




Relational Model

Database System Concepts, 5th Ed.


©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use





Chapter 2: Relational Model


- Structure of Relational Databases
- Fundamental Relational-Algebra-Operations
- Additional Relational-Algebra-Operations
- Extended Relational-Algebra-Operations
- Modification of the Database





Example of a Relation

<i>account_number</i>	<i>branch_name</i>	<i>balance</i>
A-101	Downtown	500
A-102	Perryridge	400
A-201	Brighton	900
A-215	Mianus	700
A-217	Brighton	750
A-222	Redwood	700
A-305	Round Hill	350





Basic Structure

- Formally, given sets D_1, D_2, \dots, D_n a **relation** r is a subset of

$$D_1 \times D_2 \times \dots \times D_n$$

Thus, a relation is a set of n -tuples (a_1, a_2, \dots, a_n) where each $a_i \in D_i$


- Example: If

- $customer_name = \{\text{Jones, Smith, Curry, Lindsay, ...}\}$
/* Set of all customer names */
- $customer_street = \{\text{Main, North, Park, ...}\}$ /* set of all street names*/
- $customer_city = \{\text{Harrison, Rye, Pittsfield, ...}\}$ /* set of all city names */

Then $r = \{$
 (Jones, Main, Harrison),
 (Smith, North, Rye),
 (Curry, North, Rye),
 (Lindsay, Park, Pittsfield) $\}$

is a relation over


$customer_name \times customer_street \times customer_city$





Attribute Types

- Each attribute of a relation has a name
- The set of allowed values for each attribute is called the **domain** of the attribute
- Attribute values are (normally) required to be **atomic**; that is, indivisible
 - E.g. the value of an attribute can be an account number, but cannot be a set of account numbers
- Domain is said to be atomic if all its members are atomic
- The special value *null* is a member of every domain
- The null value causes complications in the definition of many operations
 - We shall ignore the effect of null values in our main presentation and consider their effect later





Relation Schema

- A_1, A_2, \dots, A_n are *attributes*
- $R = (A_1, A_2, \dots, A_n)$ is a *relation schema*
 - Ordering of attributes is important!


Example:

$Customer_schema = (customer_name, customer_street, customer_city)$

- $r(R)$ denotes a *relation* r on the *relation schema* R

Example:


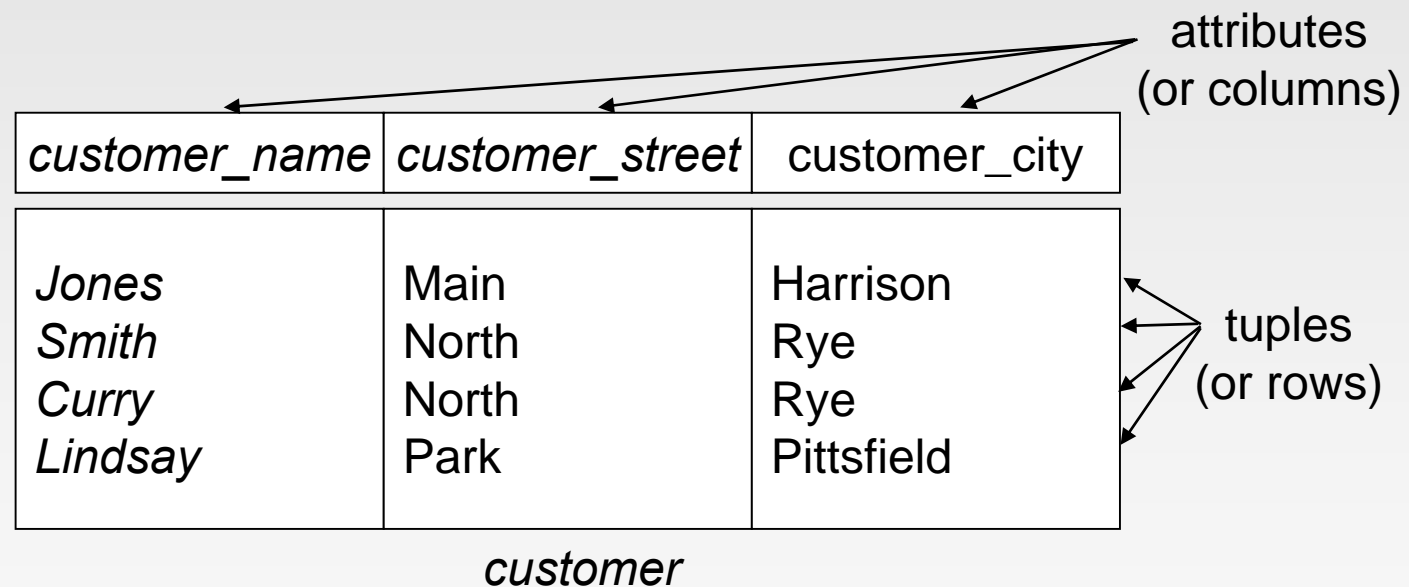
$customer (Customer_schema)$





Relation Instance

- The current values (*relation instance*) of a relation are specified by a table
- An element t of r is a *tuple*, represented by a *row* in a table






Relations are Unordered

- Order of tuples is irrelevant (tuples may be stored in an arbitrary order)
- Example: *account* relation with unordered tuples


<i>account_number</i>	<i>branch_name</i>	<i>balance</i>
A-101	Downtown	500
A-215	Mianus	700
A-102	Perryridge	400
A-305	Round Hill	350
A-201	Brighton	900
A-222	Redwood	700
A-217	Brighton	750





Database


- A database consists of multiple relations
- Information about an enterprise is broken up into parts, with each relation storing one part of the information
 - account* : stores information about accounts
 - depositor* : stores information about which customer owns which account
 - customer* : stores information about customers
- Storing all information as a single relation such as
bank(account_number, balance, customer_name, ..)
results in
 - repetition of information
 - ▶ e.g., if two customers own an account (What gets repeated?)
 - the need for null values
 - ▶ e.g., to represent a customer without an account
- Normalization theory (Chapter 7: Relational Database Design) deals with how to design relational schemas





The *customer* Relation


<i>customer_name</i>	<i>customer_street</i>	<i>customer_city</i>
Adams	Spring	Pittsfield
Brooks	Senator	Brooklyn
Curry	North	Rye
Glenn	Sand Hill	Woodside
Green	Walnut	Stamford
Hayes	Main	Harrison
Johnson	Alma	Palo Alto
Jones	Main	Harrison
Lindsay	Park	Pittsfield
Smith	North	Rye
Turner	Putnam	Stamford
Williams	Nassau	Princeton





The *depositor* Relation


<i>customer_name</i>	<i>account_number</i>
Hayes	A-102
Johnson	A-101
Johnson	A-201
Jones	A-217
Lindsay	A-222
Smith	A-215
Turner	A-305





Keys


- Let $K \subseteq R$
- K is a **superkey** of R if values for K are sufficient to identify a unique tuple of each possible relation $r(R)$
 - by “possible r ” we mean a relation r that could exist in the enterprise we are modeling.
 - Example: $\{customer_name, customer_street\}$ and $\{customer_name\}$
are both superkeys of $Customer$, if no two customers can possibly have the same name
 - ▶ In real life, an attribute such as $customer_id$ would be used instead of $customer_name$ to uniquely identify customers, but we omit it to keep our examples small, and instead assume customer names are unique.





Keys (Cont.)


- K is a **candidate key** if K is minimal
Example: $\{customer_name\}$ is a candidate key for *Customer*, since it is a superkey and no subset of it is a superkey.
- **Primary key**: a candidate key chosen as the principal means of identifying tuples within a relation
 - Should choose an attribute whose value never, or very rarely, changes.
 - E.g. email address is unique, but may change





Query Languages


- Language in which user requests information from the database.
- Categories of languages
 - Procedural
 - Non-procedural, or declarative
- “Pure” languages:
 - Relational algebra
 - Tuple relational calculus
 - Domain relational calculus
- Pure languages form underlying basis of query languages that people use.





Relational Algebra

- Procedural language
- Six basic operators
 - select: σ
 - project: Π
 - union: \cup
 - set difference: $-$
 - Cartesian product: \times
 - rename: ρ
- The operators take one or two relations as inputs and produce a new relation as a result.






Select Operation – Example

- 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





Select Operation

- Notation: $\sigma_p(r)$
- p is called the **selection predicate**
- Defined as:

$$\sigma_p(r) = \{t \mid t \in r \text{ and } p(t)\}$$

where p is a formula in *propositional calculus*:


formula := term
 term <conj> term
 (term)
term := expr
 expr <op> expr
 (expr)
expr := attribute
 constant

<conj> is one of: \wedge (**and**), \vee (**or**), \neg (**not**)

<op> is one of: =, \neq , >, \geq , <, \leq

- Example of selection:

$$\sigma_{\text{branch_name}='Perryridge'}(\text{account})$$





Project Operation – Example

■ 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

=

A	C
α	1
β	1
β	2





Project Operation


- Notation:

$$\Pi_{A_1, A_2, \dots, A_k}(r)$$

where A_1, A_2 are attribute names and r is a relation name.

- The result is defined as the relation of k columns obtained by erasing the columns that are not listed
- *Duplicate* rows removed from result, since relations are sets
- Example: To eliminate the *branch_name* attribute of *account*

$$\Pi_{\text{account_number, balance}}(\text{account})$$





Union Operation – Example

- Relations r, s :

r


A	B
α	1
α	2
β	1

s

A	B
α	2
β	3

- $r \cup s$:

A	B
α	1
α	2
β	1
β	3





Union Operation

■ Notation: $r \cup s$

■ Defined as:


$$r \cup s = \{t \mid t \in r \text{ or } t \in s\}$$

■ For $r \cup s$ to be valid.

1. r, s must have the **same arity** (same number of attributes)
2. The attribute domains must be **compatible** (example: 2nd column of r deals with the same type of values as does the 2nd column of s)

■ Example: to find all customers with either an account or a loan

$$\Pi_{customer_name}(depositor) \cup \Pi_{customer_name}(borrower)$$





Set Difference Operation – Example

- 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






Set Difference Operation

- Notation $r - s$
- Defined as:

$$r - s = \{t \mid t \in r \text{ and } t \notin s\}$$

- Set differences must be taken between **compatible** relations.
 - r and s must have the **same** arity
 - attribute domains of r and s must be compatible





Cartesian-Product Operation – Example

- Relations r, s :

r


A	B
α	1
β	2

s

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






Cartesian-Product Operation

- Notation $r \times s$
- Defined as:

$$r \times s = \{ tq \mid t \in r \text{ and } q \in s \}$$

where tq means the concatenation of tuples t and q to produce a single tuple.

- Assume that attributes of $r(R)$ and $s(S)$ are disjoint. (That is, $R \cap S = \emptyset$).
- If attributes of $r(R)$ and $s(S)$ are not disjoint, then renaming must be used.





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

r	A	B
	α	1
	β	2

s	C	D	E
	α	10	a
	β	10	a
	β	20	b
	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





Rename Operation

- Allows us to name, and therefore to refer to, the results of relational-algebra expressions.
- Allows us to refer to a relation by more than one name.
- Example of naming a relation:

$$\rho_X(E)$$


returns the expression E under the name X

- Example of naming a relation and its attributes:

If a relational-algebra expression E has arity n , then

$$\rho_{X(A_1, A_2, \dots, A_n)}(E)$$

returns the result of expression E under the name X , and with the attributes renamed to A_1, A_2, \dots, A_n .





Banking Example

branch (branch_name, branch_city, assets)

customer (customer_name, customer_street, customer_city)

account (account_number, branch_name, balance)

loan (loan_number, branch_name, amount)

depositor (customer_name, account_number)

borrower (customer_name, loan_number)





Example Queries

loan (*loan_number*, *branch_name*, *amount*)

depositor (*customer_name*, *account_number*)

borrower (*customer_name*, *loan_number*)

- Find all loans of over \$1200


$$\sigma_{amount > 1200} (loan)$$

- Find the loan number for each loan of an amount greater than \$1200

$$\Pi_{loan_number} (\sigma_{amount > 1200} (loan))$$

- Find the names of all customers who have a loan, an account, or both, from the bank

$$\Pi_{customer_name} (borrower) \cup \Pi_{customer_name} (depositor)$$





Example Queries

loan (*loan_number*, *branch_name*, *amount*)


depositor (*customer_name*, *account_number*)

borrower (*customer_name*, *loan_number*)

- Find the names of all customers who have a loan at the Perryridge branch.

$$\Pi_{customer_name} (\sigma_{branch_name="Perryridge"} (\sigma_{borrower.loan_number = loan.loan_number} (borrower \times loan)))$$

- Find the names of all customers who have a loan at the Perryridge branch but do not have an account at any branch of the bank.

$$\Pi_{customer_name} (\sigma_{branch_name = "Perryridge"} (\sigma_{borrower.loan_number = loan.loan_number} (borrower \times loan))) - \Pi_{customer_name} (depositor)$$





Example Queries

- Find the names of all customers who have a loan at the Perryridge branch.

- Query 1

$$\Pi_{\text{customer_name}} (\sigma_{\text{branch_name} = \text{"Perryridge"}} (\sigma_{\text{borrower.loan_number} = \text{loan.loan_number}} (\text{borrower} \times \text{loan})))$$

- Query 2

$$\Pi_{\text{customer_name}} (\sigma_{\text{loan.loan_number} = \text{borrower.loan_number}} (\sigma_{\text{branch_name} = \text{"Perryridge"}} (\text{loan})) \times \text{borrower}))$$





Example Queries

account (*account_number*, *branch_name*, *balance*)

- Find the largest account balance
 - Strategy:
 - ▶ Find those balances that are *not* the largest
 - Rename *account* relation as *d* so that we can compare each account balance with all others
 - ▶ Use set difference to find those account balances that were *not* found in the earlier step.
 - The query is:


$$\Pi_{balance}(account) - \Pi_{account.balance}(\sigma_{account.balance < d.balance} (account \times \rho_d (account)))$$





Formal Definition

- A basic expression in the relational algebra consists of either one of the following:
 - A relation in the database
 - A constant relation
- Let E_1 and E_2 be relational-algebra expressions; the following are all relational-algebra expressions:
 - $E_1 \cup E_2$
 - $E_1 - E_2$
 - $E_1 \times E_2$
 - $\sigma_P(E_1)$, P is a predicate on attributes in E_1
 - $\Pi_S(E_1)$, S is a list consisting of some of the attributes in E_1
 - $\rho_x(E_1)$, x is the new name for the result of E_1






Additional Operations

We define additional operations that do not add any power to the relational algebra, but that simplify common queries.


- Set intersection
- Natural join
- Division
- Assignment





Set-Intersection Operation

- Notation: $r \cap s$
- Defined as:
- $r \cap s = \{ t \mid t \in r \text{ and } t \in s \}$
- Assume:
 - r, s have the *same arity*
 - attributes of r and s are compatible
- Note: $r \cap s = r - (r - s)$





Set-Intersection Operation – Example

- 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





Natural-Join Operation

- Notation: $r \bowtie s$
- Let r and s be relations on schemas R and S respectively. Then, $r \bowtie s$ 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


- Example:

$R = (A, B, C, D)$

$S = (E, B, D)$

- Result schema = (A, B, C, D, E)
- $r \bowtie s$ is defined as:

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





Natural Join Operation – 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


r

B	D	E
1	a	α
3	a	β
1	a	γ
2	b	δ
3	b	ϵ

s

- $r \bowtie s$

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





Division Operation


- Notation: $r \div s$
- Suited to queries that include the phrase “for all”.
- Let r and s be relations on schemas R and S respectively where
 - $R = (A_1, \dots, A_m, B_1, \dots, B_n)$
 - $S = (B_1, \dots, B_n)$

The result of $r \div s$ is a relation on schema

$R - S = (A_1, \dots, A_m)$

$$r \div s = \{ t \mid t \in \Pi_{R-S}(r) \wedge \forall u \in s (tu \in r) \}$$

where tu means the concatenation of tuples t and u to produce a single tuple.





Division Operation – Example


- Relations r, s :

r	A	B
	α	1
	α	2
	α	3
	β	1
	γ	1
	δ	1
	δ	3
	δ	4
	ϵ	6
	ϵ	1
	β	2

s	B
	1
	2

- $r \div s$

A
α
β





Another Division Example


- Relations r , s :

r	A	B	C	D	E
	α	a	α	a	1
	α	a	γ	a	1
	α	a	γ	b	1
	β	a	γ	a	1
	β	a	γ	b	3
	γ	a	γ	a	1
	γ	a	γ	b	1
	γ	a	β	b	1

s	D	E
	a	1
	b	1

- $r \div s$

A	B	C
α	a	γ
γ	a	γ






Division Operation (Cont.)

- Property
 - Let $q = r \div s$
 - Then q is the largest relation satisfying $q \times s \subseteq r$
- Definition in terms of the basic algebra operation
Let $r(R)$ and $s(S)$ be relations, and let $S \subseteq R$

$$r \div s = \Pi_{R-S}(r) - \Pi_{R-S}((\Pi_{R-S}(r) \times s) - \Pi_{R-S,S}(r))$$

To see why

- $\Pi_{R-S,S}(r)$ simply reorders attributes of r
- $\Pi_{R-S}(\Pi_{R-S}(r) \times s) - \Pi_{R-S,S}(r)$ gives those tuples t in $\Pi_{R-S}(r)$ such that for some tuple $u \in s$, $tu \notin r$.





Assignment Operation


- The assignment operation (\leftarrow) provides a convenient way to express complex queries.
 - Write query as a sequential program consisting of
 - ▶ a series of assignments
 - ▶ followed by an expression whose value is displayed as a result of the query.
 - Assignment must always be made to a temporary relation variable.
- Example: Write $r \div s$ as

$$temp1 \leftarrow \Pi_{R-S}(r)$$

$$temp2 \leftarrow \Pi_{R-S}((temp1 \times s) - \Pi_{R-S,S}(r))$$

$$result = temp1 - temp2$$

- The result to the right of the \leftarrow is assigned to the relation variable on the left of the \leftarrow .
- May use variable in subsequent expressions.






Bank Example Queries

- Find the names of all customers who have a loan and an account at bank.

$$\Pi_{customer_name} (borrower) \cap \Pi_{customer_name} (depositor)$$

- Find the name of all customers who have a loan at the bank and the loan amount

$$\Pi_{customer_name, loan_number, amount} (borrower \bowtie loan)$$





Bank Example Queries

- Find all customers who have an account from at least the “Downtown” and the Uptown” branches.


- Query 1

$$\Pi_{customer_name} (\sigma_{branch_name = \text{“Downtown”}} (depositor \bowtie account)) \cap \\ \Pi_{customer_name} (\sigma_{branch_name = \text{“Uptown”}} (depositor \bowtie account))$$

- Query 2

$$\Pi_{customer_name, branch_name} (depositor \bowtie account) \\ \div \rho_{temp(branch_name)} (\{\text{“Downtown”}, \text{“Uptown”}\})$$

Note that Query 2 uses a constant relation.






Bank Example Queries

- Find all customers who have an account at all branches located in Brooklyn city.

$$\begin{aligned} & \Pi_{customer_name, branch_name} (depositor \bowtie account) \\ & \div \Pi_{branch_name} (\sigma_{branch_city = \text{"Brooklyn"}} (branch)) \end{aligned}$$





Extended Relational-Algebra Operations

- Generalized Projection
- Aggregate Functions
- Outer Join






Generalized Projection

- Extends the projection operation by allowing arithmetic functions to be used in the projection list.

$$\Pi_{F_1, F_2, \dots, F_n}(E)$$

- E is any relational-algebra expression
- Each of F_1, F_2, \dots, F_n are arithmetic expressions involving constants and attributes in the schema of E .
- Given relation $credit_info(customer_name, limit, credit_balance)$, find how much more each person can spend:

$$\Pi_{customer_name, limit - credit_balance}(credit_info)$$





Aggregate Functions and Operations

- **Aggregation function** takes a collection of values and returns a single value as a result.

avg: average value

min: minimum value

max: maximum value

sum: sum of values


count: number of values

- **Aggregate operation** in relational \mathcal{G} algebra

$$G_1, G_2, \dots, G_n \mathcal{G}_{F_1(A_1), F_2(A_2), \dots, F_n(A_n)}(E)$$

E is any relational-algebra expression

- G_1, G_2, \dots, G_n is a list of attributes on which to group (can be empty)
- Each F_i is an aggregate function
- Each A_i is an attribute name






Aggregate Operation – Example

- Relation r

A	B	C
α	α	7
α	β	7
β	β	3
β	β	10

- $\mathcal{G}_{\text{sum}(C)}(r)$

sum(C)
27






Aggregate Operation – Example

- Relation *account* grouped by *branch-name*:

<i>branch_name</i>	<i>account_number</i>	<i>balance</i>
Perryridge	A-102	400
Perryridge	A-201	900
Brighton	A-217	750
Brighton	A-215	750
Redwood	A-222	700

branch_name \mathcal{G} **sum**(*balance*) (*account*)

<i>branch_name</i>	sum (<i>balance</i>)
Perryridge	1300
Brighton	1500
Redwood	700





Aggregate Functions (Cont.)

- Result of aggregation does not have a name
 - Can use rename operation to give it a name

$$\rho_{x(\text{branch_name}, \text{sum_balance})} \left(\text{branch_name } \mathcal{G} \text{ sum}(\text{balance}) (\text{account}) \right)$$

- For convenience, we permit renaming as part of aggregate operation

$$\text{branch_name } \mathcal{G} \text{ sum}(\text{balance}) \text{ as } \text{sum_balance} (\text{account})$$




Outer Join

- An extension of the join operation that avoids loss of information.
- Example of natural join:

loan


<i>loan_number</i>	<i>branch_name</i>	<i>amount</i>
L-170	Downtown	3000
L-230	Redwood	4000
L-260	Perryridge	1700

borrower

<i>customer_name</i>	<i>loan_number</i>
Jones	L-170
Smith	L-230
Hayes	L-155

loan ⋈ *borrower*


<i>loan_number</i>	<i>branch_name</i>	<i>amount</i>	<i>customer_name</i>
L-170	Downtown	3000	Jones
L-230	Redwood	4000	Smith





Outer Join (cont.)

- Computes the join and then adds tuples from one relation that does not match tuples in the other relation to the result of the join.
- Uses *null* values:
 - *null* signifies that the value is unknown or does not exist
 - All comparisons involving *null* are (roughly speaking) **false** by definition.
 - ▶ We shall study precise meaning of comparisons with nulls later





Left Outer Join – Example

■ Left Outer Join

loan


<i>loan_number</i>	<i>branch_name</i>	<i>amount</i>
L-170	Downtown	3000
L-230	Redwood	4000
L-260	Perryridge	1700

borrower

<i>customer_name</i>	<i>loan_number</i>
Jones	L-170
Smith	L-230
Hayes	L-155

loan ⋈ *borrower*

<i>loan_number</i>	<i>branch_name</i>	<i>amount</i>	<i>customer_name</i>
L-170	Downtown	3000	Jones
L-230	Redwood	4000	Smith
L-260	Perryridge	1700	<i>null</i>





Right Outer Join – Example

■ Right Outer Join

loan


<i>loan_number</i>	<i>branch_name</i>	<i>amount</i>
L-170	Downtown	3000
L-230	Redwood	4000
L-260	Perryridge	1700

borrower

<i>customer_name</i>	<i>loan_number</i>
Jones	L-170
Smith	L-230
Hayes	L-155

loan ⋈_⊃ *borrower*

<i>loan_number</i>	<i>branch_name</i>	<i>amount</i>	<i>customer_name</i>
L-170	Downtown	3000	Jones
L-230	Redwood	4000	Smith
L-155	<i>null</i>	<i>null</i>	Hayes





Full Outer Join – Example

■ Full Outer Join

loan


<i>loan_number</i>	<i>branch_name</i>	<i>amount</i>
L-170	Downtown	3000
L-230	Redwood	4000
L-260	Perryridge	1700

borrower

<i>customer_name</i>	<i>loan_number</i>
Jones	L-170
Smith	L-230
Hayes	L-155

loan ⋈ *borrower*


<i>loan_number</i>	<i>branch_name</i>	<i>amount</i>	<i>customer_name</i>
L-170	Downtown	3000	Jones
L-230	Redwood	4000	Smith
L-260	Perryridge	1700	<i>null</i>
L-155	<i>null</i>	<i>null</i>	Hayes





Modification of the Database

- The content of the database may be modified using the following operations:
 - Deletion
 - Insertion
 - Updating
- All these operations are expressed using the assignment operator.






Deletion

- A delete request is expressed similarly to a query, except instead of displaying tuples to the user, the selected tuples are removed from the database.
- Can delete only whole tuples; cannot delete values on only particular attributes
- A deletion is expressed in relational algebra by:

$$r \leftarrow r - E$$

where r is a relation and E is a relational algebra query.





Deletion Examples

- Delete all account records in the Perryridge branch.

$account \leftarrow account - \sigma_{branch_name = \text{“Perryridge”}}(account)$

- Delete all loan records with amount in the range of 0 to 50

$loan \leftarrow loan - \sigma_{amount \geq 0 \text{ and } amount \leq 50}(loan)$

- Delete all accounts at branches located in Needham.

branch (*branch_name*, *branch_city*, *assets*)

account (*account_number*, *branch_name*, *balance*)

depositor (*customer_name*, *account_number*)


$r_1 \leftarrow \sigma_{branch_city = \text{“Needham”}}(account \bowtie branch)$

$r_2 \leftarrow \Pi_{account_number, branch_name, balance}(r_1)$

$r_3 \leftarrow \Pi_{customer_name, account_number}(r_2 \bowtie depositor)$

$account \leftarrow account - r_2$

$depositor \leftarrow depositor - r_3$





Insertion

- To insert data into a relation, we either:
 - specify a tuple to be inserted
 - write a query whose result is a set of tuples to be inserted
- In relational algebra, an insertion is expressed by:

$$r \leftarrow r \cup E$$

where r is a relation and E is a relational algebra expression.

- The insertion of a single tuple is expressed by letting E be a constant relation containing one tuple.





Insertion Examples

- Insert information in the database specifying that Smith has \$1200 in account A-973 at the Perryridge branch.

$$account \leftarrow account \cup \{("A-973", "Perryridge", 1200)\}$$
$$depositor \leftarrow depositor \cup \{("Smith", "A-973")\}$$


- Provide as a gift for all loan customers in the Perryridge branch, a \$200 savings account. Let the loan number serve as the account number for the new savings account.

account (*account_number*, *branch_name*, *balance*)

loan (*loan_number*, *branch_name*, *amount*)

depositor (*customer_name*, *account_number*)

borrower (*customer_name*, *loan_number*)

$$r_1 \leftarrow (\sigma_{branch_name = "Perryridge"}(borrower \bowtie loan))$$
$$account \leftarrow account \cup \Pi_{loan_number, branch_name, 200}(r_1)$$
$$depositor \leftarrow depositor \cup \Pi_{customer_name, loan_number}(r_1)$$





Updating

- A mechanism to change a value in a tuple without changing *all* values in the tuple
- Use the generalized projection operator to do this task

$$r \leftarrow \prod_{F_1, F_2, \dots, F_n} (r)$$

- Each F_i is either
 - the i^{th} attribute of r , if the i^{th} attribute is not updated, or,
 - if the attribute is to be updated F_i is an expression, involving only constants and the attributes of r , which gives the new value for the attribute





Update Examples


account (*account_number*, *branch_name*, *balance*)

- Make interest payments by increasing all balances by 5 percent.

$account \leftarrow \Pi_{account_number, branch_name, balance * 1.05} (account)$

- Pay all accounts with balances over \$10,000 6 percent interest and pay all others 5 percent

$account \leftarrow \Pi_{account_number, branch_name, balance * 1.06} (\sigma_{balance > 10000} (account)) \cup \Pi_{account_number, branch_name, balance * 1.05} (\sigma_{balance \leq 10000} (account))$






Views

- In some cases, it is not desirable for all users to see the entire logical model (that is, all the actual relations stored in the database.)
- Consider a person who needs to know a customer's name and loan number, but has no need to see the loan amount. This person should see a relation described, in relational algebra, by

$$\Pi_{customer_name, loan_number} (borrower \bowtie loan)$$

- A **view** provides a mechanism to hide certain data from the view of certain users.
- Any relation that is not of the conceptual model but is made visible to a user as a “virtual relation” is called a **view**.





View Definition

- A view is defined using the **create view** statement which has the form

create view *v* **as** < query expression >

where <query expression> is any legal relational algebra expression.
The view name is represented by *v*.

- Once a view is defined, the view name can be used to refer to the virtual relation that the view generates.
- When a view is created, the *query expression* is stored in the database; the expression is substituted into queries using the view.
 - So view is not the same as creating a new relation by evaluation the query expression.






Example Queries

- A view consisting of branches and their customers

create view *all_customer* as

$$\begin{aligned} & \Pi_{branch_name, customer_name} (depositor \bowtie account) \\ & \cup \\ & \Pi_{branch_name, customer_name} (borrower \bowtie loan) \end{aligned}$$


- Find all customers of the Perryridge branch

$$\Pi_{customer_name} (\sigma_{branch_name = 'Perryridge'} (all_customer))$$




Views Defined Using Other Views

- One view may be used in the expression defining another view
- A view relation v_1 is said to *depend directly on* a view relation v_2 if v_2 is used in the expression defining v_1
- A view relation v_1 is said to *depend on* view relation v_2 if either v_1 depends directly to v_2 or there is a path of dependencies from v_1 to v_2
- A view relation v is said to be *recursive* if it depends on itself.





View Expansion

- A way to define the meaning of views defined in terms of other views.
- Let view v_1 be defined by an expression e_1 that may itself contain uses of view relations.
- View expansion of an expression repeats the following replacement step:
 - repeat**
 - Find any view relation v_i in e_1
 - Replace the view relation v_i by the expression defining v_i
 - until** no more view relations are present in e_1
- As long as the view definitions are not recursive, this loop will terminate






Update of a View

- Database modifications expressed as views must be translated to modifications of the actual relations in the database.
- Consider the person who needs to see all loan data in the loan relation except amount. The view given to the person, *branch_loan*, is defined as:

create view *loan_branch* as

$$\Pi_{loan_number, branch_name}(loan)$$

- Since we allow a view name to appear wherever a relation name is allowed, the user may write:

$$loan_branch \leftarrow loan_brach \cup \{('L-37', 'Perryridge')\}$$





Tuple Relational Calculus

- A nonprocedural query language, where each query is of the form

$$\{t \mid P(t)\}$$

- It is the set of all tuples t such that predicate P is true for t
- t is a *tuple variable*, $t[A]$ denotes the value of tuple t on attribute A
- $t \in r$ denotes that tuple t is in relation r
- P is a *formula* similar to that of the predicate calculus






Predicate Calculus Formula

1. Set of attributes and constants
2. Set of comparison operators: (e.g., $<$, \leq , $=$, \neq , $>$, \geq)
3. Set of connectives: and (\wedge), or (\vee), not (\neg)
4. Implication (\Rightarrow): $x \Rightarrow y$, if x is true, then y is true

$$x \Rightarrow y \equiv \neg x \vee y$$

5. Set of quantifiers:


- ▶ $\exists t \in r (Q (t)) \equiv$ "there exists" a tuple t in relation r such that predicate $Q (t)$ is true
- ▶ $\forall t \in r (Q (t)) \equiv$ Q is true "for all" tuples t in relation r





Example Queries

- $loan (loan_number, branch_name, amount)$
- Find the $loan_number$, $branch_name$, and $amount$ for loans of over \$1200
$$\{t \mid t \in loan \wedge t[amount] > 1200\}$$
- Find the loan number for each loan of an amount greater than \$1200
$$\{t \mid \exists s \in loan (t[loan_number] = s[loan_number] \wedge s[amount] > 1200)\}$$
 - Notice that a relation on schema $(loan_number)$ is implicitly defined by the query.
- *Relation schema* of an expression is determined by either of:
 - If $t \in r$ is present in the expression, the resulting schema is of r
 - Otherwise the resulting schema is determined by all attributes of t used in the expression.
 - Note: If $t[A]$ is used more than once, the attribute A is in the relation schema just once!!!






Example Queries

- *depositor* (*customer_name*, *account_number*)
- *borrower* (*customer_name*, *loan_number*)
- Find the names of all customers having a loan, an account, or both at the bank

$$\{t \mid \exists s \in \text{borrower} (t [\text{customer_name}] = s [\text{customer_name}]) \\ \vee \exists u \in \text{depositor} (t [\text{customer_name}] = u [\text{customer_name}]) \}$$

- Find the names of all customers who have a loan and an account at the bank

$$\{t \mid \exists s \in \text{borrower} (t [\text{customer_name}] = s [\text{customer_name}]) \\ \wedge \exists u \in \text{depositor} (t [\text{customer_name}] = u [\text{customer_name}]) \}$$





Example Queries

- *loan* (*loan_number*, *branch_name*, *amount*)
- *depositor* (*customer_name*, *account_number*)
- *borrower* (*customer_name*, *loan_number*)

- Find the names of all customers having a loan at the Perryridge branch
$$\{t \mid \exists s \in \textit{borrower} (t[\textit{customer_name}] = s[\textit{customer_name}] \wedge \exists u \in \textit{loan} (u[\textit{branch_name}] = \text{"Perryridge"} \wedge u[\textit{loan_number}] = s[\textit{loan_number}])))\}$$


- Find the names of all customers who have a loan at the Perryridge branch, but no account at any branch of the bank
$$\{t \mid \exists s \in \textit{borrower} (t[\textit{customer_name}] = s[\textit{customer_name}] \wedge \exists u \in \textit{loan} (u[\textit{branch_name}] = \text{"Perryridge"} \wedge u[\textit{loan_number}] = s[\textit{loan_number}])) \wedge \neg \exists v \in \textit{depositor} (v[\textit{customer_name}] = t[\textit{customer_name}])\}$$





Example Queries


- *branch* (*branch_name*, *branch_city*, *assets*)
 - *customer* (*customer_name*, *customer_street*, *customer_city*)
 - *account* (*account_number*, *branch_name*, *balance*)
 - *loan* (*loan_number*, *branch_name*, *amount*)
 - *depositor* (*customer_name*, *account_number*)
 - *borrower* (*customer_name*, *loan_number*)
-
- Find the names of all customers having a loan at the Perryridge branch, and the cities in which they live

$$\{t \mid \exists s \in \text{loan} (s [\text{branch_name}] = \text{"Perryridge"} \\ \wedge \exists u \in \text{borrower} (u [\text{loan_number}] = s [\text{loan_number}] \\ \wedge t [\text{customer_name}] = u [\text{customer_name}]) \\ \wedge \exists v \in \text{customer} (u [\text{customer_name}] = v [\text{customer_name}] \\ \wedge t [\text{customer_city}] = v [\text{customer_city}]))) \}$$




Example Queries


- *branch* (*branch_name*, *branch_city*, *assets*)
customer (*customer_name*, *customer_street*, *customer_city*)
account (*account_number*, *branch_name*, *balance*)
loan (*loan_number*, *branch_name*, *amount*)
depositor (*customer_name*, *account_number*)
borrower (*customer_name*, *loan_number*)
- Find the names of all customers who have an account at all branches located in Brooklyn:

$$\{t \mid \exists r \in \text{customer} (t[\text{customer_name}] = r[\text{customer_name}]) \wedge$$
$$(\forall u \in \text{branch} (u[\text{branch_city}] = \text{"Brooklyn"} \Rightarrow$$
$$\exists s \in \text{depositor} (r[\text{customer_name}] = s[\text{customer_name}])$$
$$\wedge \exists w \in \text{account} (w[\text{account_number}] = s[\text{account_number}])$$
$$\wedge (w[\text{branch_name}] = u[\text{branch_name}]))))$$
$$\})$$




Safety of Expressions

- It is possible to write tuple calculus expressions that generate infinite relations.
- For example, $\{ t \mid \neg t \in r \}$ results in an infinite relation if the domain of any attribute of relation r is infinite
- To guard against the problem, we restrict the set of allowable expressions to safe expressions.
- An expression $\{ t \mid P(t) \}$ in the tuple relational calculus is *safe* if every component of t appears in one of the relations, tuples, or constants that appear in P
 - NOTE: this is more than just a syntax condition.
 - ▶ E.g. $\{ t \mid t[A] = 5 \vee \mathbf{true} \}$ is not safe – it defines an infinite set with attribute values that do not appear in any relation or tuples or constants in P .






Domain Relational Calculus

- A nonprocedural query language equivalent in power to the tuple relational calculus
- Each query is an expression of the form:

$$\{ \langle x_1, x_2, \dots, x_n \rangle \mid P(x_1, x_2, \dots, x_n) \}$$

- x_1, x_2, \dots, x_n represent domain variables
- P represents a formula similar to that of the predicate calculus






Example Queries

- *loan* (*loan_number*, *branch_name*, *amount*)
depositor (*customer_name*, *account_number*)
borrower (*customer_name*, *loan_number*)

- Find the *loan_number*, *branch_name*, and *amount* for loans of over \$1200
 - $\{ \langle l, b, a \rangle \mid \langle l, b, a \rangle \in \textit{loan} \wedge a > 1200 \}$

- Find the names of all customers who have a loan of over \$1200
 - $\{ \langle c \rangle \mid \exists l, b, a (\langle c, l \rangle \in \textit{borrower} \wedge \langle l, b, a \rangle \in \textit{loan} \wedge a > 1200) \}$


- Find the names of all customers who have a loan at the Perryridge branch and the loan amount:
 - $\{ \langle c, a \rangle \mid \exists l (\langle c, l \rangle \in \textit{borrower} \wedge \exists b (\langle l, b, a \rangle \in \textit{loan} \wedge b = \text{"Perryridge"})) \}$
 - $\{ \langle c, a \rangle \mid \exists l (\langle c, l \rangle \in \textit{borrower} \wedge \langle l, \text{"Perryridge"}, a \rangle \in \textit{loan}) \}$





Example Queries

- *branch* (*branch_name*, *branch_city*, *assets*)
customer (*customer_name*, *customer_street*, *customer_city*)
account (*account_number*, *branch_name*, *balance*)
loan (*loan_number*, *branch_name*, *amount*)
depositor (*customer_name*, *account_number*)
borrower (*customer_name*, *loan_number*)
- Find the names of all customers having a loan, an account, or both at the Perryridge branch:
 - $\{ \langle c \rangle \mid \exists l (\langle c, l \rangle \in \textit{borrower} \wedge \exists b, a (\langle l, b, a \rangle \in \textit{loan} \wedge b = \text{“Perryridge”})) \vee \exists a (\langle c, a \rangle \in \textit{depositor} \wedge \exists b, n (\langle a, b, n \rangle \in \textit{account} \wedge b = \text{“Perryridge”}))) \}$
- Find the names of all customers who have an account at all branches located in Brooklyn:
 - $\{ \langle c \rangle \mid \exists s, t (\langle c, s, t \rangle \in \textit{customer}) \wedge \forall x, y, z (\langle x, y, z \rangle \in \textit{branch} \wedge y = \text{“Brooklyn”}) \Rightarrow \exists a, b (\langle a, x, b \rangle \in \textit{account} \wedge \langle c, a \rangle \in \textit{depositor}) \}$





Safety of Expressions

The expression:

$$\{ \langle x_1, x_2, \dots, x_n \rangle \mid P(x_1, x_2, \dots, x_n) \}$$

is safe if all of the following hold:

1. All values that appear in tuples of the expression are values from *dom* (*P*) (that is, the values appear either in *P* or in a tuple of a relation mentioned in *P*).
2. For every “there exists” subformula of the form $\exists x (P_1(x))$, the subformula is true if and only if there is a value of *x* in *dom* (*P*₁) such that *P*₁(*x*) is true.
3. For every “for all” subformula of the form $\forall x (P_1(x))$, the subformula is true if and only if *P*₁(*x*) is true for all values *x* from *dom* (*P*₁).

