, AvgSal practice_emp E1,	_	
	(SELECT job, AVG FROM practice_e GROUP BY job) A	(sal) AS mp S E2	AvgSal
WHER	E E1.job = E2.job)	
GROU	P BY E1.job	job	AvgSal
		Analyst	3500.0000 -
		Clerk	1437.5000
		Manager	3158.3333
		President	6500.0000
Idea	:	President Salesman	6500.0000 1800.0000

2. Join the original practice_emp

A subquery in **FROM** returns a **relation** – used as input for another query

Subqueries in WHERE

Find employee name (or names) who earns the highest salary for each job



[more subqueries in WHERE later]

A subquery in WHERE returns a single value – to be compared to another value in a WHERE clause

Subqueries in WITH

Find employee name (or names) who earns the highest salary for each job

practice_emp job empno ename sal 7369 Smith Clerk 1200 7499 Allen Salesman 2000 7521 Ward Salesman 1650 7566 3375 Jones Manager 7654 Martin Salesman 1650 7698 Blake 3250 Manager 7782 Clark 2850 Manager 7788 Scott 3500 Analyst 7839 King President 6500 7844 Turner Salesman 1900 7876 Adams 1500 Clerk 7900 James Clerk 1350 7902 Ford 3500 Analyst 7934 Miller 1700 Clerk

WITH	Ηt	emp AS	
	(S	ELECT job, MAX(sal) AS n	naxSal
	F	ROM practice_emp	
	G	ROUP BY job)	
SELI	ЕСТ	E1.ename	
FRO	Μ	practice emp AS E1, tem	p AS T
			T
WHEI	RE	E1.sal = T.maxSal AND	ename
WHEI	RE	E1.sal = T.maxSal AND E1.job = T.job	ename Allen
WHEI	RE	E1.sal = T.maxSal AND E1.job = T.job	ename Allen Jones
WHEI	RE	E1.sal = T.maxSal AND E1.job = T.job	ename Allen Jones Scott
WHEI	RE	E1.sal = T.maxSal AND E1.job = T.job	ename Allen Jones Scott King
WHEI	RE	E1.sal = T.maxSal AND E1.job = T.job	ename Allen Jones Scott King Ford
WHEI	RE	E1.sal = T.maxSal AND E1.job = T.job	ename Allen Jones Scott King Ford Miller

A subquery in WITH clause returns a temporary relation that can be used by an associated query

[WITH -- not supported by MySQL 5.6, 5.7; work on MySQL 8.0 (GCP and CS server) and XAMPP MariaDB]

Subqueries and Set Operations



(sub-result1)

UNION

(sub-result2)

INTERSECT

[not supported by MySQL 5.6, 5.7; work on MySQL 8.0 (GCP and CS server) and local XAMPP 10.4.11-MariaDB]

(sub-result1)

INTERSECT

(sub-result2)

EXCEPT

[not supported by MySQL 5.6, 5.7; work on MySQL 8.0 (GCP and CS server) and local XAMPP 10.4.11-MariaDB]

(sub-result1)

EXCEPT

(sub-result2)

Requirements:

- Same number of columns
- Same order of columns
- Same column data types

We talked about UNION and INTERSECT. Let's consider EXCEPT

Example: Let's Solve A Problem

Use the following schema. Find IDs and names of all customers who have purchased products sold by company 7777 only. Do not list customers who have purchased from any other companies.

<pre>Product(pid, name, cid)</pre>
cid is foreign key to Company.cid
Company(<u>cid</u> , cname, city)
Customer(<u>custId</u> , name, city)
Purchase(<u>purchase_date</u> , <u>pid</u> , <u>custId</u> , quantity, price)
pid is foreign key to Product.pid,
custId is foreign key to Customer.custId

Assume each customer may purchase the same product multiple times

How should we solve this problem?

Example: Let's Solve A Problem

How should we solve this problem?

- (Find all customers who have purchased) EXCEPT
- 2 (Find all customers who have purchased from other companies, not 7777)

Use EXCEPT to Solve the Problem

- 1 (Find all customers who have purchased) EXCEPT
- (Find all customers who have purchased from other companies, not 7777)

	(sub-re	sult1)		(sub-re	esult2)			
	custid	name		custid	name		(differe	ence)
	991	Humpty		992	Dumpty		custld	name
-	992	Dumpty	_	994	Minnie		991	Humpty
_	992	Dumpty	-	995	Duh		997	Duh
-	994	Minnie		998	Duh			
_	995	Duh				I		
_	995	Duh	-					
	997	Duh						
_	998	Duh						

Use EXCEPT to Solve the Problem (2)

(Find all customers who have purchased)

Find which companies the customers have purchased. Then, find the names of the customers

SELECT T1.custId, T2.cid FROM Purchase T1 NATURAL JOIN Product T2		
GROUP BY T1.custId, T2.cid	custid	cid
	991	7777
	992	7777
	992	7778
Got all customers who	994	7778
have purchased.	995	7777
Still need to find the	995	7779
names of the customers	997	7777
	998	7779

Use EXCEPT to Solve the Problem (3)



Find which companies the customers have purchased Then, find the names of the customers

SELEC	T T3.custId, T3.name		
FROM	SELECT T1.custId, T2.cid		
	FROM Purchase T1 NATURAL JOIN Product	t T2	
	GROUP BY T1.custId. T2.cid) T	custld	name
		991	Humpty
NATUR	AL JOIN Customer T3	992	Dumpty
		992	Dumpty
		994	Minnie
		995	Duh
		995	Duh
		997	Duh
		998	Duh

Use EXCEPT to Solve the Problem (4)



Find all customers who have purchased from other companies Then, find the names of the customers



Use EXCEPT to Solve the Problem (5)



Find all customers who have purchased from other companies Then, find the names of the customers

SELEC	T T3.custId, T3.name		
FROM	(SELECT T1.custId		
	FROM Purchase T1 NATURAL JOIN Produc	t T2	
	WHERE T2.cid <> 7777	custld	name
	GROUP BY T1.custId) T	992	Dumpty
	AT TOTN Customer II2	994	Minnie
NATUR	AL JOIN CUSCOMEL IS	995	Duh
		998	Duh

Use EXCEPT to Solve the Problem (6)

	cust	ld	name
(SELECT T3.custld, T3.name	99	91	Humpty
FROM (SELECT T1.custId, T2.cid	99	92	Dumpty
FROM Purchase T1 NATURAL JOIN Product T2	99	92	Dumpty
GROUP BY T1.custId, T2.cid) T	99	94	Minnie
NATURAL JOIN Customer T3)	99	95	Duh
,	99	95	Duh
EXCEPT	99	97	Duh
	99	98	Duh
(SELECT T3.custId, T3.name		-	
(SELECT T3.custId, T3.name	cust	ld	name
(SELECT T3.custId, T3.name 2 FROM (SELECT T1.custId EDOM Durchage T1 NATURAL TOTA Droduct T2	cust	- Id 92	name Dumpty
(SELECT T3.custId, T3.name FROM (SELECT T1.custId FROM Purchase T1 NATURAL JOIN Product T2	cust 99	- 1 d 92 94	name Dumpty Minnie
(SELECT T3.custId, T3.name FROM (SELECT T1.custId FROM Purchase T1 NATURAL JOIN Product T2 WHERE T2.cid <> 7777	Cust 99 99	92 94 95	name Dumpty Minnie Duh
<pre>(SELECT T3.custId, T3.name 2 FROM (SELECT T1.custId FROM Purchase T1 NATURAL JOIN Product T2 WHERE T2.cid <> 7777 GROUP BY T1.custId) T</pre>	cust 99 99 99 99	92 94 95 98	name Dumpty Minnie Duh Duh
<pre>(SELECT T3.custId, T3.name FROM (SELECT T1.custId FROM Purchase T1 NATURAL JOIN Product T2 WHERE T2.cid <> 7777 GROUP BY T1.custId) T NATURAL JOIN Customer T3)</pre>	cust 99 99 99	92 94 95 98	name Dumpty Minnie Duh Duh
(SELECT T3.custId, T3.name FROM (SELECT T1.custId FROM Purchase T1 NATURAL JOIN Product T2 WHERE T2.cid <> 7777 GROUP BY T1.custId) T NATURAL JOIN Customer T3)	Cust 99 99 99 99 99	1 d 92 94 95 98 98	name Dumpty Minnie Duh Duh

Duh

997

Workaround for EXCEPT



Example: UNION

```
(SELECT T3.custId, T3.name
FROM (SELECT T1.custId, T2.cid
FROM Purchase T1 NATURAL JOIN Product T2
GROUP BY T1.custId, T2.cid) T
NATURAL JOIN Customer T3)
```

UNION

```
(SELECT T3.custId, T3.name
FROM (SELECT T1.custId
    FROM Purchase T1 NATURAL JOIN Product T2
    WHERE T2.cid <> 7777
    GROUP BY T1.custId) T
NATURAL JOIN Customer T3)
```



Example: INTERSECT

```
custid
                                                                  name
(SELECT T3.custId, T3.name
                                                              991
                                                                  Humpty
                                                              992
                                                                  Dumpty
FROM (SELECT T1.custId, T2.cid
                                                              992
                                                                  Dumpty
       FROM Purchase T1 NATURAL JOIN Product T2
                                                              994
                                                                  Minnie
                                                                  Duh
                                                              995
       GROUP BY T1.custId, T2.cid) T
                                                              995
                                                                  Duh
NATURAL JOIN Customer T3)
                                                              997
                                                                  Duh
                                                              998
                                                                  Duh
INTERSECT
                                                            custid
                                                                  name
                                                              992
                                                                  Dumpty
(SELECT T3.custId, T3.name
                                                                  Minnie
                                                              994
                                                              995
                                                                  Duh
FROM (SELECT T1.custId
                                                              998
                                                                  Duh
       FROM Purchase T1 NATURAL JOIN Product T2
                                                                 WHERE T2.cid <> 7777
                                                            custid
                                                                  name
       GROUP BY T1.custId) T
                                                              992
                                                                  Dumpty
                                                              994
                                                                  Minnie
NATURAL JOIN Customer T3)
                                                              995
                                                                  Duh
```

998

Duh

Workaround for INTERSECT



Wrap-Up

- Subqueries in SELECT, FROM
- Abstract immediate result using WITH
- Equivalent queries
- Intro to subqueries in WHERE
- Subqueries and set operations

Note:

- Avoid nested queries if aiming for speed
- Be careful of semantics of nested queries
 - Correlated vs. Uncorrelated

What's next?

- Subqueries in WHERE
- Existential and universal quantifiers
- Triggers and constraints