

Database Connectivity: JDBC

Database Connectivity Basics

Application-level database connectivity:

- Host language (Java, C/C++, Ruby, Python, Perl, PHP, etc)
- Target DBMS (MySQL, PostgreSQL, Oracle, MS SQL, IBM DB2, etc)
- Client — Server environment
 - Client: application program
 - Server: DBMS

General structure:

1. Load database driver/database support functionality
2. Form a SQL statement
3. Connect to the DBMS
4. Pass SQL statement to the DBMS
5. Receive result
6. Close connection

Case Study: JDBC

JDBC originally, an abbreviation for Java Database Connectivity, is the database connectivity package for Java.

Loading the database driver

First task in any JDBC application is loading the JDBC driver for the right DBMS. In newer versions of Java/JDBC, you do **not need to do this step** as long as the connector jar is in your CLASSPATH. This is done statically by the driver's class, so all you need to do is load the class into the runtime. This can be done by using the `Class.forName(String name)` method. The argument passed to the `Class.forName` method is the name of the JDBC driver for a specific DBMS server.

DBMS	Driver name
MySQL	<code>com.mysql.jdbc.driver</code>
PostgreSQL	<code>org.postgresql.Driver</code>
Oracle	<code>oracle.jdbc.OracleDriver</code>
Microsoft SQL Server	<code>com.microsoft.jdbc.sqlserver.SQLServerDriver</code>
IBM DB2	<code>COM.ibm.db2.jdbc.app.DB2Driver</code>

Example. The following code loads Oracle's JDBC driver or, if unsuccessful, reports an error.

```
try {
    Class.forName("com.mysql.jdbc.driver");
} catch (ClassNotFoundException ex) {
    System.err.println("Driver not found");
};
```

Establishing a Connection

JDBC package contains a **Connection** class representing client-server connections between the client Java applications and the DBMS servers. Note that creating a connection is a very costly operation and should only be done when needed. An instance of the **Connection** class will be created via the following driver manager call:

```
Connection conn =
    DriverManager.getConnection(connectionString, user, password);
```

Here,

connectionString is a connection URL specifying location and connection port for the database. See below for syntax.

user is the DBMS user login account.

password is the password for the DBMS account of the user **user**.

Connection String

Connection string has the following syntax:

```
<driver>:<dbms>://<server>[:<port>]/<database>[?<additionalOption>=<value>[&<anotherOption>=<value>]*]
```

The connection string for the **CoolKids** database on our class MySQL server (using the standard MySQL port (3306)) could be:

```
jdbc:mysql://csc-db0.calpoly.edu:3306/CoolKids?autoReconnect=true
```

Example. The following code establishes the connection to our MySQL server:

```
Connection conn = null;

String url =
    "jdbc:mysql://csc-db0.calpoly.edu:3306/CoolKids?autoReconnect=true";
String user = "foo";
String password= "bar";

try {
    conn = DriverManager.getConnection(url, user, password);
} catch (Exception ex) {
    System.err.println("Could not open connection");
}
```

Statements

Working with a `Connection` object within a Java program is straightforward: SQL statements are created and passed via the connection, results are received in return. There are three classes for SQL statements:

Statement: general use statement class. Used for creation and execution of SQL statements, typically, once during the run of the program.

PreparedStatement: statement class to be used in the following cases:

- a sequence of similar SQL statements, different only in values of some parameters needs to be executed;
- a single time-consuming SQL statement needs to be executed, possibly multiple times.

SQL statements represented by instances of `PreparedStatement` class are pre-compiled and thus may be more efficiently executed.

CallableStatement: statement class for execution of stored SQL (PL/SQL) procedures. A `Statement` can also be used to execute stored procedures.

Instances of each class are obtained by invoking methods (see below) from the `Connection` class.

JDBC distinguishes two types of SQL statements:

Non-Query SQL statements: All DDL and DML statements, which do NOT return relational tables.

Queries: SELECT statements and their combinations, which return relational tables.

Because queries return tables while non-queries return only exit status, different methods are used to pass these two types of SQL statements.

Class Statement

Obtaining Instances. Instances of class `Statement` can be obtained by invoking the `createStatement()` method of the `Connection` class:

```
Statement s = conn.createStatement();
```

Executing Non-Queries

Non-queries (a.k.a. updates) are executed using method `executeUpdate()` of class `Statement`. While, several method signatures exist, the method call to be used under most standard circumstances is:

```
int executeUpdate(String sql) throws SQLException
```

The method returns the number of rows affected by the update, or 0 if a DDL statement (CREATE TABLE, etc.) was executed.

For example, the following sequence executes two statements: a table `Employees` is created and a record is inserted into it.

```
String update = "CREATE TABLE Employees (" +
    "    id INT PRIMARY KEY," +
    "    Name CHAR(30)," +
    "    Salary INT" +
    ")";

try {
    System.out.println(s.executeUpdate(update));
}
```

```

        // Prints "0".

        s.executeUpdate("INSERT INTO Employees VALUES(1, 'John Smith', 30000)");
        // Prints "1".
    } catch (SQLException e) {
    }
}

```

Note, that the **Statement** instance is **reusable**. Generally speaking, in order to execute a sequence of SQL statements, you only need to create one **Statement** instance.

Executing SQL queries

Use method `executeQuery()` of class **Statement**. The method returns an instance of the class **ResultSet**, which is discussed below.

```

    try {
        String query = "SELECT * FROM Employees WHERE Name = 'Jones'";
        ResultSet results = s.executeQuery(query);
    } catch (SQLException e) {
    }
}

```

Class **PreparedStatement**

PreparedStatement should be used when the same query with possibly different parameters is to be executed multiple times in the course of a program.

Instances of **PreparedStatement** are created with an SQL statement, possibly with parameter placeholders associated with them, and that association cannot change.

Obtaining instances

Instances of class **PreparedStatement** can be obtained by invoking the `prepareStatement(String sql)` method of the **Connection** class:

```

    String sql = "INSERT INTO Employee VALUES(2, 'Bob Brown', 40000)";
    PreparedStatement ps = conn.prepareStatement(sql);

```

This code creates a prepared SQL statement associated with the SQL statement `INSERT INTO Employee VALUES(2, 'Bob Brown', 40000)`.

Parameterized Prepared Statements

The text of the SQL code for the **PreparedStatement** instance can contain `'?'` symbols: one symbol per input parameter to the query. For example, in order to create a generic parameterized INSERT statement for the **Employee** table, we can do the following:

```

    String sql = "INSERT INTO Employee VALUES(?, ?, ?)";
    PreparedStatement ps = conn.prepareStatement(sql);

```

`prepareStatement()` method parses the input SQL string and identifies locations of all parameters. Each parameter gets a number (starting with 1).

Setting Parameter Values.

PreparedStatement uses the following methods to set values for parameters (note: all methods are void and throw SQLException):

Method	Explanation
setNull(int parIndex, int jdbcType)	sets parameter parIndex to null
setBoolean(int parIndex, boolean x)	sets parameter parIndex to a boolean value x
setByte(int parIndex, byte x)	sets parameter parIndex to a byte value x
setInt(int parIndex, int x)	sets parameter parIndex to an integer value x
setLong(int parIndex, long x)	sets parameter parIndex to a long integer value x
setFloat(int parIndex, float x)	sets parameter parIndex to a floating point value x
setDouble(int parIndex, double x)	sets parameter parIndex to a double precision value x
setString(int parIndex, String x)	sets parameter parIndex to a string value x
setDate(int parIndex, java.sql.Date x)	sets parameter parIndex to a date value x
setTime(int parIndex, java.sql.Time x)	sets parameter parIndex to a time value x
clearParameters()	clears the values of all parameters

Executing Non-Query Statements

PreparedStatement class has the executeUpdate() method to execute non-query SQL statements. This method takes no input arguments (since the SQL statement is already prepared).

The following example shows how parameters are set up and prepared updates are executed:

```
try {
    // set first column value for the INSERT statement (ID)
    ps.setInt(1,3);

    // set second column value (Name)
    ps.setString(2, 'Mary Williams');

    // set third column value (Salary)
    ps.setInt(3, 45000);

    // execute INSERT INTO Employee VALUES(3, 'Mary Williams', 45000)
    ps.executeUpdate();
} catch (SQLException e) {
}
```

Executing Queries

To execute queries, use executeQuery() method, which also does not take any arguments. This method returns an instance of ResultSet.

The following code fragment shows the preparation and execution of SELECT statements which select rows of the Employee table by salary.

```
try {
    String sql = "SELECT * FROM Employee WHERE Salary > ?";
    PreparedStatement ps = conn.prepareStatement(sql);

    ps.setInt(1, 35000);

    // execute SELECT * FROM Employee WHERE Salary > 35000
    ResultSet result = ps.executeQuery();

    // clear all parameters
```

```

ps.clearParameters();

ps.setString(1,42000);

// execute SELECT * FROM Employee WHERE Salary > 42000
result = ps.executeQuery();
} catch (SQLException e) {
}

```

Working with output: Class ResultSet

Results of SELECT statements (and other SQL statements that return tables) are stored in instances of the `ResultSet` class.

An instance of a `ResultSet` maintains a *cursor* which points to the currently observed record (tuple) in the returned table. The following methods can be used to navigate a `ResultSet` object:

Method	Explanation
<code>boolean next()</code>	move cursor to the next record.
<code>boolean previous()</code>	move cursor to the previous record.
<code>boolean first()</code>	move cursor in front of the first record.
<code>boolean last()</code>	move cursor to the last row of the cursor.
<code>boolean absolute(int row)</code>	move cursor to the record number <code>row</code> .
<code>boolean relative(int rows)</code>	move cursor <code>rows</code> records from the current position.
<code>boolean isLast()</code>	<code>true</code> if the cursor is on the last row.
<code>void close()</code>	close the cursor, release JDBC resources.
<code>boolean wasNull()</code>	<code>true</code> if the last column read had <code>null</code> value
<code>int findColumn(String columnName)</code>	returns the column number given the name of the column

In addition to these methods, a collection of `get` methods is associated with `ResultSet` class. Note that the cursor for a `ResultSet` begins **BEFORE** the first record. Each `get` method retrieves one value from the current record (tuple) in the cursor. There are two families of `get` methods: one family retrieves values by column number, the other — by column name. Remember, columns indexes start at 1.

get by Column Number	get by Column Name	Explanation
<code>String getString(int colIndex)</code>	<code>String getString(String colName)</code>	retrieve a string value
<code>boolean getBoolean(int colIndex)</code>	<code>boolean getBoolean(String colName)</code>	retrieve a boolean value
<code>byte getByte(int colIndex)</code>	<code>byte getByte(String colName)</code>	retrieve a byte value
<code>short getShort(int colIndex)</code>	<code>short getShort(String colName)</code>	retrieve a short integer value
<code>int getInt(int colIndex)</code>	<code>int getInt(String colName)</code>	retrieve an integer value
<code>float getFloat(int colIndex)</code>	<code>float getFloat(String colName)</code>	retrieve a floating point value
<code>double getDouble(int colIndex)</code>	<code>double getDouble(String colName)</code>	retrieve a double precision value
<code>java.sql.Date getDate(int colIndex)</code>	<code>java.sql.Date getDate(String colName)</code>	retrieve a string value

Example. The code fragment below prints out the values of the `Name` column from the `Employee` table returned from the query.

```

Statement query = conn.createStatement();
ResultSet result =
    query.executeQuery("SELECT * FROM Employee WHERE Salary > 27000");

// original position of the cursor - before first record.
while (result.next()) {
    // 'Name' is a valid column name in the result.
    String s = result.getString("Name");
    System.out.println(s);
}

```

Types of ResultSet instances

ResultSet instances can be of one of three types:

- **TYPE_FORWARD_ONLY**: the result set is non-scrollable, the cursor can be moved using only the `next()` and `last()` methods (no methods that go back can be used) (default).
- **TYPE_SCROLL_INSENSITIVE**: the result set is scrollable, i.e., the cursor can be moved both forward (`next()`, `last()`) and **backwards** (`previous()`, `first()`). Also, the result set typically does not change in response to changes in the database.
- **TYPE_SCROLL_SENSITIVE**: the result set is scrollable, i.e., the cursor can be moved both forward (`next()`, `last()`) and **backwards** (`previous()`, `first()`). Also, the result set typically changes if the data in the underlying database changes.

In addition, the result set may, or may not be updatable. This is controlled by the **concurrency** setting:

- **CONCUR_READ_ONLY**: the result set is read-only, no programmatic updates are allowed (default).
- **CONCUR_UPDATABLE**: the result set can be updated programmatically.

The type of the **ResultSet** instance to be returned by `executeQuery()` statements can be selected at the creation time of the SQL statement object:

- **Class Statement**: default result set type set by the `createStatement()` method is **TYPE_FORWARD_ONLY**, and the concurrency setting is **CONCUR_READ_ONLY**.

To create a statement with a different type of result set use

```
createStatement(int scrollable, int concur)
```

Here, **Scrollable** is the scrollability type (one of **TYPE_FORWARD_ONLY**, **TYPE_SCROLL_INSENSITIVE**, **TYPE_SCROLL_SENSITIVE**) and **Concur** is the concurrency setting (one of **CONCUR_READ_ONLY**, **CONCUR_UPDATABLE**)¹.

¹ For the purposes of this course, we can live with the default concurrency setting, and do not even need to know about it. However, **Connection** class does not have a version of `createStatement()` that sets only the scrollability setting, hence, a brief description of the second argument is needed.