

# SQL

## Key Points

- SQL is the standard language for interacting with RDBMS
  - supported by databases from many different vendors (though there are incompatibilities due to incomplete implementations and non-standard extensions)
- bags vs sets
  - in the relational model, relations are sets
    - no duplicate tuples
    - order of tuples doesn't matter
  - in practice, tables are implemented as bags
    - order of rows doesn't matter
    - *duplicates are not automatically removed from query results*
- basic queries
  - SELECT ... FROM ... WHERE ...
    - JOIN, NATURAL JOIN
  - DISTINCT, ORDER BY
  - NULL values
- notation used for syntax in the MySQL documentation

2

## Basic Queries

```
SELECT  $a_1, a_2, \dots, a_n$   
FROM  $R$   
WHERE  $C$ 
```

- (only) table and column names are case-sensitive
- resulting table has
  - the rows from  $R$  for which  $C$  is true
  - the columns  $a_1, a_2, \dots, a_n$

## Basic Queries

```
SELECT  $a_1, a_2, \dots, a_n$   
FROM  $R$   
WHERE  $C$   
ORDER BY  $b_1, b_2, \dots, b_m$ 
```

- rows are sorted according to the specified columns just before the SELECT is applied (so  $b_1, b_2, \dots, b_m$  can involve columns other than  $a_1, a_2, \dots, a_n$ )
- default is ascending order; use DESC for descending order

```
SELECT DISTINCT  $a_1, a_2, \dots, a_n$   
FROM  $R$   
WHERE  $C$ 
```

- as the last step, remove duplicate rows (duplicates identified based only on columns  $a_1, a_2, \dots, a_n$ )

# Basic Queries

## SELECT

- can contain \* (all columns), arithmetic expressions
- can rename individual columns of the result using AS

## FROM

- can rename the whole relation using AS
- can contain R, S (cross product), R JOIN S ON C, R NATURAL JOIN S
  - JOIN keeps all columns of R and S – equivalent to FROM R, S WHERE C
  - NATURAL JOIN keeps only one copy of each join column
- qualify column name with relation name or alias to disambiguate same-named columns

## WHERE

- can contain arithmetic expressions, <, <=, >, >=, =, <>, AND, OR, NOT, IS NULL, IS NOT NULL, LIKE, NOT LIKE
  - wildcards % (0 or more), \_ (single character) allowed in LIKE/NOT LIKE patterns – used for pattern-matching within column values

# Basic Queries

## Notes on data types.

- ' (single quotes) denote string values
- DATE – 'yyyy-mm-dd'
- TIME – 'hh:mm:ss' or 'hh:mm:ss+h:mm' or 'hh:mm:ss-h:mm' (relative to GMT)
- TIMESTAMP – 'yyyy-mm-dd hh:mm:ss'

## NULL values.

- any operation applied to NULL yields NULL
- any comparison with NULL yields UNKNOWN

## Three-valued logic.

- works out as if TRUE = 1, FALSE = 0, UNKNOWN = 1/2 and AND = min, OR = max, NOT = 1-value

# Syntax Notation

- **blue** (and not italics) – type literally that
- *italics* – replace with the appropriate specifics of that kind of thing
- | - list of alternatives, choose one
- [] - optional
- {} - provides grouping when there are alternatives for a required thing (don't type the {} themselves)

```

SELECT
  [ALL | DISTINCT | DISTINCTROW ]
  [HIGH_PRIORITY]
  [STRAIGHT_JOIN]
  [SQL_SMALL_RESULT] [SQL_BIG_RESULT] [SQL_BUFFER_RESULT]
  [SQL_NO_CACHE] [SQL_CALC_FOUND_ROWS]
  select_expr [, select_expr] ...
  [into_option]
  [FROM table_references
  [PARTITION partition_list]]
  [WHERE where_condition]
  [GROUP BY {col_name | expr | position}, ... [WITH ROLLUP]]
  [HAVING where_condition]
  [WINDOW window_name AS (window_spec)
  [, window_name AS (window_spec)] ...]
  [ORDER BY {col_name | expr | position}
  [ASC | DESC], ... [WITH ROLLUP]]
  [LIMIT [{offset},] row_count | row_count OFFSET offset]
  [into_option]
  [FOR {UPDATE | SHARE}
  [OF tbl_name [, tbl_name] ...]
  [NOWAIT | SKIP LOCKED]
  | LOCK IN SHARE MODE]
  [into_option]

into_option: {
  INTO OUTFILE 'file_name'
  [CHARACTER SET charset_name]
  export_options
  | INTO DUMPFILE 'file_name'
  | INTO var_name [, var_name] ...
}
    
```

Figure 5.6 One possible database state for the COMPANY relational database schema.

Empno	Empname	Job	Mgr	Salary	Startdate	Address
1	John	Analyst		30000	1965-01-09	731 Fondren, Houston, TX
2	Franklin	Analyst	1	40000	1965-12-08	638 Voss, Houston, TX
3	Alicia	Analyst	1	25000	1968-01-19	3321 Castle, Spring, TX
4	Jennifer	Analyst	1	43000	1981-06-20	291 Berry, Bellaire, TX
5	Ramesh	Analyst	1	38000	1962-09-15	975 Fire Oak, Humble, TX
6	Joyce	Analyst	1	25000	1972-07-31	5631 Rice, Houston, TX
7	Ahmad	Analyst	1	25000	1969-03-29	980 Dallas, Houston, TX
8	James	Analyst	1	55000	1937-11-10	450 Stone, Houston, TX

  

Dname	Dnumber	Mgr_ssn	Mgr_start_date
Research	5	333445555	1968-05-22
Administration	4	987654321	1995-01-01
Headquarters	1	888665555	1981-06-19

  

Dnumber	Dlocation
1	Houston
4	Stafford
5	Bellaire
5	Sugarland
5	Houston

  

Empno	Projno	Hours
123456789	1	32.5
123456789	2	7.5
666884444	3	40.0
453453453	1	20.0
453453453	2	20.0
333445555	2	10.0
333445555	3	10.0
333445555	10	10.0
333445555	20	10.0
999887777	30	30.0
999887777	10	10.0
987987987	10	35.0
987987987	30	5.0
987654321	30	20.0
987654321	20	15.0
888665555	20	NULL

  

Pname	Pnumber	Plocation	Dnum
ProductX	1	Bellaire	5
ProductY	2	Sugarland	5
ProductZ	3	Houston	5
Computerization	10	Stafford	4
Reorganization	20	Houston	1
Newbenefits	30	Stafford	4

  

Empno	Dependent_name	Sex	Bdate	Relationship
333445555	Alice	F	1986-04-05	Daughter
333445555	Theodore	M	1983-10-25	Son
333445555	Joy	F	1958-05-03	Spouse
987987987	Abner	M	1942-02-28	Spouse
123456789	Michael	M	1988-01-04	Son
123456789	Alice	F	1988-12-30	Daughter
123456789	Elizabeth	F	1967-05-05	Spouse

```

SELECT DISTINCT Dependent_name
FROM DEPENDENT
WHERE Sex='F'

SELECT *
FROM DEPT_LOCATIONS JOIN PROJECT ON Dnumber=Dnum

SELECT *
FROM WORKS_ON JOIN PROJECT ON Pno=Pnumber
WHERE Plocation='Houston' AND Hours >= 10

SELECT Fname,Lname
FROM EMPLOYEE JOIN DEPARTMENT ON Super_ssn=Mgr_ssn
WHERE Dno=Dnumber

SELECT Fname,Lname,Dependent_name,Relationship
FROM EMPLOYEE NATURAL JOIN DEPENDENT

SELECT *
FROM DEPARTMENT NATURAL JOIN DEPT_LOCATIONS

```

0

## Other Joins

```

table_reference {[INNER | CROSS] JOIN | STRAIGHT_JOIN} table_factor [join_specification]
| table_reference {LEFT|RIGHT} [OUTER] JOIN table_reference join_specification
| table_reference NATURAL [INNER | {LEFT|RIGHT} [OUTER]] JOIN table_factor

```

- R LEFT JOIN S
  - if there is no matching row for R in S, a row with all NULLs is used as the match
  - every row of R is represented in the result
  - R RIGHT JOIN S is equivalent to S LEFT JOIN R
    - LEFT JOIN is recommended for portability
- (full) OUTER JOIN
  - if a row in either table doesn't have a match in the other, a row with all NULLs is used as the match
  - MySQL does not (directly) support a full outer join
- CROSS JOIN, INNER JOIN
  - syntactically equivalent to JOIN in MySQL (but not in standard SQL)
- R STRAIGHT\_JOIN S
  - R is always read first – gives directions to the optimizer when the query is processed

1

## Strategies for Writing Queries

- start with FROM – build a table that contains all of the rows and columns that you need
- add WHERE to pick the desired rows
  - prefer JOIN condition to WHERE condition when specifying how rows of one table combine with rows of another
- add SELECT to pick the desired columns