

# SQL: DATA DEFINITION LANGUAGE



# Database Schemas in SQL

2

- SQL is primarily a query language, for getting information from a database.
  - ▣ **Data manipulation language (DML)**
- But SQL also includes a *data-definition* component for describing database schemas.
  - ▣ **Data definition language (DDL)**

# Creating (Declaring) a Relation

3

- Simplest form is:

```
CREATE TABLE <name> (  
    <list of elements>  
);
```

- To delete a relation:

```
DROP TABLE <name>;
```

# Elements of Table Declarations

4

- Most basic element: an attribute and its type.
- The most common types are:
  - ▣ INT or INTEGER (synonyms).
  - ▣ REAL or FLOAT (synonyms).
  - ▣ CHAR( $n$ ) = fixed-length string of  $n$  characters.
  - ▣ VARCHAR( $n$ ) = variable-length string of up to  $n$  characters.

# Example: Create Table

5

```
CREATE TABLE Sells (  
    bar          CHAR(20),  
    beer        VARCHAR(20),  
    price       REAL  
);
```

# SQL Values

6

- Integers and reals are represented as you would expect.
- Strings are too, except they require single quotes.
  - ▣ Two single quotes = real quote, e.g., 'Joe' 's Bar'.
- Any value can be NULL
  - ▣ Unless attribute has NOT NULL constraint
  - ▣ E.g., price REAL not null,

# Dates and Times

7

- DATE and TIME are types in SQL.

- The form of a date value is:

DATE 'yyyy-mm-dd'

- ▣ **Example:** DATE '2007-09-30' for Sept. 30, 2007.

# Times as Values

8

- The form of a time value is:

TIME 'hh:mm:ss'

with an optional decimal point and fractions of a second following.

- **Example:** TIME '15:30:02.5' = two and a half seconds after 3:30PM.



# Declaring Keys

- An attribute or list of attributes may be declared PRIMARY KEY or UNIQUE.
- Either says that no two tuples of the relation may agree in all the attribute(s) on the list.

# Our Running Example

10

Beers(name, manf)

Bars(name, addr, license)

Drinkers(name, addr, phone)

Likes(drinker, beer)

Sells(bar, beer, price)

Frequents(drinker, bar)

- Underline = *key* (tuples cannot have the same value in all key attributes).

# Declaring Single-Attribute Keys

11

- Place PRIMARY KEY or UNIQUE after the type in the declaration of the attribute.
- Example:

```
CREATE TABLE Beers (  
    name    CHAR(20) UNIQUE,  
    manf    CHAR(20)  
);
```

# Declaring Multiattribute Keys

- A key declaration can also be another element in the list of elements of a CREATE TABLE statement.
- This form is essential if the key consists of more than one attribute.
  - ▣ May be used even for one-attribute keys.

# Example: Multiattribute Key

13

- The bar and beer together are the key for Sells:

```
CREATE TABLE Sells (  
    bar        CHAR(20),  
    beer       VARCHAR(20),  
    price      REAL,  
    PRIMARY KEY (bar, beer)  
);
```

# PRIMARY KEY vs. UNIQUE

14

1. There can be only one PRIMARY KEY for a relation, but several UNIQUE attributes.
2. No attribute of a PRIMARY KEY can ever be NULL in any tuple. But attributes declared UNIQUE may have NULL's, and there may be several tuples with NULL.

# Kinds of Constraints

15

- Keys
- Foreign-key, or referential-integrity.
- Domain constraints
  - ▣ Constrain values of a particular attribute.
- Tuple-based constraints
  - ▣ Relationship among components.
- Assertions: any SQL boolean expression

# Foreign Keys

16

- Values appearing in attributes of one relation must appear together in certain attributes of another relation.
- **Example:** in `Sells(bar, beer, price)`, we might expect that a beer value also appears in `Beers.name`



# Expressing Foreign Keys

17

- Use keyword REFERENCES, either:
  1. After an attribute (for one-attribute keys).
  2. As an element of the schema:  
FOREIGN KEY (<list of attributes>  
REFERENCES <relation> (<attributes>)
- Referenced attributes must be declared PRIMARY KEY or UNIQUE.

# Example: With Attribute

18

```
CREATE TABLE Beers (  
  name      CHAR(20) PRIMARY KEY,  
  manf      CHAR(20) );
```

```
CREATE TABLE Sells (  
  bar       CHAR(20),  
  beer      CHAR(20) REFERENCES Beers(name),  
  price     REAL );
```

# Example: As Schema Element

19

```
CREATE TABLE Beers (  
    name      CHAR(20) PRIMARY KEY,  
    manf      CHAR(20) );
```

```
CREATE TABLE Sells (  
    bar       CHAR(20),  
    beer      CHAR(20),  
    price     REAL,  
    FOREIGN KEY (beer) REFERENCES  
        Beers (name) );
```

# Enforcing Foreign-Key Constraints

- If there is a foreign-key constraint from relation  $R$  to relation  $S$ , two violations are possible:
  1. An insert or update to  $R$  introduces values not found in  $S$ .
  2. A deletion or update to  $S$  causes some tuples of  $R$  to “dangle.”

# Actions Taken --- (1)

21

- **Example:** suppose  $R = \text{Sells}$ ,  $S = \text{Beers}$ .
- An insert or update to **Sells** that introduces a nonexistent beer must be rejected.
- A deletion or update to **Beers** that removes a beer value found in some tuples of **Sells** can be handled in three ways...

# Actions Taken --- (2)

22

1. *Default* : Reject the modification.
2. *Cascade* : Make the same changes in Sells.
  - ▣ *Deleted beer*: delete Sells tuple.
  - ▣ *Updated beer*: change value in Sells.
3. *Set NULL* : Change the beer to NULL.

# Example: Cascade

23

- Delete the Bud tuple from Beers:
  - ▣ Then delete all tuples from Sells that have beer = 'Bud'.
- Update the Bud tuple by changing 'Bud' to 'Budweiser':
  - ▣ Then change all Sells tuples with beer = 'Bud' to beer = 'Budweiser'.

# Example: Set NULL

24

- Delete the Bud tuple from Beers:
  - ▣ Change all tuples of Sells that have beer = 'Bud' to have beer = NULL.
- Update the Bud tuple by changing 'Bud' to 'Budweiser':
  - ▣ Same change as for deletion.



# Choosing a Policy

- When we declare a foreign key, we may choose policies `SET NULL` or `CASCADE` independently for deletions and updates.
- Follow the foreign-key declaration by:  
`ON [UPDATE, DELETE][SET NULL CASCADE]`
- Two such clauses may be used.
- Otherwise, the default (reject) is used.

# Example: Setting Policy

26

```
CREATE TABLE Sells (  
    bar      CHAR(20),  
    beer     CHAR(20),  
    price    REAL,  
    FOREIGN KEY (beer)  
        REFERENCES Beers (name)  
        ON DELETE SET NULL  
        ON UPDATE CASCADE  
);
```

# Attribute-Based Checks

27

- Constraints on the value of a particular attribute.
- Add `CHECK(<condition>)` to the declaration for the attribute.
- The condition may use the name of the attribute, but any other relation or attribute name must be in a subquery.

# Example: Attribute-Based Check

28

```
CREATE TABLE Sells (  
  bar      CHAR(20),  
  beer     CHAR(20) CHECK ( beer IN  
    (SELECT name FROM Beers) ),  
  price    REAL CHECK ( price <= 5.00 )  
);
```

# Timing of Checks

- Attribute-based checks are performed only when a value for that attribute is inserted or updated.
  - ▣ **Example:** `CHECK (price <= 5.00)` checks every new price and rejects the modification (for that tuple) if the price is more than \$5.
  - ▣ **Example:** `CHECK (beer IN (SELECT name FROM Beers))` not checked if a beer is deleted from Beers (unlike foreign-keys).

# Tuple-Based Checks

- CHECK (<condition>) may be added as a relation-schema element.
- The condition may refer to any attribute of the relation.
  - ▣ But other attributes or relations require a subquery.
- Checked on insert or update only.

# Example: Tuple-Based Check

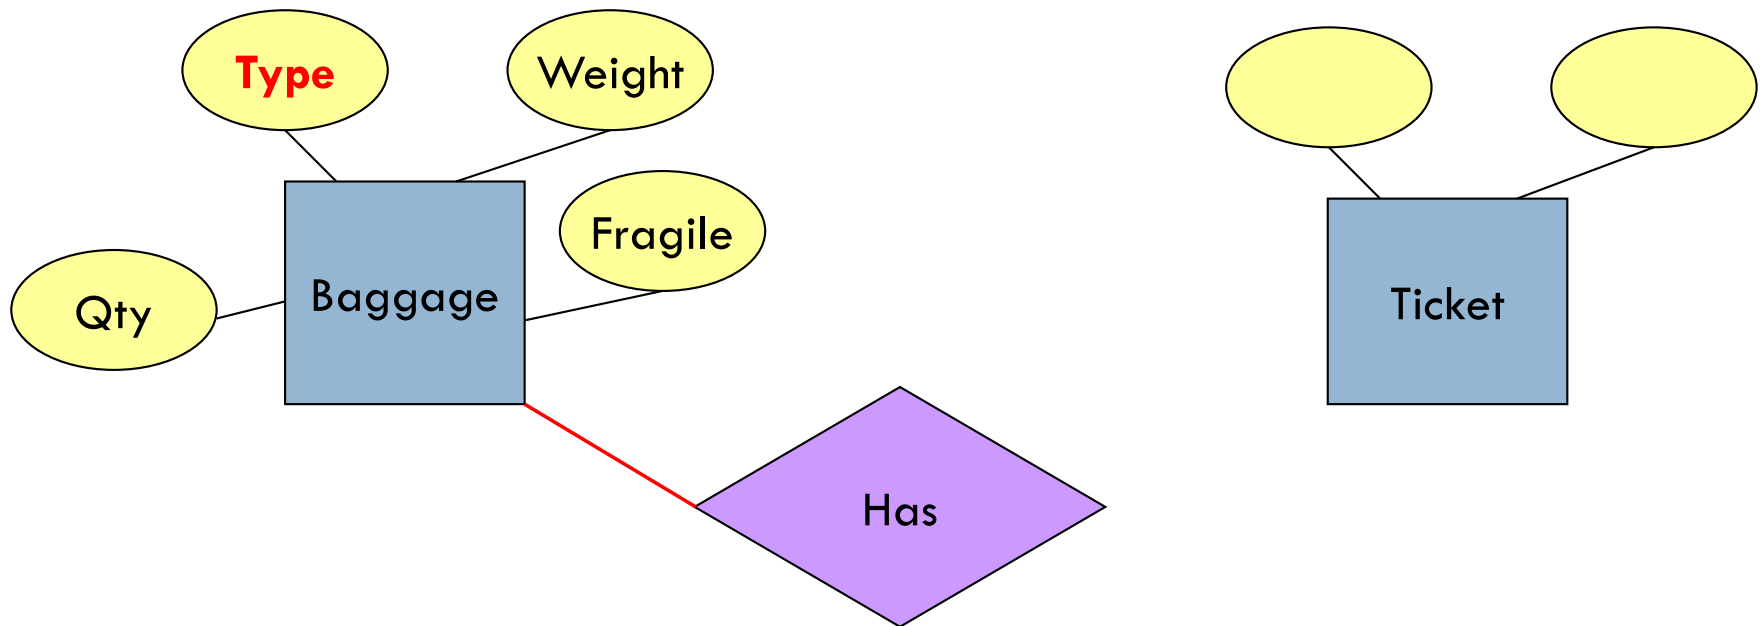
31

- Only Joe's Bar can sell beer for more than \$5:

```
CREATE TABLE Sells (  
    bar          CHAR(20),  
    beer         CHAR(20),  
    price        REAL,  
    CHECK (bar = 'Joe's Bar' OR  
           price <= 5.00)  
);
```

# Asg 1 Update: Missing attribute in Baggage

1





# INTRODUCTION TO SQL



# Why SQL?

3

- SQL is a very-high-level language.
  - ▣ Structured Query Language
  - ▣ Say “what to do” rather than “how to do it.”
  - ▣ Avoid a lot of data-manipulation details needed in procedural languages like C++ or Java.
- Database management system figures out “best” way to execute query.
  - ▣ Called “query optimization.”

# Database Schemas in SQL

4

- SQL is primarily a query language, for getting information from a database.
  - ▣ **Data manipulation language (DML)**
- But SQL also includes a *data-definition* component for describing database schemas.
  - ▣ **Data definition language (DDL)**

# Select-From-Where Statements

5

**SELECT** desired attributes

**FROM** one or more tables

**WHERE** condition about tuples of  
the tables

# Our Running Example

6

- Our SQL queries will be based on the following database schema.
  - ▣ Underline indicates key attributes.

Beers(name, manf)

Bars(name, addr, license)

Drinkers(name, addr, phone)

Likes(drinker, beer)

Sells(bar, beer, price)

Frequents(drinker, bar)

# Example

7

- Using `Beers(name, manf)`, what beers are made by Anheuser-Busch?

```
SELECT name  
FROM Beers  
WHERE manf = 'Anheuser-Busch';
```

# Result of Query

8

name

Bud
Bud Lite
Michelob
...

The answer is a relation with a single attribute, name, and tuples with the name of each beer by Anheuser-Busch, such as Bud.

# Meaning of Single-Relation Query

9

- Begin with the relation in the FROM clause.
- Apply the selection indicated by the WHERE clause.
- Apply the extended projection indicated by the SELECT clause.



# Operational Semantics - General

- Think of a *tuple variable* visiting each tuple of the relation mentioned in FROM.
- Check if the tuple assigned to the tuple variable satisfies the WHERE clause.
- If so, compute the attributes or expressions of the SELECT clause using the components of this tuple.

# Operational Semantics

11

name	manf
Bud	Anheuser-Busch

If so, include t.name  
in the result

Check if  
Anheuser-Busch

Tuple-variable  $t$   
loops over all  
tuples

# Example

12

- What beers are made by Anheuser-Busch?

```
SELECT name  
FROM Beers  
WHERE manf = 'Anheuser-Busch';
```

OR:

```
SELECT t.name  
FROM Beers t  
WHERE t.manf = 'Anheuser-Busch';
```

**Note:** these are identical queries.

## \* In SELECT clauses

13

- When there is one relation in the FROM clause, \* in the SELECT clause stands for “all attributes of this relation.”
- **Example:** Using **Beers(name, manf):**

```
SELECT *  
FROM Beers  
WHERE manf = 'Anheuser-Busch';
```

# Result of Query:

14

name	manf
Bud	Anheuser-Busch
Bud Lite	Anheuser-Busch
Michelob	Anheuser-Busch
...	...

Now, the result has each of the attributes of Beers.

# Renaming Attributes

15

- If you want the result to have different attribute names, use “AS <new name>” to rename an attribute.

- **Example:** Using `Beers(name, manf)`:

```
SELECT name AS beer, manf
FROM Beers
WHERE manf = 'Anheuser-Busch'
```

# Result of Query:

16

beer

manf

beer	manf
Bud	Anheuser-Busch
Bud Lite	Anheuser-Busch
Michelob	Anheuser-Busch
...	...

# Expressions in SELECT Clauses

17

- Any valid expression can appear as an element of a SELECT clause.

- **Example:** Using `Sells(bar, beer, price)`:

```
SELECT bar, beer,  
       price*95 AS priceInYen  
FROM Sells;
```



# Result of Query

18

bar	beer	priceInYen
Joe's	Bud	285
Sue's	Miller	342
...	...	...

# Example: Constants as Expressions

19

- Using `Likes(drinker, beer)`:

```
SELECT drinker,  
       'likes Bud' AS whoLikesBud  
FROM Likes  
WHERE beer = 'Bud';
```

# Result of Query

20

drinker	whoLikesBud
Sally	likes Bud
Fred	likes Bud
...	...

# Complex Conditions in WHERE Clause

21

- Boolean operators AND, OR, NOT.
- Comparisons =, <>, <, >, <=, >=.

# Example: Complex Condition

22

- Using `Sells(bar, beer, price)`, find the price Joe's Bar charges for Bud:

```
SELECT price
FROM Sells
WHERE bar = 'Joe''s Bar' AND
       beer = 'Bud';
```

# Patterns

23

- A condition can compare a string to a pattern by:
  - `<Attribute> LIKE <pattern>` or `<Attribute> NOT LIKE <pattern>`
- *Pattern* is a quoted string
  - `%` = “any string”;
  - `_` = “any character”.

# Example: LIKE

24

- Using `Drinkers(name, addr, phone)` find the drinkers with exchange 555:

```
SELECT name
FROM Drinkers
WHERE phone LIKE '%555-__ __ __';
```

# NULL Values

- Tuples in SQL relations can have NULL as a value for one or more components.
- Meaning depends on context. Two common cases:
  - ▣ *Missing value* : e.g., we know Joe's Bar has some address, but we don't know what it is.
  - ▣ *Inapplicable* : e.g., the value of attribute *spouse* for an unmarried person.



# Comparing NULL's to Values

- The logic of conditions in SQL is really 3-valued logic: TRUE, FALSE, UNKNOWN.
- Comparing any value (including NULL itself) with NULL yields UNKNOWN.
- A tuple is in a query answer iff the WHERE clause is TRUE (not FALSE or UNKNOWN).

# Three-Valued Logic

27

- To understand how AND, OR, and NOT work in 3-valued logic
- For TRUE result
  - ▣ OR: at least one operand must be TRUE
  - ▣ AND: both operands must be TRUE
  - ▣ NOT: operand must be FALSE
- For FALSE result
  - ▣ OR: both operands must be FALSE
  - ▣ AND: at least one operand must be FALSE
  - ▣ NOT: operand must be TRUE
- Otherwise, result is UNKNOWN

# Example

28

- From the following Sells relation:

bar	beer	price
Joe's Bar	Bud	NULL

SELECT bar

FROM Sells

WHERE price < 2.00 OR price >= 5.00;

← UNKNOWN →

← UNKNOWN →

← UNKNOWN →

# Multi-Relation Queries

- Interesting queries often combine data from more than one relation.
- We can address several relations in one query by listing them all in the FROM clause.
- Distinguish attributes of the same name by “<relation>.<attribute>” .

# Example: Joining Two Relations

30

- Using relations `Likes(drinker, beer)` and `Frequents(drinker, bar)`, find the beers liked by at least one person who frequents Joe's Bar.

```
SELECT beer
FROM Likes, Frequents
WHERE bar = 'Joe''s Bar' AND
      Frequents.drinker = Likes.drinker;
```

# Example: Joining Two Relations

31

- Alternatively can use explicit (named) tuple variables

```
SELECT beer
FROM Likes l, Frequents f
WHERE bar = 'Joe''s Bar' AND
       f.drinker = l.drinker;
```

# Formal Semantics

32

- Almost the same as for single-relation queries:
  - Start with the product of all the relations in the FROM clause.
  - Apply the selection condition from the WHERE clause.
  - Project onto the list of attributes and expressions in the SELECT clause.


# Operational Semantics

- Imagine one tuple-variable for each relation in the FROM clause.
  - ▣ These tuple-variables visit each combination of tuples, one from each relation.
- If the tuple-variables are pointing to tuples that satisfy the WHERE clause, send these tuples to the SELECT clause.



# Example

34



# Explicit Tuple-Variables

- Sometimes, a query needs to use two copies of the same relation.
- Distinguish copies by following the relation name by the name of a tuple-variable, in the FROM clause.
- It's always an option to rename relations this way, even when not essential.

# Example: Self-Join

36

- From **Beers(name, manf)**, find all pairs of beers by the same manufacturer.
  - ▣ Do not produce pairs like (Bud, Bud).
  - ▣ Do not produce the same pair twice like (Bud, Miller) and (Miller, Bud).

# Select-From-Where Statements

1

**SELECT** desired attributes

**FROM** one or more tables

**WHERE** condition about tuples of  
the tables

# Example

2

- Using `Beers(name, manf)`, what beers are made by Anheuser-Busch?

```
SELECT name  
FROM Beers  
WHERE manf = 'Anheuser-Busch';
```

# Result of Query

3

name

Bud
Bud Lite
Michelob
...

The answer is a relation with a single attribute, name, and tuples with the name of each beer by Anheuser-Busch, such as Bud.

# Operational Semantics

4

name	manf
Bud	Anheuser-Busch

If so, include t.name  
in the result

Check if  
Anheuser-Busch

Tuple-variable  $t$   
loops over all  
tuples

# Comparing NULL's to Values

5

- The logic of conditions in SQL is really 3-valued logic: TRUE, FALSE, UNKNOWN.
- Comparing any value (including NULL itself) with NULL yields UNKNOWN.
- A tuple is in a query answer iff the WHERE clause is TRUE (not FALSE or UNKNOWN).



# Three-Valued Logic

6

- To understand how AND, OR, and NOT work in 3-valued logic
- For TRUE result
  - ▣ OR: at least one operand must be TRUE
  - ▣ AND: both operands must be TRUE
  - ▣ NOT: operand must be FALSE
- For FALSE result
  - ▣ OR: both operands must be FALSE
  - ▣ AND: at least one operand must be FALSE
  - ▣ NOT: operand must be TRUE
- Otherwise, result is UNKNOWN

# Example

7

- From the following Sells relation:

bar	beer	price
Joe's Bar	Bud	NULL

SELECT bar

FROM Sells

WHERE price < 2.00 OR price >= 5.00;

← UNKNOWN →

← UNKNOWN →

← UNKNOWN →

# Multi-Relation Queries

8

- Interesting queries often combine data from more than one relation.
- We can address several relations in one query by listing them all in the FROM clause.
- Distinguish attributes of the same name by “<relation>.<attribute>” .

# Example: Joining Two Relations

9

- Using relations `Likes(drinker, beer)` and `Frequents(drinker, bar)`, find the beers liked by at least one person who frequents Joe's Bar.

```
SELECT beer
FROM Likes, Frequents
WHERE bar = 'Joe''s Bar' AND
       Frequents.drinker = Likes.drinker;
```

# Example: Joining Two Relations

10

- Alternatively can use explicit (named) tuple variables

```
SELECT beer
FROM Likes l, Frequents f
WHERE bar = 'Joe''s Bar' AND
       f.drinker = l.drinker;
```

# Formal Semantics

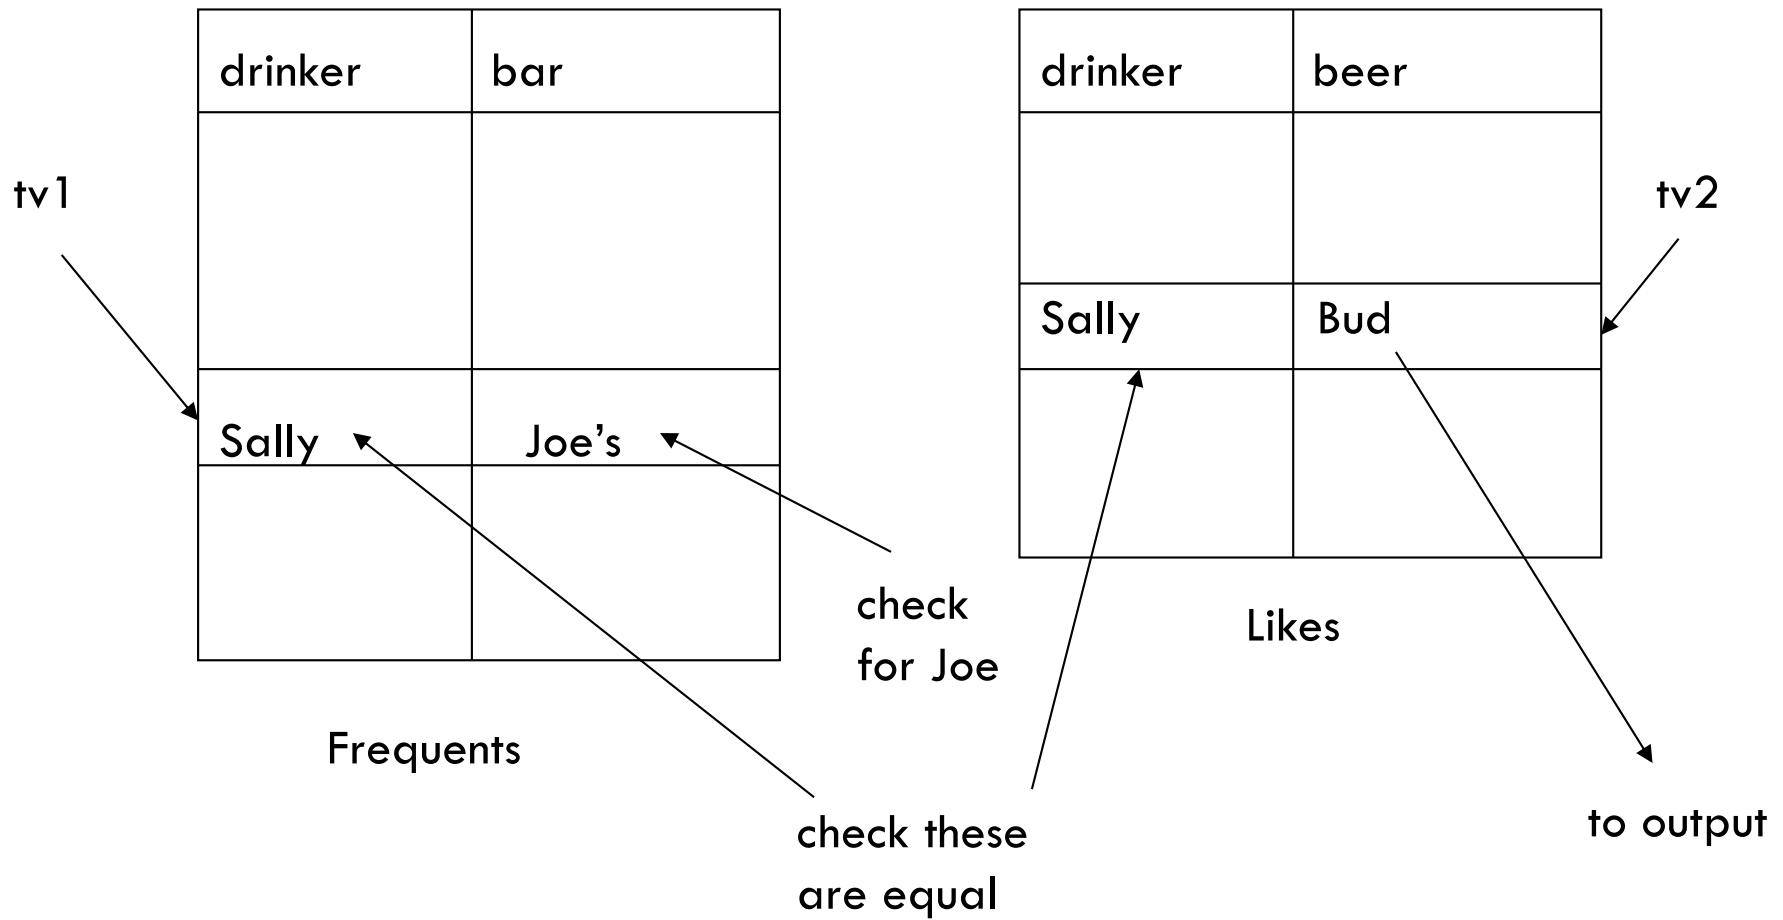
- Almost the same as for single-relation queries:
  - Start with the product of all the relations in the FROM clause.
  - Apply the selection condition from the WHERE clause.
  - Project onto the list of attributes and expressions in the SELECT clause.

# Operational Semantics

- Imagine one tuple-variable for each relation in the FROM clause.
  - ▣ These tuple-variables visit each combination of tuples, one from each relation.
- If the tuple-variables are pointing to tuples that satisfy the WHERE clause, send these tuples to the SELECT clause.

# Example

13





# Explicit Tuple-Variables

- Sometimes, a query needs to use two copies of the same relation.
- Distinguish copies by following the relation name by the name of a tuple-variable, in the FROM clause.
- It's always an option to rename relations this way, even when not essential.

# Example: Self-Join

15

- From **Beers(name, manf)**, find all pairs of beers by the same manufacturer.
  - ▣ Do not produce pairs like (Bud, Bud).
  - ▣ Do not produce the same pair twice like (Bud, Miller) and (Miller, Bud).

```
SELECT b1.name, b2.name
FROM Beers b1, Beers b2
WHERE b1.manf = b2.manf AND
      b1.name < b2.name;
```

# Subqueries

- A parenthesized SELECT-FROM-WHERE statement (*subquery*) can be used as a value in a number of places, including FROM and WHERE clauses.
- **Example:** in place of a relation in the FROM clause, we can use a subquery and then query its result.
  - Must use a tuple-variable to name tuples of the result.

# Example: Subquery in FROM

17

- Find the beers liked by at least one person who frequents Joe's Bar.

Drinkers who  
frequents Joe's Bar



```
SELECT beer
```

```
FROM Likes, (SELECT drinker  
FROM Frequents  
WHERE bar = 'Joe''s Bar') JD
```

```
WHERE Likes.drinker = JD.drinker;
```

# Subqueries often obscure queries

18

- Find the beers liked by at least one person who frequents Joe's Bar.

```
SELECT beer
FROM Likes l, Frequents f
WHERE l.drinker = f.drinker AND
      bar = 'Joe''s Bar';
```

Simple join query

# Subqueries That Return One Tuple

- If a subquery is guaranteed to produce one tuple, then the subquery can be used as a value.
  - ▣ Usually, the tuple has one component.
  - ▣ Remember SQL's 3-valued logic.

# Example: Single-Tuple Subquery

20

- Using `Sells(bar, beer, price)`, find the bars that serve Miller for the same price Joe charges for Bud.

Two queries would work:

- Find the price Joe charges for Bud.
- Find the bars that serve Miller at that price.

# Query + Subquery Solution

21

SELECT bar

FROM Sells


WHERE beer = 'Miller' AND price

- Find the price Joe charges for Bud.
- Find the bars that serve Miller at that price.

Sells(bar, beer, price)

```
= (SELECT price
   FROM Sells
   WHERE bar = 'Joe''s Bar'
        AND beer = 'Bud');
```

The price at  
which Joe  
sells Bud



What if price of Bud is NULL?



# Query + Subquery Solution

22

```
SELECT bar  
FROM Sells  
WHERE beer = 'Miller' AND
```

```
price = (SELECT price  
        FROM Sells  
        WHERE beer = 'Bud');
```

What if subquery  
returns multiple  
values?



# Recap: Conditions in WHERE Clause

23

- Boolean operators AND, OR, NOT.
- Comparisons =, <>, <, >, <=, >=.
- LIKE operator
  
- SQL includes a **between** comparison operator
- Example: Find the names of all instructors with salary between \$90,000 and \$100,000 (that is,  $\geq$  \$90,000 and  $\leq$  \$100,000)
  - **select** *name*  
**from** *instructor*  
**where** *salary* **between** 90000 **and** 100000

# Subqueries

36

- A parenthesized SELECT-FROM-WHERE statement (*subquery*) can be used as a value in a number of places, including FROM and WHERE clauses.
- **Example:** in place of a relation in the FROM clause, we can use a subquery and then query its result.
  - Must use a tuple-variable to name tuples of the result.

# Example: Subquery in FROM

37

- Find the beers liked by at least one person who frequents Joe's Bar.

Drinkers who  
frequents Joe's Bar

```
SELECT beer
```

```
FROM Likes, (SELECT drinker  
FROM Frequent  
WHERE bar = 'Joe's Bar') JD
```

```
WHERE Likes.drinker = JD.drinker;
```

# Subqueries often obscure queries

38

- Find the beers liked by at least one person who frequents Joe's Bar.

```
SELECT beer
FROM Likes l, Frequents f
WHERE l.drinker = f.drinker AND
      bar = 'Joe''s Bar';
```

Simple join query

# Subqueries That Return One Tuple

39

- If a subquery is guaranteed to produce one tuple, then the subquery can be used as a value.
  - ▣ Usually, the tuple has one component.
  - ▣ Remember SQL's 3-valued logic.

# Example: Single-Tuple Subquery

40

- Using `Sells(bar, beer, price)`, find the bars that serve Miller for the same price Joe charges for Bud.

Two queries would work:

- Find the price Joe charges for Bud.
- Find the bars that serve Miller at that price.

# Query + Subquery Solution

41

SELECT bar

FROM Sells


WHERE beer = 'Miller' AND price

- Find the price Joe charges for Bud.
- Find the bars that serve Miller at that price.

Sells(bar, beer, price)

```
= (SELECT price
   FROM Sells
   WHERE bar = 'Joe''s Bar'
        AND beer = 'Bud');
```

The price at  
which Joe  
sells Bud



What if price of Bud is NULL?



# Query + Subquery Solution

42

```
SELECT bar  
FROM Sells  
WHERE beer = 'Miller' AND
```

```
price = (SELECT price  
        FROM Sells  
        WHERE beer = 'Bud');
```

What if subquery  
returns multiple  
values?



# Recap: Conditions in WHERE Clause

43

- Boolean operators AND, OR, NOT.
- Comparisons =, <>, <, >, <=, >=.
- LIKE operator
  
- SQL includes a **between** comparison operator
- Example: Find the names of all instructors with salary between \$90,000 and \$100,000 (that is,  $\geq$  \$90,000 and  $\leq$  \$100,000)
  - **select** *name*  
**from** *instructor*  
**where** *salary* **between** 90000 **and** 100000

# The Operator ANY

- $x = \text{ANY}(\langle \text{subquery} \rangle)$  is a boolean condition that is true iff  $x$  equals at least one tuple in the subquery result.
  - ▣  $=$  could be any comparison operator.
- **Example:**  $x \geq \text{ANY}(\langle \text{subquery} \rangle)$  means  $x$  is not the uniquely smallest tuple produced by the subquery.
  - ▣ Note tuples must have one component only.

# The Operator ALL

- $x \langle \rangle \text{ALL}(\langle \text{subquery} \rangle)$  is true iff for every tuple  $t$  in the relation,  $x$  is not equal to  $t$ .
  - ▣ That is,  $x$  is not in the subquery result.
- $\langle \rangle$  can be any comparison operator.
- **Example:**  $x \geq \text{ALL}(\langle \text{subquery} \rangle)$  means there is no tuple larger than  $x$  in the subquery result.

# Example: ALL

46

- From `Sells(bar, beer, price)`, find the beer(s) sold for the highest price.

```
SELECT beer
```

```
FROM Sells
```

```
WHERE price >=
```

```
ALL( SELECT price  
      FROM Sells);
```

price from the outer  
Sells must not be  
less than any price.

# The IN Operator

47

- $\langle \text{value} \rangle \text{ IN } (\langle \text{subquery} \rangle)$  is true if and only if the  $\langle \text{value} \rangle$  is a member of the relation produced by the subquery.
  - ▣ Opposite:  $\langle \text{value} \rangle \text{ NOT IN } (\langle \text{subquery} \rangle)$ .
- IN-expressions can appear in WHERE clauses.
- WHERE col IN (value1, value2, ...)

# IN is Concise

48

- ❑ `SELECT * FROM Cartoons  
WHERE LastName IN ('Simpsons', 'Smurfs', 'Flintstones')`
  
- ❑ `SELECT * FROM Cartoons  
WHERE LastName = 'Simpsons'  
OR LastName = 'Smurfs'  
OR LastName = 'Flintstones'`

# Example: IN

49

- Using **Beers(name, manf)** and **Likes(drinker, beer)**, find the name and manufacturer of each beer that Fred likes.

```
SELECT *
```

```
FROM Beers
```

```
WHERE name IN
```

The set of  
beers Fred  
likes

```
(SELECT beer  
FROM Likes  
WHERE drinker = 'Fred');
```



# IN vs. Join

50

```
SELECT R.a  
FROM R, S  
WHERE R.b = S.b;
```

```
SELECT R.a  
FROM R  
WHERE b IN (SELECT b FROM S);
```

# IN is a Predicate About R's Tuples

51

```
SELECT a
FROM R
WHERE b IN (SELECT b FROM S);
```

Two 2's

a	b
1	2
3	4

R

b	c
2	5
2	6

S

(1,2) satisfies  
the condition;  
1 is output once.

One loop, over  
the tuples of R

# This Query Pairs Tuples from R, S

52

```
SELECT a
FROM R, S
WHERE R.b = S.b;
```

a	b
1	2
3	4

R

b	c
2	5
2	6

S

(1,2) with (2,5)  
and (1,2) with  
(2,6) both satisfy  
the condition;  
1 is output twice.

Double loop, over  
the tuples of R and S

# Query + Subquery Solution

42

```
SELECT bar  
FROM Sells  
WHERE beer = 'Miller' AND
```

```
price = (SELECT price  
        FROM Sells  
        WHERE beer = 'Bud');
```

What if subquery  
returns multiple  
values?



# The Operator ANY

- $x = \text{ANY}(\langle \text{subquery} \rangle)$  is a boolean condition that is true iff  $x$  equals at least one tuple in the subquery result.
  - ▣  $=$  could be any comparison operator.
- **Example:**  $x \geq \text{ANY}(\langle \text{subquery} \rangle)$  means  $x$  is not the uniquely smallest tuple produced by the subquery.
  - ▣ Note tuples must have one component only.

# The Operator ALL

- $x \langle \rangle \text{ALL}(\langle \text{subquery} \rangle)$  is true iff for every tuple  $t$  in the relation,  $x$  is not equal to  $t$ .
  - ▣ That is,  $x$  is not in the subquery result.
- $\langle \rangle$  can be any comparison operator.
- **Example:**  $x \geq \text{ALL}(\langle \text{subquery} \rangle)$  means there is no tuple larger than  $x$  in the subquery result.

# The IN Operator

45

- $\langle \text{value} \rangle \text{ IN } (\langle \text{subquery} \rangle)$  is true if and only if the  $\langle \text{value} \rangle$  is a member of the relation produced by the subquery.
  - ▣ Opposite:  $\langle \text{value} \rangle \text{ NOT IN } (\langle \text{subquery} \rangle)$ .
- IN-expressions can appear in WHERE clauses.
- `WHERE col IN (value1, value2, ...)`

# IN vs. Join

46

```
SELECT R.a  
FROM R, S  
WHERE R.b = S.b;
```

```
SELECT R.a  
FROM R  
WHERE b IN (SELECT b FROM S);
```



# IN is a Predicate About R's Tuples

47

```
SELECT a
FROM R
WHERE b IN (SELECT b FROM S);
```

Two 2's

a	b
1	2
3	4

R

b	c
2	5
2	6

S

(1,2) satisfies  
the condition;  
1 is output once.

One loop, over  
the tuples of R

# This Query Pairs Tuples from R, S

48

```
SELECT a
FROM R, S
WHERE R.b = S.b;
```

a	b
1	2
3	4

R

b	c
2	5
2	6

S

Double loop, over  
the tuples of R and S

(1,2) with (2,5)  
and (1,2) with  
(2,6) both satisfy  
the condition;  
1 is output twice.

# Back to our original query...

49

```
SELECT bar  
FROM Sells  
WHERE beer = 'Miller' AND  
price = (SELECT price
```

```
FROM Sells  
WHERE beer = 'Bud');
```

Use **IN()** or **= ANY()**

# Recap

50

- `IN( )` is equivalent to `= ANY( )`
- For `ANY( )`, you can use other comparison operators such as `>`, `<`,... etc, but not applicable for `IN( )`
  
- The `< >ANY` operator, however, differs from `NOT IN`:
  - ▣ `< >ANY` means `not = a, or not = b, or not = c`
  - ▣ `NOT IN` means `not = a, and not = b, and not = c.`
  - ▣ `<>ALL` means the same as `NOT IN`.

# Example: =ANY

51

## Sells

Bar	Beer	Price
Jane	Miller	3.00
Joe	Miller	4.00
Joe	Bud	3.00
Jack	Bud	4.00
Tom	Miller	4.50

```
SELECT Bar
FROM Sells
WHERE Beer = 'Miller' AND Price =
      ANY(SELECT Price
          FROM Sells
          WHERE Beer='Bud')
```

## Result

Bar
Jane
Joe

# The Exists Operator

52

- EXISTS(<subquery>) is true if and only if the subquery result is not empty.
- **Example:** From **Beers(name, manf)** , find those beers that are the unique (only) beer made by their manufacturer.

# Example: EXISTS

53

```
SELECT name
FROM Beers b1
WHERE NOT EXISTS (
```

Notice scope rule: manf refers to closest nested FROM with a relation having that attribute. (Some DBMS consider this ambiguous.)

Set of beers with the same manf as b1, but not the same beer

```
SELECT *
FROM Beers
WHERE manf = b1.manf AND
      name <> b1.name);
```

Notice the SQL “not equals” operator

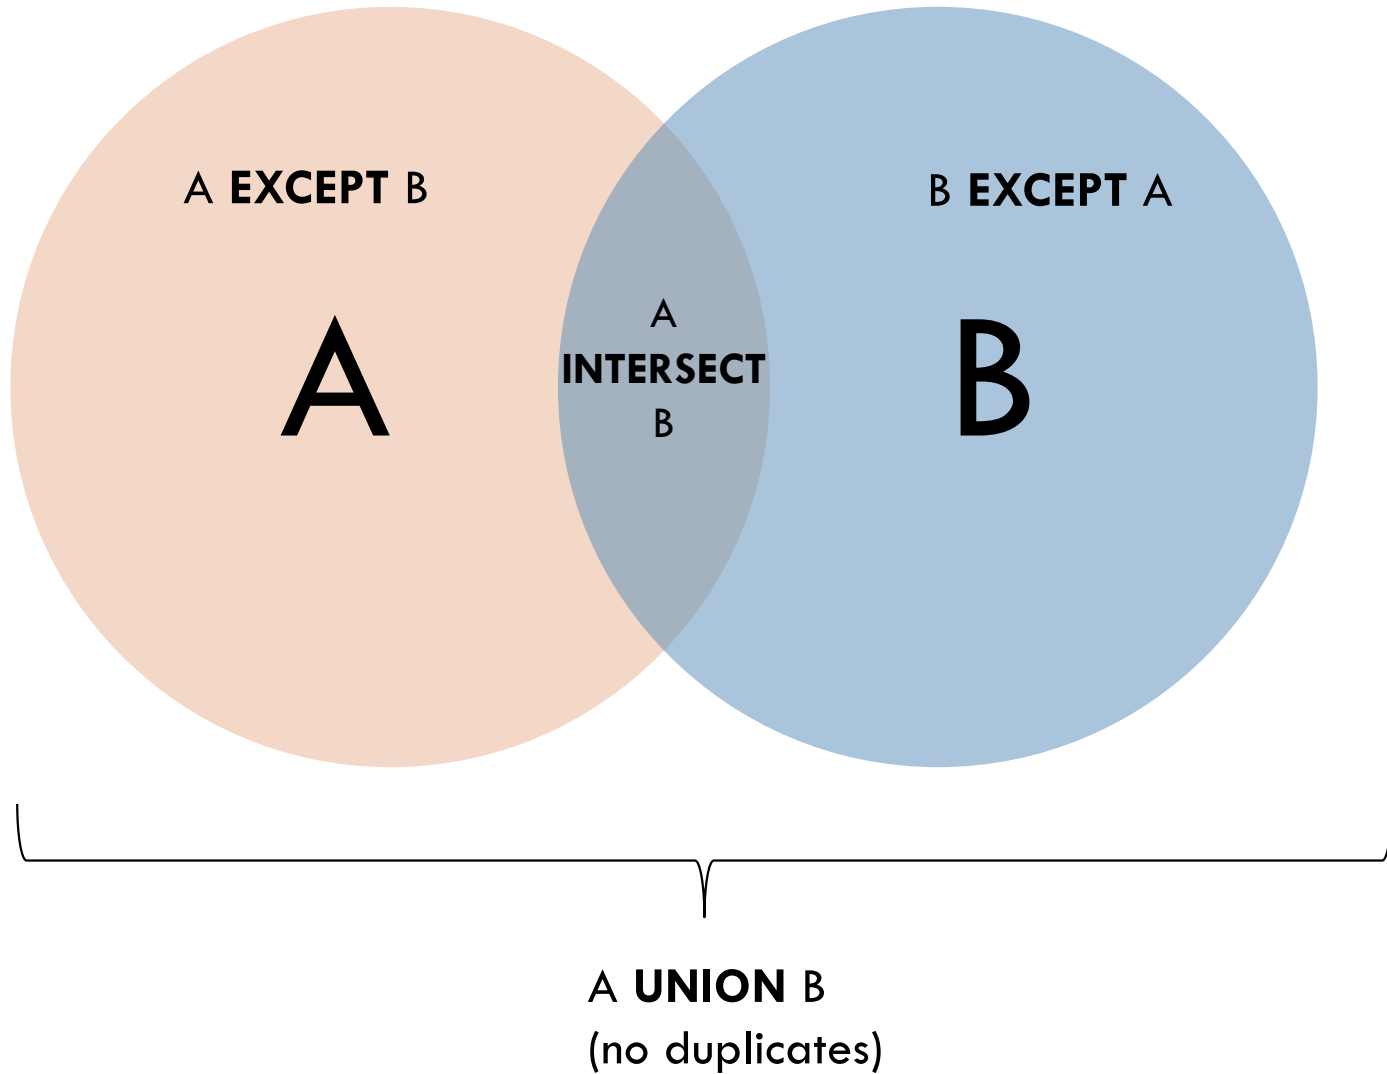
# Union, Intersection, and Difference

- Union, intersection, and difference of relations are expressed by the following forms, each involving subqueries:
  - (<subquery>) UNION (<subquery>)
  - (<subquery>) INTERSECT (<subquery>)
  - (<subquery>) EXCEPT (<subquery>)



# Visually

55



# Example: Intersection

56

- Using `Likes(drinker, beer)`, `Sells(bar, beer, price)`, and `Frequents(drinker, bar)`, find the drinkers and beers such that:
  - The drinker likes the beer, and
  - The drinker frequents at least one bar that sells the beer.

# Solution

subquery is  
really a stored  
table.

```
(SELECT * FROM Likes)
```

**INTERSECT**

```
(SELECT drinker, beer  
FROM Sells, Frequents  
WHERE Frequents.bar = Sells.bar  
);
```

The drinker frequents  
a bar that sells the  
beer.

# Union, Intersection, and Difference

1

- Union, intersection, and difference of relations are expressed by the following forms, each involving subqueries:
  - (<subquery>) UNION (<subquery>)
  - (<subquery>) INTERSECT (<subquery>)
  - (<subquery>) EXCEPT (<subquery>)

# Example: Intersection

2

- Using `Likes(drinker, beer)`, `Sells(bar, beer, price)`, and `Frequents(drinker, bar)`, find the drinkers and beers such that:
  - The drinker likes the beer, and
  - The drinker frequents at least one bar that sells the beer.

# Solution

subquery is  
really a stored  
table.

```
(SELECT * FROM Likes)
```

**INTERSECT**

```
(SELECT drinker, beer  
FROM Sells, Frequents  
WHERE Frequents.bar = Sells.bar  
);
```

The drinker frequents  
a bar that sells the  
beer.

# Bag Semantics

4

- A *bag* (or *multiset*) is like a set, but an element may appear more than once.
- **Example:**  $\{1,2,1,3\}$  is a bag.
- **Example:**  $\{1,2,3\}$  is also a bag that happens to be a set.

# Bag (Multiset) Semantics

5

- SQL primarily uses bag semantics
- The SELECT-FROM-WHERE statement uses bag semantics
  - ▣ originally for efficiency reasons
- The default for union, intersection, and difference is set semantics.
  - ▣ That is, duplicates are eliminated as the operation is applied.



# Motivation: Efficiency

6

- When doing projection, it is easier to avoid eliminating duplicates.
  - ▣ Just work tuple-at-a-time.
- For intersection or difference, it is most efficient to sort the relations first.
  - ▣ At that point you may as well eliminate the duplicates anyway.

# Controlling Duplicate Elimination

7

- Force the result to be a set by `SELECT DISTINCT . . .`
- Force the result to be a bag (i.e., don't eliminate duplicates) by `ALL`, as in  
    `. . . UNION ALL . . .`

# Example: DISTINCT

8

- From `Sells(bar, beer, price)`, find all the different prices charged for beers:

```
SELECT DISTINCT price
FROM Sells;
```

Notice that without `DISTINCT`, each price would be listed as many times as there were bar/beer pairs at that price.

# Example: ALL

9

- Using relations **Frequents(drinker, bar)** and **Likes(drinker, beer)**:
- Lists drinkers who frequent more bars than they like beers, and do so as many times as the difference of those counts.

```
(SELECT drinker FROM Frequents)  
  EXCEPT ALL  
(SELECT drinker FROM Likes);
```

# Ordering the Display of Tuples

10

- List in alphabetic order the names of all instructors  
**select** *name*  
**from** *instructor*  
**order by** *name*
- We may specify **desc** for descending order or **asc** for ascending order, for each attribute; ascending order is the default.
  - Example: **order by** *name* **desc**

# Humour

11

*SQL query walks into a bar, and approaches two tables and asks, can I join you?*



# DATABASE MODIFICATIONS



# Database Modifications

13

- A *modification* command does not return a result (as a query does), but changes the database in some way.
- Three kinds of modifications:
  1. *Insert* a tuple or tuples.
  2. *Delete* a tuple or tuples.
  3. *Update* the value(s) of an existing tuple or tuples.



# Insertion

14

- To insert a single tuple:

```
INSERT INTO <relation>
```

```
VALUES ( <list of values> );
```

- **Example:** add to **Likes(drinker, beer)** the fact that Sally likes Bud.

```
INSERT INTO Likes
```

```
VALUES ( 'Sally', 'Bud' );
```

# Specifying Attributes in INSERT

15

- We may add to the relation name a list of attributes.
- Two reasons to do so:
  1. We forget the standard order of attributes for the relation.
  2. We don't have values for all attributes, and we want the system to fill in missing components with NULL or a default value.

# Example: Specifying Attributes

16

- Another way to add the fact that Sally likes Bud to `Likes(drinker, beer)`:

```
INSERT INTO Likes (beer, drinker)
VALUES ('Bud', 'Sally');
```

# Adding Default Values

17

- In a CREATE TABLE statement, we can follow an attribute by DEFAULT and a value.
- When an inserted tuple has no value for that attribute, the default will be used.

# Example: Default Values

18

```
CREATE TABLE Drinkers (  
    name CHAR(30) PRIMARY KEY,  
    addr CHAR(50)  
        DEFAULT '123 Sesame St.',  
    phone CHAR(16)  
);
```

# Example: Default Values

19

```
INSERT INTO Drinkers (name)
VALUES ('Sally');
```

Resulting tuple:

name	address	phone
Sally	123 Sesame St	NULL

# Inserting Many Tuples

20

- We may insert the entire result of a query into a relation, using the form:

```
INSERT INTO <relation>  
( <subquery> );
```

# Example: Insert a Subquery

21

- Using `Frequents(drinker, bar)`, enter into the new relation `Buddies(name)` all of Sally's "potential buddies,"
- i.e., those drinkers who frequent at least one bar that Sally also frequents.

```
INSERT INTO Buddies
(SELECT
);
```



# Solution

22

*“Those drinkers who frequent at least one bar that Sally also frequents”*

The other  
drinker

**INSERT INTO Buddies**

```
(SELECT d2.drinker
```

```
FROM Frequents d1, Frequents d2
```

```
WHERE d1.drinker = 'Sally' AND
```

```
  d2.drinker <> 'Sally' AND
```

```
  d1.bar = d2.bar
```

```
);
```

Pairs of Drinker tuples where the first is for Sally, the second is for someone else, and the bars are the same.

# Deletion

23

- To delete tuples satisfying a condition from some relation:

```
DELETE FROM <relation>  
WHERE <condition>;
```

# Example: Deletion

24

- Delete from `Likes(drinker, beer)` the fact that Sally likes Bud:

```
DELETE FROM Likes
WHERE drinker = 'Sally' AND
      beer = 'Bud';
```

# Example: Delete all Tuples

25

- Make the relation Likes empty:

```
DELETE FROM Likes;
```

- Note no WHERE clause needed.

# Example: Delete Some Tuples


26

- Delete from **Beers(name, manf)** all beers for which there is another beer by the same manufacturer.

```
DELETE FROM Beers b
WHERE
```

```
  EXISTS (
    SELECT name
      FROM Beers
     WHERE manf = b.manf AND
           name <> b.name);
```

Beers with the same manufacturer and a different name from the name of the beer represented by tuple b.



# Semantics of Deletion --- (1)

- Suppose Anheuser-Busch makes only Bud and Bud Lite.
- Suppose we come to the tuple  $b$  for Bud first.
- The subquery is nonempty, because of the Bud Lite tuple, so we delete Bud.
- Now, when  $b$  is the tuple for Bud Lite, do we delete that tuple too?

# Semantics of Deletion --- (2)

28

- **Answer:** we do delete Bud Lite as well.
- The reason is that deletion proceeds in two stages:
  1. Mark all tuples for which the WHERE condition is satisfied.
  2. Delete the marked tuples.

# Updates

29

- To change certain attributes in certain tuples of a relation:

UPDATE <relation>

SET <list of attribute assignments>

WHERE <condition on tuples>;



# Example: Update

30

- Change drinker Fred's phone number to 555-1212:

```
UPDATE Drinkers  
SET phone = '555-1212'  
WHERE name = 'Fred';
```

# Example: Update Several Tuples

31

- Make \$4 the maximum price for beer:

```
UPDATE Sells
SET price = 4.00
WHERE price > 4.00;
```

# AGGREGATION, GROUPING & OUTER JOINS



# Aggregation

2

- SUM, AVG, COUNT, MIN, and MAX can be applied to a column in a SELECT clause to produce that aggregation on the column.
- COUNT(\*) counts the number of tuples.

# Example: Aggregation

3

- From **Sells(bar, beer, price)**, find the average price of Bud:

```
SELECT AVG(price)
FROM Sells
WHERE beer = 'Bud';
```

# Eliminating Duplicates in an Aggregation

4

- Use DISTINCT inside an aggregation.
- **Example:** find the number of *different* prices charged for Bud:

```
SELECT COUNT(DISTINCT price)
FROM Sells
WHERE beer = 'Bud';
```

# NULL's Ignored in Aggregation

5

- NULL never contributes to a sum, average, or count, and can never be the minimum or maximum of a column.
- But if all the values in a column are NULL, then the result of the aggregation is NULL.
  - ▣ **Exception:** COUNT of an empty set is 0.

# Example: Effect of NULL's

6

Sells(bar, beer, price)

```
SELECT count(*)  
FROM Sells  
WHERE beer = 'Bud';
```

The number of bars  
that sell Bud.

```
SELECT count(price)  
FROM Sells  
WHERE beer = 'Bud';
```

The number of bars  
that sell Bud at a  
known price (i.e., where  
price is not NULL)



# Example Query

7

- Find the age of the youngest employee at each rating level

```
SELECT MIN (age)
FROM Employees
WHERE rating = i
```

# Grouping

8

- We may follow a SELECT-FROM-WHERE expression by GROUP BY and a list of attributes.
- The relation that results from the SELECT-FROM-WHERE is grouped according to the values of all those attributes, and any aggregation is applied only within each group.

```
SELECT  rating, MIN(age)
FROM    Employees
GROUP BY rating
```

# Example: Grouping

9

- From `Sells(bar, beer, price)`, find the average price for each beer:

```
SELECT beer, AVG(price)
FROM Sells
GROUP BY beer;
```

beer	AVG(price)
Bud	2.33
Miller	4.55
...	...

# Example: Grouping

10

- From `Sells(bar, beer, price)` and `Frequents(drinker, bar)`, find for each drinker the average price of Bud at the bars they frequent:

```
SELECT drinker, AVG(price)
FROM Frequents, Sells
WHERE beer = 'Bud' AND
      Frequents.bar = Sells.bar
```

```
GROUP BY drinker;
```

Compute all drinker-bar-price triples for Bud.

Then group them by drinker.

# Restriction on SELECT Lists With Aggregation

- If any aggregation is used, then each element of the SELECT list must be either:
  1. Aggregated, or
  2. An attribute on the GROUP BY list.

# Illegal Query Example

12

```
SELECT bar, beer, AVG(price)
FROM Sells
GROUP BY bar
```

- But this query is illegal in SQL.
- Only one tuple output for each bar, no unique way to select which beer to output

# A Closer Look

13

```
SELECT bar, beer, AVG(price) AS avgP
FROM Sells
GROUP BY bar
```

beer

**Result**

bar	beer	avgP
Joe	?	3.50
Tom	?	3.88
Jane	?	4.00

↑  
{Bud, Miller, Coors}?

**Sells**

Bar	Beer	Price
Joe	Bud	3.00
Joe	Miller	4.00
Tom	Bud	3.50
Tom	Miller	4.25
Jane	Bud	3.25
Jane	Miller	4.75
Jane	Coors	4.00

Only one tuple output for each bar, no unique way to select which beer to output

# HAVING Clauses

14

- **HAVING** <condition> may follow a **GROUP BY** clause.
- If so, the condition applies to each group, and groups not satisfying the condition are eliminated.



# Example: HAVING

15

- From **Sells(bar, beer, price)** and **Beers(name, manf)**, find the average price of those beers that are either served in at least three bars or are manufactured by Pete's.

# Solution

Sells(bar, beer, price) and Beers(name, manf),

16


```
SELECT beer, AVG(price)
FROM Sells
GROUP BY beer
HAVING COUNT(bar) >= 3 OR
```

Beer groups with at least  
3 non-NULL bars



```
beer IN (SELECT name
FROM Beers
WHERE manf = 'Pete's');
```

Beers manu-  
factured by  
Pete's.



# Requirements on HAVING Conditions

17

- Anything goes in a subquery.
- Outside subqueries, they may refer to attributes only if they are either:
  1. A grouping attribute, or
  2. Aggregated(same condition as for SELECT clauses with aggregation).

# A Final Example

18

```
SELECT Bar, SUM(Qty) AS sumQ
FROM Sells
GROUP BY Bar
HAVING sum(Qty) > 4
```

## Result

Bar	sumQ
Tom	5
Jane	6

## Sells

Bar	Beer	Price	Qty
Joe	Bud	3.00	2
Joe	Miller	4.00	2
Tom	Bud	3.50	1
Tom	Miller	4.25	4
Jane	Bud	3.25	1
Jane	Miller	4.75	3
Jane	Coors	4.00	2

# Assignment 2

1

- Due Nov 4<sup>th</sup> at 10:00pm
  - ▣ Populate database early!
  - ▣ Practice SQL queries as prep for midterm
- This week: outerjoins, views, indexes
- Midterm material cut-off this Thursday's lecture
- Review midterm practice questions: Oct 21, 23

# Example: Grouping

2

- From `Sells(bar, beer, price)`, find the average price for each beer:

```
SELECT beer, AVG(price)
FROM Sells
GROUP BY beer;
```

beer	AVG(price)
Bud	2.33
Miller	4.55
...	...

# Restriction on SELECT Lists With Aggregation

3

- If any aggregation is used, then each element of the SELECT list must be either:
  1. Aggregated, or
  2. An attribute on the GROUP BY list.

# Illegal Query Example

4

```
SELECT bar, beer, AVG(price)
FROM Sells
GROUP BY bar
```

- But this query is illegal in SQL.
- Only one tuple output for each bar, no unique way to select which beer to output



# A Closer Look

5

```
SELECT bar, beer, AVG(price) AS avgP
FROM Sells
GROUP BY bar beer
```

**Result**

bar	beer	avgP
Joe	?	3.50
Tom	?	3.88
Jane	?	4.00

↑  
{Bud, Miller, Coors}?

**Sells**

Bar	Beer	Price
Joe	Bud	3.00
Joe	Miller	4.00
Tom	Bud	3.50
Tom	Miller	4.25
Jane	Bud	3.25
Jane	Miller	4.75
Jane	Coors	4.00

Only one tuple output for each bar, no unique way to select which beer to output

# Example: HAVING

6

- From `Sells(bar, beer, price)` and `Beers(name, manf)`, find the average price of those beers that are either served in at least three bars or are manufactured by Pete's.

# Solution

Sells(bar, beer, price) and Beers(name, manf),

7


```
SELECT beer, AVG(price)
FROM Sells
GROUP BY beer
HAVING COUNT(bar) >= 3 OR
```

Beer groups with at least  
3 non-NULL bars



```
beer IN (SELECT name
FROM Beers
WHERE manf = 'Pete's');
```

Beers manu-  
factured by  
Pete's.



# Requirements on HAVING Conditions

8

- Anything goes in a subquery.
- Outside subqueries, they may refer to attributes only if they are either:
  1. A grouping attribute, or
  2. Aggregated(same condition as for SELECT clauses with aggregation).

# Cross Product

9

- Evaluating joins involves combining two or more relations
- Given two relations,  $S$  and  $R$ , each row of  $S$  is paired with each row of  $R$
- Result schema: one attribute from each attribute of  $S$  and  $R$

# Example

10

Sells

Bar	Beer	Price
Joe	Bud	3.00
Tom	Miller	4.00
Jane	Lite	3.25

Frequents

Drinker	Bar
Aaron	Joe
Mary	Jane

Cross product,  
also known as the  
Cartesian product

Sells x Frequents

(Bar)	Beer	Price	Drinker	(Bar)
Joe	Bud	3.00	Aaron	Joe
Joe	Bud	3.00	Mary	Jane
Tom	Miller	4.00	Aaron	Joe
Tom	Miller	4.00	Mary	Jane
Jane	Lite	3.25	Aaron	Joe
Jane	Lite	3.25	Mary	Jane

```
SELECT drinker
FROM Frequents, Sells
WHERE beer = 'Bud' AND
      Frequents.bar = Sells.bar
```

**Drinker**

Aaron



# Joined Relations

- **Join operations** take two relations and return as a result another relation.
- A join operation is a Cartesian product which requires that tuples in the two relations match (under some condition). It also specifies the attributes that are present in the result of the join

# Join Operations – Example

12

- Relation *course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

- Relation *prereq*

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

- Observe that

prereq information is missing for CS-315 and  
course information is missing for CS-347



# Outer Join

- An extension of the join operation that avoids loss of information.
- Suppose you have two relations  $R$  and  $S$ . A tuple of  $R$  that has no tuple of  $S$  with which it joins is said to be *dangling*.
  - Similarly for a tuple of  $S$ .
- 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.
- Outerjoin preserves dangling tuples by padding them with NULL.

# Left Outer Join

14

□ **course**

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

■ **prereq**

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

■ **course left outer join prereq**

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_id</i>
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-315	Robotics	Comp. Sci.	3	<i>null</i>

# Right Outer Join

15

□ *course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

■ *prereq*

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

■ *course* **right outer join** *prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_id</i>
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-347	<i>null</i>	<i>null</i>	<i>null</i>	CS-101

# Full Outer Join

16

□ *course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

■ *prereq*

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

■ *course full outer join prereq*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_id</i>
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101
CS-315	Robotics	Comp. Sci.	3	<i>null</i>
CS-347	<i>null</i>	<i>null</i>	<i>null</i>	CS-101

# Inner Join

17

□ *course*

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>
BIO-301	Genetics	Biology	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3

■ *prereq*

<i>course_id</i>	<i>prereq_id</i>
BIO-301	BIO-101
CS-190	CS-101
CS-347	CS-101

- **course inner join prereq on**  
 $course.course\_id = prereq.course\_id$

<i>course_id</i>	<i>title</i>	<i>dept_name</i>	<i>credits</i>	<i>prereq_id</i>
BIO-301	Genetics	Biology	4	BIO-101
CS-190	Game Design	Comp. Sci.	4	CS-101

# Outerjoins

- R OUTER JOIN S is the core of an outerjoin expression. It is modified by:
  1. Optional NATURAL in front of OUTER.
    - Check equality on all common attributes
    - No two attributes with the same name in the output
  2. Optional ON <condition> after JOIN.
  3. Optional LEFT, RIGHT, or FULL before OUTER.
    - ◆ LEFT = pad dangling tuples of R only.
    - ◆ RIGHT = pad dangling tuples of S only.
    - ◆ FULL = pad both; this choice is the default.

# Example: Outerjoin

19

R =

A	B
1	2
4	5

S =

B	C
2	3
6	7

(1,2) joins with (2,3), but the other two tuples are dangling.

R NATURAL FULL OUTERJOIN S =

A	B	C
1	2	3
4	5	NULL
NULL	6	7

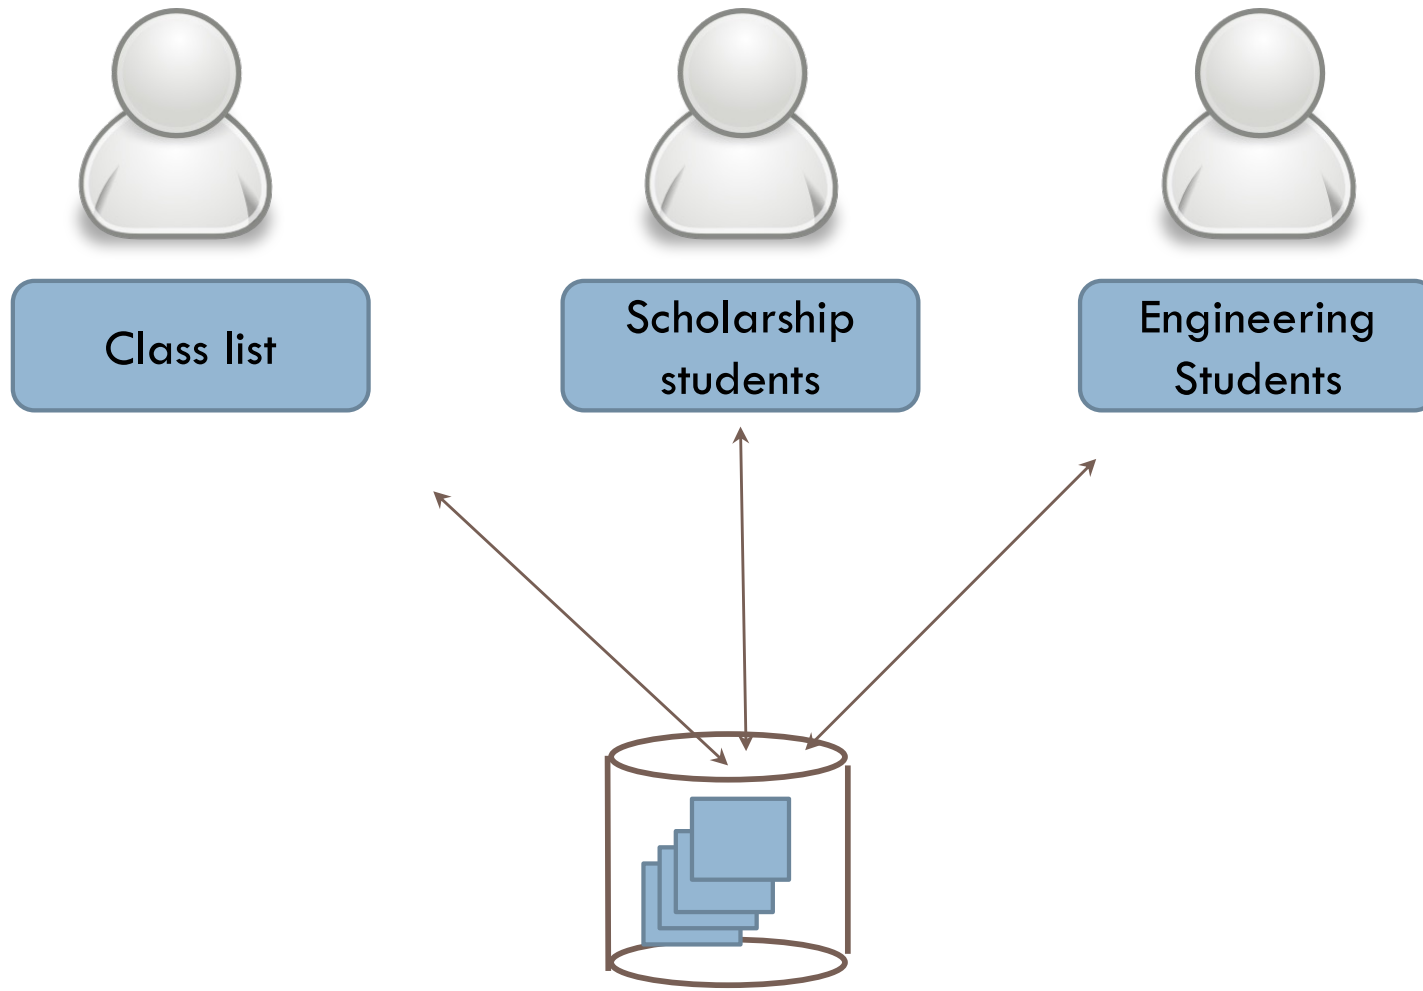
VIEWS





# Scenario

21



# Views

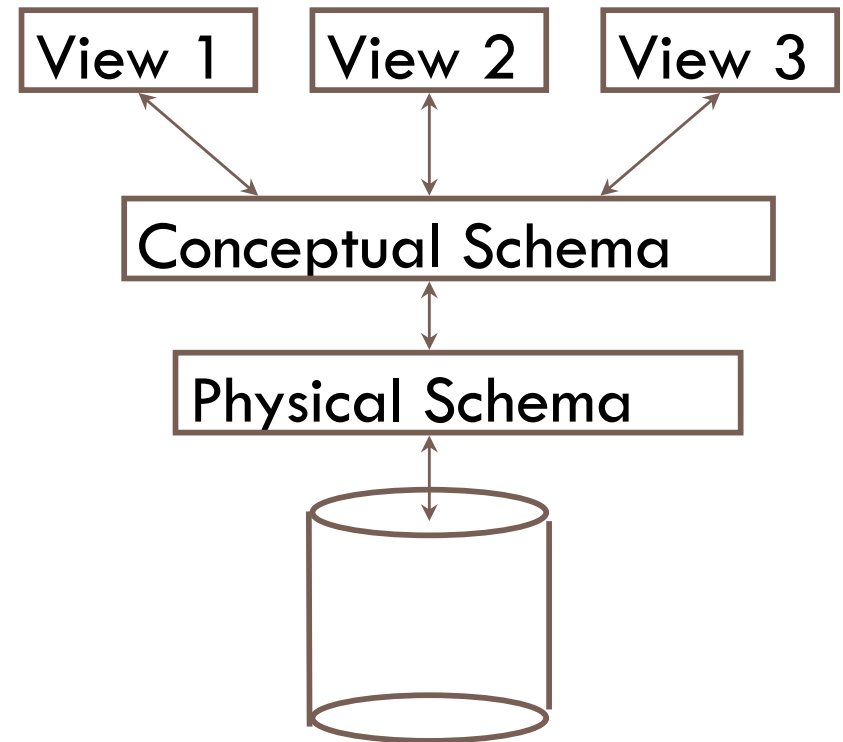
22

- In most cases, it is not desirable for all users to see the entire data instance.
- A **view** provides a mechanism to hide certain data from the view of certain users.

# Levels of Abstraction

23

- Many views, single conceptual (logical) schema and physical schema.
  - ▣ Views describe how users see the data.
  - ▣ Conceptual schema defines logical structure
  - ▣ Physical schema describes the files and indexes used.



# Views

24

- A *view* is a relation defined in terms of stored tables (called *base tables* ) and other views.
- Two kinds:
  1. *Virtual* = not stored in the database; just a query for constructing the relation.
  2. *Materialized* = actually constructed and stored.

# Declaring Views

25

- Declare by:

```
CREATE [MATERIALIZED] VIEW <name> AS <query>;
```

- A view name
- A possible list of attribute names (for example, when arithmetic operations are specified or when we want the names to be different from the attributes in the base relations)
- A query to specify the view contents
- Default is virtual.

# Example: View Definition

26

- `CanDrink(drinker, beer)` is a view “containing” the drinker-beer pairs such that the drinker frequents at least one bar that serves the beer:

```
CREATE VIEW CanDrink AS
  SELECT drinker, beer
  FROM Frequents, Sells
  WHERE Frequents.bar = Sells.bar;
```

# Example: Accessing a View

27

- Query a view as if it were a base table.
  - ▣ Also: a limited ability to modify views if it makes sense as a modification of one underlying base table.

- **Example query:**

```
SELECT beer FROM CanDrink  
WHERE drinker = 'Sally';
```

# Another Example

28

- **Example:** View Synergy has (drinker, beer, bar) triples such that the bar serves the beer, the drinker frequents the bar and likes the beer.



# Example: The View

29

```
CREATE VIEW Synergy AS
```

```
SELECT Likes.drinker, Likes.beer, Sells.bar
```

Pick one copy of  
each attribute



```
FROM Likes, Sells, Frequents
```

```
WHERE Likes.drinker = Frequents.drinker
```

```
AND Likes.beer = Sells.beer
```

```
AND Sells.bar = Frequents.bar;
```

Natural join of Likes,  
Sells, and Frequents



# Updates on Views

- Generally, it is impossible to modify a virtual view, because it doesn't exist.
- Can't we "translate" updates on views into "equivalent" updates on base tables?
  - ▣ Not always (in fact, not often)
  - ▣ Most systems prohibit most view updates
- We cannot insert into Synergy --- it is a virtual view.

# Interpreting a View Insertion

31

- But we could try to translate a (drinker, beer, bar) triple into three insertions of projected pairs, one for each of Likes, Sells, and Frequent.

# Insertion

32

```
INSERT INTO LIKES VALUES(n.drinker, n.beer);
```

```
INSERT INTO SELLS(bar, beer) VALUES(n.bar, n.beer);
```

```
INSERT INTO FREQUENTS VALUES(n.drinker, n.bar);
```

- ▣ Sells.price will have to be NULL.
- ▣ There isn't always a unique translation.

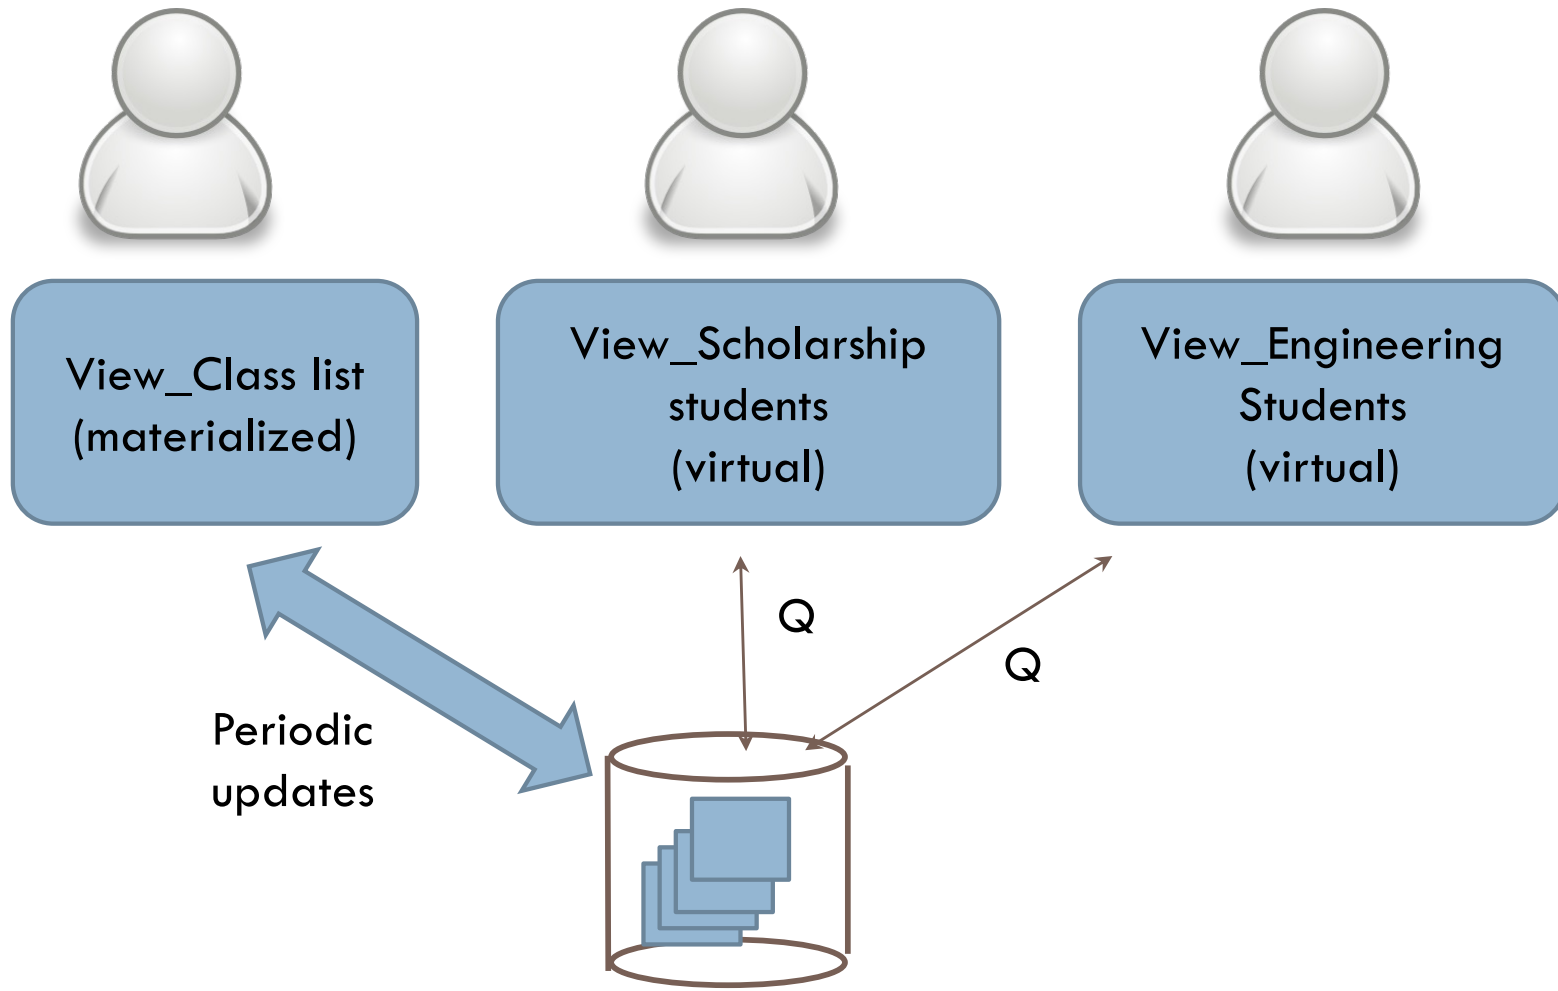
# Materialized Views

33

- *Materialized* = actually constructed and stored (keeping a temporary table)
- **Concerns:** maintaining correspondence between the base table and the view when the base table is updated
- **Strategy:** incremental update

# Example

34



# Views

1

- A *view* is a relation defined in terms of stored tables (called *base tables* ) and other views.
- Two kinds:
  1. *Virtual* = not stored in the database; just a query for constructing the relation.
  2. *Materialized* = actually constructed and stored.

# Example: View Definition

2

- **CanDrink(drinker, beer)** is a view “containing” the drinker-beer pairs such that the drinker frequents at least one bar that serves the beer:

```
CREATE VIEW CanDrink AS
  SELECT drinker, beer
  FROM Frequents, Sells
  WHERE Frequents.bar = Sells.bar;
```



# Example: Accessing a View

3

- Query a view as if it were a base table.
  - ▣ Also: a limited ability to modify views if it makes sense as a modification of one underlying base table.
- **Example query:**

```
SELECT beer FROM CanDrink  
WHERE drinker = 'Sally';
```

# Another Example

4

- **Example:** View Synergy has (drinker, beer, bar) triples such that the bar serves the beer, the drinker frequents the bar and likes the beer.

# Example: The View

5

```
CREATE VIEW Synergy AS
```

```
SELECT Likes.drinker, Likes.beer, Sells.bar
```

Pick one copy of  
each attribute



```
FROM Likes, Sells, Frequents
```

```
WHERE Likes.drinker = Frequents.drinker
```

```
AND Likes.beer = Sells.beer
```

```
AND Sells.bar = Frequents.bar;
```

Natural join of Likes,  
Sells, and Frequents



# Updates on Views

- Generally, it is impossible to modify a virtual view, because it doesn't exist.
- Can't we "translate" updates on views into "equivalent" updates on base tables?
  - ▣ Not always (in fact, not often)
  - ▣ Most systems prohibit most view updates
- We cannot insert into Synergy --- it is a virtual view.

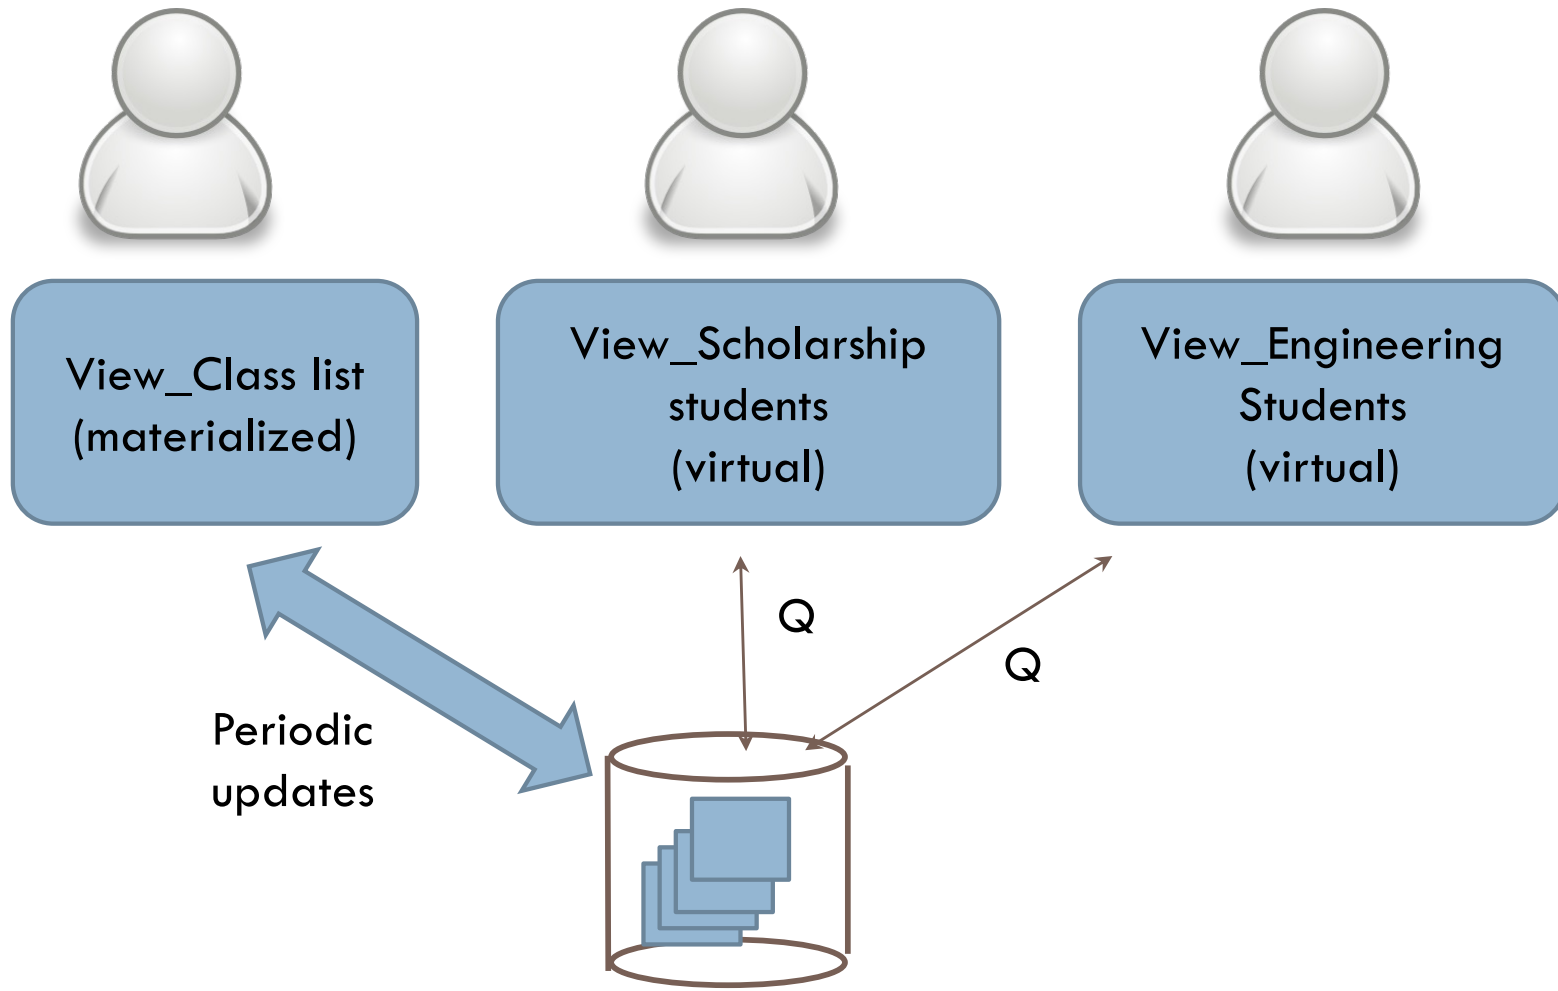
# Materialized Views

7

- *Materialized* = actually constructed and stored (keeping a temporary table)
- **Concerns:** maintaining correspondence between the base table and the view when the base table is updated
- **Strategy:** incremental update

# Example

8



# Materialized View Updates

9

- Update on a single view without aggregate operations: update may map to an update on the underlying base table (most SQL implementations)
- Views involving joins: an update *may map to an* update on the underlying base relations not always possible

# INDEXES





# Example

11

- Find the price of beers manufactured by Pete's and sold by Joe.

```
SELECT price
```

```
FROM Beers, Sells
```

```
WHERE manf = 'Pete's' AND bar = 'Joe' AND
```

```
Sells.beer = Beers.name
```

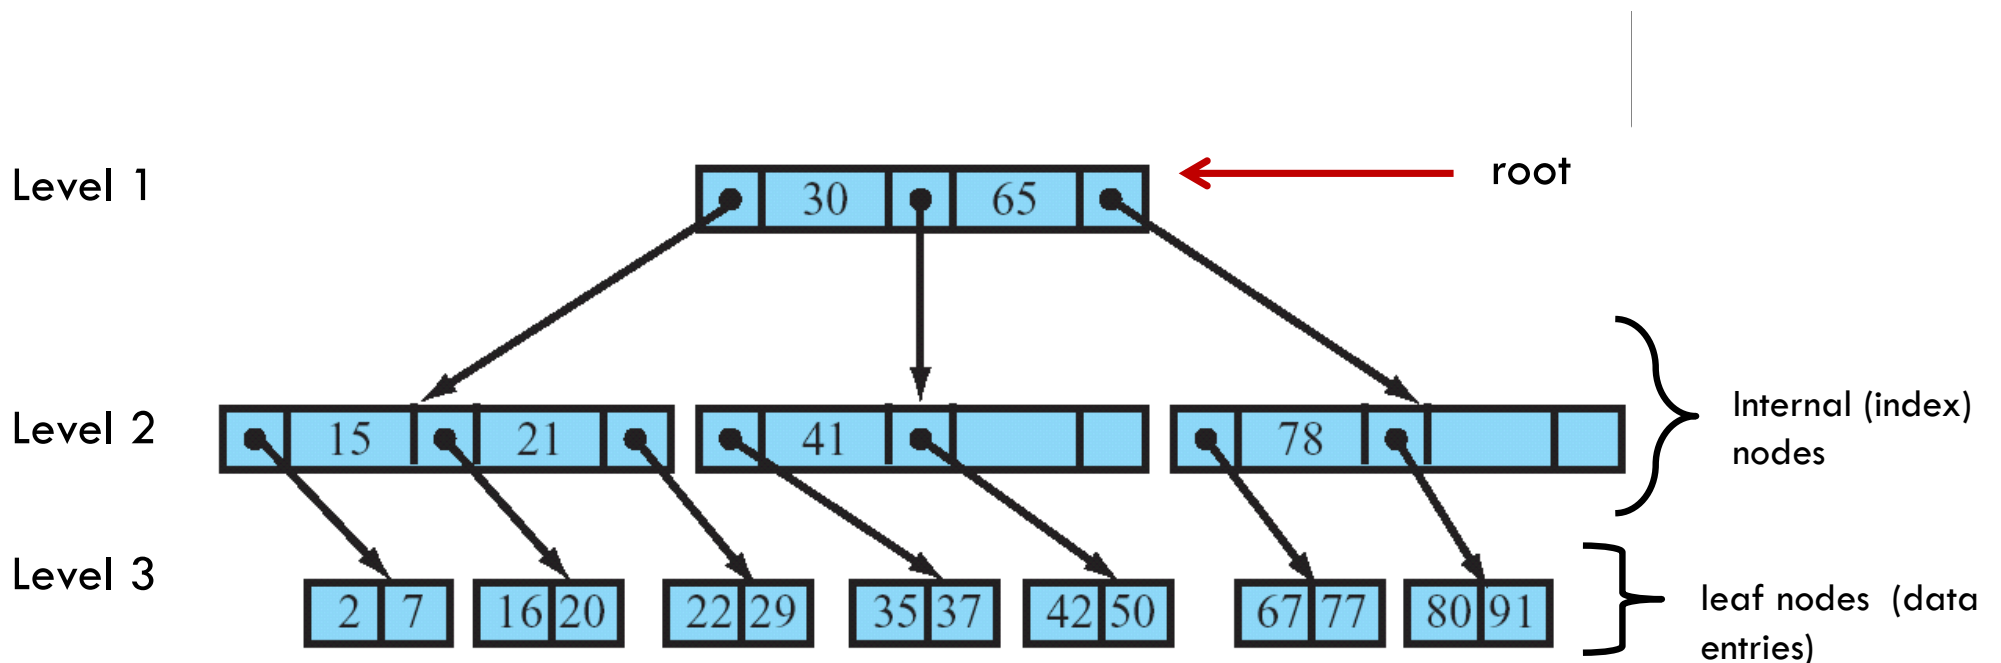
# An Index

- A data structure used to speed access to tuples of a relation, based on values of one or more attributes (“search key” fields)
- Organizes records via trees or hashing
- Given a value  $v$ , the index takes us to only those tuples that have  $v$  in the attribute(s) of the index.
- **Example:** use **BeerInd** (on manf) and **SellInd** (on bar, beer) to find the prices of beers manufactured by Pete’s and sold by Joe.

# B+ Tree Index

13



- The B+ tree structure is the most common index type in databases
- Index files can be quite large, often stored on disk, partially loaded into memory as needed
- Each node is at least 50% full



# B+ Tree Index

14


Supports equality and range-searches efficiently



Credit: Renee Miller

# Example

15




- Find 28\*? 29\*? All  $> 15^*$  and  $< 30^*$
- Insert/delete: Find data entry in leaf, then change it. Need to adjust parent sometimes.
  - And change sometimes bubbles up the tree

# Inserting a Data Entry

16

- ❑ Find correct leaf L.
- ❑ Put data entry onto L.
  - ❑ If L has enough space, done!
  - ❑ Else, must *split* L (into L and a new node L2)
    - ❑ Redistribute entries evenly, *copy up* middle key.
    - ❑ Insert index entry pointing to L2 into parent of L.
- ❑ This can happen recursively
  - ❑ To split index node, redistribute entries evenly, but *push up* middle key.
- ❑ Splits “grow” tree; root split increases height.




Insert data  
value 4

# Deleting a Data Entry

Delete value 3

17

- ❑ Start at root, find leaf  $L$  where entry belongs.
- ❑ Remove the entry.
  - ❑ If  $L$  is at least half-full, done!
  - ❑ If not,
    - ❑ Try to **re-distribute**, borrowing from sibling (adjacent node with same parent as  $L$ ).
    - ❑ If re-distribution fails, **merge**  $L$  and sibling.
- ❑ If merge occurred, must delete entry (pointing to  $L$  or sibling) from parent of  $L$ .
- ❑ Merge could propagate to root, decreasing height.




# B+ Tree: Most Widely Used Index

1

- ❖ Insert/delete at  $\log_F N$  cost; keep tree *height-balanced*. ( $F = \text{fanout}$ ,  $N = \# \text{ leaf pages}$ )
- ❖ Minimum 50% occupancy (except for root). Each node contains  $d \leq m \leq 2d$  entries. The parameter  $d$  is called the *order* of the tree.
- ❖ Node with order  $d = 2$ ,

e.g.,  $2 \leq m \leq 4$





# B+ Trees in Practice



2

- ❖ Typical order: 100.
- ❖ Typical fill-factor:  $\ln 2 = 66.5\%$  (approx)
  - ❖ average fanout  $= 2 \times 100 \times 66.5\% = 133$
- ❖ Typical capacities:
  - ❖ Height 4:  $133^4 = 312,900,721$  pages
  - ❖ Height 3:  $133^3 = 2,352,637$  pages
  
- ❖ For typical orders ( $d \sim 100-200$ ), a shallow B+ tree can accommodate very large files.

# B+ Tree Index

3

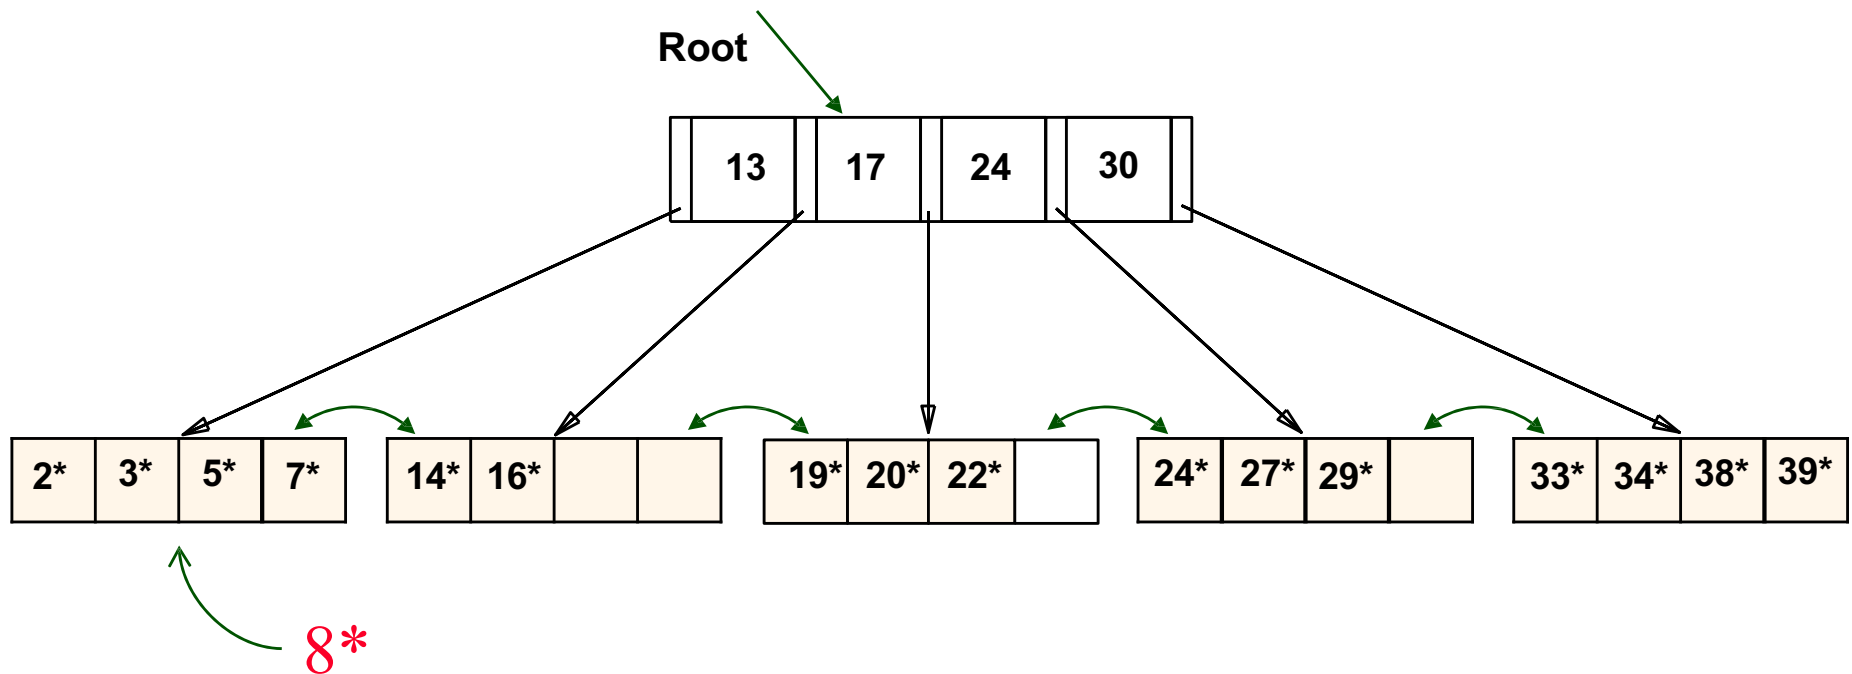
Supports equality and range-searches efficiently



Credit: Renee Miller

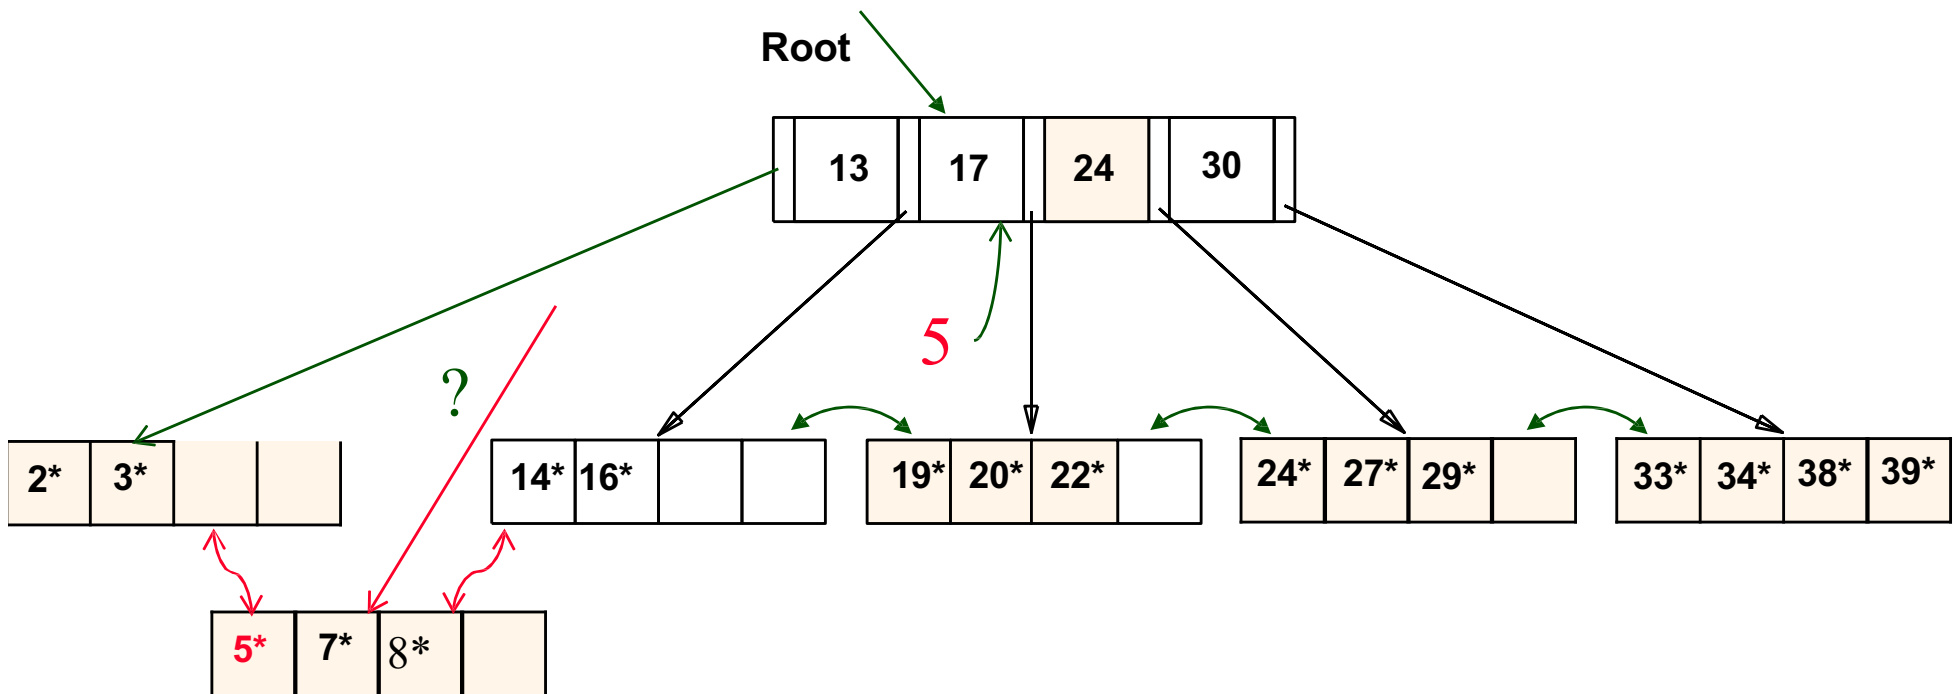
# Insertion Example

4




# Insertion Example

5



# After Inserting 8\*

6





❖ Notice that root was split, leading to increase in height.

# Copy-Up vs. Bump-Up

7

- ❖ Observe how minimum occupancy is guaranteed in both leaf and index pg splits.
- ❖ Note difference between *copy-up* and *bump-up*; Why do we handle leaf page split and index page split differently?

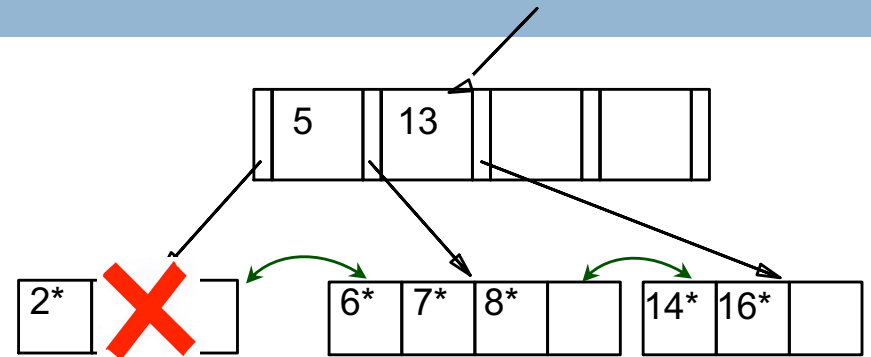


# Deleting a Data Entry

Delete value 3


8

- ❑ Start at root, find leaf  $L$  where entry belongs.
- ❑ Remove the entry.
  - ❑ If  $L$  is at least half-full, done!
  - ❑ If not,
    - ❑ Try to **re-distribute**, borrowing from sibling (adjacent node with same parent as  $L$ ).
    - ❑ If re-distribution fails, **merge**  $L$  and sibling.
- ❑ If merge occurred, must delete entry (pointing to  $L$  or sibling) from parent of  $L$ .
- ❑ Merge could propagate to root, decreasing height.



# Deleting 19\* is Straightforward

9




❖ What happens if we delete 20\* next?



# Example Tree: Deleting 19\* and 20\*

10

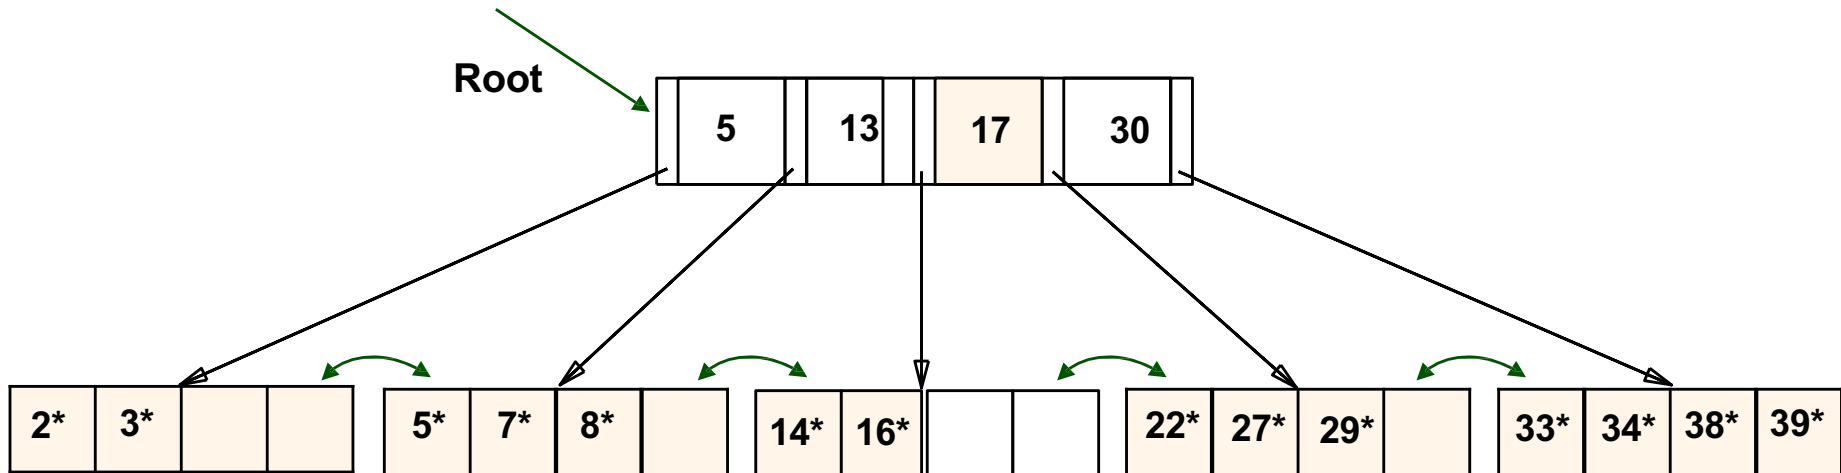



- ❖ Deleting 19\* is easy.
- ❖ Deleting 20\* is done with re-distribution. Notice how **new** middle key is **copied up**.
- ❖ What happens if we delete 24\* now?

# Deleting 24\* ...

11

- ❖ Must merge.
- ❖ Observe '*toss*' of index entry (on right), and '*pull down*' of index entry (below).




# Balanced vs. Unbalanced Trees


12

- In a balanced tree, every path from the root to a leaf node is the same length.

○ Balanced



○ Unbalanced



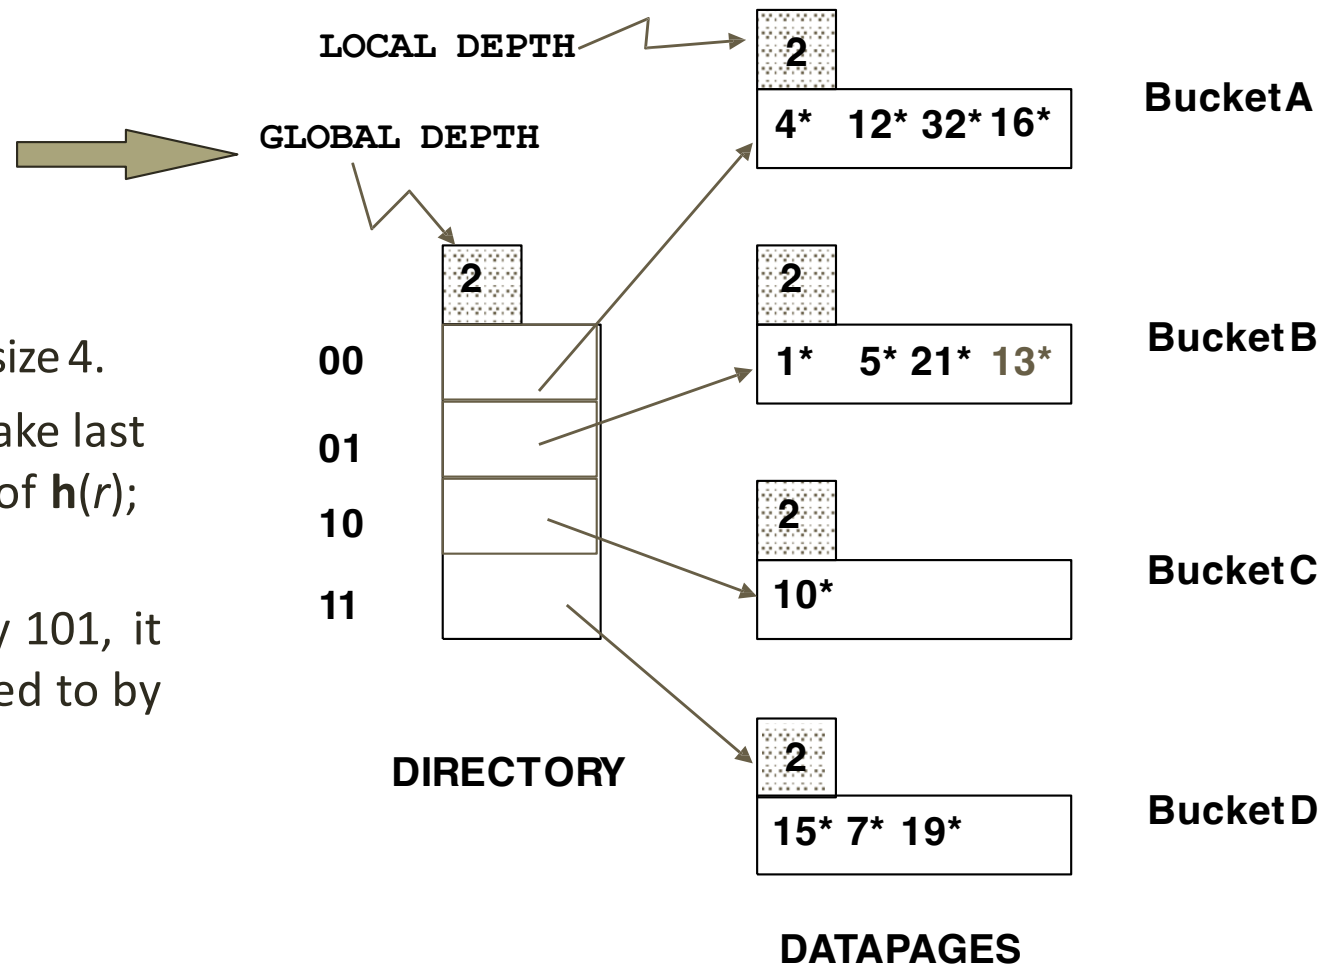
# Hash Based Indexes

13

- Good for equality searches
- Your index is a collection of *buckets* (bucket = page)
- Define a hash function,  $h$ , that maps a key to a bucket.
- Store the corresponding data in that bucket.
- Collisions
  - Multiple keys hash to the same bucket.
  - Store multiple keys in the same bucket.
- What do you do when buckets fill?
  - Chaining: link new pages(overflow pages) off the bucket.

# Example

- Directory is array of size 4.
- To find bucket for  $r$ , take last '*global depth*' # bits of  $h(r)$ ; we denote  $r$  by  $h(r)$ .
  - If  $h(r) = 5 = \text{binary } 101$ , it is in bucket pointed to by 01.



❖ **Insert:** If bucket is full, *split* it (allocate new page, re-distribute).

❖ *If necessary*, double the directory. (As we will see, splitting a bucket does not always require doubling; we can tell by comparing *global depth* with *local depth* for the split bucket.)

# Three basic alternatives for data entries in any index

15

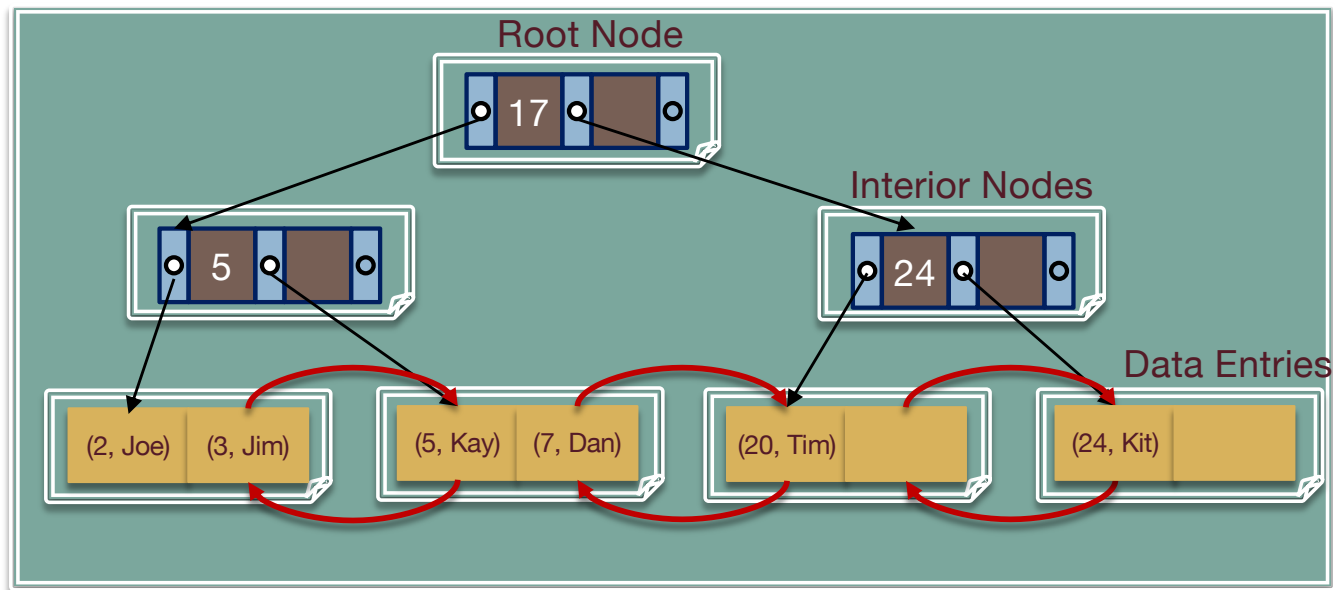
- Three basic alternatives for data entries in any index
  - ▣ Alternative 1: By Value
  - ▣ Alternative 2: By Reference
  - ▣ Alternative 3: By List of references

# Alternative 1 Index (B+ Tree)

16

- Record contents are stored in the index file
  - ▣ No need to follow pointers

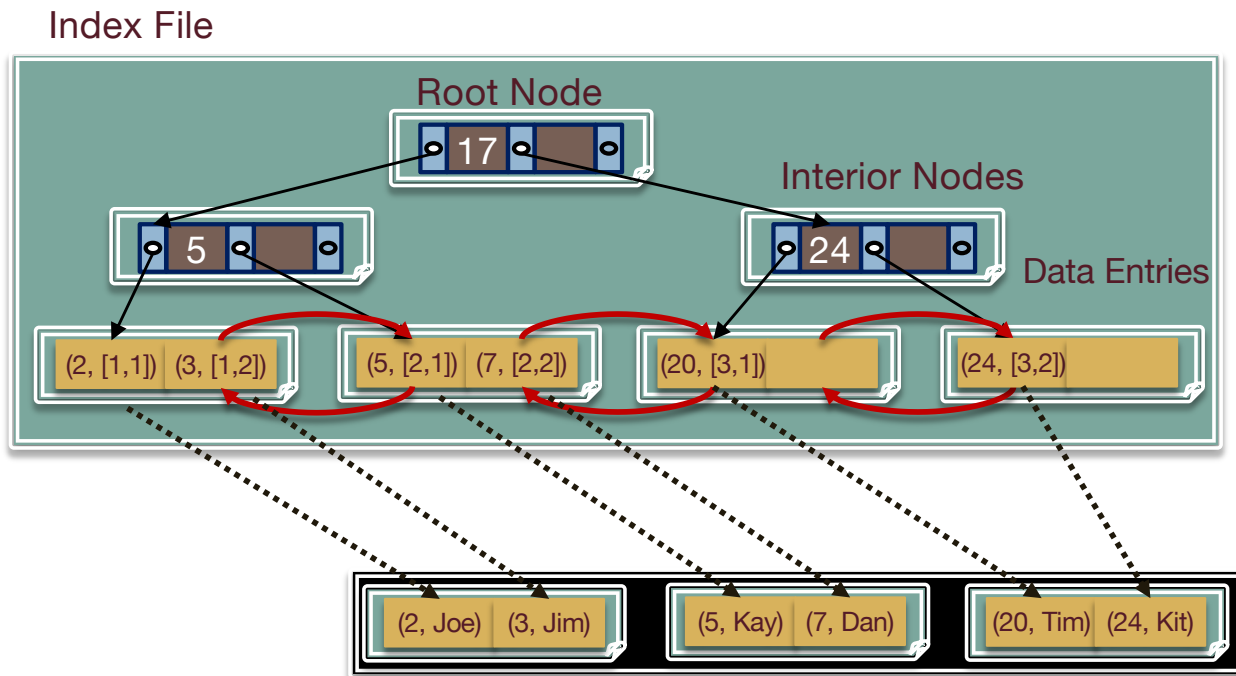
uid	name
2	Joe
3	Jim
5	Kay
7	Dan
20	Tim
24	Kit



# Alternative 2 Index

17

- Alternative 2: **By Reference**,  $\langle k, \text{rid of matching data record} \rangle$



Index Contains  
(Key, Record Id)  
Pairs

uid	name
2	Joe
3	Jim
5	Kay
7	Dan
20	Tim
24	Kit

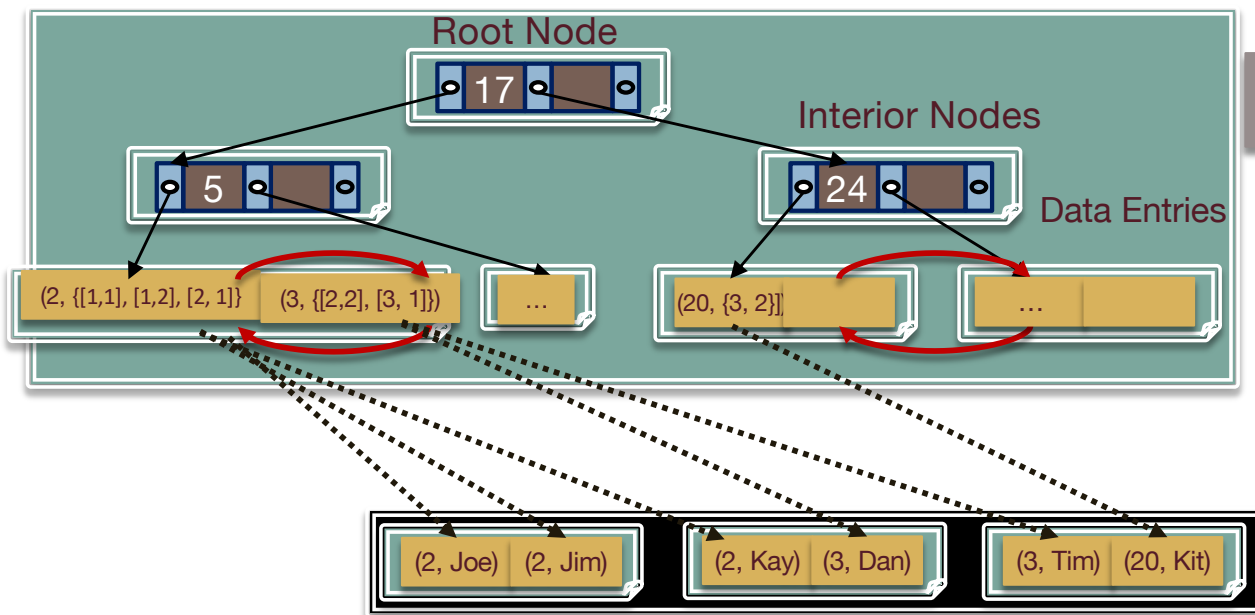


# Alternative 3 Index

18

- Alternative 3: **By List of references**,  $\langle k, \text{list of rids of matching data records} \rangle$ 
  - Alternative 3 more compact than alternative 2
    - For very large rid lists, single data entry spans multiple blocks

Index File



Index Contains  
(Key, {list of record Id}) Pairs

Key	Record Id
2	{[1,1], [1,2], [1,3]}
3	4

# Indexing By Reference

- Both Alternative 2 and Alternative 3 index data *by reference*
- By-reference is *required* to support multiple indexes per table
  - ▣ Otherwise, we would be replicating entire tuples
  - ▣ Replicating data leads to complexity when we're doing updates, so it's something we want to avoid

# Alternative 2 vs Alternative 3 Table Illustration

20

Alternative 2  
Index data entries

Key	Record Id	SSN	Last Name	First Name	Salary
Gonzalez	[3, 1]	123	Gonzalez	Amanda	\$400
Gonzalez	[3, 2]	443	Gonzalez	Joey	\$300
Gonzalez	[3, 3]	244	Gonzalez	Jose	\$140
Hong	[3, 4]	134	Hong	Sue	\$400

Alternative 3  
Index data entries

Key	Record Id
Gonzalez	[3, {1, 2, 3}]
Hong	[3,4]

