UNIVERSITÄT MANNHEIM



Heiko Paulheim

Outline

- Today
 - Overview of The SQL Query Language
 - Basic Query Structure
 - Set Operations
 - Join Operators
 - Null Values
 - Aggregate Functions
 - Nested Subqueries
- Next week
 - Data Definition
 - Data Types in SQL
 - Modifications of the database
 - Views
 - Integrity Constraints
 - Roles & Rights

Recap: Database Systems

- Users and applications interact with databases
 - By issuing queries
 - Data definition (DDL): defining, altering, deleting tables
 - Data manipulation (DML): reading from & writing to tables
- SQL is both a DDL and a DML
 - The language that most DBMS speak



History

- IBM SEQUEL language developed as part of System R project at the IBM San Jose Research Laboratory
 - Structured English QUEry Language
- Renamed Structured Query Language (SQL)
- ANSI and ISO standard SQL:
 - SQL-86
 - SQL-89
 - SQL-92
 - SQL:1999
 - SQL:2003







International Organization for Standardization

- Commercial + free systems offer most, if not all, SQL-92 features
 - plus varying feature sets from later standards and special proprietary features
 - Not all examples here may work on your particular system!

Parts of SQL: The Big Picture



Source: https://www.w3schools.in/mysql/ddl-dml-dcl/

2/28/23 Heiko Paulheim

Reading Data

- The **select** clause lists the attributes desired in the result of a query
- Example: find the names of all instructors: select name from instructor
- In relational algebra:
 - $-\prod_{name}$ (instructor)

A Note on Case Sensitivity

- SQL is completely case insensitive
 - select = SELECT = SeLeCt
- Also for names of relations and attributes
 - instructor = Instructor = INSTRUCTOR
 - name = NAME = nAmE
- Each relation / attribute can only exist once
 - Hence, two relations named *instructor* and *Instructor* would not be feasible
- Case sensitivity does *not* apply to values!
 - i.e., "Einstein" and "einstein" are different values!

Renaming Columns in a Select

- Columns can be renamed during selection
- select name, salary as payment from instructor
- In relational algebra
 - a composition of projection and renaming:

 $\rho_{payment \leftarrow salary} (\prod_{name, salary} (instructor))$

The Select Clause

- An asterisk in the select clause denotes "all attributes" select * from instructor
- An attribute can be a literal with no from clause, possibly renamed select '437'
 FOO
 select '437' as FOO
 437
- An attribute can be a literal with from clause select name, 'Instructor' as role from instructor union select name, 'Student' as role from student

name	role
Smith	Instructor
Einstein	Instructor
Johnson	Student

Duplicates

- Difference to relational algebra
 - Sets do not contain duplicates!
- SQL allows duplicates in relations as well as in query results

unless we define a constraint (see later)

- To force the elimination of duplicates, insert the keyword **distinct** after select.
- Find the department names of all instructors, and remove duplicates select distinct dept_name from instructor

Arithmetics in the Selection

- The select clause can contain arithmetic expressions involving the operation, +, –, *, and /, and operating on constants or attributes of tuples
 - Here, we leave relational algebra!
- The query

select ID, name, salary/12 from instructor

would return a relation that is the same as the *instructor* relation, except that the value of the attribute *salary* is divided by 12

- Combined with renaming:
 - **select** *ID*, *name*, *salary/12* **as** *monthly_salary*



Reading Parts of a Relation

- So far, we have always read an entire relation
- Usually, we are interested only in a small portion
- The where clause restricts which parts of the table to read
- To find all instructors in Comp. Sci. dept select name from instructor where dept_name = 'Comp. Sci.'
- In relational algebra: combination of selection and projection

 $\pi_{name}(\sigma_{dept_name} = Comp. Sci.'(r))$

Reading Parts of a Relation

 Comparison results can be combined using the logical connectives and, or, and not

```
select name
from instructor
where dept_name = 'Comp. Sci.' and salary > 90000
π<sub>name</sub>(σ<sub>dept_name</sub> = 'Comp. Sci.' ∧ salary>90000(r))
```

• Can be combined with results of arithmetic expressions

select name, salary/12 as monthly_salary
from instructor
where dept_name = 'Comp. Sci.' and monthly_salary > 7500

Searching in Texts

- So far, we have handled exact equality in selections
- Sometimes, we want to search differently
 - All books that contain "database"
 - All authors starting with "S"

- In SQL: comparing with like and two special characters:
 - _ = any arbitrary character
 - % = any number of arbitrary characters
 - masking with backslash

select ... where title like '%database%'

select ... where author like 'S%'

select ... where amount like '100\%'

most SQL engines don't check types

2/28/23 Heiko Paulheim

Reading Data from Multiple Tables

- Example: find all instructors and the courses they teach
- **select** * **from** *instructor*, *teaches*
 - this generates the *cartesian product*, i.e., instructor x teaches
 - result: generates every possible instructor teaches pair, with all attributes from both relations
- Common attributes (e.g., *ID*), the attributes in the resulting table are renamed using the relation name
 - e.g., instructor.ID, teaches.ID
- Relational algebra notation:
 - $\rho_{instructor.ID \leftarrow ID}(instructor) \ge \rho_{teaches.ID \leftarrow ID}(teaches)$



Cartesian Product

	instructor										te	ea	ches	5		
ID	name		dept_	name	sa	lary		ID		course_	id	sec	c_id	ser	nester	year
10101 12121	Srinivas Wu	san	Com Finai	p. Sci. nce	65 9(5000 0000		1010 1010)1)1	CS-101 CS-315	- >		1 1	Fa Sp	ll ring	2009 2010
22222	Inst.ID	па	me	dept_n	ame	salary	teach	ies.ID	СС	ourse_id	sec_	id	sem	ester	year	2009 2010
22222 32343	10101 10101 10101 10101 10101 10101 12121 12121 12121 12121	Srir Srir Srir Srir Srir Wu Wu Wu Wu	nivasan nivasan nivasan nivasan nivasan	Comp. Comp. Comp. Comp. Comp. Comp. Financ Financ Pinanc Pinanc	Sci. Sci. Sci. Sci. Sci. e e e	65000 65000 65000 65000 65000 65000 90000 90000 90000 90000	101 101 101 121 151 222 101 101 101 121	01 01 21 51 22 01 01 01 21		5-101 5-315 5-347 N-201 U-199 HY-101 5-101 5-315 5-347 N-201	1 1 1 1 1 1 1 1 1		Fall Spr Fall Spr Fall Fall Spr Fall Spr	l ing ing ing l l ing l ing	2009 2010 2009 2010 2010 2009 2009 2010 2009 2010	2010 2010 2009
	12121 12121	Wu Wu 	l	Financ Pinanc	e e	90000 90000	151 222	51 22	M Pł	U-199 HY-101	1		Spr Fal	ring l	2010 2009	
						•••		•		••						

2/28/23

Cartesian Products with Selection

 Find the names of all instructors who have taught some course and the course_id

select name, course_id
from instructor , teaches
where instructor.ID = teaches.ID

• Relational algebra:

 $\pi_{name, course_id}(\sigma_{instructor.ID=teaches.ID}(\rho_{instructor.ID \leftarrow ID}((instructor) \times \rho_{teaches.ID \leftarrow ID}(teaches))))$

Cartesian Product

	LI	1511	uctor								tee	acne	S		
ID	name		dept_	name	sa	lary		ID		course_	id s	ec_id	ser	nester	year
10101 12121	Srinivas Wu	san	Com Finar	p. Sci. 1ce	65 90	5000 2000		1010 1010	1 1	CS-101 CS-315	- 	1 1	Fa Sp	ll ring	2009 2010
15151 22222	Inst.ID	na	me	dept_n	ame	salary	teach	ies.ID	со	urse_id	sec_ic	l sen	nester	year	2009 2010
32343	10101 10101 10101 10101 10101 10101 10101 12121 12121 12121 12121 12121 12121 12121 	Srin Srin Srin Srin Srin Srin Wa Wa Wa Wa Wa	nivasan nivasan nivasan nivasan nivasan	Comp Comp Comp Comp Comp Comp Comp Comp	. Sci. . Sci.	65000 65000 65000 65000 65000 65000 65000 90000 90000 90000 90000 90000 90000 90000 90000	101 101 121 151 222 101 101 101 121 121 151 222 	01 01 21 51 22 01 01 01 21 51 22	CS CS FII M PI · · · · · · · · · · · · · · · · · ·	5-101 5-315 5-347 N-201 U-199 IY 101 5-101 5-315 5-347 N-201 U-199 IY 101	1 1 1 1 1 1 1 1 1 1 	Fal Sp: Fal Sp: Fal Fal Sp: Sp: Sp: Sp: Fal 	l ring ring ring l ring l ring ring l	2009 2010 2009 2010 2010 2010 2009 2009 2010 2009 2010 2010 2010 2009 	2010 2010 2009
	• • •	•••				•••			•••	•	•••		•	•••	

instructor

togchog

2/28/23

Cartesian Products with Selection

 Find the names of all instructors in the Finance department who have taught some course, together with the course_id

select name, course_id
from instructor , teaches
where instructor.ID = teaches.ID and instructor. dept_name = 'Finance'

 $\pi_{name,course_id}(\sigma_{instructor.ID=teaches.ID \land dept_name='Finance'}(\rho_{instructor.ID \leftarrow ID}(instructor) \times \rho_{teaches.ID \leftarrow ID}(teaches))))$

Cartesian Product

	iı	nstr	ructor	•							t	eac	ches	5		
ID	name		dept_	name	sa	lary		ID		course_	id	sec	_id	ser	nester	year
10101 12121 15151	Srinivas Wu	san	Com Finai	p. Sci. nce	65 90	5000 0000		1010 1010	1 1	CS-101 CS-315	5	1		Fa Sp	ll ring	2009 2010
22222	Inst.ID	na	me	dept_n	ame	salary	teach	ies.ID	со	ourse_id	sec_	_id	sem	ester	year	2009
32343	$ 10101 \\ 10101 \\ 10101 $	Srir Srir Srir	hivasan hivasan hivasan	Comp Comp	Sci. Sci. Sci	65000 65000	101 101 101	01 01 01		5 101 5 315 5 347	1 1 1 1		Fall Spr Fall	ing	2009 2010 2009	2010 2009
	10101	Srii	iivasan	Comp.	. Sci.	65000	121	21 51	FI	N-201	1		Spr.	ing	2010 2010	
	10101	Srir	nivasan	Comp	Sci.	65000	222	<u>22</u>	PI	IY 101	1		Fall	шg	2010	
		•••		•••		•••	••	••	•	••			•••		•••	
	 12121	Wu	27	Financ	ē	 90000	-101	01	CC	 5-101	-1		Fall		 2009	
	12121 12121	Wu Wu		l'inance Pinance	ie ie	90000 90000	$\begin{array}{r}101\\-101\end{array}$	01 01	Ct Ct	5-315 5-347	1 1		Spr Fall	ing	2010 2009	
	12121	Wu	• 1	Pinanc	æ	90000	121	21	FI	N-201	1		Spr	ing	2010	
	12121 12121	Wu Wu	-	Pinanc	ю :с	90000 90000	151 222	51 22		U-199 IY-101	1		Spr Fall	ing	2010	
	•••	•••				•••		•	••				•••			

2/28/23

Cartesian Product of a Table with Itself

- Find the names of all instructors who have a higher salary than some instructor in 'Comp. Sci'.
 - We need the same table twice
 - So, we have to use it under different names

select distinct T.name
from instructor as T, instructor as S
where T.salary > S.salary and S.dept_name = 'Comp. Sci.'

 $\pi_{T,name}(\sigma_{T,salary>S,salary \land S,dept_name='Comp. Sci.'}(\rho_{T}(instructor) \times \rho_{S}(instructor)))$

• What happens if we omit the **distinct** here?

Join Operations

Join operations

- take two relations
- return as new relation as their result
- A join operation
 - is a Cartesian product
 - requires that tuples in the two relations match (under some condition)
 - specifies the attributes that are present in the result of the join
- The join operations are typically used as subquery expressions in the **from** clause

Join Operations

- Recap: We have already seen a form of joins:
- A join operation
 - is a Cartesian product
 - requires that tuples in the two relations match (under some condition)
 - specifies the attributes that are present in the result of the join
- Find the names of all instructors who have taught some course and the course_id

select name, course_id
from instructor, teache
where instructor.ID = teaches.ID

- Consider the two relations below
- Desired:
 - List all courses with their prerequisites
 - Note: course CS-315 has no prerequisites

course_id	title	dept_name	credits
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

course_id	prereq_id
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

• List all courses with their prerequisites

select C.course_id, C.title, C.credits, C.dept_name, P.course_id
from course as C, prereq as P
where C.course_id = P.course_id

course_id	title	dept_name	credits	course_id	prereq_id
BIO-301	Genetics	Biology	4	BIO-301	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-190	CS-101
CS-315	Robotics	Comp. Sci.	3	CS-347	CS-101

C.course_id	C.title	C.credits	C.dept_name	P.course_id
BIO-301	Genetics	4	Biology	BIO-101
CS-190	Game Design	4	Comp. Sci.	CS-101

• List all courses with their prerequisites

select C.course_id, C.title, C.credits, C.dept_name, P.prereq_id
from course as C left outer join prereq as P
on C.course_id = P.course_id

course_id	title	dept_name	credits	course_id	prereq_id
BIO-301	Genetics	Biology	4	BIO-301	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-190	CS-101
CS-315	Robotics	Comp. Sci.	3	CS-347	CS-101

C.course_id	C.title	C.credits	C.dept_name	P.prereq_id
BIO-301	Genetics	4	Biology	BIO-101
CS-190	Game Design	4	Comp. Sci.	CS-101
CS-315	Robotics	3	Comp. Sci.	null

Join Operations

- Join type defines how tuples in each relation that do not match any tuple in the other relation (based on the join condition) are treated
 - inner join: ignore
 - outer join: fill with null values
- Join condition defines which tuples in the two relations match, and what attributes are present in the result of the join
 - explicit: on clause
 - implicit: natural keyword

Join types							
inner join							
left outer join							
right outer join							
full outer join							

Join Conditions **natural on** < predicate> **using** $(A_1, A_1, ..., A_n)$

for the moment: keyword for "a blank cell"

• List all courses with their prerequisites

select C.course_id, C.title, C.credits, C.dept_name, P.prereq_id
from course as C right outer join prereq as P
on C.course_id = P.course_id

course_id	title	dept_name	credits	course_id	prereq_id
BIO-301	Genetics	Biology	4	BIO-301	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-190	CS-101
CS-315	Robotics	Comp. Sci.	3	CS-347	CS-101

C.course_id	C.title	C.credits	C.dept_name	P.prereq_id
BIO-301	Genetics	4	Biology	BIO-101
CS-190	Game Design	4	Comp. Sci.	CS-101
CS-347	null	null	null	CS-101

• List all courses with their prerequisites

select C.course_id, C.title, C.credits, C.dept_name, P.prereq_id
from course as C full outer join prereq as P
on C.course_id = P.course_id

course_id	title	dept_name	credits	course_id	prereq_id
BIO-301	Genetics	Biology	4	BIO-301	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-190	CS-101
CS-315	Robotics	Comp. Sci.	3	CS-347	CS-101

C.course_id	C.title	C.credits	C.dept_name	P.prereq_id
BIO-301	Genetics	4	Biology	BIO-101
CS-190	Game Design	4	Comp. Sci.	CS-101
CS-347	null	null	null	CS-101
CS-315	Robotics	3	Comp. Sci.	null

Join Types at a Glance



https://www.codeproject.com/Articles/33052/Visual-Representation-of-SQL-Joins

2/28/23 Heiko Paulheim

Ordering Results

- Recap: Relational Algebra works on sets
 - i.e., it does not have orderings
- For database applications, ordering is often useful, e.g.,
 - list students ordered by names
 select id,name
 from student
 order by name
 - list instructors ordered by department first, then by name select id,name,dept_name from instructor order by dept_name, name

Limiting Results

• Find the three lecturers with the highest salaries

select id,name,salary from instructor order by salary desc limit 3;

- Note: the **desc** keyword creates a descending ordering
- **asc** also exists and creates an ascending ordering
 - also the default when not specifiying the direction

Paging with LIMIT and OFFSET

- Applications, e.g., Web applications, often offer a *paged* view
- Example:
 - Display student list on pages of 100 students
 - with navigation (next page, previous page)

select *id*,*name* from *student* order by *name* limit 100 offset 100;

- offset 100 means: skip the first 100 entries
 - i.e., this query would create the second page
- *Note:* offset should only be used with **order by**
 - otherwise, the results are not deterministic

Set Operations

- All courses that are offered in HWS 2017 and FSS 2018

 (select course_id from section where sem = 'HWS' and year = 2017) intersect
 (select course_id from section where sem = 'FSS' and year = 2018)

 π_{course_id}(σ_{sem='HWS' ∧ year=2017}(section)) ∩ π_{course_id}(σ_{sem='FSS' ∧ year=2018}(section))
- All courses that are offered in HWS 2017 but not in FSS 2018
 (select course_id from section where sem = 'HWS' and year = 2017)
 except
 (select course_id from section where sem = 'FSS' and year = 2018)
 π_{course_id}(σ_{sem='HWS' ∧ year=2017}(section)) π_{course_id}(σ_{sem='FSS' ∧ year=2018}(section))

Set Operations

- All courses that are offered in HWS 2017 or FSS 2018

 (select course_id from section where sem = 'HWS' and year = 2017) union
 (select course_id from section where sem = 'FSS' and year = 2018)

 π_{course_id}(σ_{sem='HWS' ∧ year=2017}(section)) ∪ π_{course_id}(σ_{sem='FSS' ∧ year=2018}(section))
- Alternative solution

(select course_id from section where ((sem = 'HWS' and year = 2017) or (sem = 'FSS' and year = 2018))

 $\pi_{\text{course_id}}(\sigma_{(\text{sem='HWS' } \land \text{ year=2017}) \lor (\text{sem='FSS' } \land \text{ year=2018})}(\text{section}))$

Aggregate Functions – Examples

- Find the average salary of instructors in the Computer Science department
 - select avg (salary)
 from instructor
 where dept_name= 'Comp. Sci.';
- Find the number of tuples in the course relation
 - select count (*)
 from course;
- Find the total number of instructors who teach a course in the Spring 2010 semester
 Why do we need
 - select count (distinct ID)
 from teaches
 where semester = 'Spring' and year = 2010;

Aggregate Functions with Group By

- Find the average salary of instructors in each department
 - select dept_name, avg (salary) as avg_salary
 from instructor
 group by dept_name;

ID	name	dept_name	salary
76766	Crick	Biology	72000
45565	Katz	Comp. Sci.	75000
10101	Srinivasan	Comp. Sci.	65000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000
12121	Wu	Finance	90000
76543	Singh	Finance	80000
32343	El Said	History	60000
58583	Califieri	History	62000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
22222	Einstein	Physics	95000

dept_name	avg_salary
Biology	72000
Comp. Sci.	77333
Elec. Eng.	80000
Finance	85000
History	61000
Music	40000
Physics	91000

2/28/23 Heiko Paulheim

Aggregate Functions with Group By

 Attributes in select clause outside of aggregate functions must appear in group by list

/* erroneous query */ select dept_name, ID, avg (salary) from instructor group by dept_name;



ID	name	dept_name	salary
76766	Crick	Biology	72000
45565	Katz	Comp. Sci.	75000
10101	Srinivasan	Comp. Sci.	65000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000
12121	Wu	Finance	90000
76543	Singh	Finance	80000
32343	El Said	History	60000
58583	Califieri	History	62000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
22222	Einstein	Physics	95000

dept_name	avg(salary)
Biology	72000
Comp. Sci.	77333
Elec. Eng.	80000
Finance	85000
History	61000
Music	40000
Physics	91000

Conditions on Aggregate Values

- Find the names and average salaries of all departments whose average salary is greater than 42000
 - select dept_name, avg (salary) as avg_salary from instructor group by dept_name where avg_salary > 42000;
- Problem:
 - Aggregation is performed *after* selection and projection
 - Hence, the variable *avg_salary* is not available when the **where** clause is evaluated
 - \rightarrow The above query will not work

Conditions on Aggregate Values

- Find the names and average salaries of all departments whose average salary is greater than 42000
 - select dept_name, avg (salary) as avg_salary from instructor group by dept_name having avg_salary > 42000;

performance!

- The **having** clause is evaluated *after* the aggregation
- Hence, it is different from the where clause
- Rule of thumb
 - Conditions on aggregate values can only be defined using having

NULL Values

- *null* signifies an unknown value or that a value does not exist
- It is possible for tuples to have a null value, denoted by *null*, for some of their attributes
 - can be forbidden by a **not null** constraint
 - keys can never be null!
- The result of any arithmetic expression involving *null* is *null*
- Example: 5 + *null* returns null
- The predicate is null can be used to check for null values
- Example: Find all instructors whose salary is null.

select name

from instructor where salary is null

NULL Values and Three Valued Logic

- Three values *true*, *false*, *unknown*
- Any comparison with *null* returns *unknown*
 - Example: 5 < null or null <> null or null = null
- Three-valued logic using the value *unknown*:
 - OR: (unknown or true) = true, (unknown or false) = unknown (unknown or unknown) = unknown
 - AND: (true and unknown) = unknown, (false and unknown) = false, (unknown and unknown) = unknown
 - NOT: (not unknown) = unknown
- "*P* is unknown" evaluates to true if predicate *P* evaluates to unknown
- Result of where clause predicate is treated as *false* if it evaluates to unknown

Aggregates and NULL Values

- Total all salaries
 - select sum (salary)
 from instructor
 - Above statement ignores null amounts
 - Result is null if there is no non-null amount
- All aggregate operations except count(*) ignore tuples with null values on the aggregated attributes
 ID name dept_name
- What if collection has only null values?
 - count returns 0
 - all other aggregates return null

ID	name	dept_name	salary
76766	Crick	Biology	72000
45565	Katz	Comp. Sci.	75000
10101	Srinivasan	Comp. Sci.	null
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000
12121	Wu	Finance	null
76543	Singh	Finance	80000
32343	El Said	History	60000
58583	Califieri	History	null
15151	Mozart	Music	40000
33456	Gold	Physics	87000
22222	Einstein	Physics	null

Subqueries

- SQL provides a mechanism for the nesting of subqueries. A subquery is a select-from-where expression that is nested within another query.
- The nesting can be done in the following SQL query

```
select A_1, A_2, ..., A_n
from r_1, r_2, ..., r_m
where P
```

as follows:

- A_i can be replaced be a subquery that generates a single value
- r_i can be replaced by any valid subquery
- *P* can be replaced with an expression of the form:

B <operation> (subquery)

Where *B* is an attribute and <operation> to be defined later

Subqueries in the WHERE Clause

- A common use of subqueries is to perform tests:
 - for set membership
 - for set comparisons
 - for set cardinality

Test for Set Membership

 Find courses offered this term by lectures from the biology department

```
select distinct course_id
from teaches
where semester = 'Spring' and year= 2022 and
ID in (select ID from instructor where dept_name = 'Biology');
```

• Find courses offered this term before 9 am or after 5 pm

Test for Set Membership

 Find the total number of (distinct) courses offered by instructors in the biology department

```
select count(distinct course_id )
from teaches
where semester = 'Spring' and year= 2022 and
ID in (select ID from instructor where dept_name = 'Biology');
```

- Note: in all of those cases, other (sometimes simpler) solutions are possible
 - In SQL, there are often different ways to solve a problem
 - A question of personal taste
 - But also: a question of performance...

Test for Set Membership

 Find the total number of (distinct) courses offered by instructors in the biology department

```
select count(distinct course_id )
from teaches
where semester = 'Spring' and year= 2022 and
ID in (select ID from instructor where dept name = 'Biology');
```

```
VS.
```

select count (distinct course_id)
from teaches, instructor
where teaches.ID = instructor.ID and instructor.department = 'Biology';



Set Comparison with SOME

 Find names of instructors with salary greater than that of some (at least one) instructor in the Biology department

select distinct T.name
from instructor as T, instructor as S
where T.salary > S.salary and S.dept name = 'Biology';

Same query using > some clause

select name from instructor where salary > some (select salary from instructor where dept name = 'Biology');

Set Comparison with ALL

• Find names of instructors with salary greater than that of all instructors in the Biology department

select name from instructor where salary > all (select salary from instructor where dept name = 'Biology');

 Note: we could also achieve this with MIN and MAX aggregates in the subqueries

Definition: Comparisons with SOME

• F <comp> some $r \Leftrightarrow \exists t \in r \text{ such that } (F <comp> t)$ Where <comp> can be: <, \leq , >, =, \neq



Definition: Comparisons with ALL

• F <comp> all $r \Leftrightarrow \forall t \in r \text{ (F <comp> } t)$



Existential Quantification in Subqueries

Select all courses offered this year which are taken by at least one student

```
- select course_id
from section
where semester = 'Spring' and year = 2022 and
exists (select *
from takes
where takes.course_id = section.course_id
and takes.sec_id = section.sec_id
and takes.semester = section.semester );
```

- The **exists** construct returns the value **true** if the result of the subquery is not empty
 - exists $r \Leftrightarrow r \neq \emptyset$
 - not exists $r \Leftrightarrow r = \emptyset$

Subqueries with NOT EXISTS

 Find all students who have taken all courses offered in the Biology department

- First nested query lists all courses offered in Biology
- Second nested query lists all courses a particular student took
- Note that $X Y = \emptyset \iff X \subseteq Y$
- Note: Cannot write this query using = all and its variants

Test for Duplicate Tuples

• Find all courses that were offered at most once in 2009

```
select T.course_id
from course as T
where unique (select R.course_id
from section as R
where T.course_id= R.course_id
and R.year = 2009);
```

- The **unique** construct evaluates to "true" if a given subquery contains no duplicates
- With **not unique**, we could query for courses that were offered more than once

Subqueries in the FROM Clause

- So far, we have considered subqueries in the where clause
- Find the average instructors' salaries of those departments where the average salary is greater than \$42,000."

Note that we do not need to use the having clause

– why?

Creating Temporary Relations Using WITH

• Find all departments with the maximum budget

```
with max_budget (value) as
    (select max(budget)
    from department)
select department.name
from department, max_budget
where department.budget = max budget.value;
```

 The with clause provides a way of defining a temporary relation whose definition is available only to the query in which the with clause occurs

Creating Temporary Relations Using WITH

- A more complex example involving two temporary relations:
 - Find all departments where the total salary is greater than the average of the total salary at all departments

```
with
```

dept_total (dept_name, value) as
 (select dept_name, sum(salary)
 from instructor
 group by dept_name),
dept_total_avg(value) as
 (select avg(value)
 from dept_total)
select dept_name
from dept_total, dept_total_avg
where dept_total.value > dept_total_avg.value;

Scalar Subqueries in the SELECT Part

List all departments along with the number of instructors in each department

select dept_name, (select count(*) from instructor where department.dept_name = instructor.dept_name) as num_instructors from department;

- Scalar subqueries return a single result
 - More specifically: a single *tuple*
- Runtime error if subquery returns more than one result tuple

Summary of Subqueries

- SELECT queries are the most often used part of SQL
- Their basic structure is simple, but subqueries are a powerful means to make them quite expressive

```
select A_1, A_2, ..., A_n
from r_1, r_2, ..., r_m
where P
```

- Subqueries in **select** part $(A_1, A_2, ..., A_n)$
 - Scalar subqueries (single values, like aggregates)
- Subqueries in **from** part $(r_1, r_2, ..., r_m)$
 - Temporary relations (can also be defined using with)
- Subqueries in **where** part (*P*)
 - Set comparisons, empty sets, test for duplicates
 - Universal and existential quantification

Summary: SQL SELECT at a Glance

- The tool support of SQL varies
- what we have covered here is standard SQL
 - Supported by *most* tools



Recap: The Big Picture



Source: https://www.w3schools.in/mysql/ddl-dml-dcl/

2/28/23 Heiko Paulheim

Summary and Take Aways

- SQL is a standarized language for relational databases
 - DML: Data Manipulation Language
- DML
 - Read data from tables using SELECT
- Coming Up:
 - Writing data to tables
 - Creating and changing tables
 - Rights & Roles
 - ...



Questions?

