Relational Algebra

Instructor: Lei Jiang Slides provided by Ramona Truta

Outline

- How to ask queries in the Relational Model.
- Relational Algebra Operators
 - \circ Set-theory Operators: Union, Intersection, Difference
 - \circ Selection and Projection
 - $\circ \ Renaming$
 - Cartesian Product
 - \circ Natural Join
 - \circ Division
- Relational Algebra Query Examples

Describing data: what is a data model?

• Mathematical representation of data.

PLUS

- Constraints.
- Operations.

Examples of queries

- Find titles of movies.
- Find artists who played in movies directed by Polanski.
- Find the artists who played the highest number of roles in a movie.

• • • •

What is the result of a query?

• They are relations constructed from the relations in the database.

How to ask a query?

- Query languages
 - Commercial: SQL
 - Theoretical: Relational calculus, algebra, datalog etc

Query Languages: Declarative vs Procedural

- In our queries, we ask
 - what we want to see in the output,
 - but we do not say **how** we want to get this output.
- **Declarative** query languages: they specify what is needed in the output, but do not say how to get it.
- **Procedural** languages: specify how to get the result.
 - DBMS figures out how to get the result, and gives it to the user.

Query Languages: Theoretical vs Practical

• Theoretical languages:

- Declarative: relational calculus, rule-based queries
- Procedural: relational algebra
- **Practical languages**: mix of both, but mostly one uses declarative features.

Procedural Language: Relational Algebra

- A collection of algebraic operators that
 - Are defined on relations;
 - produce a relation as result;

- and, therefore, can be combined to form complex algebraic operations.

- Operators:
 - Set operators: Union (\cup), Intersection (\cap), Difference (-);
 - Renaming ρ ;
 - Selection (σ) and Projection (π) ;
 - Join: Natural Join (\bowtie), Cartesian Product (\times), Theta Join (\bowtie_{θ})
 - Division (÷)

Set Operators: Union, Intersection, Difference

- Relations are sets, so we can apply set-theoretic operators.
- However, we want the results to be relations (that is, homogeneous sets of tuples).
- It is therefore meaningfull to only apply ∪, ∩, − to pairs of relations defined over the same attributes.

Union

• Graduates:	Number	Surname	Age
	7255	Kubrick	25
	7342	Altman	28
	7111	Polanski	30

	Number	Surname	Age
• Managers:	7155	Robinson	45
	7111	Polanski	30

• Graduates \cup Managers:

Union cont'd

• Graduates:	Number	Surname	Age
	7255	Kubrick	25
	7342	Altman	28
	7111	Polanski	30

[Number	Surname	Age
Managers:	7155	Robinson	45
	7111	Polanski	30

	Number	Surname	Age
• Graduates \cup Managers:	7255	Kubrick	25
	7342	Altman	28
	7111	Polanski	30
	7155	Robinson	45

Intersection

• Graduates:	Number	Surname	Age
	7255	Kubrick	25
	7342	Altman	28
	7111	Polanski	30

	Number	Surname	Age
Managers:	7155	Robinson	45
	7111	Polanski	30

• Graduates \cap Managers:

Intersection cont'd

• Graduates:	Number	Surname	Age
	7255	Kubrick	25
	7342	Altman	28
	7111	Polanski	30

	Number	Surname	Age
Managers:	7155	Robinson	45
	7111	Polanski	30

• Graduates () Managers:	Number	Surname	Age
• Graduates Managers.	7111	Polanski	30

Difference

• Graduates:	Number	Surname	Age
	7255	Kubrick	25
	7342	Altman	28
	7111	Polanski	30

	Number	Surname	Age
Managers:	7155	Robinson	45
·	7111	Polanski	30

• Graduates - Managers:

• Managers - Graduates:

Difference cont'd

	Number	Surname	Age
atog.	7255	Kubrick	25
1005.	7342	Altman	28
	7111	Polanski	30

• Graduates:

	Number	Surname	Age
• Managers:	7155	Robinson	45
	7111	Polanski	30

	Number	Surname	Age
• Graduates - Managers:	7255	Kubrick	25
	7342	Altman	28

• Managorg - Graduatog	Number	Surname	Age
• Hallagers – Graduates.	7155	Robinson	45

A meaningfull but impossible union



	Mother	Child
• Maternity:	Elizabeth	Charles
	Elizabeth	Andrew

• We want their union, which should be the "parent-child" relation.

– But we cannot use ${\tt Maternity} \cup {\tt Paternity}$ because they have different attributes!

• Solution: *rename* the attributes!

Renaming

- This is a unary operator which changes the attribute names for a relation without changing any values.
- Renaming eliminates the limitations associated with set operators.
- Notation: $\rho_{\texttt{new_name} \leftarrow \texttt{old_name}}(\mathbf{R})$
- Then, $\rho_{\texttt{Parent}\leftarrow\texttt{Father}}(\texttt{Paternity}) \cup \rho_{\texttt{Parent}\leftarrow\texttt{Mother}}(\texttt{Maternity})$:

Renaming cont'd

- This is a unary operator which changes the attribute names for a relation without changing any values.
- Renaming eliminates the limitations associated with set operators.
- Notation: $\rho_{\texttt{new_name} \leftarrow \texttt{old_name}}(\texttt{R})$
- Then, $\rho_{\texttt{Parent}\leftarrow\texttt{Father}}(\texttt{Paternity})\cup\rho_{\texttt{Parent}\leftarrow\texttt{Mother}}(\texttt{Maternity})$:

Parent	Child
George	Elizabeth
Philip	Charles
Charles	William
Elizabeth	Charles
Elizabeth	Andrew

Selection and Projection

- They are unary operators.
- Selection for "horizontal" decompositions;



• Projection for "vertical" decompositions.



Selection

- This is a unary operator which returns a relation
 - with the same schema as the operand,
 - but only with the tuples satisfying a condition.
- Notation: $\sigma_{cond}(R)$
- Semantics: $\sigma_{cond}(\mathbf{R}) = \{\mathbf{t} | \mathbf{t} \in \mathbf{R} \text{ st } cond(\mathbf{t})\}$

- cond is a formula in propositional calculus consisting of terms connected by \wedge,\vee,\neg

- Each term is one of:
- o <attribute> op <attribute> or
- o <attribute> op <constant>
- \circ where op is one of =, $\neq, >, \geq, <, \leq$

Selection: examples

	Number	Surname	Age
atos.	7255	Kubrick	25
ales.	7342	Altman	28
	7111	Polanski	30

• Graduates:

• $\sigma_{\texttt{Number}>7000}(\texttt{Graduates})$

Selection: examples cont'd

	Number	Surname	Age
• Graduatos:	7255	Kubrick	25
• Graduates:	7342	Altman	28
	7111	Polanski	30

• $\sigma_{\texttt{Number}>7000}(\texttt{Graduates})$

Number	Surname	Age
7255	Kubrick	25
7342	Altman	28
7111	Polanski	30

Selection: examples cont'd

	Number	Surname	Age
• Graduates:	7255	Kubrick	25
	7342	Altman	28
	7111	Polanski	30

• $\sigma_{\text{Number} < 7000}(\text{Graduates})$



• $\sigma_{\texttt{Surname} \neq '\texttt{Kubrick'} \land \texttt{Age} \geq 30}(\texttt{Graduates})$

Selection: examples cont'd

	Number	Surname	Age
• Graduatos:	7255	Kubrick	25
• Graduates:	7342	Altman	28
	7111	Polanski	30

• $\sigma_{\text{Number} < 7000}(\text{Graduates})$

Number Surname Age

• $\sigma_{\texttt{Surname} \neq '\texttt{Kubrick'} \land \texttt{Age} \geq 30}(\texttt{Graduates})$

Number	Surname	Age
7111	Polanski	30

Projection

- This is a unary operator which returns a relation
 - which includes all the tuples
 - but only a subset of the attributes of the operand.
- Notation: Given $R(X), Y \subseteq X: \pi_Y(R)$
 - only leaves attributes Y in relation R.
- Semantics: $\pi_{\mathbf{Y}}(\mathbf{R}) = \{\mathbf{t}[\mathbf{Y}] | \mathbf{t} \in \mathbf{R}\}$

Projection: examples

Number	Surname	Age
7255	Kubrick	25
7111	Polanski	30
7211	Polanski	30

• Graduates:

• $\pi_{\text{Number,Age}}(\text{Graduates})$

• $\pi_{\texttt{Surname}}(\texttt{Graduates})$

Projection: examples cont'd

	Number	Surname	Age
• Graduates:	7255	Kubrick	25
• Graduates.	7111	Polanski	30
	7211	Polanski	30

• $\pi_{\texttt{Number},\texttt{Age}}(\texttt{Graduates})$

Number	Age
7255	25
7111	30
7211	30

• $\pi_{\texttt{Surname}}(\texttt{Graduates})$

Surname
Kubrick
Polanski

Cartesian Product

- Puts together two relations
- Notation: $\mathbf{r} \times \mathbf{s}$
- Semantics: $r \times s = \{tq \mid t \in r \text{ and } q \in s\}$
 - puts together each tuple t of r with each tuple q of s
- Assume that the attributes of r and s are disjoint;
 - if they are not disjoint, in the result they have to be renamed.
- \bullet If r has n tuples and s has m tuples, then $r\times s$ has $n\times m$ tuples

Cartesian Product: example

	Employee	Department
• r:	Smith	Sales
• ſ.	Blum	Production
	Lewis	Production

	Department	Head
S:	Sales	Mori
	Production	Brown

 \bullet r \times s:

Cartesian Product: example

	Employee	Department
• r:	Smith	Sales
• []	Blum	Production
	Lewis	Production

	Department	Head
s:	Sales	Mori
	Production	Brown

	Employee	Department	Department1	Head
	Smith	Sales	Sales	Mori
	Smith	Sales	Production	Brown
S:	Blum	Production	Sales	Mori
	Blum	Production	Production	Brown
	Lewis	Production	Sales	Mori
	Lewis	Production	Production	Brown

 \bullet r \times s:

Join

- The most used operator in Relational Algebra.
- Allows us to connect data among different relations.
- Two main versions of the join:
 - Natural join: 🖂
 - Theta join: $\mathbf{r} \bowtie_{\theta} \mathbf{s} = \sigma_{\theta}(\mathbf{r} \times \mathbf{s})$

Natural Join: example

	Employee	Department
• r:	Smith	Sales
• r:	Blum	Production
	Lewis	Production



- We want a new relation, t(Employee, Department, Head), such that Head is the head of the Department where Employee works.
- One possible solution:

 $-\mathbf{r} \times \mathbf{s}$ - Is it really?

Natural Join: a better solution

	Employee	Department
	Smith Sale	
• r:	Blum	Production
	Lewis	Production
	Kifer	HR

	Department	Head
s:	Sales	Mori
	Production	Brown

• $t = r \bowtie s$:

Natural Join: a better solution cont'd

	Employee	Department
	Smith	Sales
• r:	Blum	Production
	Lewis	Production
	Kifer	HR

	Department	Head
s:	Sales	Mori
	Production	Brown

	Employee	Department	Head
r M g.	Smith	Sales	Mori
INS.	Blum	Production	Brown
	Lewis	Production	Brown

• t =

(Optional) Natural Join: it's not a new operator!

	Employee	Department	Department1	Head
	(r.Employee)	(r.Department)	(s.Department)	(s.Head)
	Smith	Sales	Sales	Mori
\mathbf{r} \times \mathbf{q}	Smith	Sales	Production	Brown
1×5 .	Blum	Production	Sales	Mori
	Blum	Production	Production	Brown
	Lewis	Production	Sales	Mori
	Lewis	Production	Production	Brown

• $\mathbf{r} \bowtie \mathbf{s} = \pi_{\text{Employee, Department, Head}}(\sigma_{r.\text{Department}=s.\text{Department}}(\mathbf{r} \times \mathbf{s}))$

- The tuples in the result are obtained by combining tuples from the operands that have equal values on the common attributes.

Natural Join: another example

	Employee	Department	Company
	Smith	Sales	Google
• r:	Blum	Production	Microsoft
	Lewis	Production	Oracle
	Kifer	HR	Pearson

	Department	Head	Company
• s:	Sales	Mori	Google
	Production	Brown	Oracle

• $t = r \bowtie s$:

Natural Join: another example cont'd

	Employee	Department	Company
	Smith	Sales	Google
• r:	Blum	Production	Microsoft
	Lewis	Production	Oracle
	Kifer	HR	Pearson

	Employee	Department	Head	Company
• $t = r \bowtie s$:	Smith	Sales	Mori	Google
	Lewis	Production	Brown	Oracle

Natural Join: the general case

- $r(A_1, A_2, \cdots, A_n, B_1, B_2, \cdots, B_m), s(A_1, A_2, \cdots, A_n, C_1, C_2, \cdots, C_p)$
- $\mathbf{r} \bowtie \mathbf{s}$ is a relation on $A_1, A_2, \cdots, A_n, B_1, B_2, \cdots, B_m, C_1, C_2, \cdots, C_p$: $\pi_{A_1, \dots, A_n, B_1, \dots, B_k, C_1, \dots, C_m}(\sigma_{\mathbf{r}, A_1 = \mathbf{s}, A_1 \land \dots \land \mathbf{r}, A_n = \mathbf{s}, A_n}(\mathbf{r} \times \mathbf{s}))$

Natural Join: definition and properties

• **Definition**:

- $\circ r(X_1), s(X_2)$
- \circ **r** \bowtie **s** is a relation on X_1X_2 (the union of the two sets):
 - $\{ \texttt{t on } \texttt{X}_1\texttt{X}_2 \mid \texttt{t}[\texttt{X}_1] \in \texttt{r} \land \texttt{t}[\texttt{X}_2] \in \texttt{s} \}$
 - or, equivalently:
 - $\{ \texttt{t on } \texttt{X}_1\texttt{X}_2 \mid \exists \texttt{t}_1 \in \texttt{r} \land \exists \texttt{t}_2 \in \texttt{s with } \texttt{t}[\texttt{X}_1] = \texttt{t}_1 \land \texttt{t}[\texttt{X}_2] = \texttt{t}_2 \}$
- Properties:
 - Commutative: $r \bowtie s = s \bowtie r$
 - Associative: $r \bowtie (s \bowtie t) = (r \bowtie s) \bowtie t$
 - Therefore, we can write n-ary joins without ambiguity:

$$r_1 \bowtie r_2 \bowtie \cdots \bowtie r_n$$

Joins can be empty

	Employee	Departmen		t
	Smith	Sales		
• • •	Blum		Production	
	Lewis	Production		
	Departmer		Head	
• 5.	HR		Morison	

t = 1	r⊠s:	Employee	Department	Heac

More on Natural Join

- $\bullet \ \mathtt{r}(\mathtt{X}_1), \mathtt{s}(\mathtt{X}_2)$
 - No assumptions have been made w.r.t. X_1, X_2
 - They could be equal, or disjoint.
- If they are equal: $r \bowtie s$ is a relation on $X_1 = X_2 = X$:
 - {t on X | $t[X] \in r \land t[X] \in s$ } =?

or, equivalently:

- {t on X |
$$\exists t_1 \in r \land \exists t_2 \in s \text{ with } t[X] = t_1 \land t[X] = t_2$$
} =?

- If they are disjoint: $\mathbf{r} \bowtie \mathbf{s}$ is a relation on $X_1 \cup X_2 = X$:
 - {t on $X_1X_2 \mid t[X_1] \in r \land t[X_2] \in s$ } =?

or, equivalently:

- {t on
$$X_1X_2 \mid \exists t_1 \in r \land \exists t_2 \in s \text{ with } t[X_1] = t_1 \land t[X_2] = t_2$$
} =?

Theta-Join

- In most cases, a Cartesian Product is meaningul only if followed by a selection.
- Theta-join:

$$\mathbf{r} \bowtie_{\theta} \mathbf{s} = \sigma_{\theta}(\mathbf{r} \times \mathbf{s})$$

• Equi-join: when the θ condition is a conjunction of equality conditions.

Interaction of relational algebra operators

- $\bullet \ \pi_{\vec{\mathbf{A}}}(\mathbf{R} \cup \mathbf{S}) = \pi_{\vec{\mathbf{A}}}(\mathbf{R}) \cup \pi_{\vec{\mathbf{A}}}(\mathbf{S})$
- $\bullet \ \sigma_{\mathsf{c}}(\mathtt{R} \cup \mathtt{S}) = \sigma_{\mathsf{c}}(\mathtt{R}) \cup \sigma_{\mathsf{c}}(\mathtt{S})$
- $\bullet \ (\mathtt{R} \cup \mathtt{S}) \times \mathtt{T} = \mathtt{R} \times \mathtt{T} \cup \mathtt{S} \times \mathtt{T}$
- $\bullet \ \mathbf{T} \times (\mathbf{R} \cup \mathbf{S}) = \mathbf{T} \times \mathbf{R} \cup \mathbf{T} \times \mathbf{S}$

Assignment Operation

- Provides a convenient way to express complex queries.
- Write a query as a sequential program consisting of
 - a series of assignments

– followed by an expression whose vales is displayed as a result of the query.

- Notation: temp \leftarrow query
- An assignment must always be made to a temporary relation variable.
 - The variable may be used in subsequent expressions.
- Example:
 - $\texttt{-temp1} \gets \texttt{r} \bowtie \texttt{s}$
 - $-\texttt{temp2} \leftarrow \pi_{\texttt{A},\texttt{B},\texttt{C}}(\texttt{temp1})$

Division Operator

- Recall the integer division operator:
 - \circ Given two integers X and Y, define Z = X/Y, so that $Z^*Y = X$.
 - \circ Can't always quite do this, so instead define Z to be the largest integer number smaller than X/Y
- Apply the same idea to the relational division operator, \div .
 - In that sense, \div is the inverse of \times

Division Operator

- Suited for queries that include "for all"
- $r(X_1), s(X_2)$, where
 - $\mathrel{\circ} X_1 = \{A_1, \cdots, A_m, B_1, \cdots, B_n\}$
 - $\circ X_2 = \{B_1, \cdots, B_n\}$
- \bullet Notation: $r \div s$
- $r \div s$ consists of those rows t such that for every row $u \in s$, the row resulting from concatenating t and u can be found in relation r, and there is no larger set of rows for which this is true.
 - The schema of $r \div s$ is $(\mathtt{A_1}, \cdots, \mathtt{A_m})$

Division Operator

- Semantics: $\mathbf{r} \div \mathbf{s} = \{ \mathbf{t} \mid \mathbf{t} \in \pi_{\mathbf{X}_1 \mathbf{X}_2}(\mathbf{r}) \land \forall \mathbf{u} \in \mathbf{s} \ (\mathbf{t}\mathbf{u} \in \mathbf{r}) \}$
 - \circ The schema of $\mathrm{r} \div \mathrm{s}$ is $(\mathtt{A_1}, \cdots, \mathtt{A_m})$

 \circ tu means the concatenation of tuple t with tuple u



 c_4

 C_5

 b_1

 a_1

 a_1

• We want $R \div S$



$$\begin{array}{c} c_3\\ c_4\\ c_5 \end{array}$$

• We want
$$R \div S$$



 a_1

 b_1



 c_4

 C_5

 b_2

 b_1

 a_1

 a_1

• Consider R:





• Consider R:

$$c_1$$

• We want $R \div S$



 C_5

 b_1

 a_1



 $|\mathbf{b}_1|$

 C_5

 a_1

• Consider R:

• We want $R \div S$



 C_5

• We want
$$R \div S$$

 a_1

	mID	title	director	year	length
	1	Shining	Kubrick	1980	146
	2	Player	Altman	1992	146
ies.	3	Chinatown	Polanski	1974	131
105.	4	Repulsion	Polanski	1965	143
	5	Star Wars IV	Lucas	1977	126
	6	American Graffiti	Lucas	1973	110
	7	Full Metal Jacket	Kubrick	1987	156

Movies:

	aID	aName
	1	Jack Nicholson
Artists:	2	Harrison Ford
	3	Stone
	4	Fisher

mID	aID	character
1	1	Jack Torrance
3	1	Jake 'J.J.' Gittes
1	3	Delbert Grady
5	2	Han Solo
6	2	Bob Falfa
5	4	Princess Leia Organa

• We want to find "actors who played in **all** Lucas's movies".

Roles:

Strategy:

 \bullet Let S be the set of all Lucas's movies:

 $\pi_{\mathrm{mID}}(\sigma_{\mathrm{director}='\mathrm{Lucas'}}(\mathrm{Movies}))$

• Let R be the set of all tuples (mID, aID) of actors (aID) playing in movies:

 $\pi_{\rm mID, aID}(\rm Roles)$

Let T be the set of actors who played in all Lucas's movies.
How do we obtain it?

Strategy:

 \bullet Let S be the set of all Lucas's movies:

 $\pi_{\mathrm{mID}}(\sigma_{\mathrm{director}='\mathrm{Lucas'}}(\mathrm{Movies}))$

• Let R be the set of all tuples (mID, aID) of actors (aID) playing in movies:

$\pi_{\mathrm{mID,aID}}(\mathrm{Roles})$

- \bullet Let T be the set of actors who played in all Lucas's movies.
 - How do we obtain it?
 - $-R \div S$
 - $-\pi_{\mathrm{mID,aID}}(\mathrm{Roles}) \div \pi_{\mathrm{mID}}(\sigma_{\mathrm{director}='\mathrm{Lucas'}}(\mathrm{Movies})$

 \bullet Recall: S is the set of all Lucas's movies:

 $\pi_{\mathrm{mID}}(\sigma_{\mathrm{director}='\mathrm{Lucas'}}(\mathrm{Movies}))$

Are the 2 queries below semantically equivalent?
– i.e., do they return the same result?

1. $\pi_{\text{mID,aID}}(\text{Roles}) \div \text{S}$

2. $\pi_{\text{aID}}(\text{Roles} \div \text{S})$

More on Relational Algebra Expressions

- A basic expression in the relational algebra consists of either one of the following:
 - \circ A relation in the database
 - \circ A constant relation

More on Relational Algebra Expressions

- Let E_1 and E_2 be RA expressions;
- the following are all RA expressions:
 - $\circ E_1 \cap E_2$
 - $\circ \ E_1 \cup E_2$
 - $\circ E_1 E_2$
 - $\circ E_1 \times E_2$
 - $\circ E_1 \bowtie E_2$
 - $\circ E_1 \div E_2$
 - \circ $\sigma_{\rm cond}(E_1)\text{,}$ where cond is a condition on attributes of E_1
 - \circ $\pi_{\mathrm{attr}}(E_1)\text{,}$ where attr is a subset of attributes in E_1
 - $\circ \ \rho_{\text{new} \leftarrow \text{old}}(\mathbf{E}_1)$