

Databases

TDA357/DIT621– LP3 2023

Lecture 13

Ana Bove

(much of the material is based on material from
both Thomas Hallgren and Jonas Duregård)

February 23rd 2023

Recall Last Lecture

- Transactions:
 - **BEGIN/COMMIT** or **ROLLBACK**;
 - ACID properties: atomicity, consistency, isolation and durability;
 - Interference problems: dirty read, non-repeatable read, phantom;
 - Isolation levels: serializable, repeatable read, read committed, read uncommitted;
- Authorisation and privileges:
 - Privileges can be granted and revoked to users/roles;
 - Most common privileges: **SELECT**, **INSERT**, **UPDATE** and **DELETE**;
- Indexes:
 - DBMS defines indexes on primary keys and unique attributes;
 - Users can define (and drop) other indexes: advantages and disadvantages on this!

Overview of Today's Lecture

- Relational algebra;
- Correspondence between **SQL** and relational algebra;
- A glance into query optimisation.

Recall (Lecture 6): Relations, Relation Schemas and Tables

- A *relation* R is a subset of the cartesian product of two or more sets T_1, T_2, \dots, T_n :

$$R \subseteq T_1 \times T_2 \times \dots \times T_n$$

- A *relation schema* $R(a_1, \dots, a_n)$ can be augmented with the domain/type of each attribute $R(a_1 : T_1, \dots, a_n : T_n)$;
- The *relation signature* of the relation R is then the corresponding cartesian product $T_1 \times \dots \times T_n$;
- Given a relation schema $R(a_1, \dots, a_n)$ with signature $T_1 \times \dots \times T_n$:
 - A *table* for the schema $R(a_1, \dots, a_n)$ is a subset of the cartesian product $T_1 \times \dots \times T_n$;
 - A *row* in the table is an element of the cartesian product $t \in T_1 \times \dots \times T_n$.

Algebra

Definition: (Wikipedia) *Algebra* is the study of mathematical symbols and the rules for manipulating these symbols.

Broad field of mathematics, going from elementary equation solving (elementary algebra, linear algebra) to the study of abstractions such as groups, rings, lattices... (abstract algebra).

In an algebra we have set of values, operations on those values, and formulas built from the values and the operators.

Example: Natural numbers with addition and multiplication form an algebra.

Example: Boolean algebra consists of the two Boolean values and the operators \vee, \wedge, \dots

Note: Observe that both the Natural numbers and Booleans are *closed* under those operations (the result of the operation is also in the set).

Relational Algebra

Definition: Relational algebra is a theory that uses algebraic structures for modeling data and defining queries on the data.

A concise mathematical notation in which we can express relations and queries.

Main advantages:

Reasoning: We can use mathematics to prove that our queries do what we intend them to do.

Simplification: Using known algebraic laws one can simplify complicated relational algebra expressions (queries).

Optimisation: Simplification can make queries faster.

SQL Query Processing

A DBMS processes a query in several steps:

Lexing: The input string is converted into a sequence of tokens.

Parsing: The sequence of tokens is converted into a(n **SQL**) syntax tree.

Type checking: The syntax tree is checked semantically.

Logical query plan generation: The syntax tree is converted into a relational algebra expression.

Optimisation: The relational algebra expression is converted into a more efficient relational algebra expression.

Which of two different queries solving the same problem is more efficient (partly) depends on what the DBMS does here.

Physical query plan generation: The more efficient expression is converted into a sequence of algorithm calls.

Execution: The physical query plan is executed and produces a result.

Back to Relational Algebra

What are the values and operations in this algebra?

The values are *relations* (the tables in the database).

This means the relation schema (name of the relation and its attributes) and the (labelled) tuples in the relation (rows in the table).

The operations are the different things we can do on a relation.

Example: Select, group by, order, ...

Each operation returns a new *relation*.

Note: When working with relational algebra, we focus on the schema, not the actual tuples.

Do we Work with Sets, Bags, Lists or Arrays?

	No order	Ordered
No duplicates	Sets	Ordered sets
Duplicates	Multisets/bags	Lists or arrays

In sets/relations: Order and duplication is irrelevant.

In tables: Order and duplication make a difference:

- **SQL** has **DISTINCT** and **ORDER BY**;
- Set operations **UNION**, **INTERSECT** and **EXCEPT** discard duplicates ...
- ... but **UNION ALL**, **INTERSECT ALL** and **EXCEPT ALL** preserve duplicates;
- Primary keys and unique constraints prevent duplicates ...
- ... but the result of a query can still contain duplicates.

Relational Algebra: Basic Notation and Correspondences

Concept	Relational algebra	Set theory	SQL
domain of attribute values		T	type
cartesian products of sets	$T_1 \times \dots \times T_n$	$\{\langle t_1, \dots, t_n \rangle \mid t_i \in T_i\}$	relation schema
relation	R	$R \subseteq T_1 \times \dots \times T_n$	table
tuple	$\{a = t_1, \dots, k = t_n\}$	$\langle t_1, \dots, t_n \rangle$	row
label	a		attribute name
component	$t.a$	$\pi_{i(a)} t, t_{i(a)}$	value of attribute
		$\{t_i \mid \langle \dots, t_i, \dots \rangle \in R\}$	column

Relational Schema in this Lecture

- Countries and their currency values:

Currencies (code, name, value)

Countries (name, abbr , capital, area, population, continent, currency)
currency → Currencies.code

- Countries with their capitals and currencies:

Capitals (country, capital)
country → Countries.name

CurrencyCodes (country, currency)
country → Countries.name

- Students and their grades:

Students (idnr, name)

Grades (student, course, grade)
student → Students.idnr
course → ...

From **SQL** to Relational Algebra: Basic Queries

Table names can be used directly.

SQL

```
SELECT projection  
FROM Table  
WHERE condition;
```



Relational Algebra

$$\pi_{\text{projection}}(\sigma_{\text{condition}}(\text{Table}))$$

π for projections of components

σ for selection of tuples/elements

Note: Here `condition` is a Boolean-expression, not an **SQL** sub-query/relational algebra expression!

Examples: Basic Queries in Relational Algebra

SQL

```
SELECT * FROM Countries;
```

Selecting rows:

```
SELECT * FROM Countries  
WHERE abbr = 'UY';
```

Projection:

```
SELECT capital FROM Countries  
WHERE abbr = 'UY';
```

Expressions:

```
SELECT capital, population/area  
FROM Countries  
WHERE abbr = 'UY';
```

Relational Algebra



```
Countries
```



```
 $\sigma_{abbr='UY'}(Countries)$ 
```



```
 $\pi_{capital}$   
 $(\sigma_{abbr='UY'}(Countries))$ 
```



```
 $\pi_{capital, population/area}$   
 $(\sigma_{abbr='UY'}(Countries))$ 
```

Note: From now on, we might omit some (,) when we break into different lines.

From SQL to Relational Algebra: Renaming

SQL

Renaming columns:

```
SELECT name AS country,  
       population/area AS density  
FROM Countries  
WHERE continent = 'EU';
```



Relational Algebra

AS becomes \rightarrow :

```
 $\pi_{name \rightarrow country, population/area \rightarrow density}$   
 $\sigma_{continent='EU'}$   
Countries
```

Renaming tables:

```
SELECT A.abbr  
FROM Countries AS A,  
WHERE A.name = 'Sweden';
```



ρ for renaming tables:

```
 $\pi_{A.abbr}(\sigma_{A.name='Sweden'}(\rho_A \text{ Countries}))$ 
```

ρ can be also be used to give a new schema (renaming tables and attributes)



```
 $\pi_{..., A.a_i, ...}$   
 $\sigma_{..., A.a_j = ...}(\rho_{A \langle a_1, a_2, ... \rangle} \text{ Table})$ 
```

From SQL to Relational Algebra: Sorting

SQL

Sorting on a
selected attribute:

```
SELECT name, capital,  
FROM Countries  
ORDER BY name;
```

Descending:

```
SELECT name, capital,  
FROM Countries  
ORDER BY name DESC;
```

Sorting on a
non-selected attribute:

```
SELECT name, capital,  
FROM Countries  
ORDER BY area;
```

Relational Algebra

τ for sorting:

$\pi_{\text{name, capital}} (\tau_{\text{name}} \text{Countries})$

$\tau_{\text{name}} (\pi_{\text{name, capital}} \text{Countries})$

$\pi_{\text{name, capital}} (\tau_{-\text{name}} \text{Countries})$

$\tau_{-\text{name}} (\pi_{\text{name, capital}} \text{Countries})$

One cannot sort by an attribute
that has already been discarded
by a projection:

$\pi_{\text{name, capital}} (\tau_{\text{area}} \text{Countries})$

From **SQL** to Relational Algebra: Duplicates

Recall: set vs. bag semantics!

SQL

```
SELECT currency,  
FROM Countries;
```

```
SELECT DISTINCT currency,  
FROM Countries;
```

Relational Algebra

Might contain duplicates:

π_{currency} Countries

δ for removing duplicates:

$\delta(\pi_{\text{currency}}$ Countries)

From **SQL** to Relational Algebra: Grouping and Aggregations

SQL

```
SELECT currency, COUNT(name),  
FROM Countries  
GROUP BY currency;
```



Relational Algebra

γ combines
SELECT and GROUP BY:

```
 $\gamma_{\text{currency}, \text{COUNT}(\text{name})}$  Countries
```

Aggregations are done in γ !
Rename to use the result elsewhere.

```
SELECT currency, SUM(population),  
FROM Countries  
GROUP BY currency  
HAVING COUNT(name) > 1;
```



```
 $\pi_{\text{currency}, \text{sum\_pop}}$   
 $\sigma_{\text{cnt} > 1}$   
 $\gamma_{\text{currency}, \text{SUM}(\text{population}) \rightarrow \text{sum\_pop},$   
 $\text{COUNT}(\text{name}) \rightarrow \text{cnt}}$   
Countries
```

Note: Both **WHERE** cond and **HAVING** cond become σ_{cond} .

From **SQL** to Relational Algebra: Cartesian Products

SQL

Full cartesian product:

```
SELECT Countries.name,  
       Currencies.name,  
FROM Countries, Currencies;
```



Relational Algebra

Not often what we need...

```
 $\pi_{\text{Countries.name, Currencies.name}}$   
 $\text{Countries} \times \text{Currencies}$ 
```

Theta join:

```
SELECT Countries.name,  
       Currencies.name,  
FROM Countries, Currencies  
WHERE currency = code;
```



```
 $\pi_{\text{Countries.name, Currencies.name}}$   
 $\sigma_{\text{currency=code}}$   
 $\text{Countries} \times \text{Currencies}$ 
```

From **SQL** to Relational Algebra: Natural and Inner Joins

SQL

Natural join:

```
SELECT capital, currency,  
FROM Capitals  
NATURAL JOIN CurrencyCodes;
```



Relational Algebra

```
 $\pi_{\text{capital, currency}}$   
 $\text{Capitals} \bowtie \text{CurrencyCodes}$ 
```

Inner join:

```
SELECT Countries.name,  
       Currencies.name,  
FROM Countries JOIN Currencies  
ON currency = code;
```



```
 $\pi_{\text{Countries.name, Currencies.name}}$   
 $\text{Countries} \bowtie_{\text{currency}=\text{code}} \text{Currencies}$ 
```

From SQL to Relational Algebra: Outer Joins

SQL

```
SELECT Countries.name,  
       Currencies.name,  
FROM Countries  
RIGHT OUTER JOIN Currencies  
ON currency = code;
```

```
SELECT Countries.name,  
       Currencies.name,  
FROM Countries  
LEFT OUTER JOIN Currencies  
ON currency = code;
```

```
SELECT Countries.name,  
       Currencies.name,  
FROM Countries  
FULL OUTER JOIN Currencies  
ON currency = code;
```

Relational Algebra

Right outer join:

$$\pi_{\text{Countries.name, Currencies.name}} \left(\text{Countries} \bowtie_{\text{currency=code}}^{\text{OR}} \text{Currencies} \right)$$

Left outer join:

$$\pi_{\text{Countries.name, Currencies.name}} \left(\text{Countries} \bowtie_{\text{currency=code}}^{\text{OL}} \text{Currencies} \right)$$

Full outer join:

$$\pi_{\text{Countries.name, Currencies.name}} \left(\text{Countries} \bowtie_{\text{currency=code}}^{\text{O}} \text{Currencies} \right)$$

What about Correlated Queries?

Consider the query:

```
SELECT name
FROM Students AS S
WHERE 4 < (SELECT AVG(grade) FROM Grades WHERE student = S.idnr)
```

The correlation needs to be replaced with a join (or cartesian product and corresponding select).

```
 $\pi_{\text{name}}(\sigma_{4 < \text{average}(\gamma_{\text{student}, \text{name}}, \text{AVG}(\text{grade}) \rightarrow \text{average}(\text{Grades} \bowtie_{\text{idnr}=\text{student}} \text{Students}))})$ 
```

From the join, group by students and name, and compute the average of the grades of each student, select those with an average of at least 4, finally project the name.

Back to Grouping in Relational Algebra

Are these the same query?

Groups by student
id *and* name;
name available
for projection


$$\begin{aligned} &\pi_{\text{name,passed}} \\ &\quad \gamma_{\text{student,name,COUNT}(*)} \rightarrow \text{passed} \\ &\quad \sigma_{\text{grade} \geq 3 \wedge \text{idnr}=\text{student}} (\text{Students} \times \text{Grades}) \end{aligned}$$

Groups *only*
by student;
needs a join to
project name


$$\begin{aligned} &\pi_{\text{name,passed}} \\ &\quad (\gamma_{\text{student,COUNT}(*)} \rightarrow \text{passed} \\ &\quad \quad (\sigma_{\text{grade} \geq 3 \wedge \text{idnr}=\text{student}} (\text{Students} \times \text{Grades}))) \\ &\quad \bowtie_{\text{student}=\text{idnr}} \text{Students} \end{aligned}$$

Note: Since we have the FD $\text{student} \rightarrow \text{name}$, these queries are the same. Otherwise they might not be since the result of grouping by $(\text{student}, \text{name})$ might be different than grouping by just student!

What about NOT IN or NOT EXISTS?

One needs to try to understand the query in terms of sets, and use set operations instead.

Example: Consider the query that selects students that have no grades:

```
SELECT idnr, name
FROM Students
WHERE idnr NOT IN (SELECT student FROM Grades)
```

We can use set difference in relational algebra to obtain this result:

$$\pi_{\text{idnr}, \text{name}}(\text{Students} \bowtie_{\text{idnr}=\text{st}} (\rho_{\text{NoGrades}\langle \text{st} \rangle}(\pi_{\text{idnr}} \text{Students} - \pi_{\text{student}} \text{Grades})))$$

The result of $(\pi_{\text{idnr}} \text{Students} - \pi_{\text{student}} \text{Grades})$ is a relation consisting of “1-tuples”. We give a new name to the information (table and attribute).

We join to retrieve the rest of the information of the students to project their name.

Relation Algebra: Set Operation

Concept	Relational algebra	Set theory	SQL
cartesian products of relations	$S \times R$	$\{\langle x, \dots, u, \dots \rangle \mid \langle x, \dots \rangle \in S, \langle u, \dots \rangle \in R\}$	FROM S, R
union with duplicates	$S \cup R$	$\{t \mid t \in S \text{ or } t \in R\}$ OBS: bags!	S UNION ALL R
union	$\delta(S \cup R)$	$\{t \mid t \in S \text{ or } t \in R\}$	S UNION R
intersection with duplicates	$S \cap R$	$\{t \mid t \in S \text{ and } t \in R\}$ OBS: bags!	S INTERSECT ALL R
intersection	$\delta(S \cap R)$	$\{t \mid t \in S \text{ and } t \in R\}$	S INTERSECT R
difference with duplicates	$S - R$	$\{t \mid t \in S \text{ and } t \notin R\}$ OBS: bags!	S EXCEPT ALL R
difference	$\delta(S - R)$	$\{t \mid t \in S \text{ and } t \notin R\}$	S EXCEPT R

Note: Recall “schemas” need to be compatible for some of the set operations to work!

Relation Algebra: Summary of Correspondences

Concept	Rel algebra	Set theory	SQL
projection	$\pi_{a,b,\dots,k} R$	$\langle t.a, t.b, \dots, t.k \mid t \in R \rangle$	SELECT a, b, ..., k
selection	$\sigma_C R$ $\sigma_C R$	$\{t \in R \mid C\}$	WHERE C HAVING C
theta join	$S \bowtie_C R = \sigma_C(S \times R)$	$\{t \in S \times R \mid C\}$	S INNER JOIN R ON C
outer join	$S \bowtie_C^O R$ $S \bowtie_C^{OL} R$ $S \bowtie_C^{OR} R$	S FULL OUTER JOIN R ON C LEFT OUTER JOIN RIGHT OUTER JOIN
natural join	$S \bowtie R$...	S NATURAL JOIN R
renaming	$a \rightarrow b$ $\rho_A R$	—	AS
new schema	$\rho_{A\langle a,\dots,k \rangle} R$	—	—
removing duplicates	δR	—	DISTINCT
sorting	$\tau_a R$	—	ORDER BY a
grouping	$\gamma_a R$	—	GROUP BY a

Example: From Problem to Relational Algebra

Select the name of all students that have passed at least 2 courses.

Students (idnr, name)
Grades (student, course, grade)
student \rightarrow Students.idnr
course \rightarrow ...

- Group first, join later: select passed courses in **Grades**, group by students and count passed course per student, now do the join, select entries where at least 2 courses are passed and project student names.

$$\pi_{\text{name}}(\sigma_{\text{passed} \geq 2 \wedge \text{idnr} = \text{student}}(\text{Students} \times \gamma_{\text{student}, \text{COUNT}(*), \rightarrow \text{passed}}(\sigma_{\text{grade} \geq 3} \text{Grades})))$$

- Join first, group later: from the join, select passed courses, group by students and names, and count passed courses per student, select entries where at least 2 courses are passed and project student names.

$$\pi_{\text{name}}(\sigma_{\text{passed} \geq 2}(\gamma_{\text{student}, \text{name}, \text{COUNT}(*), \rightarrow \text{passed}}(\sigma_{\text{grade} \geq 3 \wedge \text{idnr} = \text{student}}(\text{Students} \times \text{Grades}))))$$

Example: From SQL to Relational Algebra

A query with almost everything:

```
SELECT a1, MAX(a2) AS max
FROM T1, T2
WHERE a3 = 5
GROUP BY a1, a3
HAVING COUNT(*) > 10
ORDER BY a1 DESC;
```

A relational algebra expression for it:

$$\tau_{-a1}(\pi_{a1, \max}(\sigma_{\text{cnt} > 10}(\gamma_{a1, a3, \max(a2) \rightarrow \max, \text{COUNT}(*), \text{COUNT}(*), \text{COUNT}(*)}(\sigma_{a3=5}(T1 \times T2))))))$$

From the join, select entries where $a3 = 5$.

Group by $a1, a3$ and compute number of elements and $\max(a2)$ per group.

Select entries where the count is at least 10.

Project $a1$ and $\max(a2)$ for those entries.

Sort the result descending in $a1$.

Sanity Check (1)!

Given this schema, is the relational algebra query below correct?

Students (idnr, name)
Grades (student, course, grade)
student \rightarrow Students.idnr
course \rightarrow ...

$\pi_{\text{idnr}}(\sigma_{\text{passed} \geq 2 \wedge \text{idnr} = \text{student}}(\text{Students} \times \gamma_{\text{student}, \text{COUNT}(*), \rightarrow \text{passed}}(\sigma_{\text{grade} \geq 3} \text{Grades})))$

Let us *sanity check* the expression by computing the schema bit by bit!

- $\sigma_{\text{grade} \geq 3} \text{Grades} : (\text{student}, \text{course}, \text{grade})$
- $\gamma_{\text{student}, \text{COUNT}(*), \rightarrow \text{passed}}(\sigma_{\text{grade} \geq 3} \text{Grades}) : (\text{student}, \text{passed})$
- $\text{Students} \times \gamma_{\text{student}, \text{COUNT}(*), \rightarrow \text{passed}}(\sigma_{\text{grade} \geq 3} \text{Grades}) : (\text{idnr}, \text{name}, \text{student}, \text{passed})$
- $\sigma_{\text{passed} \geq 2 \wedge \text{idnr} = \text{student}}(\text{Students} \times \gamma_{\text{student}, \text{COUNT}(*), \rightarrow \text{passed}}(\sigma_{\text{grade} \geq 3} \text{Grades})) : (\text{idnr}, \text{name}, \text{student}, \text{passed})$
- $\pi_{\text{idnr}}(\sigma_{\text{passed} \geq 2 \wedge \text{idnr} = \text{student}}(\text{Students} \times \gamma_{\text{student}, \text{COUNT}(*), \rightarrow \text{passed}}(\sigma_{\text{grade} \geq 3} \text{Grades}))) : (\text{idnr})$

Sanity Check (2)!

Given this schema, is the relational algebra query below correct?

Students (idnr, name)
Grades (student, course, grade)
student \rightarrow Students.idnr
course \rightarrow ...

$\pi_{\text{idnr}}(\sigma_{\text{cnt} \geq 2 \wedge \text{idnr} = \text{student} \wedge \text{grade} \geq 3}(\text{Students} \times \gamma_{\text{student}, \text{COUNT}(\text{*}) \rightarrow \text{cnt}} \text{Grades}))$

Let us *sanity check* the expression by computing the schema bit by bit!

- $\gamma_{\text{student}, \text{COUNT}(\text{*}) \rightarrow \text{cnt}} \text{Grades} : (\text{student}, \text{cnt})$
- $\text{Students} \times \gamma_{\text{student}, \text{COUNT}(\text{*}) \rightarrow \text{cnt}} \text{Grades} : (\text{idnr}, \text{name}, \text{student}, \text{cnt})$
- $\sigma_{\text{cnt} \geq 2 \wedge \text{idnr} = \text{student} \wedge \text{grade} \geq 3}(\text{Students} \times \gamma_{\text{student}, \text{COUNT}(\text{*}) \rightarrow \text{cnt}} \text{Grades})$: **ERROR!**

There is no grade in the schema of the expression to be used by the σ operator in order to evaluate the condition!

Note: Make sure your expression is correct by performing a sanity check on it!

Some Algebraic Laws

These (and other) algebraic laws can be used for query simplification and/or optimisation.

Some laws generate a potentially infinite number of equivalent expressions for a query. Query optimization tries to find the best of those.

Some laws work with sets but not with bags!

$$\begin{aligned} R \bowtie S &= S \bowtie R \\ R \bowtie (S \bowtie T) &= (R \bowtie S) \bowtie T \end{aligned}$$

Set-theoretic laws:

If applicable, associativity, commutativity, distributivity, idempotence of unions, intersections, products and joins.

Repeated projection:

$$\pi_{a_1, \dots, a_n}(\pi_{b_1, \dots, b_m} R) = \pi_{a_1, \dots, a_n} R$$

b_1, \dots, b_m should be a plain projection, will not work if there is a rename that it is later projected.

a_1, \dots, a_n should be a subset of b_1, \dots, b_m for the expression to be correct.

Some Algebraic Laws (Cont.)

Some laws can dramatically reduce the number of rows in the results!

Repeated selection: $\sigma_{C_1}(\sigma_{C_2} R) = \sigma_{C_1 \wedge C_2} R$

Pushing duplicate elimination inside:

$$\begin{aligned}\delta(\sigma_C R) &= \sigma_C(\delta R) \\ \delta(R \times S) &= \delta(R) \times \delta(S)\end{aligned}$$

Pushing selection inside cartesian products:

If C only uses
attributes in R_1 :

$$\sigma_C(R_1 \times R_2) = (\sigma_C R_1) \times R_2$$

If C_1 only in R_1 and C_2
only in R_2 :

$$\sigma_{C_1 \wedge C_2}(R_1 \times R_2) = (\sigma_{C_1} R_1) \times (\sigma_{C_2} R_2)$$

Note: See section 2 in chapter 16 of the book for more algebraic laws for improving queries!

To make large relation algebra expressions more readable:

- Write them over several lines and use indentation to show the structure:

$$\begin{aligned} &\pi_{A.name} \\ &\quad (\sigma_{A.name=B.capital} \\ &\quad \quad (\rho_A \text{ Countries} \times \rho_B \text{ Countries})) \end{aligned}$$

- Split them in to several named parts:

$$\begin{aligned} R_1 &= \rho_A \text{ Countries} \times \rho_B \text{ Countries} \\ R_2 &= \sigma_{A.name=B.capital} R_1 \\ \text{Result} &= \pi_{A.name} R_2 \end{aligned}$$

Final Remarks

- A basic **SQL** query (performing some projection, selection and cartesian product) is done “altogether” and produces a single result. In relational algebra, each operation results in a *new* relation.
- It is then important to keep track of what information is available after each step (see slide 28!).

Recall the expression $(\pi_{\text{idnr}} \text{Students} - \pi_{\text{student}} \text{Grades})$ in slide 22: it results in a relation with student's id as single information; to retrieve the rest of the information of students we needed to perform a join!

- The abstract syntax tree of a relational algebra expression can help you understanding the expression.
- Translating **SQL** to relational algebra, simplifying the expression and then translating it back to **SQL** can give a not too compact query!

Overview of Next Lecture

- Quiz with recap of all course!
(Good if you have a separate device to answer the quiz than that you use to see the questions if you are on zoom.)