

Database System Internals Relational Model Review

Paul G. Allen School of Computer Science and Engineering University of Washington, Seattle

March 31, 2025

CSE 444 - Relational Model

- Lab 1 part 1 is due on Monday
- This weekend we will send a partner sign-up form
 Full lab due April 16
- HW1 out tomorrow morning, due on April 11th
 Ed announcement and Gradescope link
- 544M first paper review, due April . 20th
 - Submit via gradescope
 - (Not hard deadline)

Note

- We will go very quickly in class over the Relational Algebra and SQL
- Please review at home:
 - Read the slides that we skipped in class
 - Review material from 344 as needed

DBMS Architecture



Query Evaluation Steps Review



Database/Relation/Tuple

- A Database is collection of relations
- A Relation R is subset of $S_1 \times S_2 \times \dots \times S_n$
 - Where $\boldsymbol{S}_{\boldsymbol{i}}$ is the domain of attribute \boldsymbol{i}
 - **n** is number of attributes of the relation
 - A relation is a set of tuples
- A Tuple t is an element of $S_1 \times S_2 \times \dots \times S_n$

Other names: relation = table; tuple = row

Discussion

Rows in a relation:

Data independence!

- Ordering immaterial (a relation is a set)
- All rows are distinct set semantics
- Query answers may have duplicates bag semantics
- Columns in a tuple:
 - Ordering is significant (in theory, it shouldn't be)
 - Applications refer to columns by their names
- Domain of each column is a primitive type

Tuple.java describes a row object in SimpleDB

- Rows are the objects passed through the database
- In the same way we conceptualize RA and a series of transformations to rows, so does it work in database

Schema

Relation schema: describes column heads

- Relation name
- Name of each field (or column, or attribute)
- Domain of each field
- Degree (or arity) of relation: # attributes
- Database schema: set of all relation schemas



Relation instance: concrete table content

- Set of tuples (also called records) matching the schema
- Cardinality of relation instance: # tuples
- Database instance: set of all relation instances

What is the schema? What is the instance?

Supplier

sno	sname	scity	sstate
1	s1	city 1	WA
2	s2	city 1	WA
3	s3	city 2	MA
4	s4	city 2	MA

What is the schema? What is the instance?

Relation schema

Supplier(<u>sno: integer</u>, sname: string, scity: string, sstate: string)

Supplier

sno	sname	scity	sstate	
1	s1	city 1	WA	instance
2	s2	city 1	WA	
3	s3	city 2	MA	
4	s4	city 2	MA	

What is the schema? What is the instance?

Handled by SimpleDB **Catalog**

Relation schema

Supplier(<u>sno: integer</u>, sname: string, scity: string, sstate: string)

Supplier

sno	sname	scity	sstate
1	s1	city 1	WA
2	s2	city 1	WA
3	s3	city 2	MA
4	s4	city 2	MA



- Condition specified on a database schema
- Restricts data that can be stored in db instance
- DBMS enforces integrity constraints
 - Ensures only legal database instances exist
- Simplest form of constraint is domain constraint
 - Attribute values must come from attribute domain

- Super Key: "set of attributes that functionally determines all attributes"
- Key: Minimal super-key; a.k.a. "candidate key"
- Primary key: One minimal key can be selected as primary key

Foreign Key Constraints

• A relation can refer to a tuple in another relation

Foreign key

- Field that refers to tuples in another relation
- This field refers to the primary key of other relation

```
CREATE TABLE Part (
    pno integer,
    pname varchar(20),
    psize integer,
    pcolor varchar(20),
    PRIMARY KEY (pno)
);
```

CREATE TABLE Supply(

- sno integer,
- pno integer,
- qty integer,
- price integer
-);

CREATE TABLE Part (
pno integer,
pname varchar(20),
psize integer,
pcolor varchar(20),
PRIMARY KEY (pno)

);

CREATE TABLE Supply(

- sno integer,
- pno integer,
- qty integer,
- price integer,

```
PRIMARY KEY (sno, pno)
```

```
);
```

CREATE TABLE Part (
pno integer,
pname varchar(20),
psize integer,
pcolor varchar(20),
PRIMARY KEY (pno)

```
);
```

```
CREATE TABLE Supply (
                               CREATE TABLE Part (
 sno integer,
                                  pno integer,
                                  pname varchar(20),
 pno integer,
                                  psize integer,
 qty integer,
                                  pcolor varchar(20),
                                  PRIMARY KEY (pno)
 price integer,
                               );
 PRIMARY KEY (sno,pno),
 FOREIGN KEY (sno) REFERENCES Supplier,
 FOREIGN KEY (pno) REFERENCES Part
);
```

```
CREATE TABLE Supply(
                               CREATE TABLE Part (
 sno integer,
                                  pno integer,
                                  pname varchar(20),
 pno integer,
                                  psize integer,
 qty integer,
                                  pcolor varchar(20),
                                  PRIMARY KEY (pno)
 price integer,
                               );
 PRIMARY KEY (sno, pno),
 FOREIGN KEY (sno) REFERENCES Supplier
                         ON DELETE NO ACTION,
 FOREIGN KEY (pno) REFERENCES Part
                         ON DELETE CASCADE
```

);

Table constraints serve to express complex constraints over a single table

```
CREATE TABLE Part (
   pno integer,
   pname varchar(20),
   psize integer,
   pcolor varchar(20),
   PRIMARY KEY (pno),
   CHECK ( psize > 0 )
);
```

Note: Also possible to create constraints over many tables Alternative: use database triggers for that purpose

Relational Query Languages

Relational Query Language

Set-at-a-time:

- Query inputs and outputs are relations
- Two variants of the query language:
 - Relational algebra: specifies order of operations
 - Relational calculus / SQL: declarative

Relational Algebra

Queries specified in an operational manner

• A query gives a step-by-step procedure

Relational operators

- Take one or two relation instances as argument
- Return one relation instance as result
- Easy to compose into relational algebra expressions

Note that only two relations may be input to a join operator, to join 3+ relations we need multiple join operators.

Five Basic Relational Operators

- Selection: $\sigma_{\text{condition}}(S)$
 - Condition is Boolean combination (∧,∨) of atomic predicates (<, <=, =, ≠, >=, >)
- Projection: π_{list-of-attributes}(S)
- **Union** (∪)
- Set difference (-),
- Cross-product/cartesian product (×), Join: $\mathbb{R} \bowtie_{\theta} \mathbb{S} = \sigma_{\theta}(\mathbb{R} \times \mathbb{S})$

Selection & Projection Examples

Patient

no	name	zip	disease
1	p1	98125	flu
2	p2	98125	heart
3	р3	98120	lung
4	p4	98120	heart

$\pi_{zip,disease}(Patient)$

zip	disease
98125	flu
98125	heart
98120	lung
98120	heart

 $\sigma_{disease='heart'}(Patient)$

no	name	zip	disease
2	p2	98125	heart
4	p4	98120	heart

$$\pi_{zip} (\sigma_{disease='heart'} (Patient))$$

zip
98120
98125

Cross-Product Example

AnonPatient P

age	zip	disease
54	98125	heart
20	98120	flu

Voters V

name	age	zip
p1	54	98125
p2	20	98120

 $P \times V$

P.age	P.zip	disease	name	V.age	V.zip
54	98125	heart	p1	54	98125
54	98125	heart	p2	20	98120
20	98120	flu	p1	54	98125
20	98120	flu	p2	20	98120

Different Types of Join

• Theta-join: $R \bowtie_{\theta} S = \sigma_{\theta} (R \times S)$

- Join of R and S with a join condition $\boldsymbol{\theta}$
- Cross-product followed by selection $\boldsymbol{\theta}$

• Equijoin:
$$R \bowtie_{\theta} S = \pi_A(\sigma_{\theta}(R \times S))$$

- Join condition θ consists only of equalities
- Projection π_A drops all redundant attributes

• Natural join:
$$R \bowtie S = \pi_A (\sigma_\theta (R \times S))$$

- Equijoin
- Equality on **all** fields with same name in R and in S

Different Types of Join

Our focus in SimpleDB We have a class for the predicate θ

• Theta-join: $R \bowtie_{\theta} S = \sigma_{\theta} (R \times S)$

- Join of R and S with a join condition $\boldsymbol{\theta}$
- Cross-product followed by selection $\boldsymbol{\theta}$

• Equijoin:
$$\mathbf{R} \bowtie_{\theta} \mathbf{S} = \pi_{\mathsf{A}}(\sigma_{\theta}(\mathbf{R} \times \mathbf{S}))$$

- Join condition θ consists only of equalities
- Projection π_A drops all redundant attributes

• Natural join:
$$R \bowtie S = \pi_A (\sigma_\theta (R \times S))$$

- Equijoin
- Equality on **all** fields with same name in R and in S

AnonPatient P

age	zip	disease
50	98125	heart
20	98120	flu

Voters V

name	age	zip
p1	54	98125
p2	20	98120

$$P \bowtie_{P.zip = V.zip \text{ and } P.age = V.age} V$$

P.age	P.zip	P.disease	V.name	V.age	V.zip
20	98120	flu	p2	20	98120

AnonPatient P

age	zip	disease
50	98125	heart
19	98120	flu

Voters V

name	age	zip
p1	54	98125
p2	20	98120

Ρ	⊠P.zip = V.zip	and P.age <=	V.age + 1	and P.age >=	V.age - 1	/
---	----------------	--------------	-----------	--------------	-----------	---

P.age	P.zip	P.disease	V.name	V.age	V.zip
19	98120	flu	p2	20	98120

-

Natural Join Example

AnonPatient P

age	zip	disease
54	98125	heart
20	98120	flu

Voters V

name	age	zip
p1	54	98125
p2	20	98120

 $\mathsf{P} \bowtie \mathsf{V}$

age	zip	disease	name
54	98125	heart	p1
20	98120	flu	p2



More Joins

Outer join

- Include tuples with no matches in the output
- Use NULL values for missing attributes
- Variants
 - Left outer join
 - Right outer join
 - Full outer join

Outer Join Example

AnonPatient P

age	zip	disease
54	98125	heart
20	98120	flu
33	98120	lung

Voters V

name	age	zip
p1	54	98125
p2	20	98120

age	zip	disease	name
54	98125	heart	p1
20	98120	flu	p2
33	98120	lung	null



Logical Query Plans

Supplier(sno, sname, scity, sstate)
Supply(sno, pno, qty, price)
Part(pno, pname, psize, pcolor)

Logical Query Plans

Supplier(sno, sname, scity, sstate)
Supply(sno, pno, qty, price)
Part(pno, pname, psize, pcolor)



Q: What does this query compute?

Logical Query Plans

Supplier(sno, sname, scity, sstate)
Supply(sno, pno, qty, price)
Part(pno, pname, psize, pcolor)



Q: What does this query compute?

A: The name and city of suppliers who make any parts with size > 10

Extended Operators of RA

- Duplicate elimination (δ)
 - Since commercial DBMSs operate on multisets not sets
- Aggregate operators (γ)
 - Min, max, sum, average, count
- Grouping operators (y)
 - Partitions tuples of a relation into "groups"
 - Aggregates can then be applied to groups
- Sort operator (τ)

Structured Query Language: SQL

- Declarative query language
- Data definition language
 - Statements to create, modify tables and views
- Data manipulation language
 - Statements to issue queries, insert, delete data



Basic form: (plus many many more bells and whistles)

SELECT<attributes>FROM<one or more relations>WHERE<conditions>

Simple SQL Query

Prod	uct PName	Price	Category	Manufacturer
	Gizmo	\$19.99	Gadgets	GizmoWorks
	Powergizmo	\$29.99	Gadgets	GizmoWorks
	SingleTouch	\$149.99	Photography	Canon
	MultiTouch	\$203.99	Household	Hitachi
SELECT*FROMProduceWHEREcategories	ıct ory='Gadgets'			
	PName	Price	Category	Manufacturer
	Gizmo	\$19.99	Gadgets	GizmoWorks
"	Powergizmo	\$29.99	Gadgets	GizmoWorks

Simple SQL Query

Product	PName	Price	Category	Manufacturer
	Gizmo	\$19.99	Gadgets	GizmoWorks
	Powergizmo	\$29.99	Gadgets	GizmoWorks
	SingleTouch	\$149.99	Photography	Canon
	MultiTouch	\$203.99	Household	Hitachi

SELECTPName, Price, ManufacturerFROMProductWHEREPrice > 100

"selection" and "projection"

PName	Price	Manufacturer
SingleTouch	\$149.99	Canon
MultiTouch	\$203.99	Hitachi

Details

Case insensitive:

- Same: SELECT Select select
- Same: Product product
- Different: 'Seattle' 'seattle'
- Constants:
 - 'abc' yes
 - "abc" no

Eliminating Duplicates



Compare to:



SELECT pname, price, manufacturer FROM Product WHERE category='gizmo' AND price > 50 ORDER BY price, pname

Ties are broken by the second attribute on the ORDER BY list, etc.

Ordering is ascending, unless you specify the DESC keyword.



Product (<u>pname</u>, price, category, manufacturer) Company (<u>cname</u>, stockPrice, country)

Find all products under \$200 manufactured in Japan; return their names and prices.

SELECT	PName, Price
FROM	Product, Company
WHERE	Manufacturer=CName AND Country='Japan'
	AND Price <= 200

Quick Review of SQL

Supplier(sno, sname, scity, sstate)
Supply(sno, pno, qty, price)
Part(pno, pname, psize, pcolor)

SELECT DISTINCT z.pno, z.pname
FROM Supplier x, Supply y, Part z
WHERE x.sno = y.sno and y.pno = z.pno
and x.scity = 'Seattle' and y.price < 100</p>

What does this query compute?

Quick Review of SQL

Supplier(sno, sname, scity, sstate)
Supply(sno, pno, qty, price)
Part(pno, pname, psize, pcolor)

What about this one?

SELECT z.pname, count(*) as cnt, min(y.price) FROM Supplier x, Supply y, Part z WHERE x.sno = y.sno and y.pno = z.pno GROUP BY z.pname



SELECT	avg(price)
FROM	Product
WHERE	maker="Toyota"

SELECTcount(*)FROMProductWHEREyear > 1995

SQL supports several aggregation operations: sum, count, min, max, avg

Except count, all aggregations apply to a single attribute

Grouping and Aggregation

$$\begin{array}{ccc} \textbf{SELECT} & S \\ \textbf{FROM} & R_1, \dots, R_n \\ \textbf{WHERE} & \textbf{C1} \\ \textbf{GROUP} & \textbf{BY} & a_1, \dots, a_k \\ \textbf{HAVING} & \textbf{C2} \end{array}$$

Conceptual evaluation steps:

- 1. Evaluate FROM-WHERE, apply condition C1
- 2. Group by the attributes a_1, \ldots, a_k
- 3. Apply condition C2 to each group (may have aggregates)
- 4. Compute aggregates in S and return the result

Read more about it in the book...

From SQL to RA

Query Evaluation Steps Review



53

DBMS Architecture

Parser
Query Rewrite
Optimizer
Executor
Query Processor

From SQL to RA

Product(<u>pid</u>, name, price) Purchase(<u>pid</u>, <u>cid</u>, store) Customer(<u>cid</u>, name, city)

SELECT DISTINCT x.name, z.name FROM Product x, Purchase y, Customer z WHERE x.pid = y.pid and y.cid = z.cid and x.price > 100 and z.city = 'Seattle'

From SQL to RA



DBMS Architecture



From SQL to RA



An Equivalent Expression



DBMS Architecture



Query Execution



Query Execution



Query Evaluation Steps Review



63

