

SQL: Structured Query Language

Nested Queries

One of the most important features of SQL is that SQL `SELECT` statements can be nested within each other to produce complex queries. While some of the queries discussed below can be written without nesting, some queries require nested structure.

Return Types

In general, a `SELECT` statement returns a relational table as a result. However, there are situations, where the returned table can be cast to a more simple structure. In general, the following four cases are possible. *We note, that in all four cases, it is always possible to treat the returned result as a relational table.*

Consider, for example, the following table (underlined values are the primary key):

`University(id INT, name VARCHAR(60), city VARCHAR(15), state CHAR(2), numStudents INT)`

The return type of a `SELECT` statement may be:

1. **An atomic value.** Returned when a single column of a relation is requested in the `SELECT` clause, and the information requested is unique in the table (only one tuple is returned).

For example, the following query

```
SELECT name
FROM University
WHERE id = 1;
```

returns a single string: the `name` attribute for the university with `id = 1`.

2. **A single tuple.** Returned when the `SELECT` clause contains multiple attributes, but the information requested is unique in the table (only one tuple is returned).

For example, consider the following query:

```
SELECT *
FROM University
WHERE id = 10;
```

This query returns a single tuple, whose value for the primary key attribute is 10.

3. **A single column.** Returned when a single attribute is listed in the `SELECT` clause, and when the result contains more than one entry. This can be viewed as a set of atomic values. For example, the query below, returns the list of IDs for all universities in California.

```
SELECT id
FROM University
WHERE state = 'CA';
```

4. **A relational table.** This is the most general return type, used when the `SELECT` clause contains multiple columns, and the result of the query is multiple tuples. This return type can be viewed as a set of tuples.

The following query returns as its result the list of all California campuses and their corresponding enrollments:

```
SELECT name, numStudents  
FROM University  
WHERE state = 'CA';
```

Note: For some queries, the return type can be established statically, by analyzing constraints on the relational table. For other queries, the actual return type of the query run on some instances can be different than the statically predicted return type (e.g., the last query will return a single tuple if only one record about a California university exists in the table).

Nested Queries in the FROM clause

Since the return of any SELECT statement can be viewed as a relational table, a nested SELECT statement can be used in the FROM clause. This allows you to construct a temporary table which is useful for the purposes of the query, but does not need to be made persistent. Any attribute of the table resulting from the nested SELECT statement can be accessed from within the WHERE, GROUP BY, HAVING, and SELECT clauses of the main query, just like attributes of regular relational tables.

Syntax:

```
SELECT <select list >
FROM (SELECT query) <alias >, ...
  [WHERE <conditions> ]
  [GROUP BY <attribute list >
  [HAVING <conditions> ] ]
```

Two rules must be followed:

- The nested SELECT statement must be enclosed in parentheses.
- The result of the nested SELECT statement must be given a table alias if it is in the FROM clause and all aggregate/expression attributes in the result of the nested SELECT statement that are to be referenced in the outer query must be given column aliases.

Example. Consider the following tables (of university courses, students and student enrollments in courses):
Course(major, num, name, room, instructor, capacity)
Student(ssn, firstName, lastName, major, year)
Enrollments(courseMajor, courseNum, studentId)

Consider the following query¹:

```
SELECT C.major, C.num, room
FROM
  Course C
  JOIN (
    SELECT E.courseMajor, E.courseNum, COUNT(*) AS enrolled
    FROM
      Student S
      JOIN Enrollments E ON E.studentId = S.ssn
    WHERE S.year > 3
    GROUP BY E.courseMajor, E.courseNum
  ) S ON
  S.courseMajor = C.major
  AND S.courseNumber = C.num
WHERE S.enrolled >= capacity / 2;
```

The nested query:

```
SELECT E.courseMajor, E.courseNum, COUNT(*) AS enrolled
FROM
  Student S
  JOIN Enrollments E ON E.studentId = S.ssn
WHERE S.year > 3
GROUP BY E.courseMajor, E.courseNum
```

¹Note that some SQL style has been relaxed to conserve vertical space in the lecture notes.

Finds all seniors (`year > 3`) and joins their information with the information about the courses they are taking. This information is then grouped by course id/number, resulting in the relation that reports a list of courses and the number of senior students enrolled in each².

This information is reported back as a three-attribute relation with alias `S` to the main query, which, in turn, finds and reports the courses and the rooms of the courses, in which senior students occupy more than one half of the designated capacity.

The same query can be rewritten as:

```
SELECT
    C.major ,
    C.num,
    room
FROM
    Enrollments E
JOIN Student S ON E.studentId = S.ssn
JOIN Course C ON
    C.Major = E.courseMajor
    AND C.num = E.courseNumber
WHERE S.year > 3
GROUP BY
    C.Major ,
    E.courseNum ,
    capacity
HAVING COUNT(*) >= capacity / 2;
```

Note, that in this query, `capacity` has to be included in the `GROUP BY` clause, to allow for its use in the `HAVING` clause.

² For courses that have at least one senior student enrolled.

Nested Queries in the WHERE Clause

Nested queries that return atomic values

A SELECT statement that returns a single value can be used in any expression in the WHERE clause where the value of this type is expected.

Example Consider the following tables (a list of bank accounts and a list of transactions for the account): each transaction either adds money (After > Before) or removes money (After < Before):

Accounts(acctNo, owner, type, dateOpened, balance)

Transactions(acctNo, transactionDate, transactionType, before, after)

Let us also assume that we have a UNIQUE(owner, type) constraint on Accounts. We can write the query that finds all dates on which “John Smith” added money to his savings account as follows:

```
SELECT transactionDate
FROM Transactions T
WHERE
    T.before < T.after
    AND AcctNo = (
        SELECT acctNo
        FROM Accounts A
        WHERE
            A.owner = 'John_Smith'
            AND A.type = 'Savings'
    )
;
```

Return Value	Allowed Comparisons
numeric	<, >, <=, >=, !=, <>, =
strings	=, !=, <>, LIKE
other	=, !=, <>

Nested queries that return single tuples

SQL has a tuple (list) construct: (<Att>, <Att>, ..., <Att>). When a SELECT statement returns a single tuple, tuples formed using this construct can be compared to the returned tuple.

Example Consider the following tables:

People(id, firstName, lastName, age, gender)

Interests(firstName, lastName, interest)

The following query returns all interests of the person who has the id number 1 in the People table.

```
SELECT i.interest
FROM Interests i
WHERE (i.firstName, i.lastName) = (
    SELECT firstName, lastName
    FROM People
    WHERE id = 1
);
```

Return Value	Allowed comparisons
tuple (list)	=, !=, <>

Nested queries that return single columns

When a SELECT statement returns a single column as a result, the result can be treated as a set of atomic values.

SQL allows for a number of comparison operations between atomic values and sets:

Operation	Explanation
EXISTS <SubQuery>	true if <SubQuery> returns a non-empty table
<AtomicValue> IN <SubQuery>	true if <AtomicValue> is one of the elements in the result of <SubQuery>
<AtomicValue> <Op> ALL (<SubQuery>)	<Op> $\in \{>, <, >=, <=, =, !=, <>\}$ universal comparison
<AtomicValue> <Op> ANY (<SubQuery>)	<Op> $\in \{>, <, >=, <=, =, !=, <>\}$ existential comparison

A universal comparison operation evaluates to true iff <AtomicValue> is in the required relationship (<Op>) with **every** member of the set returned by <Expression>.

An existential comparison operation evaluates to true iff <AtomicValue> is in the required relationship (<Op>) with **at least one** member of the set returned by <Expression>.

Examples. Consider the relational tables Accounts and Transactions defined above:

Accounts(acctNo, owner, type, dateOpened, balance)

Transactions(acctNo, transactionDate, transactionType, before, after)

The following query finds the account information for all accounts that had at least one \$1000 deposit.

```
SELECT A.acctNo, A.owner, A.type
FROM Accounts A
WHERE A.acctNo IN (
    SELECT DISTINCT acctNo
    FROM Transactions
    WHERE
        after - before = 1000
        AND transactionType = 'deposit'
);
```

This query finds the information about all accounts with the maximum balance.

```
SELECT A.acctNo, A.owner, A.balance
FROM Accounts
WHERE balance >= ALL (SELECT Balance FROM Accounts);
```

Correlated Queries

In examples above, nested queries can be evaluated once for the entire “lifetime” of the full query. In some situations, though, the result of the inner query depends changes as the evaluation of the outer query proceeds. Such queries are called *correlated queries*.

Examples. The query below reports the highest balance for each type of the account.

```
SELECT DISTINCT A.type, A.balance
FROM Accounts A
WHERE A.balance >= ALL (
    SELECT Balance
    FROM Accounts AA
    WHERE AA.type = A.type
);
```

Note, that in this case, correlating subquery and the main query is unavoidable — the subquery has to select only the balances of accounts of the same type as the current account number being considered in the outer query.

The next query is one possible way of finding all accounts that had transactions on October 1, 2007.

```
SELECT A.acctNo , A.owner
FROM Accounts A
WHERE EXISTS (
    SELECT *
    FROM Transactions T
    WHERE
        T.transactionDate = DATE( '2007-10-01 ' )
        AND T.acctNo = A.acctNo
);
```

In this case the use of **EXISTS** operation naturally represents our intuition: we are not interested in the specific details of transactions that happened on October 1, 2007, rather, we are simply verifying that such transactions exist for each account.

Nested queries that return relational tables

SELECT statements that return multicolumn relational tables can also be used as subqueries. The table below shows operations which they can be used:

Expression	Allowed Operators
EXISTS (<Subquery>)	N/A
<Tuple> IN (<SubQuery>)	N/A
<Tuple> <OP> ALL (<SubQuery>)	=, <>, !=
<Tuple> <OP> ANY (<SubQuery>)	=, <>, !=

Examples. Consider the Accounts and Transactions tables discussed above.

Accounts(acctNo, owner, type, dateOpened, balance)

Transactions(acctNo, transactionDate, transactionType, before, after)

The following query finds all dates when transactions *like* (same account type, same amount) any transactions on the accounts owned by Bob Brown occurred.

```
SELECT transactionDate
FROM Transactions
WHERE (transactionType, After - Before) = ANY (
    SELECT transactionType, after - before
    FROM Transactions T, Accounts A
    WHERE
        A.AcctNo = T.AcctNo
        AND A.Owner = 'Bob_Brown'
);
```