

# Introduction to SQL (III)

# Roadmap of This Lecture

- Transactions
- Integrity Constraints
- SQL Data Types and Schemas
- Authorization
- Embedded SQL

# Transactions

- Logical unit of work – contains several sequential actions
- Atomic transaction
  - either fully executed or rolled back as if it never occurred
- Isolation from concurrent transactions
- Transactions begin implicitly
  - Ended by **commit work** or **rollback work**
- But default on most databases: each SQL statement commits automatically
  - Can turn off auto commit for a session (e.g. using API)
  - In SQL:1999, can use: **begin atomic .... end**

# Integrity Constraints

- Integrity constraints guard against accidental damage to the database, by ensuring that authorized changes to the database do not result in a loss of data consistency.
  - A checking account must have a balance greater than \$10,000.00
  - A salary of a bank employee must be at least \$4.00 an hour
  - A customer must have a (non-null) phone number

# Integrity Constraints on a Single Relation

- **not null**
- **primary key**
- **unique**
- **check (P)**, where P is a predicate

# Not Null and Unique Constraints

## ■ not null

- Declare *name* and *budget* to be **not null**

*name* **varchar(20) not null**

*budget* **numeric(12,2) not null**

## ■ unique ( $A_1, A_2, \dots, A_m$ )

- The unique specification states that the attributes

$A_1, A_2, \dots, A_m$

form a candidate key.

- Candidate keys are permitted to be null (in contrast to primary keys).

# The check clause

- **check** (P)

where P is a predicate

Example: ensure that semester value is one of fall, winter, spring or summer:

```
create table section (  
    course_id varchar (8),  
    sec_id varchar (8),  
    semester varchar (6),  
    year numeric (4,0),  
    building varchar (15),  
    room_number varchar (7),  
    time slot id varchar (4),  
    primary key (course_id, sec_id, semester, year),  
    check (semester in ('Fall', 'Winter', 'Spring', 'Summer'))  
);
```

# Referential Integrity

- Ensures that a value that appears in one relation for a given set of attributes also appears for a certain set of attributes in another relation.
  - Example: If “Biology” is a department name appearing in one of the tuples in the *instructor* relation, then there exists a tuple in the *department* relation for “Biology”.
- Let A be a set of attributes. Let R and S be two relations that contain attributes A and where A is the primary key of S. A is said to be a **foreign key** of R if for any values of A appearing in R these values also appear in S.



# Cascading Actions in Referential Integrity

- **create table** *course* (  
    *course\_id* **char**(5),  
    *title* **varchar**(20),  
    *dept\_name* **varchar**(20),  
    **primary key** (*course\_id*)  
    **foreign key** (*dept\_name*) **references** *department*)
- **create table** *course* (  
    ...  
    *dept\_name* **varchar**(20),  
    **foreign key** (*dept\_name*) **references** *department*  
        **on delete cascade**  
        **on update cascade**,  
    ...  
)
- alternative actions to cascade: **set null, set default**

# Integrity Constraint Violation During Transactions

- E.g.

```
create table person (  
  ID char(10),  
  name char(40),  
  mother char(10),  
  father char(10),  
  primary key ID,  
  foreign key father references person,  
  foreign key mother references person)
```

- How to insert the first tuple without causing constraint violation ?
  - insert father and mother of a person before inserting person
  - OR, set father and mother to null initially, update after inserting all persons (not possible if father and mother attributes declared to be **not null**)
  - OR defer constraint checking (rarely supported)

# Complex Check Clauses

- **check** (*time\_slot\_id* in (**select** *time\_slot\_id* from *time\_slot*))
  - why not use a foreign key here?
- Every section has at least one instructor teaching the section.
  - teaches(ID, course\_id, sec\_id, semester, year)
  - section(course\_id, sec\_id, semester, year)
  - how to write this?
  - check((course\_id, sec\_id, semester, year) in  
(select course\_id, sec\_id, semester, year from teaches))
- **Unfortunately: subquery in check clause not supported by pretty much any database**
  - Alternative: triggers (later)
- **create assertion** <assertion-name> **check** <predicate>;
  - Also not supported by anyone

# Index Creation

- **create table** *student*  
(*ID* **varchar** (5),  
*name* **varchar** (20) **not null**,  
*dept\_name* **varchar** (20),  
*tot\_cred* **numeric** (3,0) **default** 0,  
**primary key** (*ID*))
- **create index** *studentID\_index* **on** *student*(*ID*)
- Indices are data structures used to speed up access to records with specified values for index attributes
  - e.g. **select** \*  
    **from** *student*  
    **where** *ID* = '12345'  
can be executed by using the index to find the required record, without looking at all records of *student*  
*More on indices later.*

# Built-in Data Types in SQL

- **date:** Dates, containing a (4 digit) year, month and date
  - Example: **date** '2005-7-27'
- **time:** Time of day, in hours, minutes and seconds.
  - Example: **time** '09:00:30'      **time** '09:00:30.75'
- **timestamp:** date plus time of day
  - Example: **timestamp** '2005-7-27 09:00:30.75'
- **interval:** period of time
  - Example: **interval** '1' day
  - Subtracting a date/time/timestamp value from another gives an interval value
  - Interval values can be added to date/time/timestamp values

# User-Defined Types

- **create type** construct in SQL creates user-defined type

**create type *Dollars* as numeric (12,2) final**

- **create table *department***  
(*dept\_name* **varchar** (20),  
*building* **varchar** (15),  
*budget* *Dollars*);

No subtypes can be  
Defined from Dollar



# Domains

- **create domain** construct in SQL-92 creates user-defined domain types

```
create domain person_name char(20) not null
```

- Types and domains are similar.
  - Domains can have constraints, such as **not null**, specified on them.
  - Domains are *not* strongly typed.
- **create domain** *degree\_level* **varchar**(10)  
**constraint** *degree\_level\_test*  
**check** (**value in** ('Bachelors', 'Masters', 'Doctorate'));

# Large-Object Types

- Large objects (photos, videos, CAD files, etc.) are stored as a *large object*.
  - **blob**: binary large object -- object is a large collection of uninterpreted binary data (whose interpretation is left to an application outside of the database system)
  - **clob**: character large object -- object is a large collection of character data
  - When a query returns a large object, a pointer is returned rather than the large object itself.



# Authorization

Forms of authorization on parts of the database:

- **Read** - allows reading, but not modification of data.
- **Insert** - allows insertion of new data, but not modification of existing data.
- **Update** - allows modification, but not deletion of data.
- **Delete** - allows deletion of data.

Forms of authorization to modify the database schema

- **Index** - allows creation and deletion of indices.
- **Resources** - allows creation of new relations.
- **Alteration** - allows addition or deletion of attributes in a relation.
- **Drop** - allows deletion of relations.

# Authorization Specification in SQL

- The **grant** statement is used to confer authorization
  - grant** <privilege list>
  - on** <relation name or view name> **to** <user list>
- <user list> is:
  - a user-id
  - **public**, which allows all valid users the privilege granted
  - A role (more on this later)
- Granting a privilege on a view does not imply granting any privileges on the underlying relations.
- The grantor of the privilege must already hold the privilege on the specified item (or be the database administrator).

# Privileges in SQL

- **select**: allows read access to relation, or the ability to query using the view
  - Example: grant users  $U_1$ ,  $U_2$ , and  $U_3$  the **select** authorization on the *instructor* relation:

**grant select on *instructor* to  $U_1, U_2, U_3$**

- **insert**: the ability to insert tuples
- **update**: the ability to update using the SQL update statement
- **delete**: the ability to delete tuples.
- **all privileges**: used as a short form for all the allowable privileges

# Revoking Authorization in SQL

- The **revoke** statement is used to revoke authorization.

**revoke** <privilege list>

**on** <relation name or view name> **from** <revokee list>

- Example:

**revoke select on** *branch* **from**  $U_1, U_2, U_3$

- <privilege list> may be **all** to revoke all privileges the revokee may hold.
- If <revokee list> includes **public**, all users lose the privilege except those granted it explicitly.
- If the same privilege was granted twice to the same user by different grantors, the user may retain the privilege after the revocation.
- All privileges that depend on the privilege being revoked are also revoked.
- Question: What if the grantor and the grantee have the *same* privilege on a relation, and the *grantee* wants to revoke the privilege of the *grantor*?

# Roles

- **create role** *instructor*;
  - **grant** *instructor* **to** Amit;
- Privileges can be granted to roles:
  - **grant select on** *takes* **to** *instructor*;
- Roles can be granted to users, as well as to other roles
  - **create role** *teaching\_assistant*;
  - **grant** *teaching\_assistant* **to** *instructor*;
    - ▶ *instructor* inherits all privileges of *teaching\_assistant*
- Chain of Roles
  - **create role** *dean*;
  - **grant** *instructor* **to** *dean*;
  - **grant** *dean* **to** Satoshi;

# Authorization on Views

- **create view** *geo\_instructor* **as**  
(**select** \*  
**from** *instructor*  
**where** *dept\_name* = 'Geology');
- **grant select on** *geo\_instructor* **to** *geo\_staff*
- Suppose that a *geo-staff* member issues
  - **select** \*  
**from** *geo\_instructor*;
- Clearly the *geo-staff* should be able to issue the query?
  - Need to deal with the case where *geo-staff* does not have authorization to *instructor*

# Authorizations on Schema

- **references** privilege to create foreign key
  - **grant reference** (*dept\_name*) **on** *department* **to** Mariano;
  - why is this required?
  - Because a foreign key guarantees the existence of the value in the other table -- can perform existence check on the other table!

# Transfer of Privileges

- Transfer of privileges
  - **grant select on *department* to Amit with grant option;**
  - **revoke select on *department* from Amit, Satoshi cascade;**
  - **revoke select on *department* from Amit, Satoshi restrict;**



# Embedded SQL

- The SQL standard defines embeddings of SQL in a variety of programming languages such as C, Java, and Cobol.
- A language to which SQL queries are embedded is referred to as a **host language**, and the SQL structures permitted in the host language comprise ***embedded SQL***.
- The basic form of these languages follows that of the System R embedding of SQL into PL/I.
- **EXEC SQL** statement is used to identify embedded SQL request to the preprocessor

EXEC SQL <embedded SQL statement > END\_EXEC

Note: this varies by language (for example, the Java embedding uses  
# SQL { .... }; )

# Example Query

- From within a host language, find the ID and name of students who have completed more than the number of credits stored in variable *credit\_amount*.
- Specify the query in SQL and declare a *cursor* for it

EXEC SQL

```
declare c cursor for  
select ID, name  
from student  
where tot_cred > :credit_amount
```

END\_EXEC

# Embedded SQL (Cont.)

- The **open** statement causes the query to be evaluated

```
EXEC SQL open c END_EXEC
```

- The **fetch** statement causes the values of one tuple in the query result to be placed on host language variables.

```
EXEC SQL fetch c into :si, :sn END_EXEC
```

*si* holds the ID and *sn* holds the name

Repeated calls to **fetch** get successive tuples in the query result

- A variable called SQLSTATE in the SQL communication area (SQLCA) gets set to '02000' to indicate no more data is available
- The **close** statement causes the database system to delete the temporary relation that holds the result of the query.

```
EXEC SQL close c END_EXEC
```

Note: above details vary with language. For example, the Java embedding defines Java iterators to step through result tuples.

# Updates Through Cursors

- Can update tuples fetched by cursor by declaring that the cursor is for update

```
declare c cursor for  
  select *  
  from instructor  
  where dept_name = 'Music'  
for update
```

- To update tuple at the current location of cursor *c*

```
update instructor  
  set salary = salary + 100  
  where current of c
```

# JDBC and ODBC

- API (application-program interface) for a program to interact with a database server
- Application makes calls to
  - Connect with the database server
  - Send SQL commands to the database server
  - Fetch tuples of result one-by-one into program variables
  - SQL queries are created at runtime and hence “***dynamic SQL***”
- ODBC (Open Database Connectivity) works with C, C++, C#, and Visual Basic
  - Other API's such as ADO.NET sit on top of ODBC
- JDBC (Java Database Connectivity) works with Java

# JDBC

- **JDBC** is a Java API for communicating with database systems supporting SQL.
- JDBC supports a variety of features for querying and updating data, and for retrieving query results.
- JDBC also supports metadata retrieval, such as querying about relations present in the database and the names and types of relation attributes.
- Model for communicating with the database:
  - Open a connection
  - Create a “statement” object
  - Execute queries using the Statement object to send queries and fetch results
  - Exception mechanism to handle errors

# JDBC Code

```
public static void JDBCexample(String dbid, String userid, String
    passwd)
{
    try {
        Class.forName ("oracle.jdbc.driver.OracleDriver");
        Connection conn = DriverManager.getConnection(
            "jdbc:oracle:thin:@db.yale.edu:2000:univdb", userid,
            passwd);
        Statement stmt = conn.createStatement();
        ... Do Actual Work ....
        stmt.close();
        conn.close();
    }
    catch (SQLException sqle) {
        System.out.println("SQLException : " + sqle);
    }
}
```

# JDBC Code (Cont.)

- Update to database

```
try {  
    stmt.executeUpdate(  
        "insert into instructor values('77987', 'Kim', 'Physics', 98000)");  
} catch (SQLException sqle)  
{  
    System.out.println("Could not insert tuple. " + sqle);  
}
```

- Execute query and fetch and print results

```
ResultSet rset = stmt.executeQuery(  
    "select dept_name, avg (salary)  
    from instructor  
    group by dept_name");  
while (rset.next()) {  
    System.out.println(rset.getString("dept_name") + " " +  
        rset.getFloat(2));  
}
```



# JDBC Code Details

- Getting result fields:
  - **`rs.getString("dept_name")` and `rs.getString(1)` equivalent if `dept_name` is the first argument of select result.**
- Dealing with Null values
  - **`int a = rs.getInt("a");`  
`if (rs.isNull()) Systems.out.println("Got null value");`**

# Prepared Statement

- `PreparedStatement pStmt = conn.prepareStatement("insert into instructor values(?,?,?,?)");`  
`pStmt.setString(1, "88877");`  
`pStmt.setString(2, "Perry");`  
`pStmt.setString(3, "Finance");`  
`pStmt.setInt(4, 125000);`  
`pStmt.executeUpdate();`  
`pStmt.setString(1, "88878");`  
`pStmt.executeUpdate();`
- **WARNING:** always use prepared statements when taking an input from the user and adding it to a query
  - NEVER create a query by concatenating strings
  - `"insert into instructor values(" + ID + ", " + name + ", " + dept name + ", " + balance + ")"`
  - What if name is "D'Souza"?

# SQL Injection

- Suppose query is constructed using
  - "select \* from instructor where name = '" + name + "'"
- Suppose the user, instead of entering a name, enters:
  - X' or 'Y' = 'Y
- then the resulting statement becomes:
  - "select \* from instructor where name = '" + "X' or 'Y' = 'Y" + "'"
  - which is:
    - ▶ select \* from instructor where name = 'X' or 'Y' = 'Y'
  - User could have even used
    - ▶ X'; update instructor set salary = salary + 10000; --
- Prepared statement internally uses:  
"select \* from instructor where name = 'X\'' or \'Y\' = \'Y'"
  - **Always use prepared statements, with user inputs as parameters**

# Metadata Features

- ResultSet metadata
- E.g., after executing query to get a ResultSet rs:
  - `ResultSetMetaData rsmd = rs.getMetaData();`  
`for(int i = 1; i <= rsmd.getColumnCount(); i++) {`  
    `System.out.println(rsmd.getColumnName(i));`  
    `System.out.println(rsmd.getColumnTypeName(i));`  
}
- How is this useful?
  - Print the scheme for this relation

# Metadata (Cont)

- Database metadata

- DatabaseMetaData dbmd = conn.getMetaData();

```
ResultSet rs = dbmd.getColumns(null, "univdb", "department", "%");
```

```
// Arguments to getColumns: Catalog, Schema-pattern, Table-pattern,  
// and Column-Pattern
```

```
// Returns: One row for each column; row has a number of attributes
```

```
// such as COLUMN_NAME, TYPE_NAME
```

```
while( rs.next() ) {
```

```
    System.out.println(rs.getString("COLUMN_NAME"),  
    rs.getString("TYPE_NAME"));
```

```
}
```

- And where is this useful?

- Only those specified columns are retrieved

# Transaction Control in JDBC

- By default, each SQL statement is treated as a separate transaction that is committed automatically
  - bad idea for transactions with multiple updates
- Can turn off automatic commit on a connection
  - `conn.setAutoCommit(false);`
- Transactions must then be committed or rolled back explicitly
  - `conn.commit();` or
  - `conn.rollback();`
- `conn.setAutoCommit(true)` turns on automatic commit.

# Other JDBC Features

- Calling functions and procedures
  - `CallableStatement cStmt1 = conn.prepareCall("{? = call some function(?)})");`
  - `CallableStatement cStmt2 = conn.prepareCall("{call some procedure(?,?)})");`
- Handling large object types
  - `getBlob()` and `getClob()` that are similar to the `getString()` method, but return objects of type `Blob` and `Clob`, respectively
  - get data from these objects by `getBytes()`
  - associate an open stream with Java `Blob` or `Clob` object to update large objects
    - ▶ `blob.setBlob(int parameterIndex, InputStream inputStream).`

# SQLJ

- JDBC is overly dynamic, errors cannot be caught by compiler
- SQLJ: embedded SQL in Java

- ```
#sql iterator deptInfolter ( String dept_name, int avgSal);  
deptInfolter iter = null;  
#sql iter = { select dept_name, avg(salary) from instructor  
              group by dept name };  
while (iter.next()) {  
    String deptName = iter.dept_name();  
    int avgSal = iter.avgSal();  
    System.out.println(deptName + " " + avgSal);  
}  
iter.close();
```



# ODBC

- Open DataBase Connectivity(ODBC) standard
  - standard for application program to communicate with a database server.
  - application program interface (API) to
    - ▶ open a connection with a database,
    - ▶ send queries and updates,
    - ▶ get back results.
- Applications such as GUI, spreadsheets, etc. can use ODBC

# ODBC (Cont.)

- Each database system supporting ODBC provides a "driver" library that must be linked with the client program.
- When client program makes an ODBC API call, the code in the library communicates with the server to carry out the requested action, and fetch results.
- ODBC program first allocates an SQL environment, then a database connection handle.
- Opens database connection using `SQLConnect()`. Parameters for `SQLConnect`:
  - connection handle,
  - the server to which to connect
  - the user identifier,
  - password
- Must also specify types of arguments:
  - `SQL_NTS` denotes previous argument is a null-terminated string.

# ODBC Code

```
■ int ODBCexample()
{
    RETCODE error;
    HENV  env;  /* environment */
    HDBC  conn; /* database connection */
    SQLAllocEnv(&env);
    SQLAllocConnect(env, &conn);
    SQLConnect(conn, "db.yale.edu", SQL_NTS, "avi", SQL_NTS,
               "avipasswd", SQL_NTS);
    { .... Do actual work ... }

    SQLDisconnect(conn);
    SQLFreeConnect(conn);
    SQLFreeEnv(env);
}
```

# ODBC Code (Cont.)

- Program sends SQL commands to the database by using `SQLExecDirect`
- Result tuples are fetched using `SQLFetch()`
- `SQLBindCol()` binds C language variables to attributes of the query result
  - When a tuple is fetched, its attribute values are automatically stored in corresponding C variables.
  - Arguments to `SQLBindCol()`
    - ▶ ODBC stmt variable, attribute position in query result
    - ▶ The type conversion from SQL to C.
    - ▶ The address of the variable.
    - ▶ For variable-length types like character arrays,
      - The maximum length of the variable
      - Location to store actual length when a tuple is fetched.
      - Note: A negative value returned for the length field indicates null value
- Good programming requires checking results of every function call for errors; we have omitted most checks for brevity.

# ODBC Code (Cont.)

- Main body of program

```
char deptname[80];
float salary;
int lenOut1, lenOut2;
HSTMT stmt;
char * sqlquery = "select dept_name, sum (salary)
                  from instructor
                  group by dept_name";
SQLAllocStmt(conn, &stmt);
error = SQLExecDirect(stmt, sqlquery, SQL NTS);
if (error == SQL SUCCESS) {
    SQLBindCol(stmt, 1, SQL C CHAR, deptname , 80,
&lenOut1);
    SQLBindCol(stmt, 2, SQL C FLOAT, &salary, 0 , &lenOut2);
    while (SQLFetch(stmt) == SQL SUCCESS) {
        printf (" %s %g\n", deptname, salary);
    }
}
SQLFreeStmt(stmt, SQL DROP);
```

# ODBC Prepared Statements

## ■ Prepared Statement

- SQL statement prepared: compiled at the database
- Can have placeholders: E.g. insert into account values(?,?,?)
- Repeatedly executed with actual values for the placeholders

## ■ To prepare a statement

```
SQLPrepare(stmt, <SQL String>);
```

## ■ To bind parameters

```
SQLBindParameter(stmt, <parameter#>,  
... type information and value omitted for simplicity..)
```

## ■ To execute the statement

```
retcode = SQLExecute( stmt);
```

## ■ To avoid SQL injection security risk, do not create SQL strings directly using user input; instead use prepared statements to bind user inputs

# More ODBC Features

## ■ Metadata features

- finding all the relations in the database and
- finding the names and types of columns of a query result or a relation in the database.

## ■ By default, each SQL statement is treated as a separate transaction that is committed automatically.

- Can turn off automatic commit on a connection
  - ▶ `SQLSetConnectOption(conn, SQL_AUTOCOMMIT, 0)`
- Transactions must then be committed or rolled back explicitly by
  - ▶ `SQLTransact(conn, SQL_COMMIT)` or
  - ▶ `SQLTransact(conn, SQL_ROLLBACK)`

# ODBC Conformance Levels

- Conformance levels specify subsets of the functionality defined by the standard.
  - Core
  - Level 1 requires support for metadata querying
  - Level 2 requires ability to send and retrieve arrays of parameter values and more detailed catalog information.
- SQL Call Level Interface (CLI) standard similar to ODBC interface, but with some minor differences.



# ADO.NET

- API designed for Visual Basic .NET and C#, providing database access facilities similar to JDBC/ODBC
  - Partial example of ADO.NET code in C#  
using System, System.Data, System.Data.SqlClient;  
SqlConnection conn = new SqlConnection(  
    "Data Source=<IPaddr>, Initial Catalog=<Catalog>");  
conn.Open();  
SqlCommand cmd = new SqlCommand("select \* from students",  
    conn);  
  
SqlDataReader rdr = cmd.ExecuteReader();  
while(rdr.Read()) {  
    Console.WriteLine(rdr[0], rdr[1]); /\* Prints first 2 attributes of result\*/  
}  
rdr.Close(); conn.Close();
- Translated into ODBC calls
- Can also access non-relational data sources such as
  - OLE-DB
  - XML data
  - Entity framework