

CS422 Principles of Database Systems

Relational Algebra

Chengyu Sun
California State University, Los Angeles

Adapted from Jeffrey Ullman's lecture notes at
<http://www-db.stanford.edu/~ullman/dscb.html>

Overview

- ◆ Operators
 - Core and Extended
- ◆ Set vs. Bag
- ◆ Express constraints in relational algebra

Relational Algebra

- ◆ Basis for SQL
- ◆ Operands
 - Relations
- ◆ Operators
 - *Core*: Set operators, SPJ, rename
 - *Extended*: duplicate elimination, aggregation, grouping, sorting, outer-join

Selection

- ◆ $\sigma_C(R)$ or $\text{SELECT}_C(R)$
 - Choose *rows* of R that satisfies condition C
 - C is a boolean expression of *constants* and R's *attributes*

Selection Example

Sells

bar	beer	price
Joe's	Bud	2.50
Joe's	Miller	2.75
Sue's	Bud	2.50
Sue's	Miller	3.00

$\sigma_{\text{bar}=\text{"Joe's"}}(\text{Sells})$

bar	beer	price
Joe's	Bud	2.50
Joe's	Miller	2.75

Projection

- ◆ $\pi_L(R)$ or $\text{PROJ}_L(R)$
 - Choose *columns* of R
 - L is a list of R's attributes
 - Eliminate duplicates in the results (*set semantics*)

Projection Example

Sells

bar	beer	price
Joe's	Bud	2.50
Joe's	Miller	2.75
Sue's	Bud	2.50
Sue's	Miller	3.00

$\pi_{\text{beer,price}}(\text{Sells})$

beer	price
Bud	2.50
Miller	2.75
Miller	3.00

Product

◆ $R_1 \times R_2$

- Concatenate *each* tuple from R_1 with *each* tuple from R_2
- Also called Cross Product or Cartesian Product

Product Example

R_1

A	B
1	2
3	4

$R_1 \times R_2$

$R_1.A$	B	$R_2.A$	C
1	2	5	6
1	2	7	8
1	2	9	10
3	4	5	6
3	4	7	8
3	4	9	10

R_2

A	C
5	6
7	8
9	10

Theta Join

◆ $R_1 \bowtie_{c} R_2$ or $R_1 \text{ JOIN}_c R_2$

- $\sigma_c(R_1 \times R_2)$

Theta Join Example

Bars

name	address
Joe's	Maple St.
Sue's	River Rd.

Sells

bar	beer	price
Joe's	Bud	2.50
Joe's	Miller	2.75
Sue's	Bud	2.50
Sue's	Miller	3.00

$\text{Bars} \bowtie_{\text{Bars.name} = \text{Sells.bar}} \text{Sells}$

name	address	bar	beer	price
Joe's	Maple St.	Joe's	Bud	2.50
Joe's	Maple St.	Joe's	Miller	2.75
Sue's	River Rd.	Sue's	Bud	2.50
Sue's	River Rd.	Sue's	Miller	3.00

Natural Join

◆ $R_1 \bowtie R_2$ or $R_1 \text{ JOIN} R_2$

- Implies equality condition of the *attributes with the same name*
- Only one column from each pair of equated attributes is kept in results

Natural Join Example

Bars

bar	address
Joe's	Maple St.
Sue's	River Rd.

Sells

bar	beer	price
Joe's	Bud	2.50
Joe's	Miller	2.75
Sue's	Bud	2.50
Sue's	Miller	3.00

Bars \bowtie Sells

bar	address	beer	price
Joe's	Maple St.	Bud	2.50
Joe's	Maple St.	Miller	2.75
Sue's	River Rd.	Bud	2.50
Sue's	River Rd.	Miller	3.00

Rename

- ◆ $\rho_{S(A_1, A_2, \dots, A_n)}(R)$ or $RENAME_{S(A_1, A_2, \dots, A_n)}(R)$
 - Rename a relation and its attributes
 - Rename a relation only: $\rho_S(R)$

Rename Example

Bars

name	address
Joe's	Maple St.
Sue's	River Rd.

$\rho_{\text{Bars}(\text{bar}, \text{address})}(\text{Bars})$

bar	address
Joe's	Maple St.
Sue's	River Rd.

Set Operators

- ◆ Union: \cup
- ◆ Intersection: \cap
- ◆ Difference: $-$
- ◆ Two relations must have the same schemas
 - Same number of attributes
 - Same attribute names
 - Same attribute types

Set Operator Examples

<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>2</td><td>3</td></tr> </tbody> </table>	A	B	1	2	2	3	\cup	<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>3</td><td>4</td></tr> </tbody> </table>	A	B	1	2	3	4	\Rightarrow	<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>2</td><td>3</td></tr> <tr><td>3</td><td>4</td></tr> </tbody> </table>	A	B	1	2	2	3	3	4
A	B																							
1	2																							
2	3																							
A	B																							
1	2																							
3	4																							
A	B																							
1	2																							
2	3																							
3	4																							

<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>2</td><td>3</td></tr> </tbody> </table>	A	B	1	2	2	3	\cap	<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>3</td><td>4</td></tr> </tbody> </table>	A	B	1	2	3	4	\Rightarrow	<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> </tbody> </table>	A	B	1	2
A	B																			
1	2																			
2	3																			
A	B																			
1	2																			
3	4																			
A	B																			
1	2																			

<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>2</td><td>3</td></tr> </tbody> </table>	A	B	1	2	2	3	$-$	<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>3</td><td>4</td></tr> </tbody> </table>	A	B	1	2	3	4	\Rightarrow	<table border="1" style="border-collapse: collapse; text-align: center;"> <thead> <tr><th>A</th><th>B</th></tr> </thead> <tbody> <tr><td>2</td><td>3</td></tr> </tbody> </table>	A	B	2	3
A	B																			
1	2																			
2	3																			
A	B																			
1	2																			
3	4																			
A	B																			
2	3																			

Expressions (Queries)

- ◆ Find the name of the bars that are either on Maple St. or sells Bud for less than 3 dollars.

$\pi_{\text{bar}} \sigma_{\text{address}=\text{"Maple St." OR (beer=\text{"Bud" AND price} < 3.0)} (\rho_{\text{Bars}(\text{bar}, \text{address})}(\text{Bars}) \bowtie \text{Sells})$

◆ Relational Operator Precedence

- Unary operators
- Binary operators
- Set operators
 - Intersection
 - Union, Difference

Linear Notation for Expressions

$R(\text{bar}, \text{address}) := \text{Bars}$
 $S := R \triangleright \triangleleft \text{Sells}$
 $T := \sigma_{\text{address}=\text{"Maple St." OR (beer=\text{"Bud"} AND price < 3.0)}(S)$
 $\text{Ans} := \pi_{\text{bar}}$

Extended Algebra

- ◆ Eliminate duplicates
- ◆ Sort tuples
- ◆ Extended projection
- ◆ Grouping and aggregation
- ◆ Outerjoin

Duplicate Elimination

- ◆ $\delta(R)$ or DELTA(R)

A	B
1	2
1	2
2	3
2	3
3	4
2	3

A	B
1	2
2	3
3	4

Sorting

- ◆ $\tau_L(R)$ or TAU_L(R)
- n L is a list of R's attributes

A	B	C
1	4	2
1	3	3
1	2	3
2	3	3
3	2	4
1	2	3

A	B	C
1	4	2
1	3	3
1	2	3
1	2	3
2	3	3
3	2	4
3	2	4

Extended Projection

- ◆ $\pi_L(R)$ or PROJ_L(R)
- n Allow arithmetic expressions of R's attributes in L

A	B
1	2
3	4

A	A+B	A1
1	3	1
3	7	3

Aggregation Operators

- ◆ Aggregation operators apply to *column(s)* of a relation and produces a single result
- n SUM, AVG, COUNT, MIN, MAX

A	B
1	2
3	4

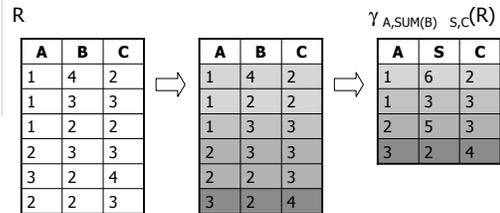
SUM(A) = 4
AVG(A) = 2
COUNT(A) = 2
MIN(A) = 1
MAX(A) = 3

Grouping

◆ $\gamma_L(R)$ or $\text{GAMMA}_L(R)$

- n L is a list of elements that are
 - w Individual attributes (*grouping attributes*)
 - w $\text{AGG}(A)$ where AGG is one of the aggregation operators and A is an attributes

Grouping Example

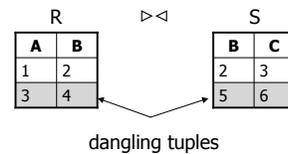


Outer Joins

- ◆ $\triangleright \circ \triangleleft$ or OUTERJOIN (*full outer join*)
- ◆ $\triangleright \circ \triangleleft_L$ or LEFTJOIN (*left outer join*)
- ◆ $\triangleright \circ \triangleleft_R$ or RIGHTJOIN (*right outer join*)

Dangling Tuples

- ◆ When two relations join, a tuple is said to be *dangling* if it does not match any tuple in the other relation.



Outer Join Examples

$R \triangleright \circ \triangleleft S$	$R \triangleright \circ \triangleleft_L S$	$R \triangleright \circ \triangleleft_R S$																																				
<table border="1" style="display: inline-table;"> <thead> <tr><th>A</th><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td><td>3</td></tr> <tr><td>3</td><td>4</td><td>NULL</td></tr> <tr><td>NULL</td><td>5</td><td>6</td></tr> </tbody> </table>	A	B	C	1	2	3	3	4	NULL	NULL	5	6	<table border="1" style="display: inline-table;"> <thead> <tr><th>A</th><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td><td>3</td></tr> <tr><td>3</td><td>4</td><td>NULL</td></tr> <tr><td>NULL</td><td>5</td><td>6</td></tr> </tbody> </table>	A	B	C	1	2	3	3	4	NULL	NULL	5	6	<table border="1" style="display: inline-table;"> <thead> <tr><th>A</th><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td><td>3</td></tr> <tr><td>NULL</td><td>5</td><td>6</td></tr> <tr><td>3</td><td>4</td><td>NULL</td></tr> </tbody> </table>	A	B	C	1	2	3	NULL	5	6	3	4	NULL
A	B	C																																				
1	2	3																																				
3	4	NULL																																				
NULL	5	6																																				
A	B	C																																				
1	2	3																																				
3	4	NULL																																				
NULL	5	6																																				
A	B	C																																				
1	2	3																																				
NULL	5	6																																				
3	4	NULL																																				

Summary of Operators

\cup	UNION	δ	DELTA
\cap	INTERSECTION	τ	TAU
$-$	DIFFERENCE	γ	GAMMA
σ_c	SELECT_c	π_L	PROJ_L
π_L	PROJ_L	$\triangleright \circ \triangleleft$	OUTERJOIN
\times	PRODUCT	$\triangleright \circ \triangleleft_L$	LEFTJOIN
$\triangleright \triangleleft$	JOIN	$\triangleright \circ \triangleleft_R$	RIGHTJOIN
$\triangleright \triangleleft_c$	JOIN_c	SUM, AVG, COUNT, MIN, MAX	
ρ	RENAME		

Set vs. Bag

- ◆ Bag (or multiset) allows *duplicates* while set does not
- ◆ SQL is a bag language
 - A relation may contain duplicate tuples
 - Only eliminate duplicate tuples when the query explicitly asks for it
 - Certain operations like projection are much more efficient on bags than sets.

Relational Operators on Bags

- ◆ Most relational operators work the same for both sets and bags – *just keep the duplicates*
- ◆ Union, Intersection, Difference

$$R = \{1, 1, 2, 3, 3, 3, 4\}, S = \{1, 2, 2, 3, 5\}$$

$$R \cup S = ??$$

$$R \cap S = ??$$

$$R - S = ??$$

Bag Laws != Set Laws

- ◆ Some laws hold for both sets and bags
 - E.g. $R \cup S = S \cup R$
- ◆ Some laws do not
 - E.g. $S \cup S = S$

Exercises - Schema

- ◆ Senators(senator_id, senator_name, state, party)
- ◆ Bills(bill_id, bill_name)
- ◆ Votes(senator_id, bill_id, vote)
 - Vote can be *Yes*, *No*, or *Abstain*

Exercises - Queries

- ◆ Find the state which senator Hillary Clinton represents.
- ◆ Find the names of the bills on which Senator Hillary Clinton has voted Yes.
- ◆ Find the number of Republican senators in the Senate.
- ◆ Find the names of the bills which have been passed, i.e. the bills that have more Yes votes than No votes.

Constraints on Relations

- ◆ Constraints
 - not null, unique, primary key, references (foreign key) ...
- ◆ Express constraints using relational algebra
 - $R = \emptyset$
 - $R \subseteq S$

Functional Dependency

- ◆ In Senators, {senator_id} → {senator_name}
- ◆ In Bills, {bill_id} → {bill_name}
- ◆ In Votes, {senator_id, bill_id} → {vote}

Other Constraints

- ◆ Referential Integrity
- ◆ NOT NULL
- ◆ Unique
- ◆ Enumeration
 - E.g. vote can only be *Yes*, *No*, or *Abstain*