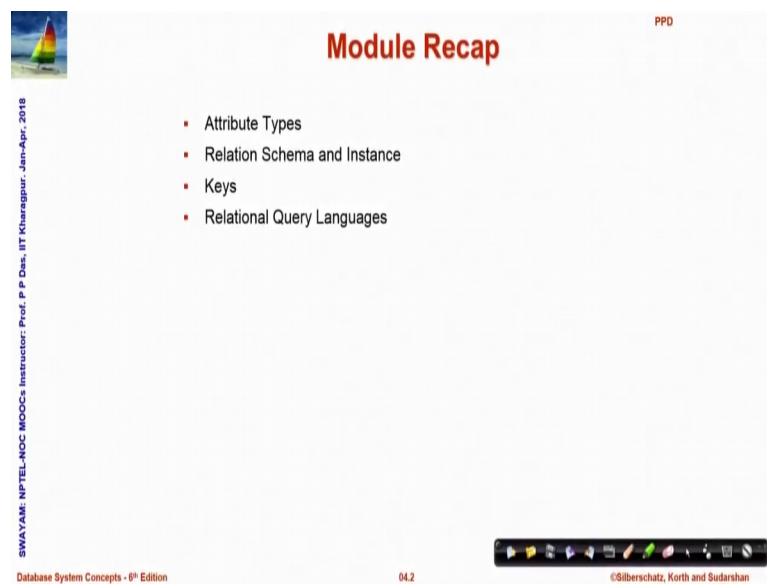


Database Management System
Prof. Partha Pratim Das
Department of Computer Science & Engineering
Indian Institute of Technology, Kharagpur

Lecture -04
Introduction to Relational Model/2

Welcome to module 5 of database management systems. In the previous module, we started discussions on introducing relational model we will conclude that in this module.

(Refer Slide Time: 00:36)



The slide is titled "Module Recap" in red text. It features a bulleted list of topics: Attribute Types, Relation Schema and Instance, Keys, and Relational Query Languages. The slide includes a small sailboat icon in the top left corner and a vertical text label on the left side that reads "SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Kharagpur, Jan-Apr, 2018". The bottom of the slide contains a footer with "Database System Concepts - 9th Edition", the slide number "04.2", and the copyright notice "©Silberschatz, Korth and Sudarshan". A small "PPD" label is in the top right corner, and a navigation bar with various icons is at the bottom right.

Module Recap

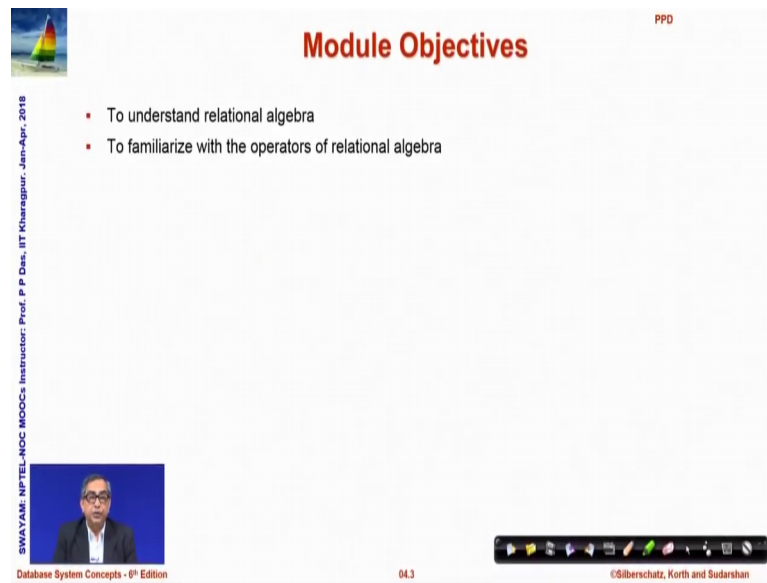
- Attribute Types
- Relation Schema and Instance
- Keys
- Relational Query Languages

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Kharagpur, Jan-Apr, 2018

Database System Concepts - 9th Edition 04.2 ©Silberschatz, Korth and Sudarshan

So, in the last module we have talked about attributes relational schemas and instances in mathematical form and very importantly we have tried to introduce discuss about the concept of keys.

(Refer Slide Time: 00:53)



PPD

Module Objectives

- To understand relational algebra
- To familiarize with the operators of relational algebra

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Kharagpur, Jan-Apr, 2018

Database System Concepts - 9th Edition

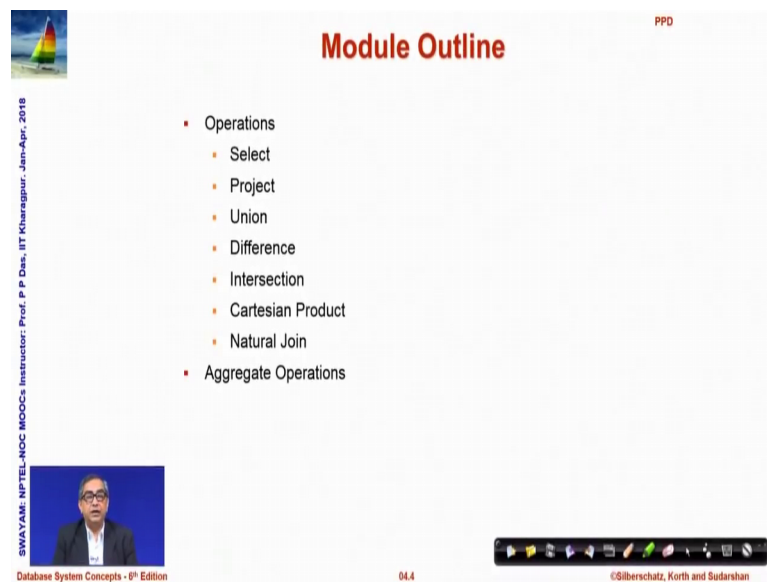
04.3

©Silberschatz, Korth and Sudarshan

This slide, titled 'Module Objectives', features a small sailboat icon in the top left corner. It lists two objectives: 'To understand relational algebra' and 'To familiarize with the operators of relational algebra'. A vertical text on the left identifies the instructor as Prof. P. P. Das from IIT Kharagpur. The bottom of the slide includes the book title 'Database System Concepts - 9th Edition', the slide number '04.3', and the copyright notice '©Silberschatz, Korth and Sudarshan'. A small video inset of the professor is visible in the bottom left, and a navigation bar is at the bottom right.

In this module we, will try to understand more on the relational algebra and familiarize with the operations of relational algebra.

(Refer Slide Time: 01:00)



PPD

Module Outline

- Operations
 - Select
 - Project
 - Union
 - Difference
 - Intersection
 - Cartesian Product
 - Natural Join
- Aggregate Operations

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Kharagpur, Jan-Apr, 2018

Database System Concepts - 9th Edition

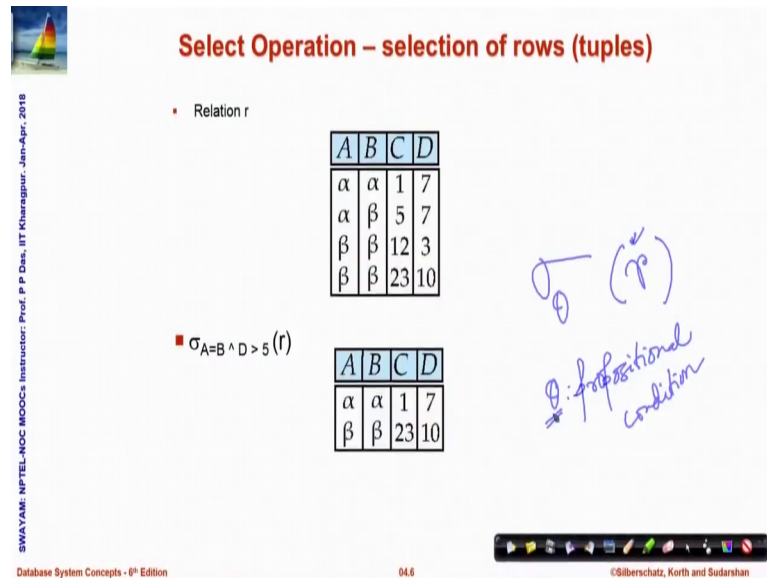
04.4

©Silberschatz, Korth and Sudarshan

This slide, titled 'Module Outline', features a small sailboat icon in the top left corner. It lists the topics to be covered: 'Operations' (which includes Select, Project, Union, Difference, Intersection, Cartesian Product, and Natural Join) and 'Aggregate Operations'. A vertical text on the left identifies the instructor as Prof. P. P. Das from IIT Kharagpur. The bottom of the slide includes the book title 'Database System Concepts - 9th Edition', the slide number '04.4', and the copyright notice '©Silberschatz, Korth and Sudarshan'. A small video inset of the professor is visible in the bottom left, and a navigation bar is at the bottom right.

So, these are the different operations that we will look at select project union and so on. Some of them are simple set theoretic operations some are newly defined operations and we look at the aggregators.

(Refer Slide Time: 01:18)



Select Operation – selection of rows (tuples)

- Relation r

A	B	C	D
α	α	1	7
α	β	5	7
β	β	12	3
β	β	23	10

$\sigma_{A=B \wedge D > 5}(r)$

A	B	C	D
α	α	1	7
β	β	23	10

Handwritten notes:
Q (r)
Q: propositional condition

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Khargapur, Jan-Apr, 2018

Database System Concepts - 9th Edition 04.6 ©Silberschatz, Korth and Sudarshan

So, relational operators, so what is a relation as we have seen already? A relation is nothing but a table. It has a set of columns and it has a set of rows or records that fill up data according to those columns. Select is an operation, which chooses a subset of rows from a relation based on a certain condition. So, it is written in terms of in relational algebra we write it with a notion of notation of sigma and following a parenthesis we put the name of the relation.

So, we say we are selecting from the relation r and then we put a condition here, which is a propositional condition. So, if for this all rows of r will be checked if a row will satisfy this condition. Then it will be included in the result, if it does not satisfy the condition, then it will not be included in the result.

(Refer Slide Time: 02:53).

Select Operation – selection of rows (tuples)

Relation r

A	B	C	D
α	α	1	7
α	β	5	7
β	β	12	3
β	β	23	10

$\theta: A=B \wedge D > 5$

$\sigma_{A=B \wedge D > 5}(r)$

A	B	C	D
α	α	1	7
β	β	23	10

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Khargpur, Jan-Apr, 2018

Database System Concepts - 9th Edition

04.6

©Silberschatz, Korth and Sudarshan

So, let us look at this example. So, our condition is sorry let this put this back. So, our condition theta is A is equal to B and D is greater than 5. So, you are saying that that any row to be selected, the value of it is A attribute should equal the value of it is B attribute and when that happens the value of it is D attribute must be greater than 5. So, you can easily if we look through this we can easily say by the first condition a equal to B you can say that this row does not satisfy this condition.

Because, A is alpha and B is beta whereas, these 3 rows satisfy because A is equal to beta. Then you again look at D we find that this D is less than 5. We say this also does not satisfy, because it fails the second condition both have to hold. So, we finally, come to that this record and this record are the selected record in the result alpha 1 7 alpha 1 7 beta beta 23 10 beta beta 23 10. In both of these records A is equal to B in both of these records D is greater than 5.

So, selection is a process is the operation which selects a subset of rows from a table, from a relation and creates a new relation. Based on a selection condition that must be true for all the rows for all the records that have been selected, but the set of columns do not change they remain the same.

(Refer Slide Time: 04:57)

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P P Das, IIT Khargpur, Jan-Apr, 2018

Select Operation – selection of rows (tuples)

- Relation r

A	B	C	D
α	α	1	7
α	β	5	7
β	β	12	3
β	β	23	10

$\sigma_{A=B \wedge D > 5}(r)$

A	B	C	D
α	α	1	7
β	β	23	10

Handwritten notes: $\sigma_{A \neq B}(x)$ and a table with columns A, B, C, D containing values α , β , 5, 7.

Database System Concepts - 9th Edition 04.6 ©Silberschatz, Korth and Sudarshan

So, for example, if I if in the contrary if we say, that sigma a not equal to B r then naturally, I will have a relation with fields A B C D and that relation will be only this row. Where that is the only row where a is not equal to b we can.

(Refer Slide Time: 05:26)

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P P Das, IIT Khargpur, Jan-Apr, 2018

Select Operation – selection of rows (tuples)

- Relation r

A	B	C	D
α	α	1	7
α	β	5	7
β	β	12	3
β	β	23	10

$\sigma_{A=B \wedge D > 5}(r)$

A	B	C	D
α	α	1	7
β	β	23	10

Handwritten notes: $\sigma_{C > 0}(y)$ and checkmarks next to the last two rows of relation r.

Database System Concepts - 9th Edition 04.6 ©Silberschatz, Korth and Sudarshan

I can have a selection saying r where C is greater than 0. Naturally, this will satisfy, this will satisfy, this will satisfy, this will satisfy. So, this whole relation would be the result of the selection. So, it is possible that it is not necessary that some rows will have to get eliminated in the result.

(Refer Slide Time: 05:54)

Select Operation – selection of rows (tuples)

Relation r

A	B	C	D
α	α	1	7
α	β	5	7
β	β	12	3
β	β	23	10

$\sigma_{A=B \wedge D > 5}(r)$

A	B	C	D
α	α	1	7
β	β	23	10

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Khargapur, Jan-Apr, 2018

Database System Concepts - 9th Edition

04.6

©Silberschatz, Korth and Sudarshan

I say that this is the selection is D is less than 1 this will fail, this will fail, this will fail this will fail. So, all of them will fail so, it is possible that the result of an operation could be either the whole relation as we saw last time or a null relation which has where none of the record with feature because, none of the record satisfy the condition. So, that is a basic select operation.

(Refer Slide Time: 06:38)

Project Operation – selection of columns (Attributes)

Relation r:

A	B	C
α	10	1
α	20	1
β	30	1
β	40	2

$\Pi_{A,C}(r)$

A	C
α	1
α	1
β	1
β	2

$\Pi_{A,C}(r) = \gamma^0$

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Khargapur, Jan-Apr, 2018

Database System Concepts - 9th Edition

04.7

©Silberschatz, Korth and Sudarshan

Let us move on look at the next one is called a projection operation. So, select chooses a subset of the rows. Projection necessarily chooses projects a set of columns from the

original relation. So, this is quite straightforward to see, it is written in terms of this notation π and then you write the columns that you want in the result of projection. So, we say this is $A \ C$. So, which means basically the column, which is not selected in the projection, you can simply forget about that simply erase it. If you erase that you get this relation and once you get that, please recall that a relation is a set and in a set every element has to be distinct. So, after erasing B this first and the second row have become identical.

So, naturally both of them cannot be there, it will have to be made distinct by erasing any one of them. And hence, they become one row in the result; obviously, I can project on any of the fields single or all the fields also I can do a projection of this $A \ B \ C$ of r of course, that means, in this case that will mean that it is a set which is equal to r anyway, but; obviously, I must have at least 1 column at least one attribute to project on. I cannot project on a null set of attributes because that does not give me a schema.

So, there will have to be some attribute one or more attribute on which I project. So, selection and projection, selection has given me the set of rows to retain and projection has given me what are the columns to retain in the result. And combining them I can do several different operations in a database table which can give me several interesting results. Before proceeding further,

Let us look into some of the typical other operations that relational algebra allows. The next one is union given 2 relations I can take a union this is nothing but a set theoretic union. The 2 relations r and s must have the same set of columns A and B because, if the columns are not same, then the union does not make sense. They because certainly if the columns are different attributes are different their types and type of data values would be different. So, they cannot be put to a same table so when 2 relations have the same set of attributes then their instances can be taken a union of.

(Refer Slide Time: 09:49)

Set difference of two relations

Relations r, s :

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

$r - s$:

A	B
α	1
β	1

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P P Das, IIT Kharagpur, Jan-Apr, 2018

Database System Concepts - 9th Edition

04.9

©Silberschatz, Korth and Sudarshan

So, all records that exist in both these relations will be put together into a single table.
So, here alpha 1 is coming here beta 1 is coming here.

(Refer Slide Time: 10:09)

Union of two relations

Relations r, s :

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

$r \cup s$:

A	B
α	1
α	2
β	1
β	3

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P P Das, IIT Kharagpur, Jan-Apr, 2018

Database System Concepts - 9th Edition

04.9

©Silberschatz, Korth and Sudarshan

In terms of relation, alpha 1 is coming here, alpha 2 is coming here, beta 1 is coming here, beta 3 is coming here and alpha 2 is coming here. You can see that this record alpha 2 exists in both the relations. And by the set theoretic notion of uniqueness, in the union they have to be unquified. So, one of them will be removed, it does not matter because

they are identical anyway. So, for relations having the same set of attributes we can simply make a union of all its records.

(Refer Slide Time: 10:57)

Set difference of two relations

• Relations r, s :

A	B
α	1
α	2
β	1

A	B
α	2
β	3

r s

• $r - s$:

A	B
α	1
β	1

Database System Concepts • 9th Edition

04.9

©Silberschatz, Korth and Sudarshan

So, other the 3rd operation is the a 4th operation is doing a set difference it is works simply as set theoretic difference. Again, the 2 relations must have the same set of attributes and I can do a difference of r minus s which mean that all tuples which exist in r , but do not exist in s will be included. So, this is included because, this is not here, but this is not included because it is in s this is included.

Because, this is this does not exist in the set s so it is the set, so you take the set r and then erase all the records like this which exist in s and you get r minus s . So, if I look into just to recap if I look into the Venn diagram, then this is these are set r minus s which belongs to r , but does not belong to s . So, there is a 4th operation that one can do with the in relational algebra.

(Refer Slide Time: 12:28)

Set intersection of two relations

Relation r , s :

A	B
α	1
α	2
β	1

r

A	B
α	2
β	3

s

$r \cap s$

A	B
α	2

Note: $r \cap s = r - (r - s)$

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Kharagpur, Jan-Apr, 2018

Database System Concepts - 9th Edition

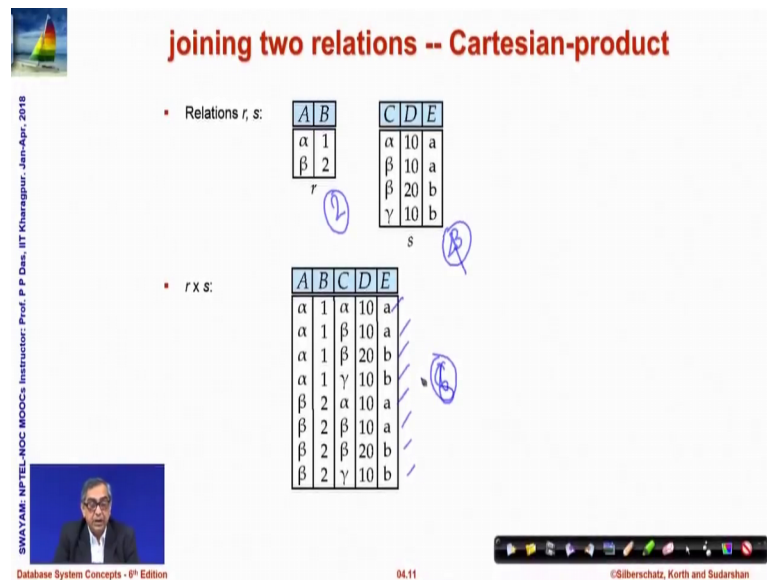
04.10

©Silberschatz, Korth and Sudarshan

Fifth is a set intersection of 2 relations. So, again the 2 relations need to have the same set of attributes. You can take their intersection, which is the record which belongs to both. It is the record that belongs to both of them and as you are aware now, set intersection actually is not a new operation. It is not a fundamental operation because, if I have r , if I have s , then this is r minus s .

Now, if I subtract this r minus s which is this set from r which is this bigger set, then what will be remaining? This is what will be remaining? So, if I subtract r minus s from r then what will remain? Is necessarily the intersection of this is r intersection s . So, set intersection is not a fundamental operation of relational algebra. But, can be used because it can be expressed in terms of set difference.

(Refer Slide Time: 14:03)



joining two relations -- Cartesian-product

- Relations r, s :

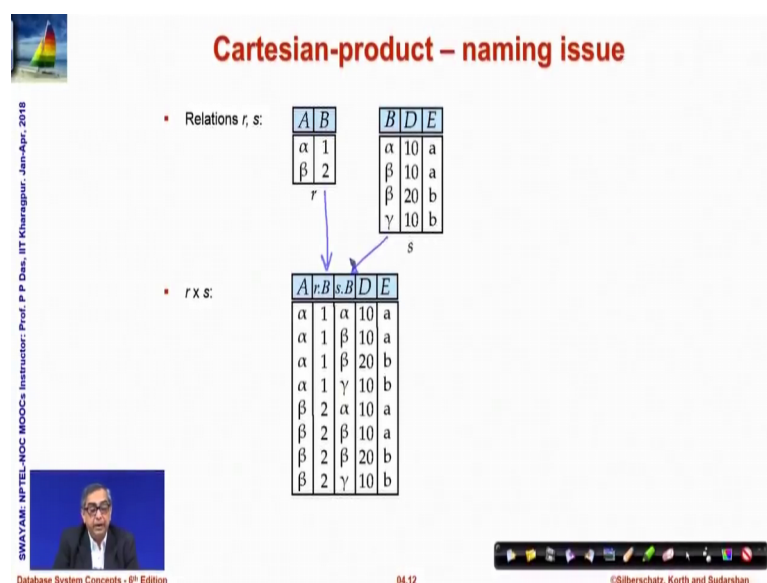
A	B
α	1
β	2

C	D	E
α	10	a
β	10	a
β	20	b
γ	10	b
- $r \times s$:

A	B	C	D	E
α	1	α	10	a
α	1	β	10	a
α	1	β	20	b
α	1	γ	10	b
β	2	α	10	a
β	2	β	10	a
β	2	β	20	b
β	2	γ	10	b

Next comes how can we join 2 different relations which have different set of attributes? So, relation r has A B and relation s has C D E. So, we can take a Cartesian product. So, taking Cartesian product is making all possible combinations. So, necessarily if since this has 2 relations and this has 3 relations. So, this will have 1 2 3 4 5 6 7 8 this has 4 relations. So, these are 8, 8 total all possible pairing of relations of r or records of r and records of s are included. So, that is the Cartesian product all possible combinations. This is this is how we can join 2 relations, but certainly what is a important is something which we will discuss shortly.

(Refer Slide Time: 15:00)



Cartesian-product – naming issue

- Relations r, s :


A	B
α	1
β	2

B	D	E
α	10	a
β	10	a
β	20	b
γ	10	b
- $r \times s$:

A	r.B	s.B	D	E
α	1	α	10	a
α	1	β	10	a
α	1	β	20	b
α	1	γ	10	b
β	2	α	10	a
β	2	β	10	a
β	2	β	20	b
β	2	γ	10	b

Now, in the Cartesian product then issue may happen because there could be attributes which are common between 2 relations. So, if you if 2 attributes are common when you take Cartesian product how do you put their name? Because as with the example show here between r and s the attribute b is common so, how do you take care of that? So, when the, such common names happen then we actually change the name of the attribute with the name of the relation. So, b coming from r will be called r dot b and s be coming from s will be called r dot s.

(Refer Slide Time: 15:46)



Renaming a Table

- Allows us to refer to a relation, (say E) by more than one name.
 $\rho_s(E)$
returns the expression E under the name X
- Relations r

A	B
α	1
β	2
- $r \times \rho_s(r)$

r.A	r.B	s.A	s.B
α	1	α	1
α	1	β	2
β	2	α	1
β	2	β	2

$r \times r$
 $r.A \quad r.B \quad r.A \quad r.B$

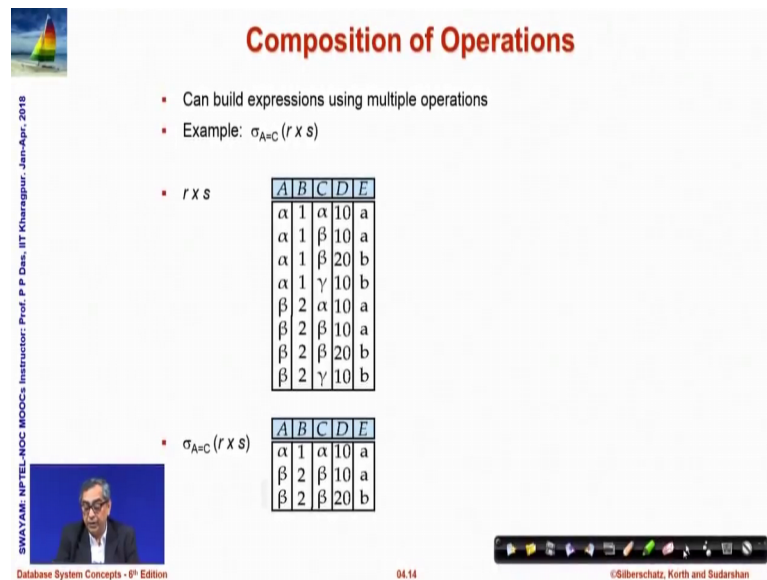
SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Kharagpur, Jan-Apr, 2018
 Database System Concepts - 9th Edition

04.13

©Silberschatz, Korth and Sudarshan

And accordingly, the relational algebra, gives you a way to rename a table and put it is name differently. So, this is given a this is given by this relation is given by this symbol rho. So, you can using a relation r you can actually give it a different name s and do that in terms of. So, with that you can actually because otherwise you cannot compute r cross r. Because, if you try to do r cross r then you will have r dot a r dot b and you will again have r dot a r dot b. So, you are using this to rename r to s and then compute this. So, renaming a table is another feature which is provided of course, it is not a fundamental operation of the algebra, but this is what makes the any kind of Cartesian product possible.

(Refer Slide Time: 16:55)



Composition of Operations

- Can build expressions using multiple operations
- Example: $\sigma_{A=C}(r \times s)$

$r \times s$

	A	B	C	D	E
α	1	α	10	a	
α	1	β	10	a	
α	1	β	20	b	
α	1	γ	10	b	
β	2	α	10	a	
β	2	β	10	a	
β	2	β	20	b	
β	2	γ	10	b	

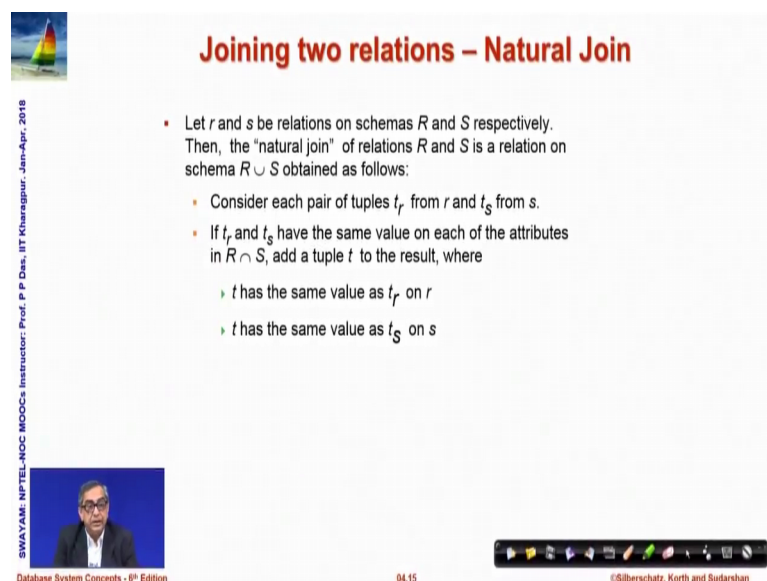
$\sigma_{A=C}(r \times s)$

	A	B	C	D	E
α	1	α	10	a	
β	2	β	10	a	
β	2	β	20	b	

Database System Concepts - 9th Edition 04.14 ©Silberschatz, Korth and Sudarshan

Finally, we can make composition of operations that if for example, what we show here is we have 2 relations r and s you have taken a Cartesian product and then we have taken a selection. So, taken a Cartesian product of r n s to produce the table as you can see the r process and then added a selection based on a equal to c based on that condition. So, all these operations can be combined in multiple different ways to give you really complex relational algebra operations.

(Refer Slide Time: 17:33)



Joining two relations – Natural Join

- Let r and s be relations on schemas R and S respectively. Then, the "natural join" of relations R and S is a relation on schema $R \cup S$ obtained as follows:
 - Consider each pair of tuples t_r from r and t_s from s.
 - If t_r and t_s have the same value on each of the attributes in $R \cap S$, add a tuple t to the result, where
 - t has the same value as t_r on r
 - t has the same value as t_s on s

Database System Concepts - 9th Edition 04.15 ©Silberschatz, Korth and Sudarshan

There is a nice operation which is a derived 1 which can be written in terms of other operations which is called a natural join which we will use very heavily.

(Refer Slide Time: 17:48)

Natural Join Example

Relations r, s:

A	B	C	D
α	1	α	a
β	2	γ	a
γ	4	β	b
α	1	γ	a
δ	2	β	b

B	D	E
1	a	α
3	a	β
1	a	γ
2	b	δ
3	b	ε

Natural Join
 $r \bowtie s$

A	B	C	D	E
α	1	α	a	α
α	1	α	a	γ
α	1	γ	a	α
α	1	γ	a	γ
δ	2	β	b	δ

$\Pi_{A, r, B, C, r, D, E}(\sigma_{r.B = s.B \wedge r.D = s.D}(r \times s))$

α 1 α a 1 α

Let me first show you an example of that, suppose I have 2 relations r and s and what is important is there are some attributes which are common between them. Now, we saw earlier that in Cartesian product, in terms of common attributes. We basically the attributes got renamed in terms of the table name, but this is not what we are looking at in relation natural joint.

What we want to say is, if an attribute is common between 2 2 tables then, while you join them, the records from 2 fields can be joined if their value on that common attribute is same. So, it is what it tries to do is it tries to make a Cartesian product of this 2 tables first take all possible combinations, but then you select only those rows where the values are identical between columns having the same name. So, for example, if you if you if you look into this row and this row. So, what will happen in the Cartesian product I will have this alpha 1 alpha a 1 a alpha.

Let me write it in smaller. So, I am doing a Cartesian product I am looking at this row I am looking at this row.

(Refer Slide Time: 19:39)

Natural Join Example

Relations r, s:

A	B	C	D
α	1	α	a
β	2	γ	a
γ	4	β	b
α	1	γ	a
δ	2	β	b

B	D	E
1	a	α
3	a	β
1	a	γ
2	b	δ
3	b	ε

Natural Join
 $r \bowtie s$

A	B	C	D	E
α	1	α	a	α
α	1	α	a	γ
α	1	γ	a	α
α	1	γ	a	γ
δ	2	β	b	δ

$\Pi_{A, r.B, C, r.D, E}(\sigma_{r.B = s.B \wedge r.D = s.D}(r \times s))$

So, I have A B C D I have B D E and alpha alpha A 1 A alpha. Now, they match on B they match on D. So, I will say this is this will get retained, but in the Cartesian product I will also have the first row of r going with the second row of s alpha 1 alpha A 3 A beta. Here, the B does not match D does not does match, but the B does not match. So, this particular entry will not go in the final result. To take the Cartesian product and you only retain those rows where the values match for the identically named attribute. That is why, so you take the Cartesian product now look at the expression the attributes common attributes are B and C B and D.

So, the B attribute is r dot B and s dot B. So, we say that in the Cartesian product r dot B must equal s dot B. The name is common; the value will have to be same similarly d is a common attribute. So, r dot D value in the r dot D and the value in the s dot D has to be same. So, based on the Cartesian product you do a selection for equality of values on attributes which are identical between the 2 column between the 2 relations between the 2 tables.

This is a final selection, as you do that you get a table where there are 2 Bs r dot Bs dot B there are 2 Ds r dot Ds dot D, but according to this selection for all at all records for all rows the value on r dot B and value on s dot B are same value on r dot D and value on r s dot B are same. Because that is how we have done the selection. So, there is no reason to keep 2 B columns or 2 D columns.

(Refer Slide Time: 22:33)

Natural Join Example

Relations r, s:

A	B	C	D
α	1	α	a
β	2	γ	a
γ	4	β	b
α	1	γ	a
δ	2	β	b

B	D	E
1	a	α
3	a	β
1	a	γ
2	b	δ
3	b	ε

Natural Join

$r \bowtie s$

A	B	C	D	E
α	1	α	a	α
α	1	α	a	γ
α	1	γ	a	α
α	1	γ	a	γ
δ	2	β	b	δ

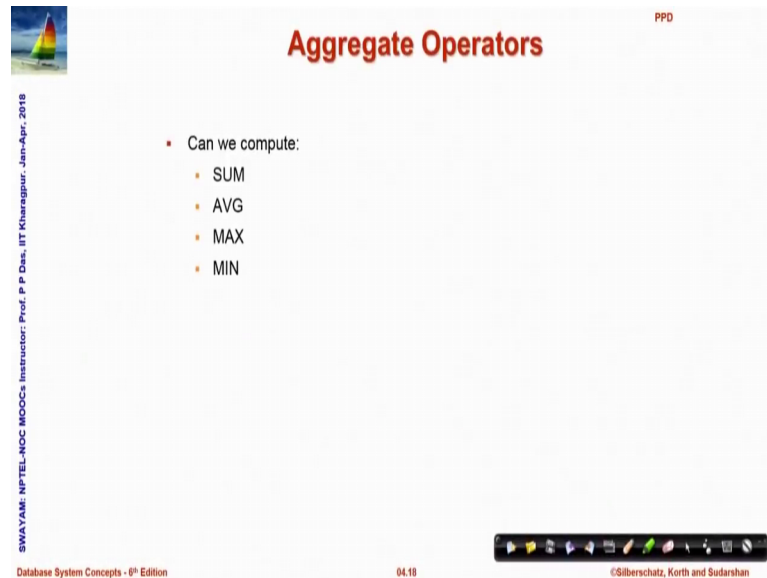
$\Pi_{A, r.B, C, r.D, E}(\sigma_{r.B = s.B \wedge r.D = s.D}(r \times s))$

Database System Concepts - 9th Edition 04.16 ©Silberschatz, Korth and Sudarshan

So now, you project based on a $r \cdot B$ $c r \cdot D$ and e which means that $s \cdot B$ and $s \cdot D$ are left out you do not project them. So, after you have done this projection, you get the final result of the natural join, which has a union of all the attributes that the 2 relations at $A B C D$ and $B D E$ union is $A B C D E$ and you have all those records whose values matched on the common attributes between relation r and relation D . So, you can say that if I now do a selection, if I now do a projection on $A B C$ or rather $A B C D$.

I will get a subset of r if I do a projection on $B D E$ I will get a subset of s . So, this is the natural join operation in relational algebra as you can see this is a derived operation because we could use the Cartesian product selection and projection to get this, but as I tell you we will see more of this when we look at the look at all these different aj query coding, but natural join is one of the most widely used most fundamental relational algebra operation that you will often need beyond selection and projection. So, these were the 6 operations and the important derived operations of relational algebra. Besides that relational algebra has some aggregation operators.

(Refer Slide Time: 24:13)



The slide is titled "Aggregate Operators" in a large, bold, red font. Below the title, there is a bulleted list of operators that can be computed. The list is as follows:

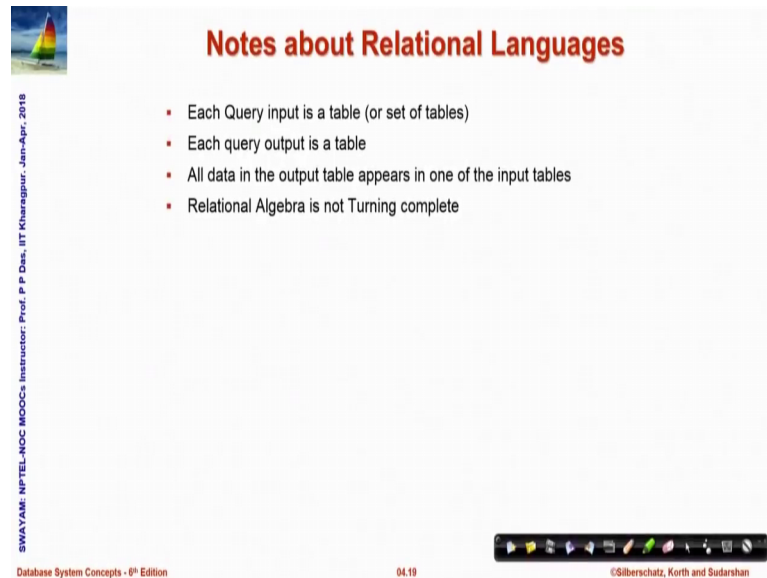
- Can we compute:
 - SUM
 - AVG
 - MAX
 - MIN

The slide also features a small sailboat icon in the top left corner. On the left side, there is vertical text: "SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Kharagpur, Jan-Apr, 2018". At the bottom, there is a footer with "Database System Concepts - 8th Edition" on the left, "04.10" in the center, and "©Silberschatz, Korth and Sudarshan" on the right. A presentation navigation bar is visible at the bottom right.

For example, given a table we could compute the sum of values on a column we could compute average of values, max of values, mean of values. So, these somehow aggregate values of multiple rows on a particular column. And therefore, these are called aggregate operators we will see when we talk about SQL. We will see how this really can be coded in SQL and used, but these are this become very convenient to use because often we will need to know.

If this is I mean these are the instructors. And so, let us see which instructor has a maximum load of courses? How many based on hours or which instructor has? What is the average load on the different instructors and so on? So, in every possible context different aggregate operators are frequently required and they are also available as operators in most of the pure as well as commercial query languages.

(Refer Slide Time: 25:18)



Notes about Relational Languages

- Each Query input is a table (or set of tables)
- Each query output is a table
- All data in the output table appears in one of the input tables
- Relational Algebra is not Turing complete

SWAYAM: NPTEL-NOC MOOCs Instructor: Prof. P. P. Das, IIT Khargpur, Jan-Apr. 2018

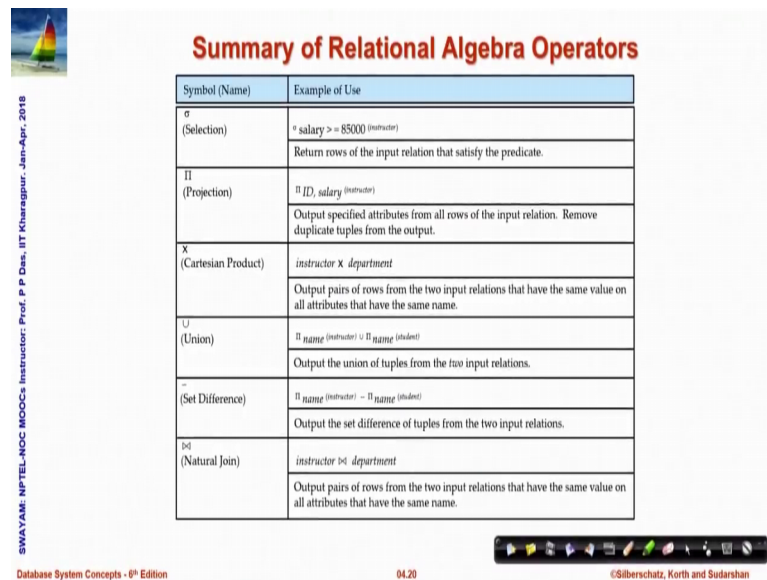
Database System Concepts - 8th Edition 04.19 ©Silberschatz, Korth and Sudarshan

Finally, to note that relational algebra in relational average every query input is a table. And the output is also a table, it is always manipulating one or more tables into a single table that is what we are, I mean in very simple terms that is the way you can look at it. And all data in the output table appears in one of the input tables that is no new data gets generated. It is basically taking, selecting, combining data from different input tables. It does not generate a new data that is that has to be that is if I see that a, we an attribute for a particular row has a value 15.

Then there must be some input table where there is a record where in that field as an attribute value 50. Otherwise, this cannot happen again relational algebra is not turing complete in the sense that, there are algorithms which cannot be coded in relational algebra. We mentioned this earlier too and that is the reason that is a foundational reason of why the SQL language commercial SQL language is based on relational algebra is not turing completely there.

So, we might need to use other programming languages along with the relational algebra coding for solving some of the application problems.

(Refer Slide Time: 26:50)



The slide is titled "Summary of Relational Algebra Operators" in red text. It features a table with two columns: "Symbol (Name)" and "Example of Use". The table lists seven operators: Selection (σ), Projection (π), Cartesian Product (\times), Union (\cup), Set Difference ($-$), and Natural Join (\bowtie). Each operator is accompanied by its symbol, a brief description, and an example query. The slide also includes a small image of a sailboat in the top left corner and a footer with the text "Database System Concepts - 8th Edition" and "©Silberschatz, Korth and Sudarshan".

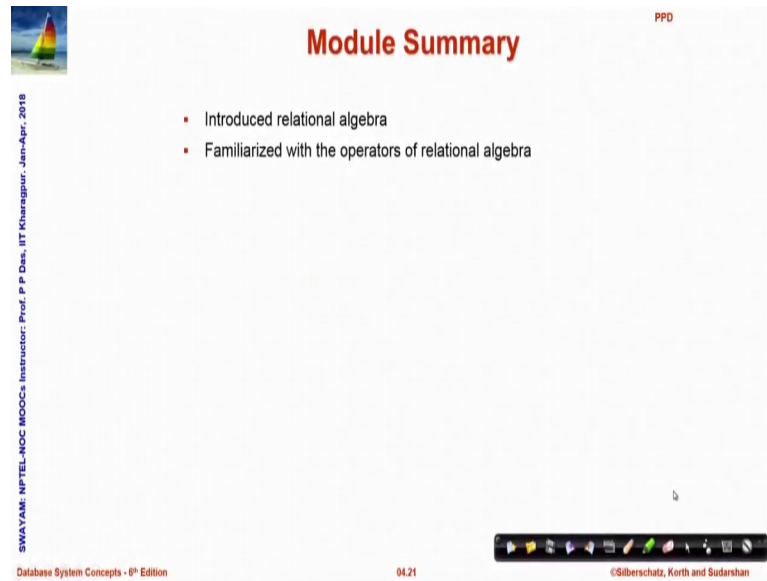
Symbol (Name)	Example of Use
σ (Selection)	$\sigma^{\text{salary} >= 85000}(\text{instructor})$ Return rows of the input relation that satisfy the predicate.
π (Projection)	$\pi_{ID, salary}(\text{instructor})$ Output specified attributes from all rows of the input relation. Remove duplicate tuples from the output.
\times (Cartesian Product)	$\text{instructor} \times \text{department}$ Output pairs of rows from the two input relations that have the same value on all attributes that have the same name.
\cup (Union)	$\pi_{name}(\text{instructor}) \cup \pi_{name}(\text{student})$ Output the union of tuples from the two input relations.
$-$ (Set Difference)	$\pi_{name}(\text{instructor}) - \pi_{name}(\text{student})$ Output the set difference of tuples from the two input relations.
\bowtie (Natural Join)	$\text{instructor} \bowtie \text{department}$ Output pairs of rows from the two input relations that have the same value on all attributes that have the same name.

To summarize this is the, so this table is what you not only should remember, but you should become an expert of in terms of the operators of relational algebra which we will start using very heavily as we start doing the query coding and processing. So, we talked about selection, which takes rows selectively we talked about projection which takes out certain columns of a table. We talked about Cartesian product of 2 relations which make all possible combined relations.

We talked about union of records from 2 tables having identical set of attributes. We talked about set difference which again is the difference of records of one relation from another, given that they have identical set of attributes. We have shown that set difference can be used to also compute set intersection, it is not in the fundamental operation, but is the derived 1 and we have shown a very interesting operation based on Cartesian product selection and projection called natural join where 2 tables can be joined based on 1 or more common attributes they have.

Now, I am sure you have already noted that if I am doing a natural join between 2 tables. Which do not have any common attribute then the result is merely the Cartesian product because the selection around the Cartesian product has no condition to select any of the you know any of the fields any of the rows separately. So, it merely turns out to be a Cartesian product.

(Refer Slide Time: 28:38)



The slide is titled "Module Summary" in a large, bold, red font. It contains two bullet points: "Introduced relational algebra" and "Familiarized with the operators of relational algebra". The slide is part of a presentation, as evidenced by the navigation icons at the bottom. The text "Database System Concepts - 9th Edition" is visible in the bottom left corner, and "©Silberschatz, Korth and Sudarshan" is in the bottom right corner. The slide number "04.21" is also present.

Module Summary

- Introduced relational algebra
- Familiarized with the operators of relational algebra

Database System Concepts - 9th Edition 04.21 ©Silberschatz, Korth and Sudarshan

So, in this module, we have introduced the relational algebra and we have familiarized ourselves with the fundamental and derived operators of relational algebra. Going forward in the next module, will take a deeper look into the relational model and start progressing towards the query design and database design.