

SQL SELECT

Database Systems Lecture 7
Natasha Alechina

In this Lecture

- SQL SELECT
 - WHERE clauses
 - SELECT from multiple tables
 - JOINS
- For more information
 - Connolly and Begg Chapter 5
 - Ullman and Widom Chapter 6.1-6.3

SQL SELECT Overview

SELECT

[DISTINCT | ALL] <column-list>

FROM <table-names>

[WHERE <condition>]

[ORDER BY <column-list>]

[GROUP BY <column-list>]

[HAVING <condition>]

- *([] - optional, | - or)*

Example Tables

Student

ID	First	Last
S103	John	Smith
S104	Mary	Jones
S105	Jane	Brown
S106	Mark	Jones
S107	John	Brown

Course

Code	Title
DBS	Database Systems
PR1	Programming 1
PR2	Programming 2
IAI	Intro to AI

Grade

ID	Code	Mark
S103	DBS	72
S103	IAI	58
S104	PR1	68
S104	IAI	65
S106	PR2	43
S107	PR1	76
S107	PR2	60
S107	IAI	35

DISTINCT and ALL

- Sometimes you end up with duplicate entries
- Using **DISTINCT** removes duplicates
- Using **ALL** retains them - this is the default

```
SELECT ALL Last  
FROM Student
```

Last
Smith
Jones
Brown
Jones
Brown

```
SELECT DISTINCT Last  
FROM Student
```

Last
Smith
Jones
Brown

WHERE Clauses

- Usually you don't want all the rows
 - A **WHERE** clause restricts the rows that are returned
 - It takes the form of a condition - only those rows that satisfy the condition are returned
- Example conditions:
 - **Mark < 40**
 - **First = 'John'**
 - **First <> 'John'**
 - **First = Last**
 - **(First = 'John') AND (Last = 'Smith')**
 - **(Mark < 40) OR (Mark > 70)**

WHERE Examples

```
SELECT * FROM Grade  
WHERE Mark >= 60
```

ID	Code	Mark
S103	DBS	72
S104	PR168	
S104	IAI	65
S107	PR176	
S107	PR260	

```
SELECT DISTINCT ID  
FROM Grade  
WHERE Mark >= 60
```

ID
S103
S104
S107

WHERE Example

- Given the table

Grade

ID	Code	Mark
S103	DBS	72
S103	IAI	
S104	PR168	
S104	IAI	
S106	PR243	
S107	PR176	
S107	PR260	
S107	IAI	

- Write an SQL query to find a list of the ID numbers and marks in IAI of students who have passed (scored 40 or higher) IAI

ID	Mark
S103	58
S104	65

One Solution

We only want the ID and Mark, not the Code
Single quotes around the string

```
SELECT ID, Mark FROM  
Grade  
WHERE (Code = 'IAI')  
AND  
(Mark >= 40)
```

We're only interested in IAI

We're looking for entries with pass marks

SELECT from Multiple Tables

- Often you need to combine information from two or more tables
- You can get the effect of a product by using
- If the tables have columns with the same name ambiguity results
- You resolve this by referencing columns with the table name

```
SELECT * FROM Table1,  
Table2...
```

```
TableName.Column
```

SELECT from Multiple Tables

```
SELECT
    First, Last, Mark
FROM Student, Grade
WHERE
    (Student.ID =
        Grade.ID) AND
    (Mark >= 40)
```

Student

ID	FirstLast	
S103	John	Smith
S104	Mary	Jones
S105	Jane	Grade
S106	Mark	
S107	John	

ID	Code	Mark
S103	DBS	72
S103	IAI	58
S104	PR168	
S104	IAI	65
S106	PR243	
S107	PR176	
S107	PR260	
S107	IAI	35

SELECT from Multiple Tables

SELECT ... FROM Student, Grade

WHERE ...

Are matched
with the first
entry from
the Student
table...

All of the
entries from
the Grade
table

And then
with the
second...

and so on

ID	FirstLastID	Code	Mark		
S103	John	Smith	S103	DBS	72
S103	John	Smith	S103	IAI 58	
S103	John	Smith	S104	PR168	
S103	John	Smith	S104	IAI 65	
S103	John	Smith	S106	PR243	
S103	John	Smith	S107	PR176	
S103	John	Smith	S107	PR260	
S103	John	Smith	S107	IAI 35	
S104	Mary	Jones	S103	DBS	72
S104	Mary	Jones	S103	IAI 58	
S104	Mary	Jones	S104	PR168	
S104	Mary	Jones	S104	IAI 65	
S104	Mary	Jones	S104	IAI 65	

SELECT from Multiple Tables

```
SELECT ... FROM Student, Grade
WHERE (Student.ID = Grade.ID) AND
```

ID	FirstLast	ID	Code	Mark		
S103	John	Smith	S103	DBS	72	
S103	John	Smith	S103	IAI 58		
S104	Mary	Jones	S104	PR168		
S104	Mary	Jones	S104	IAI 65		
S106	Mark	Jones	S106	PR243		
S107	John	Brown	S107	PR176		
S107	John	Brown	S107	PR260		
S107	John	Brown	S107	IAI 35		

Student.ID

Grade.ID

SELECT from Multiple Tables

```
SELECT ... FROM Student, Grade
WHERE (Student.ID = Grade.ID) AND (Mark >=
40)
```

ID	FirstLastID	Code	Mark		
S103	John	Smith	S103	DBS	72
S103	John	Smith	S103	IAI 58	
S104	Mary	Jones	S104	PR168	
S104	Mary	Jones	S104	IAI 65	
S106	Mark	Jones	S106	PR243	
S107	John	Brown	S107	PR176	
S107	John	Brown	S107	PR260	

SELECT from Multiple Tables

```
SELECT First, Last, Mark FROM Student, Grade  
WHERE (Student.ID = Grade.ID) AND (Mark >=  
40)
```

First	Last	Mark
John	Smith	72
John	Smith	58
Mary	Jones	68
Mary	Jones	65
Mark	Jones	43
John	Brown	76
John	Brown	60

SELECT from Multiple Tables

- When selecting from multiple tables you almost always use a **WHERE** clause to find entries with common values

```
SELECT * FROM
    Student, Grade,
    Course
WHERE
    Student.ID = Grade.ID
AND
    Course.Code =
    Grade.Code
```


SELECT from Multiple Tables

Student

Grade

Course

ID	FirstLast	ID	Code	Mark	Code	Title		
S103	John	Smith	S103	DBS	72	DBS	Database Systems	
S103	John	Smith	S103	IAI	58	IAI	Intro to AI	
S104	Mary	Jones	S104	PR168	PR1	Programming 1		
S104	Mary	Jones	S104	IAI	65	IAI	Intro to AI	
S106	Mark	Jones	S106	PR243	PR2	Programming 2		
S107	John	Brown	S107	PR176	PR1	Programming 1		
S107	John	Brown	S107	PR260	PR2	Programming 2		
S107	John	Brown	S107	IAI	35	IAI	Intro to AI	

Student.ID = Grade.ID

Course.Code = Grade.Code

JOINS

- JOINS can be used to combine tables

- There are many types of JOIN

- CROSS JOIN
 - INNER JOIN
 - NATURAL JOIN
 - OUTER JOIN

- OUTER JOINS are linked with NULLS - more later

A CROSS JOIN B

- returns all pairs of rows from A and B

A NATURAL JOIN B

- returns pairs of rows with common values for identically named columns and without duplicating columns

A INNER JOIN B

- returns pairs of rows satisfying a condition

CROSS JOIN

Student

ID	Name
123	John
124	Mary
125	Mark
126	Jane

Enrolment

ID	Code
123	DBS
124	PRG
124	DBS
126	PRG

```
SELECT * FROM
Student CROSS JOIN
Enrolment
```

ID	Name	ID	Code	
123	John	123	DBS	
124	Mary	123	DBS	
125	Mark	123	DBS	
126	Jane	123	DBS	
123	John	124	PRG	
124	Mary	124	PRG	
125	Mark	124	PRG	
126	Jane	124	PRG	
123	John	124	DBS	
124	Mary	124	DBS	

NATURAL JOIN

Student

ID	Name
123	John
124	Mary
125	Mark
126	Jane

Enrolment

ID	Code
123	DBS
124	PRG
124	DBS
126	PRG

SELECT * FROM

**Student NATURAL JOIN
Enrolment**

ID	Name	Code
123	John	DBS
124	Mary	PRG
124	Mary	DBS
126	Jane	PRG

CROSS and NATURAL JOIN

```
SELECT * FROM  
  A CROSS JOIN B
```

- is the same as

```
SELECT * FROM A, B
```

```
SELECT * FROM  
  A NATURAL JOIN B
```

- is the same as

```
SELECT A.col1,... A.coln,  
      [and all other columns  
      apart from  
      B.col1,...B.coln]
```

```
FROM A, B
```

```
WHERE A.col1 = B.col1
```

```
      AND A.col2 = B.col2
```

```
      ...AND A.coln = B.col.n
```

(this assumes that col1...
coln in A and B have
common names)

INNER JOIN

- **INNER JOINS** specify a condition which the pairs of rows satisfy

```
SELECT * FROM  
  A INNER JOIN B  
  ON <condition>
```

- Can also use

```
SELECT * FROM  
  A INNER JOIN B  
  USING  
    (col1, col2,...)
```
- Chooses rows where the given columns are equal

INNER JOIN

Student

ID	Name
123	John
124	Mary
125	Mark
126	Jane

Enrolment

ID	Code
123	DBS
124	PRG
124	DBS
126	PRG

SELECT * FROM

**Student INNER JOIN
Enrolment USING (ID)**

ID	Name	ID	Code	
123	John	123	DBS	
124	Mary	124	PRG	
124	Mary	124	DBS	
126	Jane	126	PRG	

INNER JOIN

Buyer

Name	Budget
Smith	100,000
Jones	150,000
Green	80,000

Property

Address	Price
15 High St	85,000
12 Queen St	125,000
87 Oak Row	175,000

```
SELECT * FROM
  Buyer INNER JOIN
  Property ON
    Price <= Budget
```

Name	Budget	Address	Price
Smith	100,000	15 High St	85,000
Jones	150,000	15 High St	85,000
Jones	150,000	12 Queen St	125,000

INNER JOIN

```
SELECT * FROM  
    A INNER JOIN B  
    ON <condition>
```

- is the same as

```
SELECT * FROM A, B  
    WHERE <condition>
```

```
SELECT * FROM  
    A INNER JOIN B  
    USING(col1, col2, ...)
```

- is the same as

```
SELECT * FROM A, B  
    WHERE A.col1 = B.col1  
        AND A.col2 = B.col2  
        AND ...
```

JOINS vs WHERE Clauses

- JOINS (so far) are not needed
 - You can have the same effect by selecting from multiple tables with an appropriate WHERE clause
 - So should you use JOINS or not?
- Yes, because
 - They often lead to concise queries
 - NATURAL JOINS are very common
- No, because
 - Support for JOINS varies a fair bit among SQL dialects

Writing Queries

- When writing queries
 - There are often many ways to write the query
 - You should worry about being correct, clear, and concise in that order
 - Don't worry about being clever or efficient
- Most DBMSs have query optimisers
 - These take a user's query and figure out how to efficiently execute it
 - A simple query is easier to optimise
 - We'll look at some ways to improve efficiency later

This Lecture in Exams

Track

cID	Num	Title	Time	aID
1	1	Violent	2391	
1	2	Every Girl	4101	
1	3	Breather	2171	
1	4	Part of Me	2791	
2	1	Star	3621	
2	2	Teaboy	4172	

CD

cID	Title	Price
1	Mix	9.99
2	Compilation	12.99

Artist

aID	Name
1	Stellar
2	Cloudboy

This Lecture in Exams

Find a list of all the CD titles.

(1 mark)

Find a list of the titles of tracks that are more than 300 seconds long.

(2 marks)

Find a list of the names of those artists who have a track on the CD with the title "Compilation".

(4 marks)

Next Lecture

- More SQL SELECT
 - Aliases
 - 'Self-joins'
 - Subqueries
 - IN, EXISTS, ANY, ALL
- For more information
 - Connolly and Begg Chapter 5
 - Ullman and Widom Chapter 6