



University of
Southampton

Relational Algebra

COMP3211 Advanced Databases

Dr Heather Packer – hp3@ecs.soton.ac.uk

What is a Relational Algebra?

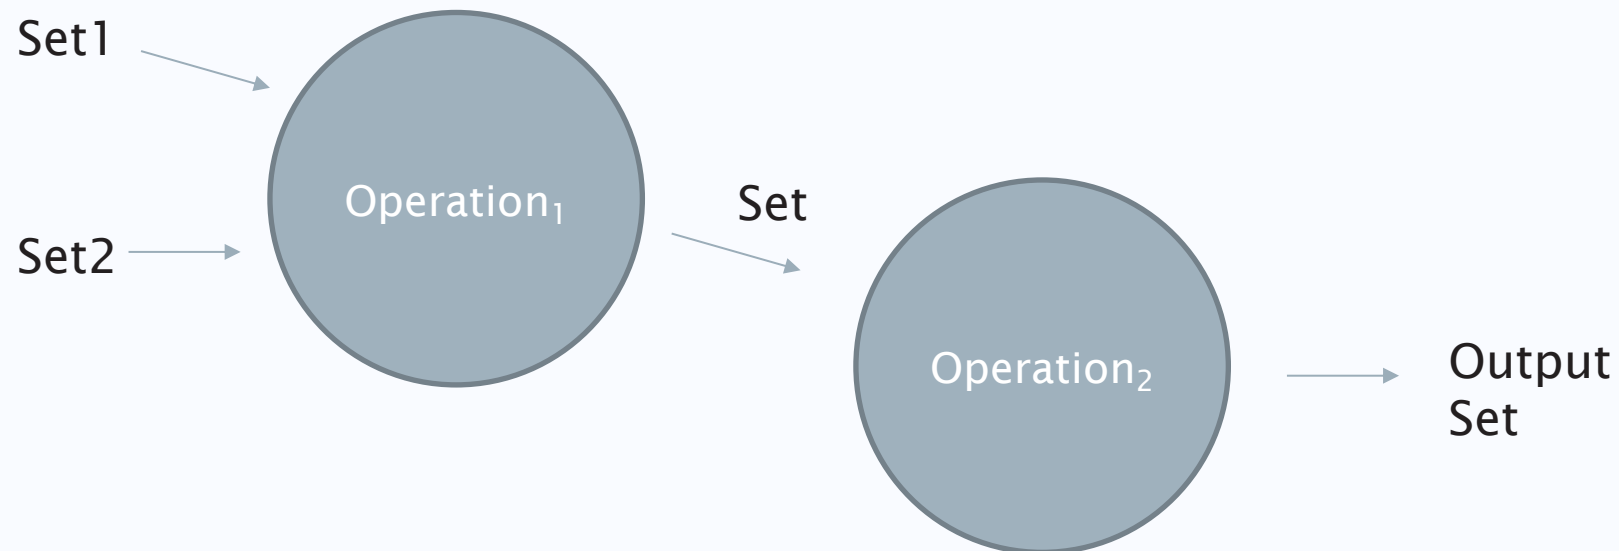
- An Algebra is a mathematical system consisting of:
 - Operands – variables or values from which new values can be constructed
 - Operators – symbols denoting procedures that construct new values from the given values
- A Relational Algebra
 - Operands – relations or variables that represent relations
 - Operators – common things that can be performed on relations

Relational Algebra

- Set Operations
 - \cup Set union
 - \cap Set intersection
 - $-$ Set difference
 - \times Cartesian product
 - \div Set Division
- Relational Database Specific Operations
 - σ Selection
 - π Projection
 - \bowtie Join
- Set Functions
 - sum
 - avg
 - count
 - any
 - max
 - min

Relational Algebra

- The input of a relational algebra operator is one or more relations
- The result of an operation is always a relation
- Necessary to use relational algebra in a cascaded manner



What is a relation?

Relations as Subset of Cartesian Products of Sets

Set S_1 (Student IDs)



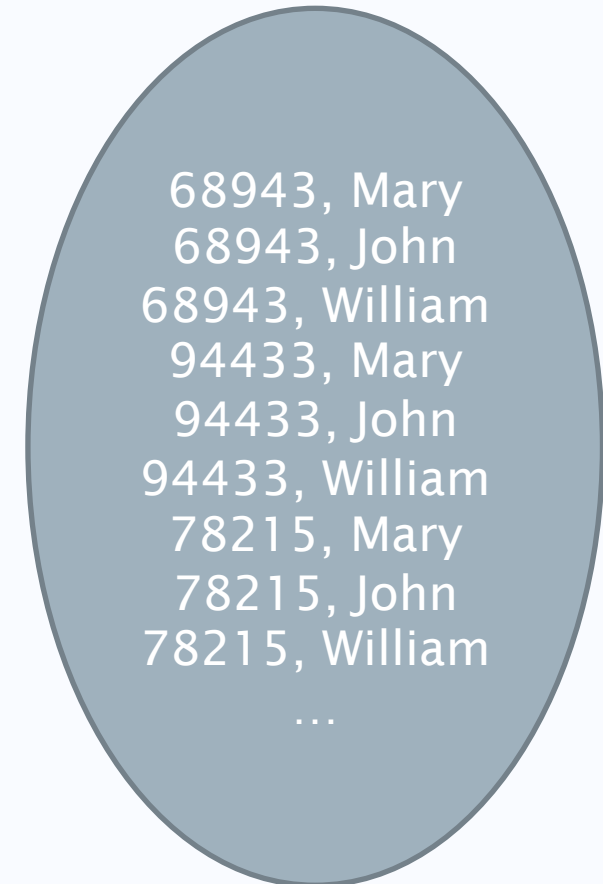
Set S_2 (Student Name)



x

$S_1 \times S_2 =$

=

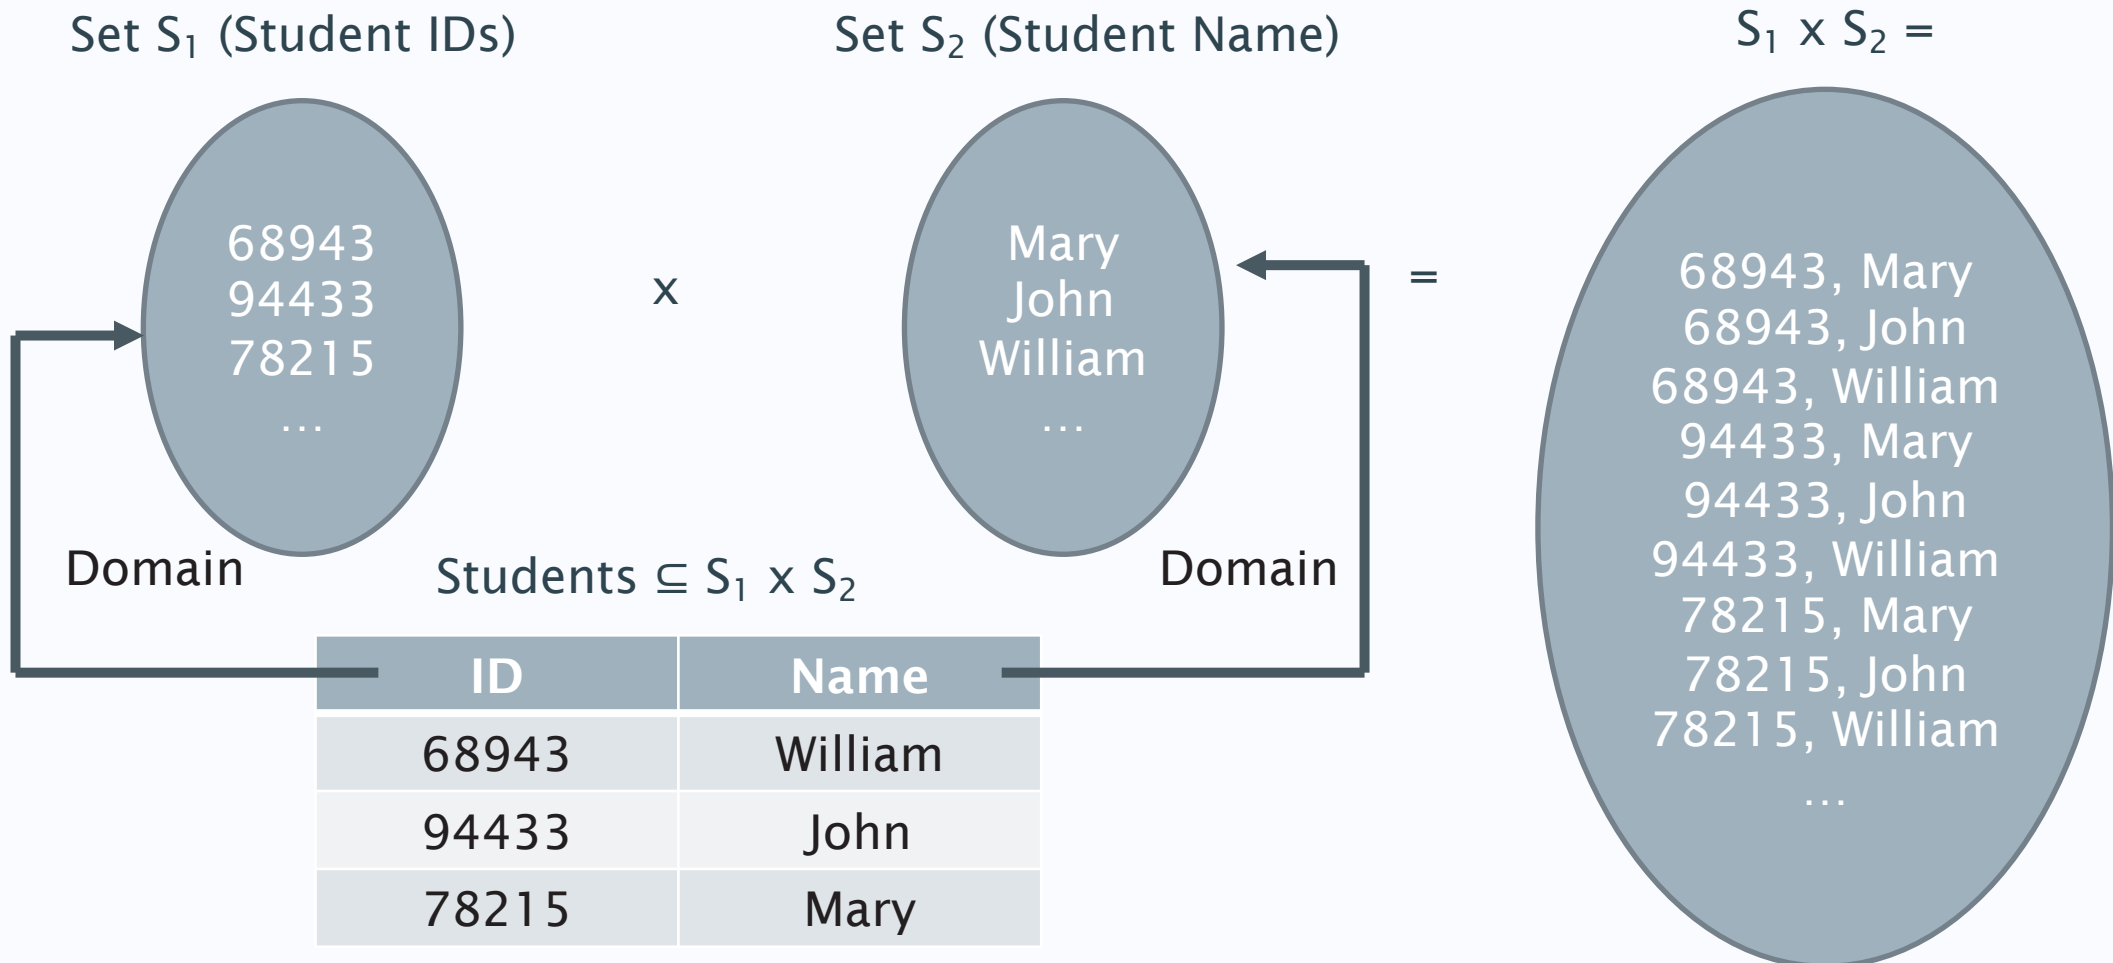


Students $\subseteq S_1 \times S_2$

tuple

ID	Name
68943	William
94433	John
78215	Mary

Relations as Subset of Cartesian Products of Sets



Tuples and attributes

- $R \subseteq D_1 \times D_2 \times \dots \times D_k$ is a set of k-tuples
- R can be represented in a table with k columns
- Values from domain D_i are the only values allowed in the i^{th} column
- The columns in relational data models have names called attributes

Students

ID	Name	DeptID
68943	William	ECS
94433	John	ECS

In the relation Students,
attributes are: ID, Name, DeptID

Properties of Relations

A relation $R(A_1, \dots, A_k)$ has the following properties:

1. Each row represents a k-tuple of R
2. The values of an attribute are all from the same domain
3. Each attribute of each tuple in a relation contains a single atomic value
4. The ordering of rows is immaterial (relations are just sets)
5. All rows are distinct (relations are just sets)
6. Named perspective: the semantics of each column is conveyed its name
7. Named perspective: the ordering of the attributes is not significant
8. Unnamed perspective: the ordering of attributes is significant (we access columns by their positions)

Relational Operators

Relational Algebra

- Set Operations
 - \cup Set union
 - $-$ Set difference
 - \times Cartesian product
 - \cap Set intersection
- Relational Specific Operations
 - ρ Renaming
 - σ Selection
 - π Projection
 - \bowtie Join

Set Operations

Union

- The union of two relations is the set of tuples where each tuple appears in either set

$R \cup S$

B	C
1	2
3	4

U

B	C
4	4
5	6

=

B	C
1	2
3	4
4	4
5	6

- $R \cup S = \{(a_1, \dots, a_k) : (a_1, \dots, a_k) \text{ is in } R \text{ or } (a_1, \dots, a_k) \text{ is in } S\}$
- Note: no duplicate tuples

Union

- Some versions of the algebra allow the attributes of R and S to have different names

R			S			R \cup S	
B	C		D	E	=	F	G
1	2	U	4	4		1	2
3	4		5	6		3	4
						4	4
						5	6

- The columns in the result are given new names

Difference

- The difference between R and S is all the tuples that appear in R but not in S

$$R - S$$

B	C
1	2
3	4

 -

B	C
3	4
5	6

 =

B	C
1	2

- $R - S = \{(a_1, \dots, a_k) : (a_1, \dots, a_k) \text{ is in } R \text{ and } (a_1, \dots, a_k) \text{ is not in } S\}$

Cartesian Product

R X S

R			S		=	R X S			
A	B	X	B	C		A	R.B	S.B	C
X	2		sda	1		X	2	sda	1
Y	3		bfd	0		Y	3	sda	1
Z	3					Z	3	sda	1
						X	2	bfd	0
						Y	3	bfd	0
						Z	3	bfd	0

- Notice the clashing attribute names are renamed

Cartesian Product

- Cartesian Product of relations $R \times S$, given R and S having a m -ary and n -ary relation respectively, is calculated by:
 - Every tuple of R is paired with every tuple of S
 - The tuples are concatenated together to give a $(m+n)$ -ary tuple
- Note that the resulting size of the the relation is the product of the size of each relation.
 - The Cartesian Product can get large

- Concretely:

$$R \times S = \{(a_1, \dots, a_m, b_1, \dots, b_n) : (a_1, \dots, a_m) \text{ is in } R \text{ and } (b_1, \dots, b_n) \text{ is in } S\}$$

$$|R \times S| = |R| \times |S|$$

Set Operation Rules

- Union
 - Commutativity **does** hold because order within a set is unimportant : $R \cup S = S \cup R$
 - Associativity **does** hold: $R \cup (S \cup T) = (R \cup S) \cup T$
- Difference:
 - Commutativity **does not** hold for Difference: $R - S \neq S - R$ (except if $R=S$)
 - Associativity **does not** hold for Difference: $R - (S - T) \neq (R - S) - T$
- Cartesian Product:
 - Note: Ordering of attributes is important here
 - Commutativity **does not** hold: $R \times S \neq S \times R$
 - Associativity **does** hold: $R \times (S \times T) = (R \times S) \times T$
 - Distributivity across Union **does** hold: $R \times (S \cup T) = (R \times S) \cup (R \times T)$

Derived Operations

- Some operations can be derived by other relational algebra operations
- Intersection:

$$R \cap S = R - (R - S)$$

R	S		R - S																	
<table border="1" style="border-collapse: collapse; width: 100%;"> <thead> <tr><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>3</td><td>4</td></tr> </tbody> </table>	B	C	1	2	3	4	-	<table border="1" style="border-collapse: collapse; width: 100%;"> <thead> <tr><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>3</td><td>4</td></tr> <tr><td>5</td><td>6</td></tr> </tbody> </table>	B	C	3	4	5	6	=	<table border="1" style="border-collapse: collapse; width: 100%;"> <thead> <tr><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> </tbody> </table>	B	C	1	2
B	C																			
1	2																			
3	4																			
B	C																			
3	4																			
5	6																			
B	C																			
1	2																			
R	R - S		R ∩ S																	
<table border="1" style="border-collapse: collapse; width: 100%;"> <thead> <tr><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> <tr><td>3</td><td>4</td></tr> </tbody> </table>	B	C	1	2	3	4	-	<table border="1" style="border-collapse: collapse; width: 100%;"> <thead> <tr><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>1</td><td>2</td></tr> </tbody> </table>	B	C	1	2	=	<table border="1" style="border-collapse: collapse; width: 100%;"> <thead> <tr><th>B</th><th>C</th></tr> </thead> <tbody> <tr><td>3</td><td>4</td></tr> </tbody> </table>	B	C	3	4		
B	C																			
1	2																			
3	4																			
B	C																			
1	2																			
B	C																			
3	4																			

Relational Specific Operations

Renaming Operator ρ

- The renaming operator takes one relation and changes attribute names

$$\rho_{B / A} (S)$$

- Returns a relation with schema identical to S but the attribute name A has been replaced by B
- Rename more than one attributes using ‘,’ in the subscript
- E.g. let S with schema $S(A, B, C, D, E)$

$$\rho_{X/B, Y/D} (S)$$

- Creates a relation with schema $S(A, X, C, Y, E)$

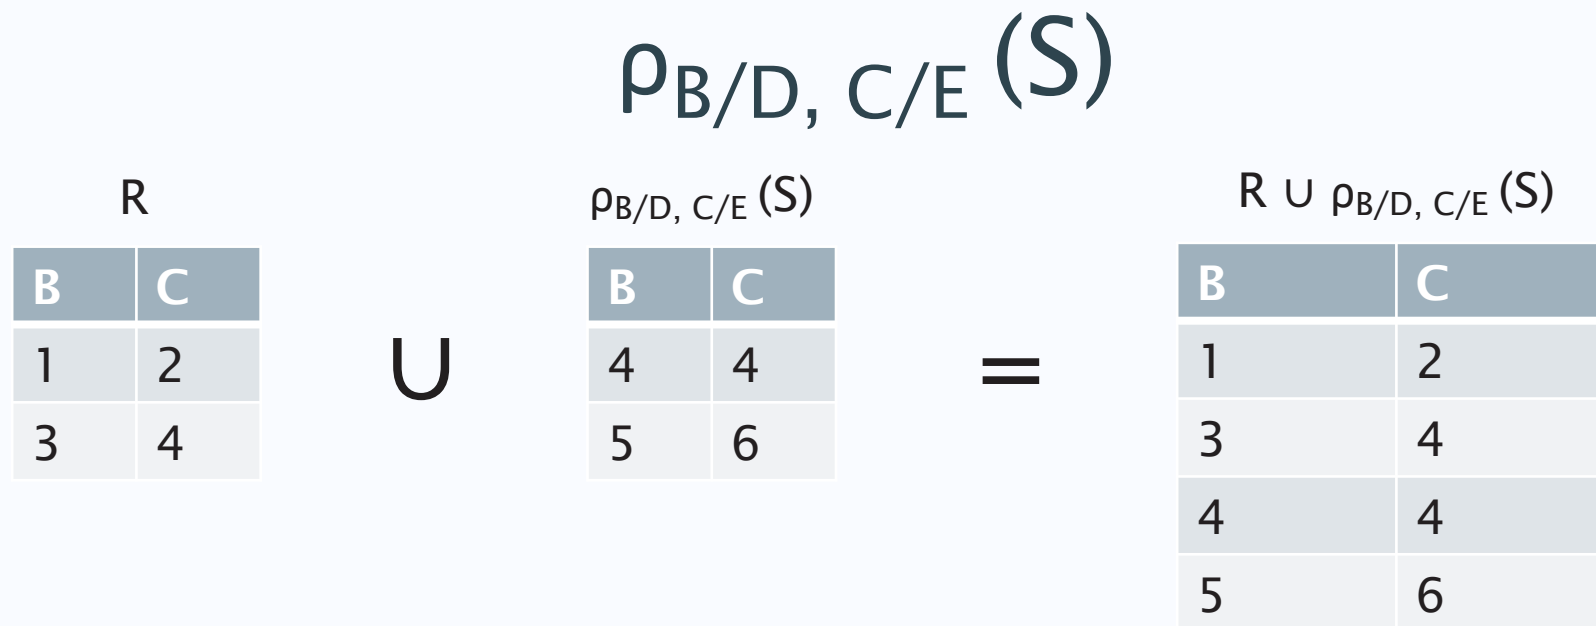
Renaming Operator ρ Example

- If the relational algebra variant needs same attribute names for union
- Then use the renaming operator

R			S		
B	C		D	E	
1	2	U	4	4	=
3	4		5	6	

Renaming Operator ρ Example

- If the relational algebra variant needs same attribute names for union
- Then use the renaming operator



Projection Operator π

- The Projection Operator removes changes what attributes appear in the relation

$$\pi_L (R)$$

- It removes all columns whose attributes do not appear in the list L
- Columns may be re-arranged according to the order in the list
- Any duplicate rows are also eliminated

Projection Operator π Example

Sells

Shop	Food	Price	Units
Union	Apples	0.50	2
Union	Bananas	0.80	4
Co-op	Apples	0.50	5
Co-op	Peaches	0.75	3
Costa	Bananas	0.90	1
Costa	Peaches	1.10	1

$\pi_{\text{Food,Price}}(\text{Sells})$

Food	Price
Apples	0.50
Bananas	0.80
Peaches	0.75
Bananas	0.90
Peaches	1.10

Note: only one entry for apples

Extended Projection

- Some algebras extend projection to allow arbitrary expressions involving attributes

$$\pi_{B+C \rightarrow A}(R)$$

- Arithmetic on attributes eg $B + C \rightarrow A$
- Allows duplicate occurrences of same attribute eg:

$$\pi_{B,B} \left(\begin{array}{|c|c|} \hline B & C \\ \hline 1 & 2 \\ \hline 3 & 4 \\ \hline \end{array} \right) = \begin{array}{|c|c|} \hline B1 & B2 \\ \hline 1 & 1 \\ \hline 3 & 3 \\ \hline \end{array}$$

Extended Projection Example

Sells

Shop	Food	Price	Units
Union	Apples	0.50	2
Union	Bananas	0.80	4
Co-op	Apples	0.50	5
Co-op	Peaches	0.75	3
Costa	Bananas	0.90	1
Costa	Peaches	1.10	1

$\pi_{\text{Food, Price/Units} \rightarrow \text{Cost}}$ (Sells)

Food	Cost
Apples	0.25
Bananas	0.20
Apples	0.10
Peaches	0.25
Bananas	0.90
Peaches	1.10

Selection Operator

- The Selection Operator returns a subset of the relation where the tuples satisfy a predicate

$$\sigma_{\Theta}(R)$$

- R is a relation and Θ is a condition or predicate
- The condition Θ is an expression built from:
 - Comparison operators =, <, >, \neq , \leq , \geq applied to operands that are constants or attribute names (or positions)
 - The Boolean logic operators \wedge , \vee , \neg applied to basic clauses.

Selection Operator σ Predicates

- Example predicates:
 - Price > 0.70
 - Shop = “Co-op”
 - (Shop = “Co-op”) \wedge (Units > 1)
- We reference columns via attribute names or via their “component number” (position) by using \$
 - Shop = “Co-op”
 - \$1 = “Co-op”
 - \$4 > 1

Selection Operator σ Example

Sells

Shop	Food	Price	Units
Union	Apples	0.50	2
Union	Bananas	0.80	4
Co-op	Apples	0.50	5
Co-op	Peaches	0.75	3
Costa	Bananas	0.90	1
Costa	Peaches	1.10	1

$\sigma_{\text{Food} = \text{"Bananas"}} (\sigma_{\text{Units} > 1} (\text{Sells}))$

Shop	Food	Price	Units
Union	Bananas	0.80	4

This is equivalent to

$\sigma_{\text{Food} = \text{"Bananas"}} \wedge \text{Units} > 1 (\text{Sells})$

Algebraic Laws for the Selection Operation

- Algebraic laws can be useful in query optimization
 - $\sigma_{\theta_1}(\sigma_{\theta_2}(R)) = \sigma_{\theta_1 \wedge \theta_2}(R)$
 - Commutative: $\sigma_{\theta_1}(\sigma_{\theta_2}(R)) = \sigma_{\theta_2}(\sigma_{\theta_1}(R))$
 - $\sigma_{\theta}(R \times S) = \sigma_{\theta}(R) \times S$
 - if θ mentions only attributes of R
 - Cartesian product is expensive so reducing the size of R is beneficial

Sets vs. Multisets

Sets vs. Multisets (Bags)

- A Multiset or Bag is like a Set but elements may appear more than once
- Example:
 - $R = \{1, 3, 3, 6, 6, 6, 7\}$, $S = \{1, 1, 6, 6, 7\}$ are both multisets
- Union of two multisets:
 - $\{1, 3, 3, 6, 6, 6, 7\} \cup \{1, 1, 6, 6, 7\} = \{1, 1, 1, 3, 3, 6, 6, 6, 6, 6, 7, 7\}$
- Difference between two multisets:
 - $\{1, 3, 3, 6, 6, 6, 7\} - \{1, 1, 6, 6, 7\} = \{3, 3, 6\}$
- Cartesian product:
 - $\{1, 3, 3\} \times \{1, 1, 6\} = \{\langle 1, 1 \rangle, \langle 1, 1 \rangle, \langle 1, 6 \rangle, \langle 3, 1 \rangle, \langle 3, 1 \rangle, \langle 3, 6 \rangle, \langle 3, 1 \rangle, \langle 3, 1 \rangle, \langle 3, 6 \rangle\}$

Operations on Multisets (Bags)

- Given $\mu(x, B)$, defined as the number of occurrences of x in multiset B
- Union $R \cup S$
 - $\mu(t, R \cup S) = \mu(t, R) + \mu(t, S)$ for all t in R and S
- Difference $R - S$
 - $\mu(t, R - S) = \max\{\mu(t, R) - \mu(t, S), 0\}$ for all t in R and S
- Intersection $R \cap S$
 - $\mu(t, R \cap S) = \min\{\mu(t, R), \mu(t, S)\}$ for all t in R and S
- Cartesian Product:
 - $\mu(tt', R \times S) = \mu(t, R) * \mu(t', S)$ for all t in R and t' in S

Operations on Multisets (Bags)

- Projection:
 - If R is a multiset, then $\pi_x(R)$ is also a multiset
 - If R is a set, then $\pi_x(R)$ **may be** a multiset
- Selection:
 - If R is a multiset, then the selection $\sigma_\theta(R)$ may be a set
 - If R is a set, then the selection $\sigma_\theta(R)$ is a set

Multisets in SQL

- SQL is multiset based
- Efficiency:
 - Duplicate elimination may take quadratic time
 - For example, after a projection
- Necessity:
 - Eliminating duplicates might result in information loss/errors (e.g., in computing averages)
- How to eliminate duplicates in SQL? Use the DISTINCT keyword

```
SELECT DISTINCT <attribute list>
```

```
FROM <relation list>
```

**Next Lecture:
Relational Algebra 2**