IJISET - International Journal of Innovative Science, Engineering & Technology, Vol. 1 Issue 3, May 2014.
www.ijiset.com
ISSN 2348 - 7968

MIDP based J2ME driver for accessing MySQL from mobile devices

P Ravi Kiran & Y K Sundara Krishna
Department of Computer Science, Krishna University, Machilipatnam, A.P

Abstract: Now-a-days mobile has become a part of our life. We use it for multipurpose like calling, messaging, multimedia as well as client for Information Systems (IS).Mobiles support Java Mobile Edition(J2ME) and .NET micro framework. To access Database we use Java's JDBC driver as a middleware that handles data communication. Mobile clients need a middleware to handle data communication like Java's JDBC Driver but this interface is missing in JavaME. JDBC driver cannot applicable for MIDP applications because of incompatibility and more size. Most existing approaches like syncML or Hotsync concentrate on synchronization aspects or use an additional middleware. In this paper we present a MIDP based JavaME driver for accessing MySQL from mobile devices similar to JDBC embedded in our MIDP applications without a middleware.

1. Introduction

Nowadays ubiquitous, nomadic, and pervasive computing is not a longer vision but reality. Devices become smaller and easier to carry around. Accessing world wide information via wireless links is possible almost anytime and everywhere. Hence, cell phones or smart phones are no longer voice communication devices but small footprinted information system (IS) clients. Even though most current mobile devices provide an environment for third party applications most prominently Java ME's Connected Limited Device Configuration (CLDC) in combination with the Mobile Information Device Profile (MIDP) – there is a lack of standardized software for combining IS and mobile devices. Most existing approaches like SyncML or Hotsync concentrate on synchronization aspects or use an additional middleware. A direct access to data from a database server like MySQL is not possible. However, due to the recent developments in the areas of mobile hardware and wireless networks, this approach becomes more and more reasonable. Java desktop applications often use an JDBC driver for this purpose. In order to fill the gap we implemented an MIDP-based Java ME driver for MySQL similar to JDBC that allows direct communication of MIDP applications to MySQL servers without a middleware. The remainder of this paper is structured as follows. Section 2 includes a brief introduction to the development of mobile applications with Java ME and illustrates the overall architecture of the driver. Section 3 describes how the driver can be used from within an MIDP application. Section 4 explains the limitations of our driver. Section 5 briefly introduces MIDP-Client, a prototype MySQL client that uses the MIDP driver. Section 6 summarizes the paper and shows open issues that will be researched in the future.

2. Overall Architecture

In this section we present the driver architecture. As this requires some background information into Java ME, we would like to first introduce Java ME before giving a detailed architectural overview. We assume that the reader knows a little bit about programming with Java.

2.1 Java SE vs. Java ME delevopment

Java ME development differs significantly from development with Java SE. In Java SE, all clients have identical class libraries and virtual machines given they use the same Java version as they are standardized by Sun. With JavaSE and especially with the Java EE extension Sun tries to fullfill the needs of every application developer. Java ME on the other hand is based on the concept of minimal building blocks. First, so called configurations define a basic runtime environment. This environment only defines "a minimum complement or the 'lowest common de- nominator' of Java technology", is "possibly incomplete for real target devices" and "shall not define any optional fea- tures". This results in execution environments with less than 80 classes for the popular CLDC in version 1.1. In comparison, Java SE has more packages than CLDC has classes. Second, so called profiles, such as MIDP, define additional libraries on top of a configuration for specific device categories or purposes. Profiles adhere to the same general design principles as configurations but can differentiate between obligatory, optional and recommended features. Third, a number of standard APIs from the JSR process define additional features like encryption, content handling, Bluetooth access and more. They are again bundled in JSR plattform definition standards (for example JSRs 751, 1852 and 2483) to simplify development (as adherence to a plattform definition implies the support for all bundled standards). For more information on the differences in the Java versions we refere to Suns Java website http://java.sun.com. As almost no two mobile devices are equal, each manufacture must implement his own version of the Java virtual machine on his devices. By doing so the manufacture can choose which standards and features should be implemented. This creates a problem because each Java ME implementation differs in some aspects. Additionally, each implementation comes with certain restrictions. For example, many virtual machines limit the maximum Jar file size and ignore everything above this limit. (A limit of about 64kB was and is very popular, especially on low-end devices.)

2.2 Preliminary considerations

Before we started the development we set ourselves four design goals:

1. Keep the driver API as near to the JDBC specification as possible

2. Keep the Jar file size below 32kB – half the popular 64kB limit to leave enough space for the application

3. Only implement required features

4. Keep the implementation code as simple and performant as possible These goals were mostly achieved. Our current development version provides database access sufficient for most applications in just 27kB. (In comparison, the MySQL JConnector JDBC driver has more than 500kB.) On the other hand we had to cut short on some aspects like parametrized queries and meta data usage.

2.2.1 Analysis

In order to communicate successful with a database server, a number of problems must be solved:

• Packet sequence control (available packet types and their expected order and content)

• Packet assembly and dissembly (order and type of packet fields and their expected or allowed values)

• Packet field encoding and decoding (types of fields and how are they encoded)

• Protocol flags (what are the protocol options and their meanings, which combinations are allowed or expected)

• Password encryption algorithm (how is the password encrypted during authentication)

• MySQL data types and their structure (how is database content transfered to the client)

• Database meta data (what kind of database meta data is available and must or should be analysed)

Below we give some information on some of the problems we encountered and their solution.

2.2.2 Conversion between Java and MySQL data types

The CLDC runtime environment only implements basic types but not extended SQL data types as provided by the Java SE java.sql package nor the large data types provided by the java.math package. This meant from the beginning that not every MySQL data type could be natively supported or without loss of usability. To standardize our conversation approach, we set up a number of rules:

• All MySQL data types are converted in their nearest Java counterpart if possible

• Signed and unsigned numbers are handled the same way

• MySQL date and time data types are converted to the Java Date type only if it is possible with minimal effort, otherwise process as string

Finally, process all inconvertible data types as string (i.e. MySQL *DECIMAL*, *ENUM* and similar)

As the MySQL client/server protocol uses two different data encoding schemes, namely 'converted to string' and 'binary as stored', we had to implement two decoding routines. This results in a few conversion problems as we could not match both implementations perfectly. Fortunately one can work around this problem by avoiding problematic data types. In Section 4 we give additional information on this delicate topic.

2.2.3 MySQL client/server protocol documentation

Even though MySQL has a very good end user documentation, the kind of in-depth information we needed is not readily available, especially not background information on flags and field. This meant we had to look for documentation at a number of different places, including source code and captured packets. Unfortunately, the information we gathered was sketchy at some places and even wrong at others. For example, one can not use the *CONNECT WITH DB* option of the authentication process as described in the MySQL Internals documentation but must specify the database afterwards. This fact is stated nowhere but was discovered during out tests and solved through packet analysis.

2.2.4 Application developer API

As stated before, our goal is an API very similar to the JDBC one. This meant identifying the necessary functionality and methods as well as making some design decisions. For example, we decided to drop all interfaces and to implement the functionality described directly as classes, using the same method and class names. In the end we implemented three JDBC interfaces, namely *Connection*, *Statement* and *ResultSet*. Additionally we implemented a *Query* class to provide missing functionality for parametrized statements.

2.3 Class overview

Figure 1 shows the basic class diagram of the driver. As one can see, the driver consists of just eight classes (plus four helper classes and exceptions not shown). Below follows some detailed information on the tasks of each class.

2.3.1 Buffer

The *Buffer* class is responsible for encoding and decoding packet fields as well as for the conversion between

MySQL and Java data types. The large number of different field and encoding types in the client/server protocol results in

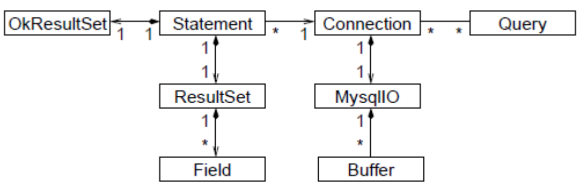

Figure 1. Class Overview

a large number of methods, making this class the largest one. Its inner workings are inspired by the Buffer class of the MySQL JConnector and uses an automatically growing byte array. The aim with this class was to keep memory consumption small and to reduce garbage collection as it is quite processing intensive. Reusing a byte array as much as possible is a good way to do this. Only in case the array is to small to hold all data, a new, bigger one is created and the data is copied over. In addition, by making the array larger then actually needed, the growing only happens infrequently, thus keeping the effort small. Internal pointers keep track of the actually used array space and the 'reserve'; internal boundary checks perform the necessary pointer adjustments.

2.3.2 MysqlIO

The MysqlIO class handles the communication with the database server. It contains methods to assemble and dissemble all packet types, but does not have any clue about the packet sequence. Most other classes can access an instance of this class and use its methods to perform their tasks. The *MysqlIO* class uses two *Buffer* instances, one for sending and one for receiving, to which it has exclusive access, ensuring strict task separation between the classes.

2.3.3 Connection

The *Connection* class is very similar to the *Connection* interface of JDBC. It owns an instance of *MysqlIO* and usesit to provide connection specific methods like opening and closing a connection and changing the database. It also works as factory for *Statement* and *Query* objects.

2.3.4 Statement

The *Statement* class is very similar to the JDBC *Statement* interface. It provides methods to execute database queries and fetch the the result. For this it implements the packet sequence logic necessary, but relies on the instance of the *MysqlIO* class hold by the *Connection* class factory for d oing all packet processing.

2.3.5 Query

The *Query* class provides basic functionality for parametrized queries. It uses a string buffer to replace occurences of question marks in the query string with the desired parameters. Because MIDP does not provide any regular expression capabilities, the algorithm works strictly linear, does not consider already processed or inserted parts of the query string and does not offer any escape capabilities.

2.3.6 Field, ResultSet and OKResultSet

These three classes manage the retrieved database data. The *OKResultSet* is an simple query information storage and is solely used by the *Statement* class. It is not directly available to the application developer but must be accessed through methods provided by the *Statement* class. The *ResultSet* class performs the same job as the JDBC *ResultSet* interface, providing exactly the same row-pointer based access methods. It uses an array of *Field* class instances to store and process all column specific data. In fact, the *ResultSet* does only act as a facade to the *Field* class, managing the row dimension of the database result set. The *Field* class provides column wise storage for database result sets and meta data. A large number of simple methods provide access to specific column- and meta data. It is only used internally by the *ResultSet*.

2.3.7 Helper classes

Three additional helper classes provide a number of static methods for common tasks not really part of the driver (like string operations). An additional *Constants* class contains all the necessary constants. Finally, there is one *SQLException* class used throughout the driver.

3 Using MIDP

3.1 Device requirements

EveryMIDP 2.0 compatible device should be able to use the driver when the following, additional requirements are fulfilled.

- The mobile device must support socket connections (optional in MIDP 2.1)
- The mobile device must support JSR 177 (needed for MySQL authentication via SHA-1)
- The device should have at least one megabyte of free heap memory (depending on the implementing application)

3.2 Simple usage scenario

Since the driver API is a very similar to JDBC, a developer familiar with JDBC will not have any problems using our driver. And even developers new to database APIs will find our driver easy to use as it always follows four steps.

1. Create a database connection
2. Create and execute a database statement
3. Process the result set
4. Close the connection

Steps two and three can be repeated in case more than one query must be executed.

3.3 Main differences to JDBC

There are a few important usage differences to JDBC we would like to point out. First, the driver does not use the JDBC style connection URL but method parameters for simplicity and performance reasons. Second, the *Statement* object can be reused thus improving garbage collection. Third, it is not possible to execute multiple statements at the same time as they use the same *MysqlIO* class instance and thus share buffers.

4 Limitations

Because of the previously stated restrictions of the CLDC/ MIDP runtime, we had to compromise and drop a number of features available in the normalMySQL JDBC driver. Following is an overview of the limitations and their cause.

4.1 No support for pre 4.1 MySQL server

Themost significant restrictions of our driver is themissing support for pre 4.1 MySQL server. This decision was caused by two facts. First, version 4.1 of theMySQL server introduced an enhanced version of the client/server protocol. This new protocol differs significantly from the previous versions, providing a new authentication mechanism using SHA-1 and extended meta data support. These changes had a price, namely the loss of its backward compatibility. Hence two protocol implementationswould be necessary to support pre 4.1 server as well as current ones. This would have required hundreds of additional lines of code at the prize of an increased `jar` file size. Second, the extended support time frame for all pre 4.1 MySQL server will end at the end of 2008. This made the likelihood of encountering an unsupported server version unlikely at best.

4.2 Limited character set support

The CLDC specification states that the "Character information is based on the Unicode Standard, version 3.0. However, since the full character tables required for Unicode support can be excessively large for devices with tight memory budgets, by default the character property and case conversion facilities in CLDC assume the presence of ISO Latin-1 range of characters only. More specifically, implementations must provide support for character properties and case conversions for characters in the '*Basic Latin*' and *Latin-1 Supplement*' blocks of Unicode 3.0. Other Unicode character blocks may be supported as necessary. For this reason, character conversation capabilities are very limited in CLDC. In order to keep our driver small and fast, we decided to limit all string conversions to the ISO Latin-1 range and to use the standard conversations provided by the String class instead. This means a number of restrictions for both the MySQL server and the mobile device. First, the server must use Latin-1 for all communication with the client and second the client device must use the Latin-1 character set as default. Also, the target database should use the Latin-1 character set but this is not necessary if the returned strings are post processed by the application.

4.3 No support for prepared statements

In the current implementation the driver does not support prepared statements. This has two reasons. First, the implementation code is quite complex, requiring the implementation of at least five additional packet types. Second,
the protocol documentation in the MySQL Internals wiki is still 'tentative' and incomplete.

4.4 Limited data type support

The client/server protocol uses two different encoding schemas: All standard result sets are returned as a character string – non-character types are converted into a string while results from prepared statements and default column values are returned as binary [9]. To keep the implementation size small, we did not match both conversation routines perfectly. This is especially true for the date and time data types. In the current implementation state this problem is not very serious as we use the binary conversion only at one place – decoding the default column value. But in future implementations with support for prepared statements this might be a problem. Additionally, we recommend to use only typeswhich can be converted directly by Java (i.e. signed integers, floatingpoint numbers, character strings and timestamps) as these types can be used without complication.

4.5 No transaction handling

Due to `jar` size considerations, we droped any built in support for transaction handling as we deemed it not necessary for mobile devices. The fact that MySQL just uses standard SQL commands for transaction handling (SQL *START TRANSACTION*, *COMMIT* and *ROLLBACK*) and no custom command codes (as it does for the *USE* command), made this decision easy as an application developer can add his own transactional logic if needed.

ISSN 2348 - 7968

4.6 Limited SQL support

The client/server protocol uses special packets for some SQL commands, for example for the SQL *USE* but also for the MySQL specific administration commands. Currently, with the exception of *USE*, the driver does not implement these special packets, they are hence not supported.

5 MIDP-Client

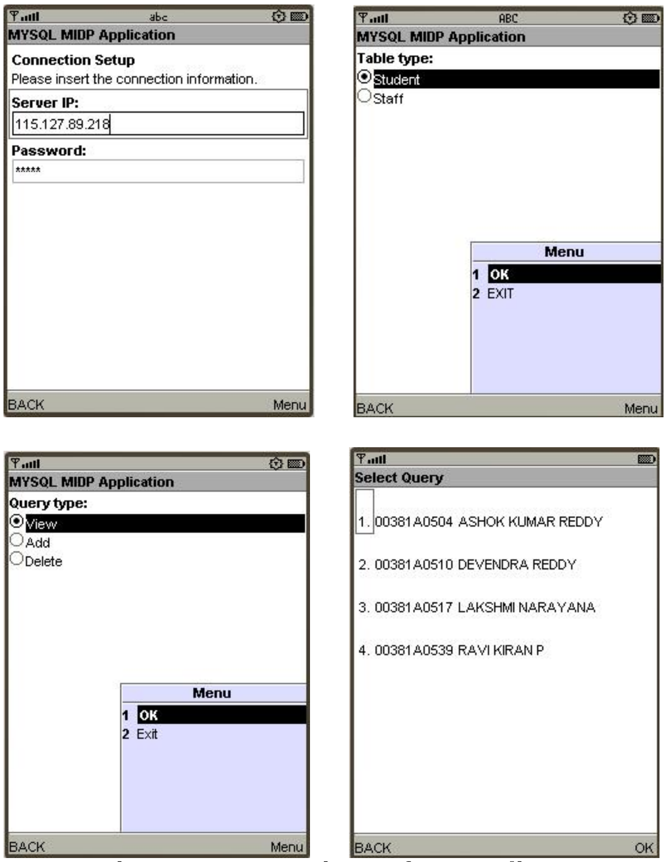

Figure 2. Screenshots of MIDP-client

Figure 2 shows four screenshots of the MIDP-client, a MIDP based prototype MySQL client application that has been implemented for demonstration purposes. After inserting the connection information (server address, password) the client offers alternative options to the user. However, this step helps reducing the query typing by the user via the cell phones' keyboard. It results in the third screen that allows to complete the query. The next step is to submit the query. Due to security reasons each java application authomatically requests confirmation before establishing any internet connection. The screen of this request, that is not shown in Figure 2, is followed by a screen that displays the query results. This query result is the result of the direct communication between our client and a MySQL server via the MIDP driver. Due to the limited display size we decided to show only the two attributes of each tuple and to number the tuples. However, this function managed by the application only and *no* feature of the MIDP driver.

6 Summary

In this paper we presented our implementation of a MIDP-based Java ME driver for MySQL. Similar to JDBC for Java SE it allows direct communication of MIDP applications to MySQL servers without a middle-ware. Hence, it is possible to directly connect a mobile device like a cell phone to aMySQL server without using a middle-ware. We discussed implementation details as well as how to use the driver. Furthermore we described existing limitations which mostly result from the restrictions of Java ME and MIDP. The driver is only the first step in the direction of a full MySQL support for mobile devices. Wireless transmissions are expensive and relatively slow.

References

1. Hagen H¨opfner and J¨org Schad and Sebastian Wendland and Essam Mansour *MIDP: An JDBC driver for accessing MySQL from mobile devices*

2. A. Caracas¸, I. Ion, M. Ion, and H. H¨opfner. *Towards Javabased Data Caching for Mobile Information System Clients.*

3. H. H¨opfner, J. Schad, S. Wendland, and E. Mansour. *MIDP and MIDP-Client: Direct Access to MySQL Databases from Cell Phones*, Mar. 2009. BTW-DEMO.

4. IBM Corporation. *IBM DB2 Everyplace Application and Development Guide Version 8.2*, Aug. 2004.

5. IBM Corporation. *IBM DB2 Everyplace Sync Server Administration Guide Version 8.2*, Aug. 2004.

6. I. Ion, A. Caracas¸, and H. H¨opfner. *MTrainSchedule: Combining Web Services and Data Caching on Mobile Devices.* Datenbank-Spektrum, 21:51–53, May 2007.

7. MySQL AB. *MySQL Internals ClientServer Protocol*, 2008. Retrieved on January 18, 2008 from `http://forge.mysql.com/wiki/MySQL Internals ClientServer Protocol`.

8. MySQL AB. *MySQL Lifecycle Policy*, 2008. Retrieved on February 21, 2008 from `http://www.mysql.com/about/legal/mysql lifecycle policy.pdf`.

9. Oracle Corporation. *Oracle Database Lite, Developer'sGuide 10g (10.0.0)*, June 2004.

10. Oracle Corporation. *Oracle Database Lite, SQL Reference 10g (10.0.0)*, June 2004.

11. Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California, 95054 U.S.A. *Connected Limited Device Configuration(CLDC) Specification*, 1.1 edition, Mar. 2003.

12. Sun Microsystems, Inc. and Motorola, Inc. *Mobile Information Device Profile Specification*, 2.1 edition, May 2006.

13. Sybase Inc. `http://www.sybase.com/ianywhere/ aproducts`, 2008.