

# SQL: The Query Language

Fall 2019

Life is just a bowl of queries.

-Anon

1

## Relational Query Languages

- **A major strength of the relational model: supports simple, powerful *querying* of data.**
- **Query languages can be divided into two parts**
  - Data Manipulation Language (DML)
    - Allows for queries and updates
  - Data Definition Language (DDL)
    - Define and modify schema (at all 3 levels)
    - Permits database tables to be created or deleted. It also defines indexes (keys), specifies links between tables, and imposes constraints between tables
- **The DBMS is responsible for efficient evaluation.**
  - The key: precise semantics for relational queries.
  - Allows the optimizer to extensively re-order operations, and still ensure that the answer does not change.
  - Internal cost model drives use of indexes and choice of access paths and physical operators.

2

# The SQL Query Language

- The most widely used relational query language.
- Originally IBM, then ANSI in 1986
- **Current standard is SQL-2008**
  - 2003 was last major update: XML, window functions, sequences, auto-generated IDs.
  - Not fully supported yet
- SQL-1999 Introduced "Object-Relational" concepts.
  - Also not fully supported yet.
- **SQL92 is a basic subset**
  - Most systems support at least this
- PostgreSQL has some "unique" aspects (as do most systems).
- SQL is not synonymous with Microsoft's "SQL Server"

3

## The SQL DML

- Single-table queries are straightforward.
- To find all 18 year old students, we can write:

```
SELECT *  
  FROM Students  
 WHERE age=18
```

```
SELECT *  
  FROM Students S  
 WHERE S.age=18
```

```
SELECT *  
  FROM Students  
 WHERE Students.age=18
```

4

# The SQL DML

- Single-table queries are straightforward.
- To find all 18 year old students, we can write:

```
SELECT *  
FROM Students S  
WHERE S.age=18
```

- To find just names and logins, replace the first line:  
SELECT S.name, S.login

5

## Querying Multiple Relations

- Can specify a join over two tables as follows:

```
SELECT S.name, E.cid  
FROM Students S, Enrolled E  
WHERE S.sid=E.sid AND E.grade='B'
```

sid	cid	grade
53831	Carnatic101	C
53831	Reggae203	B
53650	Topology112	A
53666	History105	B

sid	name	login	age	gpa
53666	Jones	jones@cs	18	3.4
53688	Smith	smith@ee	18	3.2

result =

S.name	E.cid
Jones	History105

Note: obviously no referential integrity constraints have been used here.

6

## Basic SQL Query

SELECT	[DISTINCT] <i>target-list</i>
FROM	<i>relation-list</i>
WHERE	<i>qualification</i>

- *relation-list*: A list of relation names
  - possibly with a *range-variable* after each name
- *target-list*: A list of attributes of tables in *relation-list*
- *qualification*: Comparisons combined using AND, OR and NOT.
  - Comparisons are Attr *op* const or Attr1 *op* Attr2, where *op* is one of = < > ≤ ≥
- *DISTINCT*: optional keyword indicating that the answer should not contain duplicates.
  - In SQL SELECT, the default is that duplicates are not eliminated! (Result is called a “multiset”)

## Query Semantics

- Semantics of an SQL query are defined in terms of the following conceptual evaluation strategy:
  1. do FROM clause: compute *cross-product* of tables (e.g., Students and Enrolled).
  2. do WHERE clause: Check conditions, discard tuples that fail. (i.e., “selection”).
  3. do SELECT clause: Delete unwanted fields. (i.e., “projection”).
  4. If DISTINCT specified, eliminate duplicate rows.

Probably the least efficient way to compute a query!

- An optimizer will find more efficient strategies to get the *same answer*.

## Step 1 – Cross Product





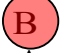

S.sid	S.name	S.login	S.age	S.gpa	E.sid	E.cid	E.grade
53666	Jones	jones@cs	18	3.4	53831	Carnatic101	C
53666	Jones	jones@cs	18	3.4	53832	Reggae203	B
53666	Jones	jones@cs	18	3.4	53650	Topology112	A
53666	Jones	jones@cs	18	3.4	53666	History105	B
53688	Smith	smith@ee	18	3.2	53831	Carnatic101	C
53688	Smith	smith@ee	18	3.2	53831	Reggae203	B
53688	Smith	smith@ee	18	3.2	53650	Topology112	A
53688	Smith	smith@ee	18	3.2	53666	History105	B

sid	cid	grade
53831	Carnatic101	C
53831	Reggae203	B
53650	Topology112	A
53666	History105	B

sid	name	login	age	gpa
53666	Jones	jones@cs	18	3.4
53688	Smith	smith@ee	18	3.2

9

## Step 2) Discard tuples that fail predicate

S.sid	S.name	S.login	S.age	S.gpa	E.sid	E.cid	E.grade
53666	Jones	jones@cs	18	3.4	53831	Carnatic101	C
53666	Jones	jones@cs	18	3.4	53832	Reggae203	
53666	Jones	jones@cs	18	3.4	53650	Topology112	A
	Jones	jones@cs	18	3.4		History105	
53688	Smith	smith@ee	18	3.2	53831	Carnatic101	C
53688	Smith	smith@ee	18	3.2	53831	Reggae203	
53688	Smith	smith@ee	18	3.2	53650	Topology112	A
53688	Smith	smith@ee	18	3.2	53666	History105	

```
SELECT S.name, E.cid
FROM Students S, Enrolled E
WHERE S.sid=E.sid AND E.grade='B' 10
```

## Step 3) Discard Unwanted Columns

S.sid	S.name	S.login	S.age	S.gpa	E.sid	E.cid	E.grade
53666	Jones	jones@cs	18	3.4	53831	Carnatic101	C
53666	Jones	jones@cs	18	3.4	53832	Reggae203	B
53666	Jones	jones@cs	18	3.4	53650	Topology112	A
53666	Jones	jones@cs	18	3.4	53666	History105	B
53688	Smith	smith@ee	18	3.2	53831	Carnatic101	C
53688	Smith	smith@ee	18	3.2	53831	Reggae203	B
53688	Smith	smith@ee	18	3.2	53650	Topology112	A
53688	Smith	smith@ee	18	3.2	53666	History105	B

```
SELECT S.name, E.cid
FROM Students S, Enrolled E
WHERE S.sid=E.sid AND E.grade='B' 11
```

## Null Values

- Field values in a tuple are sometimes **unknown** (e.g., a rating has not been assigned) or **inapplicable** (e.g., no spouse's name).
  - SQL provides a special value **null** for such situations.
- The presence of **null** complicates many issues. E.g.:
  - Special operators needed to check if value is/is not **null**.
  - Is  $rating > 8$  true or false when  $rating$  is equal to **null**?  
What about **AND**, **OR** and **NOT** connectives?
  - We need a **3-valued logic** (true, false and **unknown**).
  - Meaning of constructs must be defined carefully. (e.g., WHERE clause eliminates rows that don't evaluate to true.)
  - New operators (in particular, **outer joins**) possible/needed.

# Null Values – 3 Valued Logic

(null > 0) is null  
 (null + 1) is null  
 (null = 0) is null  
 null AND true is null

AND	T	F	Null
T	T	F	Null
F	F	F	F
NULL	Null	F	Null

OR	T	F	Null
T	T	T	T
F	T	F	Null
NULL	T	Null	Null

13

## Now the Details

We will use these instances of relations in our examples.

*Reserves*

<u>sid</u>	<u>bid</u>	<u>day</u>
22	101	10/10/96
95	103	11/12/96

*Sailors*

<u>sid</u>	sname	rating	age
22	Dustin	7	45.0
31	Lubber	8	55.5
95	Bob	3	63.5

*Boats*

<u>bid</u>	bname	color
101	Interlake	blue
102	Interlake	red
103	Clipper	green
104	Marine	red

14

## Example Schemas (in SQL DDL)

<u>sid</u>	sname	rating	age
22	Dustin	7	45.0
31	Lubber	8	55.5
95	Bob	3	63.5

```
CREATE TABLE Sailors (sid INTEGER, sname  
CHAR(20), rating INTEGER, age REAL,  
PRIMARY KEY sid)
```

Consider the use  
Of VARCHAR instead

15

## Example Schemas (in SQL DDL)

<u>bid</u>	bname	color	
101	Interlake	blue	
102	Interlake	red	
103	Clipper	green	
104	Marine	red	

```
CREATE TABLE Boats (bid INTEGER, bname CHAR (20),  
color CHAR(10)  
PRIMARY KEY bid)
```

16



## Example Schemas (in SQL DDL)

<u>sid</u>	<u>bid</u>	<u>day</u>
22	101	10/10/96
95	103	11/12/96

```
CREATE TABLE Reserves (sid INTEGER, bid INTEGER,  
    day DATE,  
    PRIMARY KEY (sid, bid, day),  
    FOREIGN KEY sid REFERENCES Sailors,  
    FOREIGN KEY bid REFERENCES Boats)
```

17

## Another Join Query

```
SELECT sname  
FROM Sailors, Reserves  
WHERE Sailors.sid=Reserves.sid  
AND bid=103
```

(sid)	sname	rating	age	(sid)	bid	day
22	dustin	7	45.0	22	101	10/10/96
22	dustin	7	45.0	58	103	11/12/96
31	lubber	8	55.5	22	101	10/10/96
31	lubber	8	55.5	58	103	11/12/96
95	Bob	3	63.5	22	101	10/10/96
95	Bob	3	63.5	95	103	11/12/96

18

## Some Notes on Range Variables

- Can associate "range variables" with the tables in the FROM clause.
  - saves writing, makes queries easier to understand
- Needed when ambiguity could arise.
  - for example, if same table used multiple times in same FROM (called a "self-join")

```
SELECT sname
FROM Sailors, Reserves
WHERE Sailors.sid=Reserves.sid AND bid=103
```

Can be  
rewritten using  
range variables as:

```
SELECT S.sname
FROM Sailors S, Reserves R
WHERE S.sid=R.sid AND bid=103
```

## More Notes

- Here's an example where range variables are required (self-join example):

```
SELECT x.sname, x.age, y.sname, y.age
FROM Sailors x, Sailors y
WHERE x.age > y.age
```

- Note that target list can be replaced by "\*" if you don't want to do a projection:

```
SELECT *
FROM Sailors x
WHERE x.age > 20
```

## Find sailors who've reserved at least one boat

```
SELECT S.sid
FROM Sailors S, Reserves R
WHERE S.sid=R.sid
```

- Would adding DISTINCT to this query make a difference?
- What is the effect of replacing *S.sid* by *S.sname* in the SELECT clause?
  - Would adding DISTINCT to this variant of the query make a difference?

21

## Expressions

- Can use arithmetic expressions in SELECT clause (plus other operations we'll discuss later)
- Use **AS** to provide column names

```
SELECT S.age, S.age-5 AS age1, 2*S.age AS age2
FROM Sailors S
WHERE S.sname = 'dustin'
```

- Can also have expressions in WHERE clause:

```
SELECT S1.sname AS name1, S2.sname AS name2
FROM Sailors S1, Sailors S2
WHERE 2*S1.rating = S2.rating - 1
```

22

## String operations

- SQL also supports some string operations
- “LIKE” is used for string matching.

```
SELECT  S.age, age1=S.age-5, 2*S.age AS age2
FROM    sailors S
WHERE   S.sname LIKE 'B_%B'
```

`\_` stands for any one character and `%` stands for 0 or more arbitrary characters.

23

Find sid's of sailors who've reserved a red **or** a green boat

```
SELECT DISTINCT R.sid
FROM Boats B,Reserves R
WHERE R.bid=B.bid AND
(B.color='red'OR B.color='green')
```

Vs.

(note:  
UNION  
eliminates  
duplicates  
by default.  
Override w/  
UNION ALL)

```
SELECT  R.sid
FROM Boats B, Reserves R
WHERE R.bid=B.bid AND B.color='red'
UNION  SELECT R.sid
FROM Boats B, Reserves R
WHERE  R.bid=B.bid AND
        B.color='green'
```

24

Find sid's of sailors who've reserved a red **and** a green boat

- If we simply replace **OR** by **AND** in the previous query, we get the wrong answer. (Why?)

```
SELECT R1.sid
FROM Boats B1, Reserves R1,
      Boats B2, Reserves R2
WHERE R1.sid=R2.sid
      AND R1.bid=B1.bid
      AND R2.bid=B2.bid
      AND (B1.color='red' AND B2.color='green')
```

25

## AND Continued...

- **INTERSECT:**
  - Can be used to compute the intersection of any two *union-compatible* sets of tuples.
- **EXCEPT**
  - (sometimes called MINUS)
  - Included in the SQL/92 standard, but **many** systems (including MySQL) don't support them.

Key field!

```
SELECT S.sid
FROM Sailors S, Boats B,
      Reserves R
WHERE S.sid=R.sid
      AND R.bid=B.bid
      AND B.color='red'
INTERSECT
SELECT S.sid
FROM Sailors S, Boats B,
      Reserves R
WHERE S.sid=R.sid
      AND R.bid=B.bid
      AND B.color='green'
```

26

## Nested Queries

- **Powerful feature of SQL: WHERE clause can itself contain an SQL query!**

– Actually, so can FROM and HAVING clauses.

*Names of sailors who've reserved boat #103:*

```
SELECT S.sname
FROM Sailors S
WHERE S.sid IN (SELECT R.sid
                FROM Reserves R
                WHERE R.bid=103)
```

- **To find sailors who've *not* reserved #103, use NOT IN.**
- **To understand semantics of nested queries:**
  - *think* of a nested loops evaluation: *For each Sailors tuple, check the qualification by computing the subquery.*

## Nested Queries with Correlation

*Find names of sailors who've reserved boat #103:*

```
SELECT S.sname
FROM Sailors S
WHERE EXISTS (SELECT *
              FROM Reserves R
              WHERE R.bid=103 AND S.sid=R.sid)
```

Tests whether the set is nonempty

- **EXISTS** is another set comparison operator, like **IN**.
- Can also specify **NOT EXISTS**
- If **UNIQUE** is used, and \* is replaced by *R.bid*, finds sailors with at most one reservation for boat #103.
  - **UNIQUE** checks for duplicate tuples in a subquery;
- **Subquery must be recomputed for each Sailors tuple.**
  - Think of subquery as a function call that runs a query!

## More on Set-Comparison Operators

- We've already seen **IN**, **EXISTS** and **UNIQUE**. Can also use **NOT IN**, **NOT EXISTS** and **NOT UNIQUE**.
- Also available: **op ANY**, **op ALL**
- Find sailors whose rating is greater than that of some sailor called Horatio:

```
SELECT *
FROM sailors S
WHERE S.rating > ANY
      (SELECT S2.rating
       FROM sailors S2
       WHERE S2.sname='Horatio')
```

29

## Rewriting INTERSECT Queries Using IN

*Find sid's of sailors who've reserved both a red and a green boat:*

```
SELECT R.sid
FROM Boats B, Reserves R
WHERE R.bid=B.bid
      AND B.color='red'
      AND R.sid IN (SELECT R2.sid
                   FROM Boats B2, Reserves R2
                   WHERE R2.bid=B2.bid
                   AND B2.color='green')
```

- Similarly, **EXCEPT** queries re-written using **NOT IN**.
- How would you change this to find *names* (not *sids*) of Sailors who've reserved both red and green boats?

# Division in SQL

Find names of sailors who've reserved all boats.

```
SELECT S.sname  Sailors S such that ...
FROM Sailors S
WHERE NOT EXISTS (SELECT B.bid
there is no boat B
                  FROM Boats B
                  WHERE NOT EXISTS (SELECT R.bid
that doesn't have ...
                                   FROM Reserves R
                                   WHERE R.bid=B.bid
                                   AND R.sid=S.sid))
a Reserves tuple showing S reserved B
```

Recall Exists Tests whether the set is nonempty

31

## Division Operations in SQL (1)

Find names of sailors who've reserved all boat:

```
SELECT S.sname
FROM Sailors S
WHERE NOT EXIST ((SELECT B.bid
                  FROM Boats B) EXCEPT
                (SELECT R.bid
                  FROM Reserves R
                  WHERE R.sid = S.sid ))
```

The sailor reserved all boats

Boats not reserved by the sailor

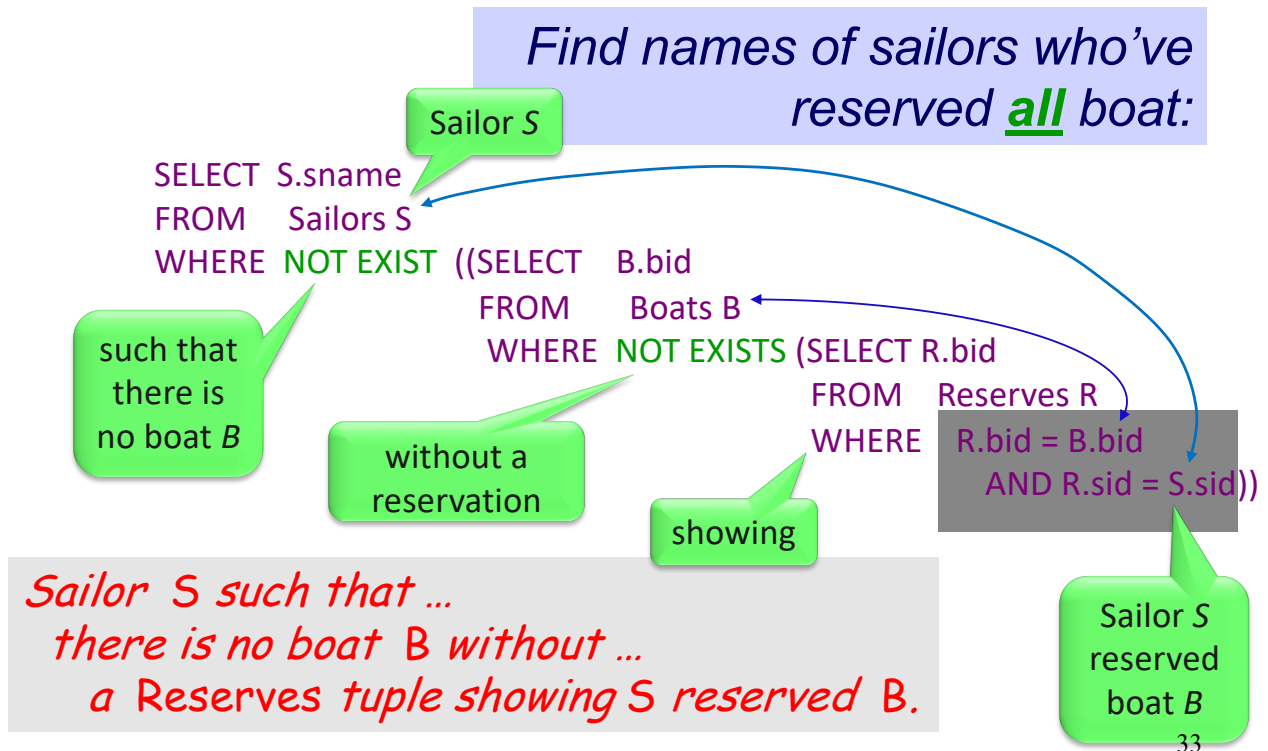
All boats

All boats reserved by the sailor

32



## Division Operations in SQL (2)



## SQL Operators

- **BETWEEN**
- **NOT BETWEEN**
- **IN**
- **UNION [DISTINCT | ALL]**
- **EXCEPT**
  - Not Supported in MySQL

```
SELECT * FROM Products
WHERE (Price BETWEEN 10 AND 20)
AND NOT CategoryID IN (1,2,3);
```

## More on Set-Comparison Operators

- **op ANY, op ALL**, where **op**: >, <, =, ≠, ≥, ≤

*Find sailors whose rating is greater than that of some sailor called Horatio*

```
SELECT *
FROM sailors S
WHERE S.rating > ANY (SELECT S2.rating
                     FROM sailors S2
                     WHERE S2.sname='Horatio')
```

*>ANY means greater than at least one value  
>ALL means greater than every value*

35

## ARGMAX?

- **The sailor with the highest rating**  
– what about ties for highest?!

```
SELECT *
FROM sailors S
WHERE S.rating >= ALL
      (SELECT S2.rating
       FROM sailors S2)
```

```
SELECT *
FROM sailors S
WHERE S.rating =
      (SELECT MAX(S2.rating)
       FROM sailors S2)
```

```
SELECT *
FROM sailors S
ORDER BY rating DESC
LIMIT 1;
```

36


# Joins

```
SELECT (column_list)
FROM table_name
  [INNER | {LEFT | RIGHT | FULL } OUTER] JOIN table_name
  ON qualification_list
WHERE ...
```

**Explicit join semantics needed unless it is an INNER join (INNER is default)**

37

## Inner Join



**Selects records that have matching values in both tables.**

```
SELECT s.sid, s.name, r.bid
FROM Sailors s INNER JOIN Reserves r
ON s.sid = r.sid
```

**Returns only those sailors who have reserved boats**

**SQL-92 also allows:**

```
SELECT s.sid, s.name, r.bid
FROM Sailors s NATURAL JOIN Reserves r
```

**"NATURAL" means equi-join for each pair of attributes with the same name (may need to rename with "AS")**

38

```
SELECT s.sid, s.name, r.bid
FROM Sailors s INNER JOIN Reserves r
ON s.sid = r.sid
```

<u>sid</u>	sname	rating	age
22	Dustin	7	45.0
31	Lubber	8	55.5
95	Bob	3	63.5

<u>sid</u>	<u>bid</u>	<u>day</u>
22	101	10/10/96
95	103	11/12/96

s.sid	s.name	r.bid
22	Dustin	101
95	Bob	103

39

## Outer Joins

*S1*

<u>sid</u>	sname	rating	age
22	dustin	7	45.0
31	lubber	8	55.5
58	rusty	10	35.0

No "sid = 31"

*R1*

<u>sid</u>	<u>bid</u>	<u>day</u>
22	101	10/10/96
58	103	11/12/96


*S1* ⋈ *R1*

sid	sname	rating	age	bid	day
22	dustin	7	45.0	101	10/10/96
31	lubber	8	55.55	null	null
58	rusty	10	35.0	103	11/12/96

No match  
in *R1*

40

# Left Outer Join



**Left Outer Join returns all matched rows, plus all unmatched rows from the table on the left of the join clause**

**(use nulls in fields of non-matching tuples)**

```
SELECT s.sid, s.name, r.bid  
FROM Sailors s LEFT OUTER JOIN Reserves r  
ON s.sid = r.sid
```

**Returns all sailors & information on whether they have reserved boats**

41

```
SELECT s.sid, s.name, r.bid  
FROM Sailors s LEFT OUTER JOIN Reserves r  
ON s.sid = r.sid
```


<u>sid</u>	sname	rating	age
22	Dustin	7	45.0
31	Lubber	8	55.5
95	Bob	3	63.5

<u>sid</u>	<u>bid</u>	<u>day</u>
22	101	10/10/96
95	103	11/12/96

s.sid	s.name	r.bid
22	Dustin	101
95	Bob	103
31	Lubber	<i>null</i>

42

# Right Outer Join



**Right Outer Join returns all matched rows, plus all unmatched rows from the table on the right of the join clause**

```
SELECT r.sid, b.bid, b.name
FROM Reserves r RIGHT OUTER JOIN Boats b
ON r.bid = b.bid
```

**Returns all boats & information on which ones are reserved.**

43

```
SELECT r.sid, b.bid, b.name
FROM Reserves r RIGHT OUTER JOIN Boats b
ON r.bid = b.bid
```

<u>sid</u>	<u>bid</u>	<u>day</u>
22	101	10/10/96
95	103	11/12/96


<u>bid</u>	bname	color
101	Interlake	blue
102	Interlake	red
103	Clipper	green
104	Marine	red

r.sid	b.bid	b.name
22	101	Interlake
<i>null</i>	102	Interlake
95	103	Clipper
<i>null</i>	104	Marine

44

# Full Outer Join

FULL OUTER JOIN



**Full Outer Join returns all (matched or unmatched) rows from the tables on both sides of the join clause**

```
SELECT r.sid, b.bid, b.name
FROM Sailors s FULL OUTER JOIN Boats b
ON s.sname = b.bname
```

45

```
SELECT s.sid, s.sname, b.bid, b.name
FROM Sailors s FULL OUTER JOIN Boats b
ON s.sname = b.bname
```

<u>sid</u>	sname	rating	age
22	Dustin	7	45.0
31	Lubber	8	55.5
95	Bob	3	63.5

<b>bid</b>	<b>bname</b>	<b>color</b>
101	Interlake	blue
105	Lubber	purple

<b>sid</b>	<b>sname</b>	<b>bid</b>	<b>bname</b>
22	Dustin	<i>null</i>	<i>null</i>
31	Lubber	105	Lubber
95	Bob	<i>null</i>	<i>null</i>
<i>null</i>	<i>null</i>	101	Interlake

46

# Aggregate Functions

*Significant extension of relational algebra*

<b>COUNT (*)</b>	The number of rows in the relation
<b>COUNT ([DISTINCT] A)</b>	The number of (unique) values in the A column
<b>SUM ([DISTINCT] A)</b>	The sum of all (unique) values in the A column
<b>AVG ([DISTINCT] A)</b>	The average of all (unique) values in the A column
<b>MAX (A)</b>	The maximum value in the A column
<b>MIN (A)</b>	The minimum value in the A column

## Aggregate Operators

```
SELECT COUNT (*)
FROM Sailors S
```

Count the number of sailors

Find the names of sailors with the highest rating

```
SELECT AVG (S.age)
FROM Sailors S
WHERE S.rating=10
```

```
SELECT S.sname
```

Find the average age

Count the number of distinct ratings of sailors called "Bob"

Find the average of the distinct ages of sailors with a rating of 10

```
SELECT COUNT (DISTINCT S.rating)
FROM Sailors S
WHERE S.sname='Bob'
```

```
SELECT AVG (DISTINCT S.age)
FROM Sailors S
WHERE S.rating=10
```



# Aggregate Operators

```
SELECT COUNT (*)  
FROM Sailors S
```

```
SELECT AVG (S.age)  
FROM Sailors S  
WHERE S.rating=10
```

```
SELECT S.sname  
FROM Sailors S  
WHERE S.rating= (SELECT MAX(S2.rating)  
FROM Sailors S2)
```

```
SELECT COUNT (DISTINCT S.rating)  
FROM Sailors S  
WHERE S.sname='Bob'
```

```
SELECT AVG (DISTINCT S.age)  
FROM Sailors S  
WHERE S.rating=10
```

50

## Summary Aggregate Functions...

Avg ()

Count ()

First ()

Last ()

Max ()

Min ()

SQL

Sum ()

Group By

Having

Ucase ()

Lcase ()

Mid ()

Len ()

Round ()

Now ()

Format ()

```
SELECT column_name,  
aggregate_function(column_name)  
FROM table_name  
WHERE column_name operator value  
GROUP BY column_name  
HAVING  
aggregate_function(column_name)  
operator value
```

# Restriction on SELECT Lists With Aggregation

- **If any aggregation is used, then each element of the SELECT list must be either:**
  1. Aggregated, or
  2. An attribute on the GROUP BY list.

51

## Illegal Query Example

- **You might think you could find the bar that sells Bud the cheapest by:**

```
SELECT bar, MIN(price)
FROM Sells
WHERE beer = 'Bud' ;
```

- **But this query is illegal in SQL.**
- **Why?**

52

# Find name and age of the oldest sailor(s)



Only aggregate operations allowed

```
SELECT S.sname, MAX (S.age)
FROM Sailors S
```

## Find name and age of the oldest sailor(s)

- **The first query is incorrect!**
- **Third query equivalent to second query**
  - allowed in SQL/92 standard, but not supported in some systems.


```
SELECT S.sname, MAX (S.age)
FROM Sailors S
```

```
SELECT S.sname, S.age
FROM Sailors S
WHERE S.age =
      (SELECT MAX (S2.age)
       FROM Sailors S2)
```


```
SELECT S.sname, S.age
FROM Sailors S
WHERE (SELECT MAX (S2.age)
       FROM Sailors S2)
      = S.age
```

# GROUP BY and HAVING (1)

- So far, we've applied aggregate operators to all (qualifying) tuples.



- Sometimes, we want to apply them to each of several *groups* of tuples.



55

# GROUP BY and HAVING (2)

**Consider: Find the age of the youngest sailor for each rating level.** /\* *Min(age) for multiple groups*

- If we know that rating values go from 1 to 10, we can write 10 queries that look like this:

For  $i = 1, 2, \dots, 10$ :  $\left\{ \begin{array}{l} \text{SELECT MIN (S.age)} \\ \text{FROM Sailors S} \\ \text{WHERE S.rating} = i \end{array} \right.$

- In general, we don't know how many rating levels exist, and what the rating values for these levels are !

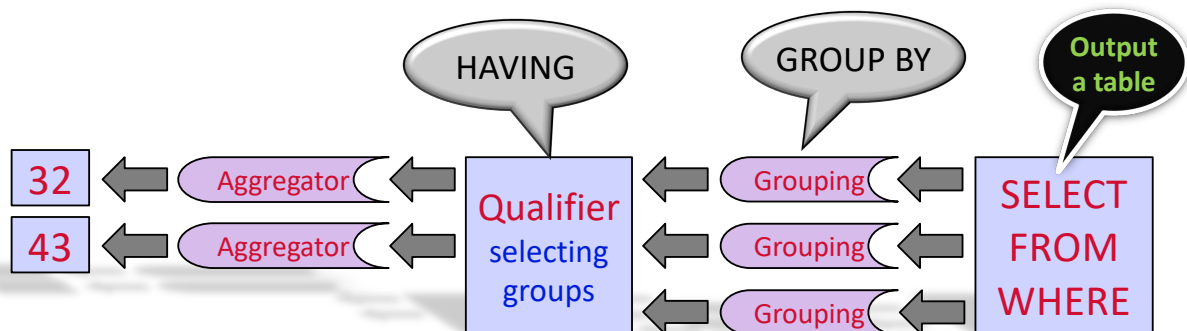
56

# Queries With GROUP BY and HAVING

```

SELECT      [DISTINCT] target-list
FROM        relation-list
WHERE       qualification
GROUP BY    grouping-list
HAVING      group-qualification
    
```

MIN(Attribute)



57

## Queries With GROUP BY

- To generate values for a column based on groups of rows, use **aggregate** functions in **SELECT** statements with the **GROUP BY** clause

```

SELECT      [DISTINCT] target-list
FROM        relation-list
[WHERE      qualification]
GROUP BY    grouping-list
    
```

The *target-list* contains

- (i) list of column names &
  - (ii) terms with aggregate operations (e.g., MIN (*S.age*)).
- column name list (i) can contain only attributes from the *grouping-list*.

58

## Group By Examples

For each rating, find the average age of the sailors

```
SELECT S.rating, AVG (S.age)
FROM Sailors S
GROUP BY S.rating
```

For each rating find the age of the youngest sailor with age  $\geq 18$

```
SELECT S.rating, MIN (S.age)
FROM Sailors S
WHERE S.age >= 18
GROUP BY S.rating
```

59

## Conceptual Evaluation

```
SELECT      [DISTINCT] target-list
FROM        relation-list
WHERE       qualification
GROUP BY   grouping-list
HAVING     group-qualification
```

1. The cross-product of *relation-list* is computed
2. Tuples that fail *qualification* are discarded
3. 'Unnecessary' fields are deleted
4. The remaining tuples are partitioned into groups by the value of attributes in *grouping-list*.
5. The *group-qualification* is then applied to eliminate some groups
6. One answer tuple is generated per qualifying group

60

**Find the age of the youngest sailor with age  $\geq 18$ , for each rating with at least 2 such sailors**

```

SELECT S.rating, MIN (S.age)
FROM Sailors S
WHERE S.age >= 18
GROUP BY S.rating
HAVING COUNT (*) > 1

```

Input relation

sid	sname	rating	age
22	dustin	7	45.0
31	lubber	8	55.5
71	zorba	10	16.0
64	horatio	7	35.0
29	brutus	1	33.0
58	rusty	10	35.0

rating	age
7	35.0
7	45.0
7	35.0
8	55.5
10	35.0

Answer

Only one group satisfies HAVING

Only S.rating and S.age are mentioned in SELECT

## "GROUP BY and HAVING" Examples

Find the age of the youngest sailor with age  $\geq 18$

```

SELECT MIN (S.age)
FROM Sailors S
WHERE S.age >= 18

```

```

SELECT S.rating, MIN (S.age)
FROM Sailors S
WHERE S.age >= 18
GROUP BY S.rating

```

Find the age of the youngest sailor with age  $\geq 18$ , **for each rating**

Find the age of the youngest sailor with age  $\geq 18$ , **for each rating with at least 2 such sailors**

```

SELECT S.rating, MIN (S.age)
FROM Sailors S
WHERE S.age >= 18
GROUP BY S.rating
HAVING COUNT (*) > 1

```

For each red boat, find the number of reservations for this boat

```
SELECT B.bid, COUNT (*) AS scout
FROM Boats B, Reserves R
WHERE R.bid=B.bid AND B.color='red'
GROUP BY B.bid
```

3) Count the number of reservations for each red-boat group

1) Find all reservations for red boats

2) Group the reservations for red boats according to bid

63

For each red boat, find the number of reservations for this boat

Grouping over a join of two relations

*B.color* is not in the grouping-list

```
SELECT B.bid, COUNT (*) AS scout
FROM Boats B, Reserves R
WHERE R.bid=B.bid
GROUP BY B.bid
HAVING B.color='red'
```

**Illegal!**

Note: HAVING clause is to select groups !

64



**Find the age of the youngest sailor older than 18, for each rating level that has at least 2 sailors**

```
SELECT  S.rating, MIN (S.age)
FROM    Sailors S
WHERE   S.age > 18
GROUP BY S.rating
HAVING  1 < (SELECT COUNT (*)
```

Replacing this by  
"HAVING COUNT(\*) > 1"

```
FROM    Sailors S2
WHERE   S.rating = S2.rating)
```

- Shows HAVING clause can also contain a subquery.
- We can use S.rating inside the nested subquery because it has a single value for the current group of sailors.
- What if HAVING clause is replaced by "HAVING COUNT(\*) >1"
  - Find the age of the youngest sailor older than 18, for each rating level that has at least two such sailors. /\* see next page

65

**Find the age of the youngest sailor older than 18, for each rating level that has at least 2 sailors**

```
SELECT  S.rating, MIN (S.age)
FROM    Sailors S
WHERE   S.age > 18
GROUP BY S.rating
```

```
HAVING  1 < (SELECT COUNT (*)
FROM    Sailors S2
WHERE   S.rating = S2.rating)
```

At least 2 sailors

Counting including sailors younger than 18

"age" is not mentioned in this subquery

```
SELECT  S.rating, MIN (S.age)
FROM    Sailors S
WHERE   S.age > 18
GROUP BY S.rating
HAVING  COUNT (*) > 1
```

Counting include only adult sailors

At least 2 such sailors, i.e., older than 18

66

**Find those ratings for which the average age is the minimum over all ratings**

```
SELECT S.rating
FROM Sailors S
WHERE S.age = (SELECT MIN (AVG (S2.age))
              FROM Sailors S2)
```

*Aggregate operations cannot be nested*

67

**Find those ratings for which the average age is the minimum over all ratings**

Correct solution (in SQL/92):

```
SELECT Temp.rating, Temp.avgage
FROM (SELECT S.rating, AVG (S.age) AS avgage
      FROM Sailors S
      GROUP BY S.rating) AS Temp
WHERE Temp.avgage = (SELECT MIN (Temp.avgage)
                    FROM Temp)
```

*Output this rating group and its average age*

*Find average age for each rating group. This table has two columns*

*Average age for some rating group*

*Minimum over all ratings*

68

# Null Values

- Field values in a tuple are sometimes
  - *unknown* (e.g., a rating has not been assigned),  
or
  - *inapplicable* (e.g., no spouse's name).
- SQL provides a special value *null* for such situations.

69

# Null Values

The presence of *null* complicates many issues:

- Special operators needed, e.g., **IS NULL** to test if a value is *null*.
- Is *rating* > 8 true or false when *rating* is equal to *null*? *null*
- What about **AND**, **OR** and **NOT**? Need a 3-valued logic (*true*, *false*, and *unknown*), e.g., (*unknown* OR *false*) = *unknown*.
- Meaning of constructs must be defined carefully, e.g., **WHERE** clause eliminates rows that don't evaluate to true.
  - Null + 5 = null; but SUM (null, 5) = 5. (nulls can cause some unexpected behavior)
- New operators (in particular, *outer joins*) possible/needed.

70

# Views

```
CREATE VIEW view_name  
AS select_statement
```

Makes development simpler  
Often used for security  
Not instantiated - makes updates tricky

```
CREATE VIEW Reds  
AS SELECT B.bid, COUNT (*) AS scout  
FROM Boats B, Reserves R  
WHERE R.bid=B.bid AND B.color='red'  
GROUP BY B.bid
```

71

```
CREATE VIEW Reds  
AS SELECT B.bid, COUNT (*) AS scout  
FROM Boats B, Reserves R  
WHERE R.bid=B.bid AND B.color='red'  
GROUP BY B.bid
```

b.bid	scout
102	1

Reds

72

## Sailors who have reserved all boats

Sailors

sid	sname	rating	age
1	Frodo	7	22
2	Bilbo	2	39
3	Sam	8	27

Boats

bid	bname	color
101	Nina	red
102	Pinta	blue
103	Santa Maria	red

Reserves

sid	bid	day
1	102	9/12
2	102	9/12
2	101	9/14
1	102	9/10
2	103	9/13

```

SELECT S.name
FROM Sailors S, reserves R
WHERE S.sid = R.sid
GROUP BY S.name, S.sid
HAVING COUNT(DISTINCT R.bid) =
    (Select COUNT (*) FROM Boats)
    
```

sname	sid	bid
Frodo	1	102
Bilbo	2	101
Bilbo	2	102
Frodo	1	102
Bilbo	2	103

sname	sid	count
Frodo	1	1
Bilbo	2	3

count
3

sname	sid	bid
Frodo	1	102,102
Bilbo	2	101, 102, 103

## Two more important topics

- Constraints
- SQL embedded in other languages

# Integrity Constraints

- **IC conditions that every legal instance of a relation must satisfy.**
  - Inserts/deletes/updates that violate ICs are disallowed.
  - Can ensure application semantics (e.g., sid is a key),
  - ...or prevent inconsistencies (e.g., sname has to be a string, age must be < 200)
- **Types of IC's: Domain constraints, primary key constraints, foreign key constraints, general constraints.**
  - Domain constraints: Field values must be of right type. Always enforced.
  - Primary key and foreign key constraints: coming right up.

75

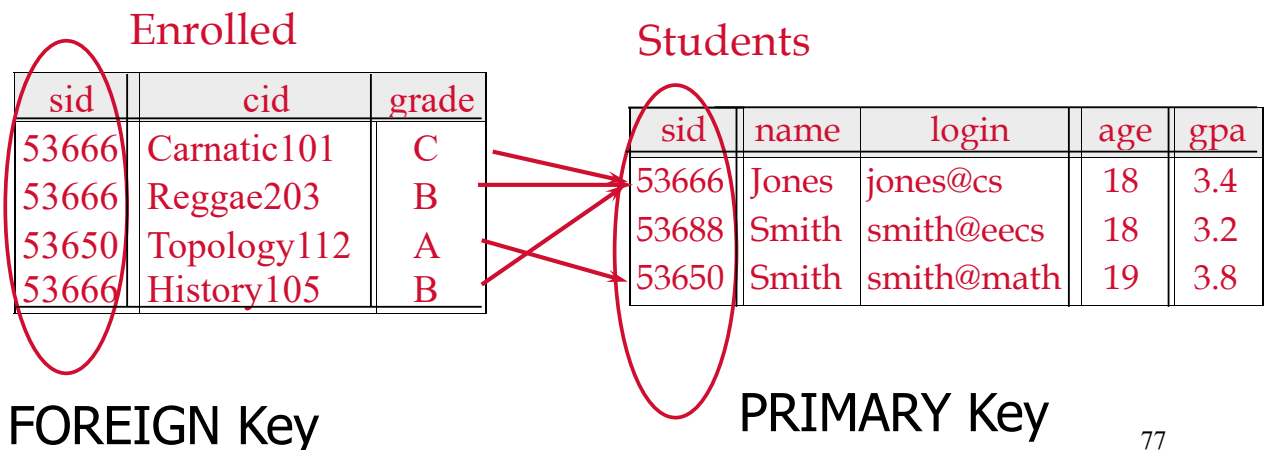
## Where do ICs Come From?

- **Semantics of the real world!**
- **Note:**
  - We can check IC violation in a DB instance
  - We can NEVER infer that an IC is true by looking at an instance.
    - An IC is a statement about all possible instances!
  - From example, we know name is not a key, but the assertion that sid is a key is given to us.
- **Key and foreign key ICs are the most common**
- **More general ICs supported too.**

76

# Keys

- **Keys are a way to associate tuples in different relations**
- **Keys are one form of IC**



77

## Primary Keys

- **A set of fields is a **superkey** if:**
  - No two distinct tuples can have same values in all key fields
- **A set of fields is a **key** for a relation if :**
  - It is a superkey
  - No subset of the fields is a superkey
- **what if >1 key for a relation?**
  - One of the keys is chosen (by DBA) to be the **primary key**. Other keys are called **candidate keys**.
- **E.g.**
  - sid is a key for Students.
  - What about name?
  - The set {sid, gpa} is a superkey.

78

## Primary and Candidate Keys

- Possibly many ***candidate keys*** (specified using **UNIQUE**), one of which is chosen as the ***primary key***.

- Keys must be used carefully!

```
CREATE TABLE Enrolled1 (sid CHAR(20),  
cid CHAR(20),  
grade CHAR(2),  
PRIMARY KEY (sid,cid)) vs. CREATE TABLE Enrolled2  
(sid CHAR(20),  
cid CHAR(20),  
grade CHAR(2),  
PRIMARY KEY (sid),  
UNIQUE (cid, grade))
```

*“For a given student and course, there is a single grade.”*

79

## Primary and Candidate Keys

```
CREATE TABLE Enrolled1 (sid CHAR(20),  
cid CHAR(20),  
grade CHAR(2),  
PRIMARY KEY (sid,cid)) vs. CREATE TABLE Enrolled2  
(sid CHAR(20),  
cid CHAR(20),  
grade CHAR(2),  
PRIMARY KEY (sid),  
UNIQUE (cid, grade))
```

```
INSERT INTO enrolled1 VALUES ('1234', 'cs186', 'A+');  
INSERT INTO enrolled1 VALUES ('1234', 'cs186', 'F');  
INSERT INTO enrolled1 VALUES ('1234', 'cs61C', 'A+');
```

```
INSERT INTO enrolled2 VALUES ('1234', 'cs186', 'A+');  
INSERT INTO enrolled2 VALUES ('1234', 'cs186', 'F');  
INSERT INTO enrolled2 VALUES ('1234', 'cs61C', 'A+');  
INSERT INTO enrolled2 VALUES ('4567', 'cs186', 'A+');
```

*“For a given student and course, there is a single grade.”*

80



# Primary and Candidate Keys

```
CREATE TABLE Enrolled1
(sid CHAR(20),
 cid CHAR(20),
 grade CHAR(2),
 PRIMARY KEY (sid,cid));
```

VS.

```
CREATE TABLE Enrolled2
(sid CHAR(20),
 cid CHAR(20),
 grade CHAR(2),
 PRIMARY KEY (sid),
 UNIQUE (cid, grade));
```

```
INSERT INTO enrolled1 VALUES ('1234', 'cs186', 'A+');
INSERT INTO enrolled1 VALUES ('1234', 'cs186', 'F');
INSERT INTO enrolled1 VALUES ('1234', 'cs61C', 'A+');
```

```
INSERT INTO enrolled2 VALUES ('1234', 'cs186', 'A+');
INSERT INTO enrolled2 VALUES ('1234', 'cs186', 'F');
INSERT INTO enrolled2 VALUES ('1234', 'cs61C', 'A+');
INSERT INTO enrolled2 VALUES ('4567', 'cs186', 'A+');
```

*“Students can take only one course, and no two students in a course receive the same grade.”*

81

## Foreign Keys, Referential Integrity

- **Foreign key**: a “logical pointer”
  - Set of fields in a tuple in one relation that ‘refer’ to a tuple in another relation.
  - Reference to *primary key* of the other relation.
- **All foreign key constraints enforced?**
  - referential integrity!
  - i.e., no dangling references.

82

# Foreign Keys in SQL

- **E.g. Only students listed in the Students relation should be allowed to enroll for courses.**

– *sid* is a foreign key referring to **Students**:

```
CREATE TABLE Enrolled
(sid CHAR(20),cid CHAR(20),grade CHAR(2),
PRIMARY KEY (sid,cid),
FOREIGN KEY (sid) REFERENCES Students);
```

Enrolled

sid	cid	grade
53666	Carnatic101	C
53666	Reggae203	B
53650	Topology112	A
53666	History105	B
<del>11111</del>	<del>English102</del>	<del>A</del>

Students

sid	name	login	age	gpa
53666	Jones	jones@cs	18	3.4
53688	Smith	smith@eecs	18	3.2
53650	Smith	smith@math	19	3.8

83

## Enforcing Referential Integrity

- ***sid* in Enrolled: foreign key referencing Students.**
- **Scenarios:**
  - Insert Enrolled tuple with non-existent student id?
  - Delete a Students tuple?
    - Also delete Enrolled tuples that refer to it? (CASCADE)
    - Disallow if referred to? (NO ACTION)
    - Set sid in referring Enrolled tups to a *default* value? (SET DEFAULT)
    - Set sid in referring Enrolled tuples to *null*, denoting '*unknown*' or '*inapplicable*'. (SET NULL)
- **Similar issues arise if primary key of Students tuple is updated.**

84

85

## General Constraints

- Useful when more general ICs than keys are involved.
- Can use queries to express constraint.
- Checked on insert or update.
- Constraints can be named.

```
CREATE TABLE Sailors
  ( sid INTEGER,
    sname CHAR(10),
    rating INTEGER,
    age REAL,
    PRIMARY KEY (sid),
    CHECK ( rating >= 1
           AND rating <= 10 ))
```

```
CREATE TABLE Reserves
  ( sname CHAR(10),
    bid INTEGER,
    day DATE,
    PRIMARY KEY (bid,day),
    CONSTRAINT noInterlakeRes
    CHECK ('Interlake' <>
           ( SELECT B.bname
             FROM Boats B
             WHERE B.bid=bid)))
```

86

## Constraints Over Multiple Relations

```
CREATE TABLE Sailors
  ( sid INTEGER,
    sname CHAR(10),
    rating INTEGER,
    age REAL,
    PRIMARY KEY (sid),
    CHECK
  ( (SELECT COUNT (S.sid) FROM Sailors S)
    + (SELECT COUNT (B.bid) FROM
      Boats B) < 100 )
```

*Number of boats  
plus number of  
sailors is < 100*

# Constraints Over Multiple Relations

- Awkward and wrong!
  - Only checks sailors!

```
CREATE TABLE Sailors
( sid INTEGER,
  sname CHAR(10),
  rating INTEGER,
  age REAL,
  PRIMARY KEY (sid),
  CHECK
```

*Number of boats  
plus number of  
sailors is < 100*

- ASSERTION is the right solution; not associated with either table.
  - Unfortunately, not supported in many DBMS.
  - *Triggers* are another solution.

```
( (SELECT COUNT (S.sid) FROM Sailors S)
+ (SELECT COUNT (B.bid) FROM
  Boats B) < 100 )
```

```
CREATE ASSERTION smallClub
CHECK
( (SELECT COUNT (S.sid) FROM Sailors S)
+ (SELECT COUNT (B.bid)
  FROM Boats B) < 100 )
```

## Two more important topics

### • ~~Constraints~~

- **SQL embedded in other languages**

# Writing Applications with SQL

- **SQL is not a general purpose programming language.**
  - + Tailored for data retrieval and manipulation
  - + Relatively easy to optimize and parallelize
  - Can't write entire apps in SQL alone
- **Options:**
  - Make the query language “Turing complete”
    - Avoids the “impedance mismatch”
    - makes “simple” relational language complex
  - Allow SQL to be embedded in regular programming languages.
  - Q: What needs to be solved to make the latter approach work?

89

## Cursors

- **Can declare a cursor on a relation or query**
- **Can *open* a cursor**
- **Can repeatedly *fetch* a tuple (moving the cursor)**
- **Special return value when all tuples have been retrieved.**
- **ORDER BY allows control over the order tuples are returned.**
  - Fields in ORDER BY clause must also appear in SELECT clause.
- **LIMIT controls the number of rows returned (good fit w/ORDER BY)**
- **Can also modify/delete tuple pointed to by a cursor**
  - A “non-relational” way to get a handle to a particular tuple

90

# Database APIs

- **A library with database calls (API)**
  - special objects/methods
  - passes SQL strings from language, presents **result sets** in a language-friendly way
  - *ODBC* a C/C++ standard started on Windows
  - *JDBC* a Java equivalent
  - Most scripting languages have similar things
    - E.g. in Ruby there's the "pg" gem for Postgres
- **ODBC/JCDB try to be DBMS-neutral**
  - at least try to hide distinctions across different DBMSs

91

## Summary

- Relational model has **well-defined query semantics**
- SQL provides functionality close to basic relational model
  - (some differences in duplicate handling, null values, set operators, ...)*
- Typically, many ways to write a query
  - **DBMS figures out a fast way to execute a query, regardless of how it is written.**

92

# ADVANCED EXAMPLES

93

## Getting Serious

- **Two “fancy” queries for different applications**
  - Clustering Coefficient for Social Network graphs
  - Medians for “robust” estimates of the central value

94

# Serious SQL: Social Nets Example

```
-- An undirected friend graph. Store each link once
CREATE TABLE Friends(
    fromID integer,
    toID integer,
    since date,
    PRIMARY KEY (fromID, toID),
    FOREIGN KEY (fromID) REFERENCES Users,
    FOREIGN KEY (toID) REFERENCES Users,
    CHECK (fromID < toID));

-- Return both directions
CREATE VIEW BothFriends AS
    SELECT * FROM Friends
    UNION ALL
    SELECT F.toID AS fromID, F.fromID AS toID, F.since
    FROM Friends F;
```

95

## 6 degrees of friends

```
SELECT F1.fromID, F5.toID
    FROM BothFriends F1, BothFriends F2, BothFriends F3,
         BothFriends F4, BothFriends F5
    WHERE F1.toID = F2.fromID
        AND F2.toID = F3.fromID
        AND F3.toID = F4.fromID
        AND F4.toID = F5.fromID;
```



# Clustering Coefficient of a Node


$$C_i = 2|\{e_{jk}\}| / k_i(k_i-1)$$

- where:
  - $k_i$  is the number of neighbors of node  $i$
  - $e_{jk}$  is an edge between nodes  $j$  and  $k$  neighbors of  $i$ , ( $j < k$ ). (A triangle!)
- I.e. Cliquishness: the fraction of your friends that are friends with each other!
- Clustering Coefficient of a graph is the average CC of all nodes.

97

## In SQL

$$C_i = 2|\{e_{jk}\}| / k_i(k_i-1)$$




```
CREATE VIEW NEIGHBOR_CNT AS
SELECT fromID AS nodeID, count(*) AS friend_cnt
  FROM BothFriends
  GROUP BY nodeID;

CREATE VIEW TRIANGLES AS
SELECT F1.toID as root, F1.fromID AS friend1,
       F2.fromID AS friend2
  FROM BothFriends F1, BothFriends F2, Friends F3
 WHERE F1.toID = F2.toID      /* Both point to root */
    AND F1.fromID = F3.fromID /* Same origin as F1 */
    AND F3.toID = F2.fromID   /* points to origin of F2 */
 ;
```

98

# In SQL



$$C_i = 2|\{e_{jk}\}| / k_i(k_i-1)$$

```
CREATE VIEW NEIGHBOR_EDGE_CNT AS
SELECT root, COUNT(*) as cnt FROM TRIANGLES
GROUP BY root;
```

```
CREATE VIEW CC_PER_NODE AS
SELECT NE.root, 2.0*NE.cnt /
      (N.friend_cnt*(N.friend_cnt-1)) AS CC
FROM NEIGHBOR_EDGE_CNT NE, NEIGHBOR_CNT N
WHERE NE.root = N.nodeID;
```

```
SELECT AVG(cc) FROM CC_PER_NODE;
```

99

# Median

- **Given  $n$  values in sorted order, the one at position  $n/2$** 
  - Assumes an odd # of items
  - For an even #, can take the lower of the middle 2
- **A much more “robust” statistic than average**
  - Q: Suppose you want the mean to be 1,000,000. What fraction of values do you have to corrupt?
  - Q2: Suppose you want the median to be 1,000,000. Same question.
  - This is called the *breakdown point* of a statistic.
  - Important for dealing with data *outliers*
    - E.g. dirty data
    - Even with real data: “overfitting”

100

# Median in SQL

```
SELECT c AS median FROM T
WHERE
  (SELECT COUNT(*) from T AS T1
   WHERE T1.c < T.c)
=
  (SELECT COUNT(*) from T AS T2
   WHERE T2.c > T.c);
```

101

# Median in SQL

```
SELECT c AS median FROM T
WHERE
  (SELECT COUNT(*) from T AS T1
   WHERE T1.c < T.c)
=
  (SELECT COUNT(*) from T AS T2
   WHERE T2.c > T.c);
```

102

# Faster Median in SQL

```
SELECT x.c as median
  FROM T x, T y
  GROUP BY x.c
  HAVING
    SUM(CASE WHEN y.c <= x.c THEN 1 ELSE 0 END)
    >= (COUNT(*)+1)/2
  AND
    SUM(CASE WHEN y.c >= x.c THEN 1 ELSE 0 END)
    >= (COUNT(*)/2)+1
```

Why faster?

Note: handles even # of items!

103

## EXAMPLES

# Essential SQL Statements

<code>CREATE DATABASE database_name</code>	<code>DROP DATABASE database_name</code>
<code>CREATE TABLE table_name</code> ( column_name1 data_type, column_name2 data_type, column_name3 data_type, ... )	<code>DROP TABLE table_name</code> <code>DELETE FROM table_name</code> <code>DELETE * FROM table_name</code>
<code>CREATE INDEX index_name</code> ON table_name (column_name)	<code>DROP INDEX index_name</code>
<code>CREATE UNIQUE INDEX index_name</code> ON table_name (column_name)	
<code>INSERT INTO table_name</code> VALUES (value1, value2,....) <code>INSERT INTO table_name</code> (column1, column2,...) VALUES (value1, value2,....)	<code>DELETE FROM table_name</code> <code>WHERE some_column=some_value</code>

## More SQL Statements

<code>ALTER TABLE table_name</code> ADD column_name datatype	<code>ALTER TABLE table_name</code> <code>DROP COLUMN column_name</code>
<code>SELECT column_name AS column_alias</code> FROM table_name	
<code>SELECT column_name</code> FROM table_name AS table_alias	
<code>SELECT column_name(s)</code> FROM table_name WHERE column_name BETWEEN value1 AND value2	
<code>SELECT column_name(s)</code> FROM table_name WHERE condition AND OR condition	

# Other SQL Statements

- **AUTO INCREMENT Field**
- **SELECT INTO**
  - Selects data from one table and inserts it into a new table
- **LIMIT**
  - Specify the number of records to return
- **CREATE VIEW**
  - Create a virtual table based on the result-set of an SQL statement
- **TRUNCATE TABLE**
  - Delete all table contents

```
CREATE VIEW view_name AS      SELECT column_name(s)      SELECT column_name(s)
SELECT column_name(s)        INTO newtable [IN        FROM table_name
FROM table_name              externaldb]              LIMIT number;
WHERE condition              FROM table1;
```

107

## Example: Relation (Table)

Row/Tuple/Record

Column/Attribute/Field


	name	birth	gpa	grad
	Anderson	1987-10-22	3.9	2009
	Jones	1990-4-16	2.4	2012
	Hernandez	1989-8-12	3.1	2011
	Chen	1990-2-4	3.2	2011

Column Types → VARCHAR(30)    DATE    FLOAT    INT

108

## Example: Primary Key

Unique For Each Row



id	name	birth	gpa	grad
14	Anderson	1987-10-22	3.9	2009
38	Jones	1990-4-16	2.4	2012
77	Hernandez	1989-8-12	3.1	2011
104	Chen	1990-2-4	3.2	2011

INT    VARCHAR(30)    DATE    FLOAT    INT

109

## Basic Table Operations

```
CREATE TABLE students (  
    id INT AUTO_INCREMENT,  
    name VARCHAR(30),  
    birth DATE,  
    gpa FLOAT,  
    grad INT,  
    PRIMARY KEY(id));  
  
INSERT INTO students(name, birth, gpa, grad)  
    VALUES ('Anderson', '1987-10-22', 3.9, 2009);  
INSERT INTO students(name, birth, gpa, grad)  
    VALUES ('Jones', '1990-4-16', 2.4, 2012);  
  
DELETE FROM students WHERE name='Anderson';  
  
DROP TABLE students;
```

110

## Query: Display Entire Table

id	name	birth	gpa	grad
1	Anderson	1987-10-22	3.9	2009
2	Jones	1990-4-16	2.4	2012
3	Hernandez	1989-8-12	3.1	2011
4	Chen	1990-2-4	3.2	2011

```
SELECT * FROM students;
```

```
+-----+-----+-----+-----+-----+
| id | name       | birth       | gpa  | grad  |
+-----+-----+-----+-----+-----+
|  1 | Anderson  | 1987-10-22 | 3.9  | 2009 |
|  2 | Jones     | 1990-04-16 | 2.4  | 2012 |
|  3 | Hernandez | 1989-08-12 | 3.1  | 2011 |
|  4 | Chen      | 1990-02-04 | 3.2  | 2011 |
+-----+-----+-----+-----+-----+
```

111

## Query: Select Columns name and gpa

id	name	birth	gpa	grad
1	Anderson	1987-10-22	3.9	2009
2	Jones	1990-4-16	2.4	2012
3	Hernandez	1989-8-12	3.1	2011
4	Chen	1990-2-4	3.2	2011

```
SELECT name, gpa
FROM students;
```

```
+-----+-----+
| name       | gpa  |
+-----+-----+
| Anderson  | 3.9  |
| Jones     | 2.4  |
| Hernandez | 3.1  |
| Chen      | 3.2  |
+-----+-----+
```

112



## Query: Filter Rows

id	name	birth	gpa	grad
1	Anderson	1987-10-22	3.9	2009
2	Jones	1990-4-16	2.4	2012
3	Hernandez	1989-8-12	3.1	2011
4	Chen	1990-2-4	3.2	2011

```
SELECT name, gpa FROM students
WHERE gpa > 3.0;
```

```
+-----+-----+
| name      | gpa  |
+-----+-----+
| Anderson  | 3.9  |
| Hernandez | 3.1  |
| Chen      | 3.2  |
+-----+-----+
```

113

## Query: Sort Output

- The **ORDER BY** keyword is used to sort the result-set by one or more columns

```
SELECT column_name, column_name
FROM table_name
ORDER BY column_name ASC | DESC, column_name ASC | DESC;
```

114

## Query: Sort Output by gpa

id	name	birth	gpa	grad
1	Anderson	1987-10-22	3.9	2009
2	Jones	1990-4-16	2.4	2012
3	Hernandez	1989-8-12	3.1	2011
4	Chen	1990-2-4	3.2	2011

```
SELECT gpa, name, grad FROM students
WHERE gpa > 3.0
ORDER BY gpa DESC;
```

```
+-----+-----+-----+
| gpa  | name      | grad  |
+-----+-----+-----+
| 3.9  | Anderson  | 2009  |
| 3.2  | Chen      | 2011  |
| 3.1  | Hernandez | 2011  |
+-----+-----+-----+
```

115

## Update Value(s)

- The **UPDATE** statement is used to update existing records in a table.

```
UPDATE table_name
SET column1=value1,column2=value2,...
WHERE some_column=some_value;
```

116

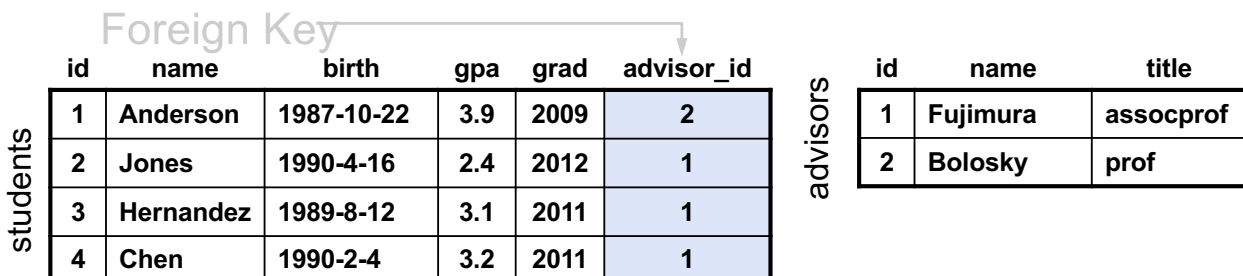
# Update Value(s)

id	name	birth	gpa	grad
1	Anderson	1987-10-22	3.9	2009
2	Jones	1990-4-16	2.4	2012
3	Hernandez	1989-8-12	3.1	2011
4	Chen	1990-2-4	3.2	2011

```
UPDATE students
SET gpa = 2.6, grad = 2013
WHERE id = 2
```

id	name	birth	gpa	grad
1	Anderson	1987-10-22	3.9	2009
2	Jones	1990-4-16	2.6	2013
3	Hernandez	1989-8-12	3.1	2011
4	Chen	1990-2-4	3.2	2011

117



```
SELECT s.name, s.gpa
FROM students s, advisors p
WHERE s.advisor_id = p.id AND p.name = 'Fujimura';
```

s.id	s.name	s.birth	s.gpa	s.grad	s.advisor_id	p.id	p.name	p.title
1	Anderson	1987-10-22	3.9	2009	2	1	Fujimura	assocprof
1	Anderson	1987-10-22	3.9	2009	2	2	Bolosky	prof
2	Jones	1990-4-16	2.4	2012	1	1	Fujimura	assocprof
2	Jones	1990-4-16	2.4	2012	1	2	Bolosky	prof
3	Hernandez	1989-8-12	3.1	2011	1	1	Fujimura	assocprof
3	Hernandez	1989-8-12	3.1	2011	1	2	Bolosky	prof
4	Chen	1990-2-4	3.2	2011	1	1	Fujimura	assocprof
4	Chen	1990-2-4	3.2	2011	1	2	Bolosky	prof

Slide 118

id	name	birth	gpa	grad	advisor_id
1	Anderson	1987-10-22	3.9	2009	2
2	Jones	1990-4-16	2.4	2012	1
3	Hernandez	1989-8-12	3.1	2011	1
4	Chen	1990-2-4	3.2	2011	1

students

id	name	title
1	Fujimura	assocprof
2	Bolosky	prof

advisors

```
SELECT s.name, s.gpa
FROM students s, advisors p
WHERE s.advisor_id = p.id AND p.name = 'Fujimura';
```

s.id	s.name	s.birth	s.gpa	s.grad	s.advisor_id	p.id	p.name	p.title
1	Anderson	1987-10-22	3.9	2009	2	1	Fujimura	assocprof
1	Anderson	1987-10-22	3.9	2009	2	2	Bolosky	prof
2	Jones	1990-4-16	2.4	2012	1	1	Fujimura	assocprof
2	Jones	1990-4-16	2.4	2012	1	2	Bolosky	prof
3	Hernandez	1989-8-12	3.1	2011	1	1	Fujimura	assocprof
3	Hernandez	1989-8-12	3.1	2011	1	2	Bolosky	prof
4	Chen	1990-2-4	3.2	2011	1	1	Fujimura	assocprof
4	Chen	1990-2-4	3.2	2011	1	2	Bolosky	prof

Slide 119

id	name	birth	gpa	grad	advisor_id
1	Anderson	1987-10-22	3.9	2009	2
2	Jones	1990-4-16	2.4	2012	1
3	Hernandez	1989-8-12	3.1	2011	1
4	Chen	1990-2-4	3.2	2011	1

students

id	name	title
1	Fujimura	assocprof
2	Bolosky	prof

advisors

```
SELECT s.name, s.gpa
FROM students s, advisors p
WHERE s.advisor_id = p.id AND p.name = 'Fujimura';
```

```
+-----+-----+
| name      | gpa  |
+-----+-----+
| Jones     | 2.4  |
| Hernandez | 3.1  |
| Chen      | 3.2  |
+-----+-----+
```

	id	name	birth	gpa	grad
students	1	Anderson	1987-10-22	3.9	2009
	2	Jones	1990-4-16	2.4	2012
	3	Hernandez	1989-8-12	3.1	2011
	4	Chen	1990-2-4	3.2	2011

	id	number	name	quarter
courses	1	CS142	Web stuff	Winter 2009
	2	ART101	Finger painting	Fall 2008
	3	ART101	Finger painting	Winter 2009
	4	PE204	Mud wrestling	Winter 2009

	course_id	student_id
courses_students	1	1
	3	1
	4	1
	1	2
	2	2
	1	3
	2	4
	4	4

```
SELECT s.name, c.quarter
FROM students s, courses c, courses_students cs
WHERE c.id = cs.course_id AND s.id = cs.student_id
AND c.number = 'ART101';
```

```
+-----+-----+
| name      | quarter    |
+-----+-----+
| Jones     | Fall 2008  |
| Chen      | Fall 2008  |
| Anderson  | Winter 2009|
+-----+-----+
```

Slide 121

## Back to Our Running Example ...

```
CREATE TABLE Sailors (
  sid INTEGER,
  sname CHAR(20),
  rating INTEGER,
  age REAL,
  PRIMARY KEY sid);
```

```
CREATE TABLE Boats (
  bid INTEGER,
  bname CHAR(20),
  color CHAR(10)
  PRIMARY KEY bid);
```

```
CREATE TABLE Reserves (
  sid INTEGER,
  bid INTEGER,
  day DATE,
  PRIMARY KEY (sid, bid, day),
  FOREIGN KEY sid REFERENCES Sailors,
  FOREIGN KEY bid REFERENCES Boats);
```

sid	sname	rating	age
1	Fred	7	22
2	Jim	2	39
3	Nancy	8	27

bid	bname	color
101	Nina	red
102	Pinta	blue
103	Santa Maria	red

sid	bid	day
1	102	9/12
2	102	9/13

# Back to Our Running Example ...

*Reserves*

<u>sid</u>	<u>bid</u>	<u>day</u>
22	101	10/10/96
95	103	11/12/96

*Sailors*

<u>sid</u>	sname	rating	age
22	Dustin	7	45.0
31	Lubber	8	55.5
95	Bob	3	63.5

*Boats*

<u>bid</u>	bname	color	
101	Interlake	blue	
102	Interlake	red	
103	Clipper	green	
104	Marine	red	

123

## APPENDIX

## DDL – Create Table

- **CREATE TABLE *table\_name***  
( { ***column\_name data\_type***  
[ **DEFAULT *default\_expr*** ] [ ***column\_constraint*** [, ... ] ] | ***table\_constraint*** } [, ... ] )
- **Data Types (PostgreSQL) include:**
  - character(n) – fixed-length character string
  - character varying(n) – variable-length character string
  - smallint, integer, bigint, numeric, real, double precision
  - date, time, timestamp, ...
  - serial - unique ID for indexing and cross reference
  - ...
- **PostgreSQL also allows OIDs and other “system types”, arrays, inheritance, rules...**
  - conformance to the SQL-1999 standard is variable.

125

## Constraints

- **Recall that the schema defines the legal instances of the relations.**
- **Data types are a way to limit the kind of data that can be stored in a table, but they are often insufficient.**
  - e.g., prices must be positive values
  - uniqueness, referential integrity, etc.
- **Can specify constraints on individual columns or on tables.**

126

# Column constraints

```
[ CONSTRAINT constraint_name ]  
  { NOT NULL | NULL | UNIQUE | PRIMARY KEY |  
  CHECK (expression) |  
  REFERENCES reftable [ ( refcolumn ) ] [ ON  
  DELETE action ] [ ON UPDATE action ] }
```

**primary key = unique + not null; also used as default target for references. (can have at most 1)**

**expression must produce a boolean result and reference that column's value only.**

**references is for foreign keys; action is one of:  
NO ACTION, CASCADE, SET NULL, SET  
DEFAULT**

127

# Table constraints

- CREATE TABLE *table\_name*  
( { *column\_name data\_type* [ DEFAULT  
*default\_expr* ] [ *column\_constraint* [, ... ] ] |  
***table\_constraint*** } [, ... ] )

Table Constraints:

- [ CONSTRAINT *constraint\_name* ]  
{ UNIQUE ( *column\_name* [, ... ] ) |  
PRIMARY KEY ( *column\_name* [, ... ] ) |  
CHECK ( *expression* ) |  
FOREIGN KEY ( *column\_name* [, ... ] )  
REFERENCES *reftable* [ ( *refcolumn* [, ... ] ) ] [ ON  
DELETE *action* ] [ ON UPDATE *action* ] }

128



## Create Table (Examples)

```
CREATE TABLE films (  
  code      CHAR(5) PRIMARY KEY,  
  title     VARCHAR(40),  
  did       DECIMAL(3),  
  date_prod DATE,  
  kind      VARCHAR(10),  
  CONSTRAINT production UNIQUE(date_prod)  
  FOREIGN KEY did REFERENCES distributors  
  ON DELETE NO ACTION  
);  
CREATE TABLE distributors (  
  did    DECIMAL(3) PRIMARY KEY,  
  name   VARCHAR(40)  
  CONSTRAINT con1 CHECK (did > 100 AND name <> '')  
);
```

## Other DDL Statements

- **Alter Table**
  - use to add/remove columns, constraints, rename things ...
- **Drop Table**
  - Compare to “Delete \* From Table”
- **Create/Drop View**
- **Create/Drop Index**
- **Grant/Revoke privileges**
  - SQL has an authorization model for saying who can read/modify/delete etc. data and who can grant and revoke privileges!