

Lecture 4: Updates, Views, Constraints and Other Fun Features

Monday, April 19th, 2004

1

Agenda

- Nulls and outer joins
- Creating and updating schemas
- Views: updating and reusing them
- Constraints
- Programming with SQL
- Relational algebra

2

Null Values and Outer Joins

Explicit joins in SQL:
Product(name, category)
Purchase(prodName, store)

```
SELECT Productname, Purchase store
FROM Product JOIN Purchase ON
      Productname = Purchase prodName
```

Same as:

```
SELECT Productname, Purchase store
FROM Product, Purchase
WHERE Productname = Purchase prodName
```

But Products that never sold will be lost!

3

Null Values and Outer Joins

Left outer joins in SQL:
Product(name, category)
Purchase(prodName, store)

```
SELECT Productname, Purchase store
FROM Product LEFT OUTER JOIN Purchase ON
      Productname = Purchase prodName
```

4

Product

| Name | Category |
|----------|----------|
| Gizmo | gadget |
| Camera | Photo |
| OneClick | Photo |

Purchase

| ProdName | Store |
|----------|-------|
| Gizmo | Wiz |
| Camera | Ritz |
| Camera | Wiz |

| Name | Store |
|----------|-------|
| Gizmo | Wiz |
| Camera | Ritz |
| Camera | Wiz |
| OneClick | NULL |

5

Outer Joins

- Left outer join:
 - Include the left tuple even if there's no match
- Right outer join:
 - Include the right tuple even if there's no match
- Full outer join:
 - Include the both left and right tuples even if there's no match

6

Modifying the Database

Three kinds of modifications

- Insertions
- Deletions
- Updates

Sometimes they are all called "updates"

7

Insertions

General form :

```
INSERT INTO R (A1, ..., An) VALUES (v1, ..., vn)
```

Example: Insert a new purchase to the database:

```
INSERT INTO Purchase (buyer, seller, product, store)
VALUES ('Joe', 'Fred', 'wake-up-clock-espresso-machine',
'The Shaper Image')
```

Missing attribute fill NULL.

May drop attribute names if give them in order.

8

Insertions

```
INSERT INTO PRODUCT (name)
SELECT DISTINCT Purchase.product
FROM Purchase
WHERE Purchase.date > "10/26/01"
```

The query replaces the VALUES keyword.
Here we insert many tuples into PRODUCT

9

Insertion: an Example

```
Product (name, listPrice, category)
Purchase (prodName, buyerName, price)
```

prodName is foreign key in ProductName

Suppose database got corrupted and we need to fix it:

Product

| name | listPrice | category |
|-------|-----------|----------|
| gim o | 100 | gadgets |

Purchase

| prodName | buyerName | price |
|----------|-----------|-------|
| cam era | John | 200 |
| gim o | Smith | 80 |
| cam era | Smith | 225 |

Task: insert in Product all prodNames from Purchase

10

Insertion: an Example

```
INSERT INTO Product (name)
SELECT DISTINCT prodName
FROM Purchase
WHERE prodName NOT IN (SELECT name FROM Product)
```

| name | listPrice | category |
|---------|-----------|----------|
| gim o | 100 | Gadgets |
| cam era | - | - |

11

Insertion: an Example

```
INSERT INTO Product (name, listPrice)
SELECT DISTINCT prodName, price
FROM Purchase
WHERE prodName NOT IN (SELECT name FROM Product)
```

| name | listPrice | category |
|------------|-----------|----------|
| gim o | 100 | Gadgets |
| cam era | 200 | - |
| cam era ?? | 225 ?? | - |

← Depends on the implementation

Deletions

Example:

```
DELETE FROM PURCHASE
WHERE seller= 'Jeb' AND
      product= 'Brooklyn Bridge'
```

Factoid about SQL: there is no way to delete only a single occurrence of a tuple that appears twice in a relation.

13

Updates

Example:

```
UPDATE PRODUCT
SET price = price/2
WHERE Productname IN
  (SELECT product
   FROM Purchase
   WHERE Date = 'Oct, 25, 1999');
```

14

Data Definition in SQL

So far we have seen the Data Manipulation Language, DML
Next: Data Definition Language (DDL)

Data types:

Defines the types.

Data definition: defining the schema.

- Create tables
- Delete tables
- Modify table schema

Indexes: to improve performance

15

Data Types in SQL

- Characters:
 - CHAR (20) - fixed length
 - VARCHAR (40) - variable length
- Numbers:
 - INT, REAL plus variations
- Times and dates:
 - DATE, DATETIME (SQL Server only)
- To reuse domains:
CREATE DOMAIN address AS
VARCHAR (55)

16

Creating Tables

Example:

```
CREATE TABLE Person (
  name          VARCHAR (30),
  social-security-number INT,
  age           SHORTINT,
  city          VARCHAR (30),
  gender        BIT (1),
  Birthdate     DATE
);
```

17

Deleting or Modifying a Table

Deleting:

Example: `DROP Person;` Exercise with care !!

Altering: (adding or removing an attribute).

Example:

```
ALTER TABLE Person
ADD phone CHAR (16);

ALTER TABLE Person
DROP age;
```

What happens when you make changes to the schema?

18

Default Values

Specifying default values:

```
CREATE TABLE Person (  
  name          VARCHAR(30),  
  social-security-number INT,  
  age           SHORTINT DEFAULT 100,  
  city          VARCHAR(30) DEFAULT 'Seattle',  
  gender        CHAR(1) DEFAULT '?',  
  Birthdate     DATE
```

The default of defaults: NULL

19

Indexes

REALLY important to speed up query processing time.

Suppose we have a relation

Person (name, age, city)


```
SELECT *  
FROM Person  
WHERE name = "Smith"
```

Sequential scan of the file Person may take long

20

Indexes

- Create an index on name:



| | | | | | |
|------|-------|---------|-----|-------|-----|
| Adam | Betty | Charles | ... | Smith | ... |
|------|-------|---------|-----|-------|-----|

- B+ trees have fan-out of 100s; max 4 levels!

21

Creating Indexes

Syntax:

```
CREATE INDEX nameIndex ON Person (name)
```

22

Creating Indexes

Indexes can be created on more than one attribute:

Example:

```
CREATE INDEX doubleIndex ON  
Person (age, city)
```

Helps in:

```
SELECT *  
FROM Person  
WHERE age = 55 AND city = "Seattle"
```

But not in:

```
SELECT *  
FROM Person  
WHERE city = "Seattle"
```

23

Creating Indexes

Indexes can be useful in range queries too:

```
CREATE INDEX ageIndex ON Person (age)
```

B+ trees help in:

```
SELECT *  
FROM Person  
WHERE age > 25 AND age < 28
```

Why not create indexes on everything?

24

Agenda

Nulls and outerjoins

Creating and updating schemas

- Views: updating and reusing them
- Constraints
- Programming with SQL
- Relational algebra

25

Defining Views

Views are relations, except that they are not physically stored.

For presenting different information to different users

Employee (ssn, name, department, project, salary)

```
CREATE VIEW Developers AS
SELECT name, project
FROM Employee
WHERE department = 'Development'
```

Payroll has access to Employee, others only to Developers

26

A Different View

Person (name, city)

Purchase (buyer, seller, product, store)

Product (name, maker, category)

```
CREATE VIEW Seattle-view AS
SELECT buyer, seller, product, store
FROM Person, Purchase
WHERE Person.city = 'Seattle' AND
      Person.name = Purchase.buyer
```

We have a new virtual table:

Seattle-view (buyer, seller, product, store)

27

Using a View

We can later use the view:

```
SELECT name, store
FROM Seattle-view, Product
WHERE Seattle-view.product = Product.name AND
      Product.category = 'shoes'
```

28

What Happens When We Query a View?

```
SELECT name, Seattle-view.store
FROM Seattle-view, Product
WHERE Seattle-view.product = Product.name AND
      Product.category = 'shoes'
```



```
SELECT name, Purchase.store
FROM Person, Purchase, Product
WHERE Person.city = 'Seattle' AND
      Person.name = Purchase.buyer AND
      Purchase.product = Product.name AND
      Product.category = 'shoes'
```

29

Types of Views

- Virtual views:
 - Used in databases
 - Computed on-demand - slower at runtime
 - Always up to date
- Materialized views
 - Used in data warehouses
 - Precomputed offline - faster at runtime
 - May have stale data

30

Updating Views

How can I insert a tuple into a table that doesn't exist?

Employee (ssn, name, department, project, salary)

```
CREATE VIEW Developers AS
SELECT name, project
FROM Employee
WHERE department = "Development"
```

If we make the following insertion:

```
INSERT INTO Developers
VALUES ("Joe", "Optimizer")
```

It becomes:

```
INSERT INTO Employee
VALUES (NULL, "Joe", "Development", "Optimizer", NULL)
```

31

Non-Updatable Views

```
CREATE VIEW Seattle-view AS
SELECT seller, product, store
FROM Person, Purchase
WHERE Person.city = "Seattle" AND
Person.name = Purchase.buyer
```

How can we add the following tuple to the view?

("Joe", "ShoeModel12345", "NewWest")

We need to add "Joe" to Person first, but we don't have all its attributes

32

Answering Queries Using Views

- What if we want to use a set of views to answer a query.
- Why?
 - The obvious reason...
 - Answering queries over web data sources.
- Very cool stuff! (i.e., I did a lot of research on this).

33

Reusing a Materialized View

- Suppose I have only the result of SeattleView:

```
SELECT buyer, seller, product, store
FROM Person, Purchase
WHERE Person.city = 'Seattle' AND
Person.name = Purchase.buyer
```
- and I want to answer the query

```
SELECT buyer, seller
FROM Person, Purchase
WHERE Person.city = 'Seattle' AND
Person.name = Purchase.buyer AND
Purchase.product = 'gizmo'.
```

Then, I can rewrite the query using the view.

34

Query Rewriting Using Views

Rewritten query:

```
SELECT buyer, seller
FROM SeattleView
WHERE product = 'gizmo'
```

Original query:

```
SELECT buyer, seller
FROM Person, Purchase
WHERE Person.city = 'Seattle' AND
Person.name = Purchase.buyer AND
Purchase.product = 'gizmo'.
```

35

Another Example

- I still have only the result of SeattleView:

```
SELECT buyer, seller, product, store
FROM Person, Purchase
WHERE Person.city = 'Seattle' AND
Person.name = Purchase.buyer
```
- but I want to answer the query

```
SELECT buyer, seller
FROM Person, Purchase
WHERE Person.city = 'Seattle' AND
Person.name = Purchase.buyer AND
Person.Phone LIKE '206 543 %'.
```

36

And Now ?

- I still have only the result of Seattle view :

```
SELECT buyer, seller, product, store
FROM Person, Purchase, Product
WHERE Person.city = 'Seattle' AND
      Person.pernam e = Purchase.buyer AND
      Purchase.product = Product.name
```

- but I want to answer the query

```
SELECT buyer, seller
FROM Person, Purchase
WHERE Person.city = 'Seattle' AND
      Person.pernam e = Purchase.buyer.
```

37

And Now ?

- I still have only the result of:

```
SELECT seller, buyer, Sum (Price)
FROM Purchase
WHERE Purchase.store = 'The Bon'
Group By seller, buyer
```

- but I want to answer the query

```
SELECT seller, Sum (Price)
FROM Purchase
WHERE Person.store = 'The Bon'
Group By seller
```

And what if it's the other way around?

38

Finally...

- I still have only the result of:

```
SELECT seller, buyer, Count(*)
FROM Purchase
WHERE Purchase.store = 'The Bon'
Group By seller, buyer
```

- but I want to answer the query

```
SELECT seller, Count(*)
FROM Purchase
WHERE Person.store = 'The Bon'
Group By seller
```

39

The General Problem

- Given a set of views V_1, \dots, V_n , and a query Q , can we answer Q using only the answers to V_1, \dots, V_n ?
- Why do we care?
 - We can answer queries more efficiently.
 - We can query data sources on the WWW in a principled manner.
- Many, many papers on this problem.
- The best performing algorithm: The MiniCon Algorithm, (Pottinger & (Ha)Levy, 2000).⁴⁰

Querying the WWW

- Assume a virtual schema of the WWW, e.g.,
 - Course (number, university, title, prof, quarter)
- Every data source on the web contains the answer to a view over the virtual schema:

```
UW database: SELECT number, title, prof
FROM Course
WHERE univ = 'UW' AND quarter = '2/02'
Stanford database: SELECT number, title, prof, quarter
FROM Course
WHERE univ = 'Stanford'
```

User query: find all professors who teach "database systems"

41

Agenda

- Nulls and outer-joins
- Creating and updating schemas
- Views: updating and reusing them
- Constraints
- Programming with SQL
- Relational algebra

42



Constraints in SQL

- A constraint = a property that we'd like our database to hold
- The system will enforce the constraint by taking some actions:
 - forbid an update
 - or perform compensating updates

43

Constraints in SQL

Constraints in SQL:

- Keys, foreign keys 
- Attribute-level constraints
- Tuple-level constraints
- Global constraints: assertions 

The more complex the constraint, the harder it is to check and to enforce

44

Keys

```
CREATE TABLE Product (
  name CHAR (30) PRIMARY KEY ,
  category VARCHAR (20))
```

OR :

```
CREATE TABLE Product (
  name CHAR (30),
  category VARCHAR (20)
  PRIMARY KEY (name))
```

45

Keys with Multiple Attributes

```
CREATE TABLE Product (
  name CHAR (30),
  category VARCHAR (20),
  price INT,
  PRIMARY KEY (name, category))
```

46

Other Keys

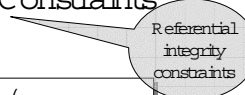
```
CREATE TABLE Product (
  productID CHAR (10),
  name CHAR (30),
  category VARCHAR (20),
  price INT,
  PRIMARY KEY (productID),
  UNIQUE (name, category))
```

There is at most one PRIMARY KEY ;
there can be many UNIQUE

47

Foreign Key Constraints

```
CREATE TABLE Purchase (
  prodName CHAR (30)
  REFERENCES Product (name),
  date DATE TIME)
```



prodName is a foreign key to Product (name)
name must be a key in Product

48

| Product | | Purchase | |
|----------|----------|----------|-------|
| Name | Category | ProdName | Store |
| Gizmo | gadget | Gizmo | Wiz |
| Camera | Photo | Camera | Ritz |
| OneClick | Photo | Camera | Wiz |

49

Foreign Key Constraints

- OR

```

CREATE TABLE Purchase (
  prodName CHAR (30),
  category VARCHAR (20),
  date DATE TIME,
  FOREIGN KEY (prodName, category)
  REFERENCES Product(prodName, category)

```

- (name, category) must be a PRIMARY KEY

50

What happens during updates ?

Types of updates:

- In Purchase: insert/update
- In Product: delete/update

| Product | | Purchase | |
|----------|----------|----------|-------|
| Name | Category | ProdName | Store |
| Gizmo | gadget | Gizmo | Wiz |
| Camera | Photo | Camera | Ritz |
| OneClick | Photo | Camera | Wiz |

51

What happens during updates ?

- SQL has three policies for maintaining referential integrity:
- Reject violating modifications (default)
- Cascade: after a delete/update do a delete/update
- Set-null set foreign-key field to NULL

READING ASSIGNMENT: 7.1.5, 7.1.6

52

Constraints on Attributes and Tuples

- Constraints on attributes:
 - NOT NULL — obvious meaning...
 - CHECK condition — any condition !
- Constraints on tuples
 - CHECK condition

53

What is the difference from Foreign-Key ?

```

CREATE TABLE Purchase (
  prodName CHAR (30)
  CHECK (prodName IN
  SELECT ProductName
  FROM Product),
  date DATE TIME NOT NULL)

```

54

General Assertions

```
CREATE ASSERTION myAssert CHECK
NOT EXISTS (
  SELECT Productname
  FROM Product, Purchase
  WHERE Productname = Purchase.prodName
  GROUP BY Productname
  HAVING count(*) > 200)
```

55

Final Comments on Constraints

- Can give them names, and alter later
 - Read in the book.
- We need to understand exactly when they are checked
- We need to understand exactly what actions are taken if they fail

56

Triggers

Enable the database programmer to specify:

- when to check a constraint,
- what exactly to do.

A trigger has 3 parts:

- An event (e.g., update to an attribute)
- A condition (e.g., a query to check)
- An action (deletion, update, insertion)

When the event happens, the system will check the constraint, and if satisfied, will perform the action.

NOTE: triggers may cause cascading effects.
Database vendors did not wait for standards with triggers!

57

Elements of Triggers (in SQL3)

- Timing of action execution: before, after or instead of triggering event
- The action can refer to both the old and new state of the database.
- Update events may specify a particular column or set of columns.
- A condition is specified with a WHEN clause.
- The action can be performed either for
 - once for every tuple, or
 - once for all the tuples that are changed by the database operation.

58

Example: Row Level Trigger

```
CREATE TRIGGER NoLowerPrices
AFTER UPDATE OF price ON Product
REFERENCING
  OLD AS OldTuple
  NEW AS NewTuple
WHEN (OldTuple.price > NewTuple.price)
UPDATE Product
SET price = OldTuple.price
WHERE name = NewTuple.name
FOR EACH ROW
```

59

Statement Level Trigger

```
CREATE TRIGGER average-price-preserve
INSTEAD OF UPDATE OF price ON Product
REFERENCING
  OLD_TABLE AS OldStuff
  NEW_TABLE AS NewStuff
WHEN (1000 <
  (SELECT AVG (price)
   FROM ((Product EXCEPT OldStuff) UNION NewStuff)))
DELETE FROM Product
WHERE (name, price, company) IN OldStuff;
INSERT INTO Product
(SELECT * FROM NewStuff)
```

60

Bad Things Can Happen

```
CREATE TRIGGER Bad-trigger
AFTER UPDATE OF price IN Product
REFERENCING OLD AS OldTuple
NEW AS NewTuple

WHEN (NewTuple.price > 50)

UPDATE Product
SET price = NewTuple.price * 2
WHERE name = NewTuple.name

FOR EACH ROW
```

61

Agenda

- Nulls and outerjoins
- Creating and updating schemas
- Views: updating and reusing them
- Constraints
- Programming with SQL
- Relational algebra

62

Embedded SQL

- direct SQL (= ad-hoc SQL) is rarely used
- in practice: SQL is embedded in some application code
- SQL code is identified by special syntax

63

Impedance Mismatch

- Example: SQL in C:
 - C uses int, char[.], pointers, etc
 - SQL uses tables
- Impedance mismatch = incompatible types

64

The Impedance Mismatch Problem


Why not use only one language?

- Forgetting SQL: "we can quickly dispense with this idea" [textbook, pg. 351].
- SQL cannot do everything that the host language can do.

Solution: use cursors

65

Programs with Embedded SQL



66

Interface: SQL / Host Language

Values get passed through shared variables.

Columns precede shared variables when they occur within the SQL statements.

EXEC SQL; precedes every SQL statement in the host language.

The variable SQLSTATE provides error messages and status reports (e.g., "00000" says that the operation completed with no problem).

```
EXEC SQL BEGIN DECLARE SECTION ;
char productName[30];
EXEC SQL END DECLARE SECTION ;
```

67

Example

Product (pname, price, quantity, maker)
Purchase (buyer, seller, store, pname)
Company (cname, city)
Person (name, phone, city)

68

Using Shared Variables

```
void simpleInsert() {
    EXEC SQL BEGIN DECLARE SECTION ;
    char n[20], c[30]; /* product-name, company-name */
    int p, q; /* price, quantity */
    char SQLSTATE[6];
    EXEC SQL END DECLARE SECTION ;

    /* get values for name, price and company somehow */

    EXEC SQL INSERT INTO Product(pname, price, quantity, maker)
    VALUES (n, p, q, c);
}
```

69

Single-Row Select Statements

```
int getPrice(char *name) {
    EXEC SQL BEGIN DECLARE SECTION ;
    char n[20];
    int p;
    char SQLSTATE[6];
    EXEC SQL END DECLARE SECTION ;

    strcpy(n, name); /* copy name to local variable */

    EXEC SQL SELECT price INTO p
    FROM Product
    WHERE Productname = n;

    return p;
}
```

70

Cursors

1. Declare the cursor
2. Open the cursor
3. Fetch tuples one by one
4. Close the cursor

71

Cursors

```
void productXML() {
    EXEC SQL BEGIN DECLARE SECTION ;
    char n[20], c[30];
    int p, q;
    char SQLSTATE[6];
    EXEC SQL END DECLARE SECTION ;

    EXEC SQL DECLARE cscursor FOR
    SELECT pname, price, quantity, maker
    FROM Product;

    EXEC SQL OPEN csc;
```

72

Cursors

```

printf("<allProducts>\n");
while (1) {
    EXEC SQL FETCH FROM cns INTO n, p, q, c;
    if (NO_MORE_TUPLES) break;
    printf(" <product>\n");
    printf("    <name>    %s </name>\n", n);
    printf("    <price>   %d </price>\n", p);
    printf("    <quantity> %d </quantity>\n", q);
    printf("    <maker>   %s </maker>\n", c);
    printf(" </product>\n");
}
EXEC SQL CLOSE cns;
printf("</allProducts>\n");
}

```

73

- What is NO_MORE_TUPLES ?

```
#define NO_MORE_TUPLES !($?cmp(SQLSTATE,'02000'))
```

74

More on Cursors

- cursors can modify a relation as well as read it.
- We can determine the order in which the cursor will get tuples by the ORDER BY keyword in the SQL query.
- Cursors can be protected against changes to the underlying relations.
- The cursor can be a scrolling one: can go forward, backward +n, -n, Abs(n), Abs(-n).

75

Agenda

- Nulls and outerjoins
- Creating and updating schemas
- Views: updating and reusing them
- Constraints
- Programming with SQL
- Relational algebra

76

Relational Algebra

- Formalism for creating new relations from existing ones
- Its place in the big picture:



SQL,
relational calculus

Relational algebra
Relational bag algebra

77

Relational Algebra

- Five operators:
 - Union: ∪
 - Difference: -
 - Selection: σ
 - Projection: π
 - Cartesian Product: ×
- Derived or auxiliary operators:
 - Intersection, complement
 - Joins (natural, equi-join, theta join, semi-join)
 - Renaming: ρ

78

1. Union and 2. Difference

- $R1 \cup R2$
- Example:
 - ActiveEmployees \cup RetiredEmployees
- $R1 - R2$
- Example:
 - AllEmployees - RetiredEmployees

79

What about Intersection ?

- It is a derived operator
- $R1 \cap R2 = R1 - (R1 - R2)$
- Also expressed as a join (will see later)
- Example
 - UnionizedEmployees \cap RetiredEmployees

80

3. Selection

- Returns all tuples which satisfy a condition
- Notation: $\sigma_c(R)$
- Examples
 - $\sigma_{\text{salary} > 40000}(\text{Employee})$
 - $\sigma_{\text{name} = 'Alice'}(\text{Employee})$
- The condition c can be $=, <, \neq, >, \neq, <>$

81

Selection Example

| Employee | | | |
|-----------|-------|--------------|--------|
| SSN | Name | DepartmentID | Salary |
| 999999999 | John | 1 | 30,000 |
| 777777777 | Tony | 1 | 32,000 |
| 888888888 | Alice | 2 | 45,000 |

Find all employees with salary more than \$40,000.

$\sigma_{\text{Salary} > 40000}(\text{Employee})$

| $\sigma_{\text{Salary} > 40000}(\text{Employee})$ | | | |
|---|-------|--------------|--------|
| SSN | Name | DepartmentID | Salary |
| 888888888 | Alice | 2 | 45,000 |

82

4. Projection

- Eliminates columns, then removes duplicates
- Notation: $\rho_{A_1, \dots, A_n}(R)$
- Example: project social-security number and names:
 - $\rho_{\text{SSN}, \text{Name}}(\text{Employee})$
 - Output schema: Answer(SSN, Name)

83

Projection Example

| Employee | | | |
|-----------|-------|--------------|--------|
| SSN | Name | DepartmentID | Salary |
| 999999999 | John | 1 | 30,000 |
| 777777777 | Tony | 1 | 32,000 |
| 888888888 | Alice | 2 | 45,000 |

$\rho_{\text{SSN}, \text{Name}}(\text{Employee})$

| $\rho_{\text{SSN}, \text{Name}}(\text{Employee})$ | |
|---|-------|
| SSN | Name |
| 999999999 | John |
| 777777777 | Tony |
| 888888888 | Alice |

84

5. Cartesian Product

- Each tuple in R1 with each tuple in R2
- Notation: $R1 \cdot R2$
- Example:
 - Employee · Dependents
- Very rare in practice; mainly used to express joins

85

Cartesian Product Example

| Employee | |
|----------|-----------|
| Name | SSN |
| John | 999999999 |
| Tony | 777777777 |

| Dependents | |
|-------------|-------|
| EmployeeSSN | Dname |
| 999999999 | Emily |
| 777777777 | Joe |

| Employee x Dependents | | | |
|-----------------------|-----------|-------------|-------|
| Name | SSN | EmployeeSSN | Dname |
| John | 999999999 | 999999999 | Emily |
| John | 999999999 | 777777777 | Joe |
| Tony | 777777777 | 999999999 | Emily |
| Tony | 777777777 | 777777777 | Joe |

86

Relational Algebra

- Five operators:
 - Union: \cup
 - Difference: $-$
 - Selection: σ
 - Projection: π
 - Cartesian Product: \cdot
- Derived or auxiliary operators:
 - Intersection, complement
 - Joins (natural, equi-join, theta join, semi-join)
 - Renaming: ρ

87

Renaming

- Changes the schema, not the instance
- Notation: $\rho_{B_1, \dots, B_n}(R)$
- Example:
 - $\rho_{\text{LastName, SocSocNo}}(\text{Employee})$
 - Output schema: Answer(LastName, SocSocNo)

88

Renaming Example

| Employee | |
|----------|-----------|
| Name | SSN |
| John | 999999999 |
| Tony | 777777777 |

| $\rho_{\text{LastName, SocSocNo}}(\text{Employee})$ | |
|---|-----------|
| LastName | SocSocNo |
| John | 999999999 |
| Tony | 777777777 |

89

Natural Join

- Notation: $R1 \bowtie R2$
- Meaning: $R1 \bowtie R2 = \pi_A(\sigma_C(R1 \cdot R2))$
- Where:
 - The selection σ_C checks equality of all common attributes
 - The projection eliminates the duplicate common attributes

90

NaturalJoin Example

| Employee | |
|----------|-----------|
| Name | SSN |
| John | 999999999 |
| Tony | 777777777 |

| Dependents | |
|------------|-------|
| SSN | Dname |
| 999999999 | Emily |
| 777777777 | Joe |

Employee \bowtie Dependents =

$P_{Name, SSN, Dname} (S_{SSN=SSN2} (Employee \times_{SSN2, Dname} Dependents))$

| Name | SSN | Dname |
|------|-----------|-------|
| John | 999999999 | Emily |
| Tony | 777777777 | Joe |

91

NaturalJoin

- $R =$

| A | B |
|---|---|
| X | Y |
| X | Z |
| Y | Z |
| Z | V |
- $S =$

| B | C |
|---|---|
| Z | U |
| V | W |
| Z | V |
- $R \bowtie S =$

| A | B | C |
|---|---|---|
| X | Z | U |
| X | Z | V |
| Y | Z | U |
| Y | Z | V |
| Z | V | W |

92

NaturalJoin

- Given the schema $R(A, B, C, D), S(A, C, E)$, what is the schema of $R \bowtie S$?
- Given $R(A, B, C), S(D, E)$, what is $R \bowtie S$?
- Given $R(A, B), S(A, B)$, what is $R \bowtie S$?

93

Theta Join

- A join that involves a predicate
- $R1 \bowtie_q R2 = s_q (R1 \cdot R2)$
- Here q can be any condition

94

Eq-join

- A theta join where q is an equality
- $R1 \bowtie_{A=B} R2 = s_{A=B} (R1 \cdot R2)$
- Example:
 - Employee $\bowtie_{SSN=SSN}$ Dependents
- Most useful join in practice

95


Semijoin

- $R \ltimes S = P_{A_1, \dots, A_n} (R \bowtie S)$
- Where A_1, \dots, A_n are the attributes in R
- Example:
 - Employee \ltimes Dependents

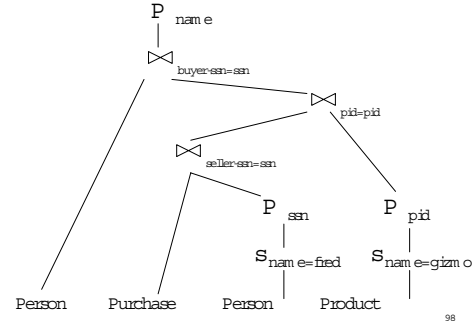
96

Semijoins in Distributed Databases

- Semijoins are used in distributed databases



Complex RA Expressions



Operations on Bags

A bag = a set with repeated elements

All operations need to be defined carefully on bags

- $\{a,b,b,c\} \cup \{a,b,b,b,e,f,f\} = \{a,a,b,b,b,b,b,c,e,f,f\}$
- $\{a,b,b,b,c,c\} - \{b,c,c,c,d\} = \{a,b,b,d\}$
- $s_c(R)$: preserve the number of occurrences
- $P_A(R)$: no duplicate elimination
- Cartesian product, join: no duplicate elimination

Important! Relational Engines work on bags, not sets!

Reading assignment: 5.3 - 5.4

99

Finally: RA has Limitations!

- Cannot compute "transitive closure"

| Name1 | Name2 | Relationship |
|-------|-------|--------------|
| Fred | Mary | Father |
| Mary | Joe | Cousin |
| Mary | Bill | Spouse |
| Nancy | Lou | Sister |

- Find all direct and indirect relatives of Fred
- Cannot express in RA !!! Need to write C program

100