

The BetaDBC Library – Reference Manual and Tutorial

**Mjølner Informatics Report
MIA 99–43
March 2004**

Copyright © 1999–2004 Mjølner Informatics.

All rights reserved.

No part of this document may be copied or distributed
without the prior written permission of Mjølner Informatics

Table of Contents

1 Introduction.....	1
2 How to use the BetaDBC interface.....	2
2.1 Introduction.....	2
2.2 Setting up your system.....	2
2.2.1 On Windows.....	2
2.2.2 On Unix.....	2
2.3 Initializing a database.....	3
2.3.1 Frontbase.....	3
2.3.2 MySQL.....	3
2.4 Writing the Beta code.....	4
2.4.1 Including an interface.....	4
2.4.2 Opening a connection.....	4
2.4.3 SQL Statements.....	4
2.4.4 Declaring Shared Variables.....	5
2.4.5 Using shared variables.....	5
2.4.6 Results.....	6
2.4.7 Specifics for the native Oracle interface.....	6
2.4.8 Specifics for the native Frontbase interface.....	7
2.4.9 Specifics for the native MySQL interface.....	7
3 BetaDBC Interface Tutorial.....	9
3.1 Introduction.....	9
3.2 Creating a Database.....	9
3.3 Querying and Retrieving from the Database.....	9
3.4 Executesqlfile.....	11
3.5 Embedded SQL – Using Shared Variables.....	11
3.6 Embedded SQL – Fetching Results.....	12
3.7 An Ad-hoc Query Evaluator.....	13
4 Text Files for the Tutorial.....	15
4.1 createmoviedbtables.txt.....	15
4.2 deletemoviedbtables.txt.....	15
4.3 populatemoviedbtables.txt.....	15
5 Transactions.....	17
6 Scrolling Cursors.....	18
7 Datatypes.....	19
7.1 Types in Oracle and Beta.....	19
7.2 Types in FrontBase and Beta.....	19
7.3 Types in MySQL and Beta.....	20
8 References.....	22
9.1 Betadbc Interface.....	23
9.2 Scrollingresultset Interface.....	33
9.3 Tables Interface.....	35

Table of Contents

9.4 Transactions Interface.....	36
Index.....	37
A.....	37
B.....	37
C.....	37
D.....	37
F.....	37
H.....	38
I.....	38
K.....	38
L.....	38
N.....	38
O.....	38
P.....	38
R.....	38
S.....	38
T.....	38

1 Introduction

This document describes the BetaDBC (Beta DataBase Connectivity) library for communicating with relational databases using SQL (Structured Query Language, [Date 93]). The library implements the pattern Connection used to model connections to relational databases. This pattern contains patterns for querying and manipulating relational databases. Additional support for transactions may be added by including the fragment Transactions. Scrolling cursors are implemented in the fragment Scrollinresultset.

Author: Klaus Marius Hansen

Heavily modified and updated May 2001 by Karsten Strandgaard Jorgensen (karsten@mjolner.dk)
If you come across any errors or inaccuracies in this manual please let us know: [Error Reporting](#)

2 How to use the BetaDBC interface

2.1 Introduction

This section is a step-by-step guide to get you started using BetaDBC.

The first time you connect to a database from a Beta program, it is a 3 step process:

- Set up your system
- Initialize database and users
- Write some Beta code

Step 1 is necessary once per platform, step 2 is necessary once per database/user and step 3 is always necessary :)

2.2 Setting up your system

This may be the most troublesome of the 3 steps, but on the bright side you may not have to do it at all. Perhaps your systems administrator has already done it. If you are stuck with the job then proceed as follows:

2.2.1 On Windows

First: Install your database manager according to the installation guidelines that came with it.

Then:

- If your Beta system uses Microsoft tools (as opposed to GNU tools) and your database manager is Oracle or Frontbase, you are in luck. No further setup is required.
- If your Beta system uses Microsoft tools and your database manager is MySQL you must do the following:
Copy the file `mysql\lib\opt\libmySQL.lib`
into the directory `beta\lib\nti_ms\`
- Otherwise you must connect to your database manager via ODBC. A driver for that purpose is either included in the standard database manager installation (example: Microsoft Access), or must be installed separately (example: MySQL).
To set up the driver, click Start > Settings > Control Panel > ODBC Data Sources (32-bit)
Click the tab 'User DNS'
Look at the list 'User Data Sources'
If you do not see a driver for your database manager, you need to download and install one.
Click the driver of your database manager and click 'Configure'
You can now adjust various parameters of the data source depending on the driver. In particular, you can probably set the name of the data source as well as a user name and password. These settings are used when you open a connection in a Beta program.

2.2.2 On Unix

First: Install your database manager according to the installation guidelines that came with it.

Then:

- If your database manager is supported by a native BetaDBC library – currently Oracle, Frontbase and MySQL – you are in luck. No further setup is required.
- Otherwise you must connect to your database manager via ODBC. You will need a .odbc.ini file in your home directory. Create such a file if it does not exist. Suppose a driver for the PostgreSQL database manager is located in /usr/local/lib/libcliPG.so. If you want to connect to the database 'mydb' on the database server 'delirium' insert the following in your .odbc.ini file ('ReadOnly = 0' tells the driver that you want to manipulate the 'mydb' database)

```
[mydatasource]
# data source containing the 'mydb' db on delirium
Driver = /usr/local/lib/libcliPG.so
Database = mydb
Servername = delirium
ReadOnly = 0
```

This defines a data source named 'mydatasource', that you may use BetaDBC to connect to. In the example 'ReadOnly' is a driver specific attribute of the data source 'mydatasource'. See the documentation for the drivers used for definitions of driver specific attributes.

2.3 Initializing a database

2.3.1 Frontbase

sql92 is a small tool included in the FrontBase installation. It is used for creating databases and database users. By default it is located in FrontBase/bin/sql92

A small example:

start sql92 and type the following:

```
create database firstdb;
connect to firstdb user _system;
create user firstuser;
disconnect current;
connect to firstdb user firstuser;
disconnect current;
```

2.3.2 MySQL

When the MySQL database manager is up and running it has by default only one user called 'root' whose password is empty. New users are created using eg. GRANT statements as follows:

Start MySQL as user 'root' on the database 'mysql':

```
./bin/mysql -u root mysql
```

Create a user called 'firstusr' whose password is 'firstpswd':

```
mysql> GRANT ALL PRIVILEGES ON *.* TO firstusr@localhost IDENTIFIED BY 'firstpswd' WITH
mysql> GRANT ALL PRIVILEGES ON *.* TO firstusr@ "%" IDENTIFIED BY 'firstpswd' WITH GRANT
```

Additional users can be created in a similar fashion. Now create a database for the user 'firstusr':

```
./bin/mysqladmin -u firstusr -p create firstdb
```

This creates a database 'firstdb' after prompting for the password. The user 'firstusr' can now log on to the database 'firstdb':

```
./bin/mysql -u firstusr -p firstdb
```

Tables can now be created using standard SQL statements:

```
mysql> CREATE TABLE test (name CHAR(200) NOT NULL, KEY index_name (name(10)));
```

2.4 Writing the Beta code

2.4.1 Including an interface

You must always INCLUDE the relevant interface

If you use ODBC: INCLUDE ' ~beta/betadbc/betadbc_odbc '
If you use native Oracle: INCLUDE ' ~beta/betadbc/betadbc_oracle '
If you use native Frontbase: INCLUDE ' ~beta/betadbc/betadbc_frontbase '
If you use native MySql: INCLUDE ' ~beta/betadbc/betadbc_mysql '

2.4.2 Opening a connection

You must first create an instance of the connection pattern.

```
sqlCon: @Connection;  
stmt: @sqlCon.directSqlStatement;
```

Calling the open method on a connection...

```
('access','marius',none)->sqlCon.open
```

...then makes it possible to communicate with the data source. The statement opens a connection to the data source name "access" for the user "marius" without specifying a password.

2.4.3 SQL Statements

Data manipulation and definition is done using the SQL statement pattern directSQLStatement. To use a directSQLStatement `stmt`, one must first open it and then associate it with an SQL statement as in

```
stmt.open;  
'SELECT title, length FROM Movie WHERE studioName = \'Disney\''->stmt;
```

or, equivalently

```
'SELECT title, length FROM Movie WHERE studioName = \'Disney\''->stmt.open;
```

Invoking `stmt.execute` will then cause the SQL statement to be executed at the database. A statement should be closed after use: `stmt.close`

2.4.4 Declaring Shared Variables

Definition of shared variables are done via the "declare..." methods. E.g. to declare a text `studioName` as a shared variable named "`studioName`" use

```
studioName: @text;
...
do 'studioName'->declareText
  (# set::(# do value->studioName #);
   get::(# do studioName[]->value[] #)
  #)
```

2.4.5 Using shared variables

The contents of a `directSQLStatement` can be any SQL statement with embedded shared variables and/or markers, i.e. a `directSQLStatement` `stmt` may be initialized as in

```
'SELECT title,length FROM Movie WHERE studioName = :studioName AND year = %i'->stmt;
```

Here "`studioName`" is the name of a shared variable declared as shown above. The value of the `%i` marker may be set using the `i` pattern in `directSQLStatement`. Now suppose that the following statements have been executed

```
'Disney'->studioName;
1990->stmt.i;
```

When executing the SQL statement the SQL contents of the statement will then conceptually be

```
SELECT title,length FROM Movie WHERE studioName = 'Disney' AND year = 1990
```

i.e., before sending an SQL statement to a database the embedded shared variables and the markers are, conceptually, substituted for their current values. After execution the contents of the statement can be changed, the markers can be reset or the statement can be closed.

Also, embedded shared variables and markers can be named as in

```
'SELECT title,length FROM Movie WHERE studioName = :studioName AND year = theYear%i'->stmt;
```

In this case, the `%i` marker is named "`theYear`" and the statement may be used as

```
'Disney'->studioName; ('theYear',1990)->stmt.setIntegerByName
```

Please note that the use of shared variables is only supported in the ODBC interface. Suggested workaround for the other interfaces: build your SQL statements dynamically using `append` or `putformat`.

2.4.6 Results

Executing an SQL statement stmt will yield an instance of resultSet (in the following rs is a reference to a resultSet):

```
rs: ^connection.resultSet
...
do stmt.execute->rs[]
```

Given a resultSet, the scan method iterates over the tuples in the resultSet. There are three distinct ways to control the scan. First, one may simply execute

```
rs.scan(# do ... #)
```

In the do-part of the scan one may then refer to the values of the columns in the result. This is done sequentially by referring to the markers of the current result. If, e.g., rs was retrieved as shown above,

```
rs.scan(# do current.s -> putline; current.i->putint; newline #)
```

will scan over the results in the resultSet and print the values of their columns on the screen.

Second, one may enter a string when evaluating a scan pattern as in:

```
':title %i'->rs.scan(# do title->putline; current.i->putint; newline #)
```

Here title is a shared variable named "title". This statement prints the same as above but by providing an input string it is here specified that the first column of each result should be assigned to the shared variable "title" and that the second column of each result is an integer that will be fetched via the i marker. In general one may in this way specify how each column of a result should be treated.

The two ways of scanning shown above may, in some circumstances, be problematic in that they assume a specific ordering of columns in the results. Therefore, the third way of doing a scan names the columns in the resultSet, as in e.g.:

```
'length%i title:title'->rs.scan
  (# do title->putline; current.i->putint; newline #)
```

In this way the order of the columns in the result may be changed from "title, length" to "length, title" without any problems for the last way of scanning.

2.4.7 Specifics for the native Oracle interface

Please note that the use of shared variables is only supported in the ODBC interface. Suggested workaround for the other interfaces: build your SQL statements dynamically using append or putformat.

For obscure reasons the datatype boolean is not supported by Oracle. Furthermore, the use of time/date datatypes is discouraged, and therefore sparingly supported by the Oracle interface. Oracle recommends storing dates as text.

A small example using Oracle:

The user `scott` with password `tiger` wants to connect to the default server (hence the empty text in the opening of the connection):

```
ORIGIN '~beta/betadbc/betadbc-oracle';
-- program: Descriptor --
(# 
  oraCon: @Connection;
  stmt: @oraCon.DirectSQLStatement;

  do
    ('','scott','tiger')->oraCon.open;

    'create table Students("Name" varchar(100), "Age" integer);'->stmt.open;
    stmt.execute;
    stmt.close;
    'commit;'->stmt.open;
    stmt.execute;
    stmt.close;
    oraCon.close;
#)
```

2.4.8 Specifics for the native Frontbase interface

Please note that the use of shared variables is only supported in the ODBC interface. Suggested workaround for the other interfaces: build your SQL statements dynamically using `append` or `putformat`.

Selecting a host:

The Beta program must select the host on which the FrontBase server is running before opening a connection. Selecting the host is done by using `selecthost`

A small example using Frontbase:

```
ORIGIN '~beta/betadbc/betadbc-frontbase';
-- program: Descriptor --
(# 
  frontbaseCon: @Connection;
  stmt: @frontbaseCon.DirectSQLStatement;

  do
    'brage.mjolner.dk' -> frontbaseCon.selecthost;
    ('firstdb','firstuser','')->frontbaseCon.open;

    'create table Students("Name" varchar(100), "Age" integer);'->stmt.open;
    stmt.execute;
    stmt.close;
    frontbaseCon.close;
#)
```

2.4.9 Specifics for the native MySQL interface

Please note that the use of shared variables is only supported in the ODBC interface. Suggested workaround for the other interfaces: build your SQL statements dynamically using `append` or `putformat`.

Selecting a database:

To utilize the MySQL interface in Beta you must select a database by using `selectdb` after

opening a connection, as shown in this excerpt from the adhoc.bet demo:

```
ORIGIN '~beta/betadbc/betadbc-mysql';
INCLUDE  '~beta/basiclib/formatio'
          '~beta/basiclib/numberio'  '~beta/basiclib/file';
-- program: Descriptor --
(#  
  sqlCon: @Connection; stmt: @sqlCon.DirectSQLStatement;
  res: ^sqlCon.resultSet
do
  'hi there'->putline;
  (2->arguments,3->arguments,4->arguments)->sqlCon.open;
  'firstdb'->sqlCon.selectdb;
...  
)
```

adhoc is then invoked as follows:

```
./adhoc vertigo firstusr firstpswd
```

Where vertigo is the host on which the MySQL server is running. It is possible to use 'localhost'.

3 BetaDBC Interface Tutorial

3.1 Introduction

The sample programs shown in this tutorial may be found in the tutorial directory accompanying BetaDBC. The examples use a database schema and examples from [Ullman 97]. It supposes that the reader is familiar with basic SQL and focuses on teaching the essentials of using BetaDBC. All examples take (up to) three command line arguments: a data source, a user name, and a password. If the specific interface used permits, the data source, user name and/or password may be omitted.

Although this tutorial is based on the ODBC interface, it should be quite useful for all BetaDBC users.

In the following it is assumed that a suitable data source named "mysource" has been created.

3.2 Creating a Database

Check this section with the new setup

The next step will then be to create and insert values into a database. Let's use the following sample database schema

```
Movie (title, year, length, inColor, studioName, producerCNo)
StarsIn (movieTitle, movieYear, starName)
MovieStar (name, address, gender, birthdate)
MovieExec (name, address, certNo, netWorth)
Studio (name, address, presCNo)
```

A series of SQL statements creating the database schema may be found in `createmoviebttables.txt`. The text files used in this section are shown in the section "Text Files".

Create the tables corresponding to this schema by running the `executesqlfile` program also found in the tutorials directory:

```
.../tutorial> ./executesqlfile createmovie.txt mysource marius
```

How the `executesqlfile` program is implemented will be discussed later. You may now insert some values in the database by running

```
.../tutorial> ./executesqlfile moviedbttables.txt mysource marius
```

The tables may later be deleted by running

```
.../tutorial> ./executesqlfile dropmovies.txt mysource marius
```

3.3 Querying and Retrieving from the Database

This section will introduce the basics of BetaDBC: connecting to data sources, executing simple queries and retrieving the results.

Consider the simple SQL statement

```
SELECT *
FROM Movie
WHERE studioName = 'Disney' AND year = 1990;
```

An application "simple" that uses the BetaDBC–ODBC interface, executes the above query and retrieves the result may look like:

```
ORIGIN '~beta/betadbc/betadbc-odbc';
-- program: Descriptor --
(# 
    sqlCon: @connection;
    stmt: @sqlCon.directSqlStatement;
    rs: ^sqlCon.resultSet
do
    (2->arguments,3->arguments,4->arguments)->sqlCon.open;
    'SELECT title, length FROM Movie WHERE studioName = \'Disney\' AND year = 1990'
        ->stmt.open;
    stmt.execute->rs[];
    rs.scan
        (# do
            'title: '->puttext;
            current.s->putline;
            'length: '->puttext;
            current.i->putint;
            newline;
        #);
    stmt.close;
    sqlCon.close;
#)
```

The program starts out by declaring a connection, a directSQLStatement belonging to that connection and a resultSet belonging to that connection. The connection is used in order to connect to a data source in the first line of the program:

```
(2->arguments,3->arguments,4->arguments)->sqlCon.open;
```

Connection's open method takes as arguments a name of the connection, a username and a password. Thus an invocation of the program like

```
.../tutorial> ./simple mysource marius foobar
```

means that the first statement will be an attempt to connect the user "marius" with password "foobar" to the data source named "mysource". If this succeeds

```
'SELECT title, length FROM Movie WHERE studioName = ''Disney'' AND year = 1990'
->stmt.open;
```

will open the directSQLStatement "stmt" and set its content to the query we want to execute. Executing the query yields a resultSet holding a cursor for the result

```
stmt.execute->rs[];
```

The resultSet may then be scanned. During the scan 'current' will hold a reference to a tuple in the

resultSet. The values of this result may then be accessed consecutively by using the marker attributes

```
rs.scan
(#
do
    'title: '->puttext;
    current.s->putline;
    'length: '->puttext;
    current.i->putint;
    newline
#);
```

Finally, in order to free resources, the directSQLStatement and the connection are closed.

3.4 Executesqlfile

The simple scheme presented in the last section can now be used for implementing the executesqlfile program. The executeLoop shows how to reuse an SQLStatement by simply replacing it's textual contents

```
executeTxt[]->sqlCon.stmt
```

3.5 Embedded SQL – Using Shared Variables

Using shared variables makes it possible to use the values of BETA objects in place of a concrete value in SQL statements. Since no preprocessor is used by BetaDBC, it is necessary to declare shared variables imperatively, as in

```
sqlCon:@connection
    studioName:@text;
do ...;
    'studioName'
        ->sqlCon.declareText
            (# set:: (# do value->studioName #);
            get:: (# do studioName[]->value[] #)
            #);
    ...
    ...;
```

Here a shared text variable named "studioName" is declared. The set pattern is final bound to describe how the shared variable's value is to be set. get is final bound to describe how the value of the shared variable is to be fetched.

Then, using embedded SQL syntax, one may use shared variables in SQL statements:

```
stmt:@sqlCon.directSQLStatement;
do ...;
    'INSERT INTO Studio(name, address) VALUES (:studioName, :studioAddr)'
        ->stmt.open;
    ...
    ...;
```

This means that when executing stmt, ":studioName" and ":studioAddr" will (conceptually) be replaced by the values of the BETA text variables "studioName" and "studioAddr", and the resulting SQL statement will then be executed.

Using BetaDBC it is possible to declare most commonly used objects as shared variables (i.e., boolean, integer, real, text, date and time). The figure below shows a full program that will execute the statement above. "stmt.getExpanded" returns in a text how the SQL statement would look if it was executed at that point.

```
ORIGIN '~beta/betadbc/betadbc-odbc';
-- program: Descriptor --
(# sqlCon: @connection;
  stmt: @sqlCon.directSqlStatement;
  studioName,studioAddr: @text
do (2->arguments,3->arguments,4->arguments)->sqlCon.open;
  'studioName'
  ->sqlCon.declareText
    (# set:: (# do value->studioName #);
      get:: (# do studioName[]->value[] #)
    #);
  'studioaddr'
  ->sqlCon.declareText
    (# set:: (# do value->studioAddr #);
      get:: (# do studioAddr[]->value[] #)
    #);
  'Input a studio name: '->puttext;
getLine->studioName.puttext;
' and address: '->puttext;
getLine->studioAddr.puttext;
'INSERT INTO Studio(name, address) VALUES (:studioName, :studioAddr)'
  ->stmt.open;
stmt.getExpanded->putline;
stmt.close;
sqlCon.close
#)
```

3.6 Embedded SQL – Fetching Results

Suppose that we are executing a statement that return a result. Embedded SQL can then also be used to fetch results directly into shared variables. In BetaDBC this is done through the use of the scan pattern. Suppose we are executing

```
'SELECT MovieExec.name, netWorth FROM Studio, MovieExec
 WHERE presCNo = certNo AND Studio.name = :studioName'
```

Then,

```
' :presName :presNetWorth' -> (stmt.execute).scan(# ... #)
```

will cause the first column in each result tuple to be assigned to the shared integer variable "presName", and the second column to "presNetWorth". The full code is shown below:

```
ORIGIN '~beta/betadbc/betadbc-odbc';
INCLUDE '~beta/basiclib/formatio';
-- program: Descriptor --
(# sqlCon: @connection;
  stmt: @sqlCon.directSqlStatement;
  studioName,presName: @text;
  presNetWorth: @integer
do 'studioName'
  ->sqlCon.declareText
    (# set:: (# do value->studioName #);
```

```

        get::  (# do studioName[]->value[] #)
    #);
'presName'
->sqlCon.declareText
(# set:: (# do value->presName #);
get:: (# do presName[]->value[] #)
#);
'presNetWorth'
->sqlCon.declareInteger
(# set:: (# do value->presNetWorth #);
get:: (# do presNetWorth->value #)
#);
(2->arguments,3->arguments,4->arguments)->sqlCon.open;
'Input a studio name: '->puttext;
getLine->studioName.puttext;
'SELECT MovieExec.name, netWorth FROM Studio, MovieExec WHERE presCNo = certNo AND St
->stmt.open;
stmt.getExpanded->putline;
':presName :presNetWorth'
->(stmt.execute).scan
(
do 'The net worth of the president %s \nof %s is %i $\n'
->putFormat
    (# do presName[]->s; studioName[]->s; presNetWorth->i #)
#);
stmt.close;
sqlCon.close
#)

```

3.7 An Ad-hoc Query Evaluator

We now have most of the building blocks to create an ad-hoc query evaluator, i.e., a program that connects to a data source and in a loop prompts for SQL statements that are to be executed on this data source. The following implements such a program.

As long as the user inputs anything but an empty line this input is sent to the data source as an SQL statement:

```

getline->stmt;
(if (stmt).empty then leave L if);
stmt.execute->res[];

```

If successful, the result is examined. First the column information of the resultSet is extracted:

```

(for j: res.columnCount repeat
'%s: %s\t'->putFormat
(
do (j->res.getColumn).name[]->s;
(j->res.getColumn).dataTypeName[]->s
#)
for);

```

Each resultSet has a columnCount yielding the number of columns in the resultSet. For each column, information such as name and dataTypeName (a DBMS specific datatype name) may be retrieved.

If the columnCount is non-zero, the results are fetched:

```
(if res.columnCount > 0 then
    res.scan
    (#
        do (for i: res.columnCount repeat
            (if (i->res.getColumn).DataType##)
                // text## then
                current.s->puttext
                // integerObject## then
                current.i->putint
                // realObject## then
                current.f->putreal
                // booleanObject## then
                (if current.b then
                    'true'->puttext;
                    else
                    'false'->puttext
                    if)
                // time## then
                current.t->puttime
                else
                    'Unknown data type!!!'->puttext
                    if);
                    '\t'->puttext
                for);
                newline
            #)
        else
            'DML/DDL statement executed successfully!'->putLine
    if)
```

Again the information about columns in the resultSet is used. By evaluating

```
(i->res.getColumn).DataType##
```

the BETA pattern corresponding to the SQL datatype in column i is found.

4 Text Files for the Tutorial

4.1 createmoviedtables.txt

```
CREATE TABLE MovieStar (
    name CHAR(30),
    address VARCHAR(255),
    gender CHAR(1),
    birthdate INTEGER
)

CREATE TABLE Movie (
    title VARCHAR(255),
    year INTEGER,
    length INTEGER,
    inColor CHAR(1),
    studioName CHAR(50),
    producerCNo INTEGER
)

CREATE TABLE StarsIn (
    movieTitle VARCHAR(255),
    movieYear INTEGER,
    starName CHAR(30)
)

CREATE TABLE MovieExec (
    name CHAR(30),
    address VARCHAR(255),
    certNo INTEGER,
    netWorth INTEGER
)

CREATE TABLE Studio (
    name CHAR(50),
    address VARCHAR(255),
    presCNo INTEGER
)
```

4.2 deletemoviedtables.txt

```
DROP TABLE MovieStar
DROP TABLE Movie
DROP TABLE StarsIn
DROP TABLE MovieExec
DROP TABLE Studio
```

4.3 populatemoviedtables.txt

```
INSERT INTO Movie VALUES ('Star Wars', 1990, 120, 'y', 'Disney', 1)
INSERT INTO Movie VALUES ('Pretty Woman', 1990, 119, 'y', 'Disney', 999)
INSERT INTO Studio VALUES ('Disney', 'Drive Rd.', 1)
```

The BetaDBC Library – Reference Manual and Tutorial

```
INSERT INTO Studio VALUES ('Paramount', 'Drive Way', 2)  
INSERT INTO MovieExec VALUES ('Marius', 'Home', 1, 3000000)  
INSERT INTO MovieExec VALUES ('Lucas', 'Home', 2, 2000000)
```

5 Transactions

Many data sources support the use of transactions. In order to use this capability from BetaDBC, the fragment transactions must be included. An outline of the interface of this fragment is shown below. The full interface may be seen in the interface section.

```
transactionsSupported:  
    (# isSupported: @boolean  
     ...  
     exit isSupported  
     #);  
autoCommitMode:  
    (# autoCommit: @boolean  
     enter (# enter autoCommit ... #)  
     exit (# ... exit autoCommit #)  
     #);  
readUncommitted: integerValue (# ... #);  
readCommitted: integerValue (# ... #);  
repeatableRead: integerValue (# ... #);  
serializable: integerValue (# ... #);  
transactionLevelSupported: booleanValue  
    (# level: @integer  
     enter level  
     ...  
     #);  
transactionLevel:  
    (# level: @integer  
     enter (# enter level ... #)  
     exit (# ... exit level #)  
     #);  
commit: (# ... #);  
rollBack: (# ... #)
```

The scope of a transaction is a whole connection including all statements allocated in it. The default is that every execution of an SQL statement starts a new transaction and automatically commits the effects of this statement after the statement has completed. This auto commit mode may be changed to manual commit mode by evaluating

```
false->sqlCon.autoCommitMode
```

where sqlCon is an instance of a connection. Note that this is only meaningful if transactionsSupported evaluates to true. In manual commit mode a series of database manipulations may be committed by executing commit. Equivalently a series of database manipulations may be aborted by executing rollBack.

Four transaction isolation levels (as defined in the SQL standard) are available, namely readUncommitted, readCommitted, repeatableRead, and serializable. Whether a transaction isolation level, such as serializable, is supported in a given connection may be checked by evaluating e.g.

```
serializable->sqlCon.transactionLevelSupported
```

6 Scrolling Cursors

Commonly, a result set is scanned from the beginning to the end. Sometimes, however, other orders of fetching of results are desirable. Scrolling result sets support this.

To use scrolling cursors and result sets, include the fragment `scrollingResultSet`. Applications can investigate whether a connection supports scrolling cursors by calling `scrollingCursorsSupported` on a connection.

Suppose the connection `con` supports scrolling cursors. Then the declaration of a statement `stmt` as

```
stmt: @con.directSQLStatement
  (#  
    cursorType:::  
      (# scrollable::: (# do true->value #) #);  
    resultSetType::: con.scrollingResultSet  
  #)
```

will create a scrolling result set whenever `stmt` is executed:

```
rs: ^connection.scrollingResultSet;  
...  
do ...; stmt.execute->rs[]; ...
```

The result set, `rs`, may now be used to scroll through the result set using the methods shown below.

```
scrollingResultSet: resultSet
  (#  
    fetchFirst: ...;  
    fetchLast: ...;  
    fetchNext: ...;  
    fetchPrior: ...;  
    fetchRelative: ...;  
    fetchAbsolute: ...;  
    scanReverse: ...;  
    ...  
  #)
```

`fetchLast`, for example, will fetch the last tuple in the result set and `scanReverse` will scan through the result set backwards.

Please note that it is not possible to use both `fetch` and the methods declared on `scrollingResultSet` on the same result set.

7 Datatypes

7.1 Types in Oracle and Beta

The following table shows the possible types of data in Oracle, and the type to which they are mapped in Beta:

Oracle type	Beta type
Varchar	Text
Number	RealObject
Integer	IntegerObject
Float	RealObject
String	Text
Varnum	IntegerObject
Long	IntegerObject
RowID	IntegerObject
Date	Time
VarRaw	<i>Unsupported type</i>
Raw	<i>Unsupported type</i>
Long Raw	<i>Unsupported type</i>
Unsigned Int	IntegerObject
Long Varchar	Text
Long VarRaw	<i>Unsupported type</i>
Char	Text
CharZ	Text
RowID	<i>Unsupported type</i>
Named Datatype	<i>Unsupported type</i>
Ref	<i>Unsupported type</i>
CLOB	<i>Unsupported type</i>
BLOB	<i>Unsupported type</i>
Binary File	<i>Unsupported type</i>
OCI String	<i>Unsupported type</i>
OCI Date	<i>Unsupported type</i>

7.2 Types in FrontBase and Beta

The following table shows the possible types of data in FrontBase, and the type to which they are mapped in Beta:

--	--

FrontBase type	Beta type
Boolean	booleanObject
Integer	integerObject
SmallInt	integerObject
Float	realObject
Real	realObject
Double precision	realObject
Numeric	realObject
Decimal	realObject
Char	text
VarChar	text
Bit	<i>Unsupported type</i>
TimeTZ	<i>Unsupported type</i>
YearMonth	<i>Unsupported type</i>
DayTime	<i>Unsupported type</i>
CLOB	<i>Unsupported type</i>
BLOB	<i>Unsupported type</i>

7.3 Types in MySQL and Beta

The following table shows the possible types of data in MySQL, and the type to which they are mapped in Beta:

MySQL type	Beta type
Decimal	integerObject
Tiny	charObject
Short	integerObject
Long	integerObject
Float	realObject
Double	realObject
VarString	text
String	text
Null	<i>Unsupported type</i>
Timestamp	<i>Unsupported type</i>
Longlong	text
Int24	<i>Unsupported type</i>
Date	time

Enum	<i>Unsupported type</i>
MediumBLOB	<i>Unsupported type</i>
LongBLOB	<i>Unsupported type</i>
BLOB	<i>Unsupported type</i>

8 References

[Date 93] Date, C.J. and Darwen, H., A Guide to the SQL Standard, Addison Wesley, Reading, MA, 1993.

[Geiger 95] Geiger, K. Inside ODBC, Microsoft Press, 1995.

[Ullman 97] Ullman, J.D., Widom, J., A First Course in Database Systems, Prentice Hall International, 1997.

9.1 Betadbc Interface

```
ORIGIN '~beta/basiclib/betaenv';
INCLUDE '~beta/basiclib/timedate';
-- lib: Attributes --
Connection: (* A connection to a relational DBMS *)
  (# 
    <<SLOT ConnectionLib:Attributes>>;
    declareVar:
      (* Used to declare a wrapper of a BETA object so that
       * the object may be used in SQLStatements *)
      (# 
        <<SLOT DeclareVarLib:Attributes>>;
        varName: ^text;
        set:<
          (* Furtherbind this to set the value of the object that you wrap *)
          object;
        get:<
          (* Furtherbind this to get the value of
           * the object that you wrap *) object;
        setAsText:<
          (* Furtherbind in concrete subpatterns to set
           * the wrapped object's value to the value 'value' *)
          (# value: ^text
            enter value[]
            ...
            #);
        getAsText:<
          (* Furtherbind in concrete subpatterns to get
           * the wrapped object's value as a text *)
          (# value: ^text
            ...
            exit value[]
            #)
        enter varName[]
        ...
        exit varName[]
        #);
      declareInteger: (* Declares a wrapper around an integer *) declareVar
      (# 
        <<SLOT DeclareIntegerLib:Attributes>>;
        set:< (# value: @integer
          enter value
          ...
          #);
        get:< (# value: @integer do INNER exit value #);
        setAsText:< (# ... #);
        getAsText:<
          (# 
            ...
            #)
        #);
      declareReal: (* Declares a wrapper around a real *) declareVar
      (# 
        <<SLOT DeclareRealLib:Attributes>>;
        set:<
          (# value: @real
            enter value
            do INNER
            #);
        get:<
          (# value: @real
            do INNER
            exit value
```

```

        #);
setAsText::<
  (#
  ...
  #);
getAsText::<  (# ... #)
#);
declareText:
(* Declares a wrapper around a text *) declareVar
(# 
<<SLOT DeclareTextLib:Attributes>>;
set::<
  (# value: ^text
enter value[]
do INNER
#);
get::<
  (# value: ^text
do INNER
exit value[]
#);
setAsText::<
  (#
  ...
  #);
getAsText::<  (# ... #)
#);
declareBoolean:
(* Declares a wrapper around a boolean *) declareVar
(# 
<<SLOT DeclareBooleanLib:Attributes>>;
set::<
  (# value: @boolean
enter value
do INNER
#);
get::<
  (# value: @boolean
do INNER
exit value
#);
setAsText::<
  (#
  ...
  #);
getAsText::<  (# ... #)
#);
declareTime: (* Declares a wrapper around a 'time' *) declareVar
(# 
<<SLOT DeclareTimeLib:Attributes>>;
set::<
  (# value: @time
enter value
do INNER
#);
get::<
  (# value: @time
do INNER
exit value
#);
setAsText::<
  (#
  ...
  #);
getAsText::<  (# ... #)
#);

```

```

declareDate:<
(* Declares a wrapper around a 'date'. Only the year, month and day
 * attributes of the time value object are taken into account
 *) declareTime (# <<SLOT DeclareDateLib:Attributes>> #);
declareClock:<
(* Declares a wrapper around a 'time'. Only the hour, minute and sec
 * attributes of the time value object are taken into account *)
declareTime (# <<SLOT DeclareClockLib:Attributes>> #);
formatTime:<
(* Called by BetaDBC to format 'time' values
 * Furtherbind this if a date format other than
 * 'YYYY-MM-DD HH:MM:SS' is needed *)
(# t: @time; value: ^text
enter t
...
exit value[]
#);
formatDate:<
(* Called by BetaDBC to format 'date' values
 * Furtherbind this if a date format other than
 * 'YYYY-MM-DD' is needed *)
(# d: @time; value: ^text
enter d
...
exit value[]
#);
formatClock:<
(* Called by BetaDBC to format 'clock' values
 * Furtherbind this if a clock format other than
 * 'HH:MM:SS' is needed *)
(# c: @time; value: ^text
enter c
...
exit value[]
#);
SQLStatement:<
(* The abstract superpattern for all SQL statements.
 * To use an instance of a subclass of SQLStatement:
 *   1. Open the statement by calling 'open',
 *      1a. specifying an enter parameter to, or
 *      1b. setting the statement's SQL contents by evaluating the enter part.
 *   2. Call 'execute'.
 *   3. Go to 1b., if necessary
 *   4. Call 'close'
*)
(# <<SLOT SQLStatementLib:Attributes>>;
cursorType:<
(* The type of cursor that is created when
 * an SQLStatement is executed *)
(# 
    insensitive:<
        (* Should be furtherbound to set value to 'true'
         * if the cursor should be insensitive to concurrent changes *)
        booleanValue;
    readOnly:<
        (* Furtherbind to to set value to 'false' if modification by cursor is needed *)
        booleanValue (# do true->value; INNER #);
    scrollable:<
        (*
            Furtherbind this if the resulting resultSet should be scrollable
        *)
        booleanValue;
    name:<
        (* Furtherbind this to set the name of the cursor.
         * If setName is not furtherbound a default name will be generated. *)
)

```

```

        (# value: ^text do INNER exit value[] #)
    #);
resultSetType:<
(* The type of resultSet that will be opened
 * when executing this(SQLStatement) *) resultSet;
execute:< (* Executes this(SQLStatement) *)
    (# res: ^resultSetType
    ...
    exit res[]
    #);
open:
(* Opens this(SQLStatement).
 * An SQLStatement must be opened before use *)
    (# value: ^text
    enter value[]
    ...
    #);
close:
(* Closes this(SQLStatement).
 * Call close when done with this(SQLStatement) *)
    (# ... #);
SQLStatementException:<
    BetaDBCEException (# do INNER #);
SQLStatementNotification:< BetaDBCNotification (# do INNER #);
get:< (# t: ^text ... exit t[] #);
set:<
    (# t: ^text
    enter t[]
    ...
    #);
private: @...
enter set
do INNER
exit get
#);
directSQLStatement:
(* Use this statement type if a statement will be executed at most
 * a few times. The contents may contain variable placeholders in the form
 * :varname
 * designating that the shared variable named 'varname'
 * will be bound to that place. In addition, also
 * %b for booleans
 * %c for clocks
 * %d for dates
 * %f for reals
 * %i for integers
 * %s for texts
 * %t for time
 * may be used when setting the contents of this(directSQLStatement)
 * Furthermore, placeholders may be named as in
 * aname%f or aname:varname
 * signifying that the value of the %f or :varname placeholder may be set/get
 * by calling setByName/getByName
 * with "aname" as parameter
*) SQLStatement
(# <<SLOT DirectSQLStatementLib:Attributes>>;
currentMarker:
(* The current marker decides which non-variable placeholder will be set
 * if one of the subpatterns of "marker" is used
*)
(# 
    set:
        (# 
            enter no
            ...

```

```

    #);
get:
  (# ...
   exit no
  #);
  no: @integer
enter set
exit get
#);
marker:
(* Abstract superpattern for markers.
 * Advances currentMarker by 1
 *) (# t: ^text ... #);
b:
(* Set placeholder number 'currentMarker'
 * to the boolean 'value' *) marker
(# value: @boolean
enter value
...
#);
c:
(* Set placeholder number 'currentMarker'
 * to the clock 'value' *) marker
(# value: @time
enter value
...
#);
d:
(* Set placeholder number 'currentMarker'
 * to the date 'value' *) marker
(# value: @time
enter value
...
#);
f:
(* Set placeholder number 'currentMarker'
 * to the real 'value' *) marker
(# value: @real
enter value
...
#);
i:
(* Set placeholder number 'currentMarker'
 * to the integer 'value' *) marker
(# value: @integer
enter value
...
#);
s:
(* Set placeholder number 'currentMarker'
 * to the text 'value' *) marker
(# value: ^text
enter value[])
...
#);
t:
(* Set placeholder number 'currentMarker'
 * to the time 'value' *) marker
(# value: @time
enter value
...
#);
setByName:
(* Abstract superpattern for patterns that
 * sets the value of a named placeholder *)

```

```

(#
  name: ^text;
  t: ^text;
  nameNotFound:< exception
  (# ...
  #)
enter name[ ]
...
#);
setBooleanByName:
(* Set placeholder named 'name'
 * to the boolean 'value' *) setByName
(# value: @boolean
enter value
...
#);
setClockByName:
(* Set placeholder named 'name'
 * to the clock 'value' *) setByName
(# value: @time
enter value
...
#);
setDateByName:
(* Set placeholder named 'name'           * to the date 'value' *)
setByName
(# value: @time
enter value
...
#);
setFloatByName:
(* Set placeholder named 'name'
 * to the real 'value' *) setByName
(# value: @real
enter value
...
#);
setIntegerByName:
(* Set placeholder named 'name'
 * to the integer 'value' *) setByName
(# value: @integer
enter value
...
#);
setTextByName:
(* Set placeholder named 'name'
 * to the text 'value' *) setByName
(# value: ^text
enter value[ ]
...
#);
 setTimeByName:
(* Set placeholder named 'name'
 * to the time 'value' *) setByName
(# value: @time
enter value
...
#);
execute:::
  (# ...
  #);
execDirectException:< BetaDBCException (# do INNER #);
execDirectNotification:< BetaDBCNotification (# do INNER #);
private: @...;
```

```

set::<
  (#  

   varNotDeclared:<  

    exception  

     (# name: ^text  

      enter name[]  

      ...  

      #)  

    ...  

    #);  

  getExpanded:  

  (* Get the contents of this(directSQLStatement)  

   * as it would appear if the statement was executed now *)  

  (# value: ^text  

   ...  

   exit value[]  

   #)  

  do INNER  

  #);  

preparedSQLStatement:  

(* Use this statement type if a statement will be executed multiple  

 * times with different bindings.  

 * ONLY IMPLEMENTED AS A DIRECT STATEMENT  

*) directSQLStatement  

  (# <<SLOT preparedSQLStatementLib:Attributes>> do INNER #);  

resultSet:  

(* A result of an SQLStatement.  

 * If columnCount <> 0 then the resultSet is can be scanned or fetched from.  

 * The tuples of a resultSet can be read at most once.  

*)  

(#  

  <<SLOT resultSetLib:Attributes>>;  

columnCount:  

  (* The number of columns in this(resultSet) *) integerValue  

  (# ... #);  

tupleCount:  

  (* If the statement that created this(resultSet) was a  

   * 1. INSERT, DELETE or UPDATE statement  

   *      tupleCount yield the number of tuples  

   *      affected by the statement  

   * 2. SELECT statement  

   *      tupleCount MAY (i.e. does not in all circumstances) yield  

   *      the number of rows selected  

*) integerValue (# ... #);  

column:  

(* A column in this(resultSet) *)  

(#  

  name: ^Text;  

  no: @Integer;  

  dataType:  

    (* The BETA pattern corresponding to the SQL datatype  

     * for this column. If the SQL datatype is DATE, TIME or  

     * TIMESTAMP then the corresponding BETA pattern will  

     * be time  

     *) ##Object;  

  dataTypeName: (* DBMS specific type name *) ^Text;  

  dataTypeNo: (* ODBC specific numbering of SQL datatypes *)  

    @integer;  

  nullable: @Boolean  

#);  

getColumn: (* Gets the column number 'i' in this(resultSet) *)  

  (# i: @Integer; res: ^column  

  enter i  

  ...  

  exit res[]  

#);

```

```

getColumnName:
(* Gets the column designated by 'name' in this(resultSet) *)
(
    name: ^text;
    res: ^column;
    nameNotFound:< exception
        (# ... #)
enter name[]
...
exit res[]
#);
cursorName:
(* Gets the name of the cursor that points to this(resultSet) *)
(# value: ^text
...
exit value[]
#);
tuple:
(* A row in this(resultSet). If this(resultSet) has been set
 * with a value that contained non-variable placeholders,
 * the values may be retrieved by using the markers below
 *)
(
    <<SLOT ResultLib:Attributes>>;
marker:
(* Gets the value of a non-variable placeholder and advances
 * the placeholder currently referred to
 *) (# ... #);
b: marker
(# value: @boolean
...
exit value
#);
c: marker
(# value: @time
...
exit value
#);
d: marker
(# value: @time
...
exit value
#);
f: marker
(# value: @real
...
exit value
#);
i: marker
(# value: @integer
...
exit value
#);
s: marker
(# value: ^text
...
exit value[]
#);
t: marker
(# value: @time
...
exit value
#);
private: @...
#);
EOT:

```

```

(* If current = EOT then no more tuples are
 * available in this(resultSet) *)
  (# value: ^tuple ... exit value[] #);
set:
(* May called before using fetch on this(resultSet)
 * The pattern entered may either
 * contain
 * 1. Named columns of the form
 *      'name1:var1 name2:%i name3:var2'
 *      which means that the column named 'name1'('name3')
 *      is bound to the variable named 'var1'('var2') and the
 *      column named 'name2' may be retrieved from the
 *      current tuple using the 'i'-marker,
 * or
 * 2. Consecutive columns of the form
 *      ':var1 %i :var2'
 *      which means that the first (third) column in each result
 *      is bound to the variable named 'var1'('var3') and that
 *      the value of the second column may be retrieved from the
 *      current tuple using the 'i'-marker
*)
(# varNotDeclared:<
(* Raised if variable used in the entered
 * pattern was not found *) exception
  (# name: ^text
  enter name[])
  ...
#);
columnNotFound:<
(* Raised if this is a named scan and
 * named column was not found *) exception
  (# name: ^text
  enter name[])
  ...
#);
pattern: ^text
enter pattern[]
...
#);
fetch:
(* Get the current tuple of the relation over which the resultSet ranges.
 * If no more tuples exist, 'current' will be EOT
 * Advance current tuple if possible *)
  (# result: ^tuple
  ...
exit result[]
#);
scan:
(* Scans over the tuples of this(resultSet) starting from the current tuple.
 * A text pattern may be entered. See 'set' for details.
*)
(# current: ^tuple;
varNotDeclared:<
(* Raised if variable used in the entered
 * pattern was not found *) exception
  (# name: ^text
  enter name[])
  ...
#);
columnNotFound:<
(* Raised if this is a named scan and
 * named column was not found *) exception
  (# name: ^text
  enter name[])

```

```

    ...
    #);
    pattern: ^text
enter pattern[]

...
#);
resultSetException:< BetaDBCEception (# do INNER #);
resultSetNotification:< BetaDBCNotification (# do INNER #);
private: @...
#);
open:<
(* Opens this(connection). The name of the connection
 * to be opened must be supplied.
 * Supplying userName and/or password is voluntary.
*)
(# 
    name: ^text;
    userName: ^text;
    password: ^text;
    openConnectionException:< BetaDBCEception (# do INNER #);
    openConnectionNotification:< BetaDBCNotification (# do INNER #)
enter (name[],userName[],password[])
...
#);
close:<
(# 
    closeException:< BetaDBCEception
        (# do INNER #);
    closeNotification:< BetaDBCNotification
        (# do INNER #)
...
#);
connectionException:< BetaDBCEception (# do INNER #);
connectionNotification:< BetaDBCNotification (# do INNER #);
private: @...
#);
BetaDBCEception:
(* Low level interface for catching exceptions.
 * A general exception message is supplied in msg,
 * SQL states and native error codes in SQLState,
 * NativeError in a comma-separated list.
*) Exception
(# 
    SQLState: @text;
    NativeError: @text;
    KnownErrorMsg: ^text;
    HandleType: @integer;
    Handle: @integer
enter (KnownErrorMsg[],HandleType,Handle)
...
#);
BetaDBCNotification: BetaDBCEception
(# do true->continue; INNER #)

```

9.2 Scrollingresultset Interface

```
ORIGIN 'betadbc';
BODY 'private/scrollingresultSetbody';
-- connectionLib: Attributes --
scrollingResultSet:
(* A scrolling resultSet supports relative and absolute
 * positioning in the tuples of a resultSet.
 * Fetch and scan *cannot* be used on a 'scrollingResultSet' if
 * the methods defined in scrollingResultSet has been used on
 * a result.
*) resultSet
(#
  fetchFirst:
  (* Fetches the first tuple in the resultSet.
   * Advances current tuple, if possible.
   * Current will be EOT if the resultSet is empty *)
  (# result: ^tuple
   ...
   exit result[]
   #);
  fetchLast:
  (* Fetches the last tuple in the resultSet.
   * Current will be EOT if the resultSet is empty *)
  (# result: ^tuple
   ...
   exit result[]
   #);
  fetchNext:
  (* Fetch next tuple in the resultSet.
   * Advance current tuple.
   * Equivalent to 'fetch'
   *)
  (# result: ^tuple
   ...
   exit result[]
   #);
  fetchPrior:
  (* Fetch prior tuple in the resultSet
   * Deadvance current tuple.
   *)
  (# result: ^tuple
   ...
   exit result[]
   #);
  fetchRelative:
  (* Fetches tuple number 'increment' counting from the current tuple
   * If increment is > 0, the current tuple is advanced, if increment is < 0,
   * the current tuple is deadvanced
   * Current will be EOT if the operation is not meaningful *)
  (# increment: @integer; result: ^tuple
  enter increment
  ...
  exit result[]
  #);
  fetchAbsolute:
  (* Positions current at absolute position 'index'
   * If index > 0 then index is counted from start.
   * If index < 0 then index is counted from end,
   *)
  (# index: @integer; result: ^tuple
  enter index
  ...
  exit result[]
```

```

#);
scanReverse:
(* Like 'scan' except that tuples are scanned in reverse order *)
(#  

    current: ^tuple;  

    varNotDeclared:<  

    (* Raised if variable used in the entered  

     * pattern was not found *) exception  

    (# name: ^text  

    enter name[]  

    ...  

    #);  

    columnNotFound:<  

    (* Raised if this is a named scan and  

     * named column was not found *) exception  

    (# name: ^text  

    enter name[]  

    ...  

    #);  

    pattern: ^text  

    enter pattern[]  

    ...  

    #);  

scrollingResultSetException:< BetaDBCException;  

scrollingResultSetNotification:< BetaDBCNotification  

#);
scrollingCursorsSupported: booleanValue  

(* Returns true if this(connection) supports scrolling cursors *)  

(# ... #)

```

9.3 Tables Interface

```
ORIGIN 'betadbc';
BODY 'private/tablesbody';
-- connectionLib: Attributes --
table: (* A table in a relational database *)
  (# CatalogName,SchemaName,TableName,TableType: ^text #);
scanTables:
(* Scans the tables in this(connection)
 * invoking INNER for each table found *)
  (# current: ^table ... #)
```

9.4 Transactions Interface

```
ORIGIN 'betadbc';
BODY 'private/transactionsbody';
-- connectionLib: Attributes --
transactionsSupported:
(* Returns true iff this(connection) supports transactions *)
  (# isSupported: @boolean
   ...
   exit isSupported
   #);
autoCommitMode:
(* In autoCommitMode, every database operation is a transaction
 * that is committed when performed. Disabling autoCommitMode is
 * only meaningful when this(connection) supports transactions.
 * The default in BetaDBC is autoCommitMode
*)
  (# autoCommit: @boolean
   enter (# enter autoCommit ... #)
   exit
   (#
   ...
   exit
   autoCommit
   #)
   #);
readUncommitted: (* Allows dirty reads, nonrepeatable reads, and phantoms *)
  integerValue (# ... #);
readCommitted:
(* Disallows dirty reads, but allows nonrepeatable reads and phantoms *)
  integerValue (# ... #);
repeatableRead:
(* Disallows dirty reads and nonrepeatable reads, but allows phantoms *)
  integerValue (# ... #);
serializable:
(* Disallows dirty reads, nonrepeatable reads, and phantoms *) integerValue
  (# ... #);
transactionLevelSupported: booleanValue
(* Check whether level is supported by this(connection) *)
  (# level: (* One of the above defined transaction levels *) @integer
   enter level
   ...
   #);
transactionLevel:
(* Sets the transaction level of the current transaction in
 * this(connection). Serializable is the default
 *)
  (# level: (* One of the above defined transaction levels *) @integer
   enter (# enter level ... #)
   exit
   (#
   ...
   exit level
   #)
   #);
commit:
(* Commits all transactions pertaining to this connection.
 * Use transactionsSupported to check whether this is meaningful *)
  (# ... #);
rollBack:
(* Rolls back all transactions pertaining to this(connection)
 * Use transactionsSupported to check whether this is meaningful.
 *)
  (# ... #)
```

Index

The entries in the alphabetic index consists of selected words and symbols from the body files of this manual – these are in **bold** font – as well as the identifiers defined in the public interfaces of the libraries – set in regular font.

In the manual, the entries, which can be found in the index are typeset like this. This can help localizing the identifier, when the link from the index if followed – especially in the case where the browser does not scroll the line to the top, e.g. because there is less than a page of text left.

In the small table of letters and symbols below, each entry links directly to the section of the index containing entries starting with the corresponding letter or symbol.

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

A

[autoCommit](#)

[autoCommitMode](#)

[autoCommit](#)

B

[BetaDBCEception](#)
[Handle](#)
[HandleType](#)

[KnownErrorMsg](#)
[NativeError](#)
[SQLState](#)

[BetaDBCNotification](#)

C

[CatalogName](#)
[close](#)
[commit](#)
[Connection](#)
[close](#)
[connectionException](#)
[connectionNotification](#)
[declareBoolean](#)
[declareClock](#)

[declareDate](#)
[declareInteger](#)
[declareReal](#)
[declareText](#)
[declareTime](#)
[declareVar](#)
[directSQLStatement](#)
[formatClock](#)
[formatDate](#)

[formatTime](#)
[open](#)
[preparedSQLStatement](#)
[private](#)
[resultSet](#)
[SQLStatement](#)
[connectionException](#)
[connectionNotification](#)
[current](#)

D

[declareBoolean](#)
[declareClock](#)
[declareDate](#)

[declareInteger](#)
[declareReal](#)
[declareText](#)

[declareTime](#)
[declareVar](#)
[directSQLStatement](#)

F

[fetchAbsolute](#)
[fetchFirst](#)
[fetchLast](#)

[fetchNext](#)
[fetchPrior](#)
[fetchRelative](#)

[formatClock](#)
[formatDate](#)
[formatTime](#)

H

Handle	HandleType
--------	------------

I

Introduction	isSupported
--------------	-------------

K

KnownErrorMsg	
---------------	--

L

level [2]	
-----------	--

N

NativeError	
-------------	--

O

open	
------	--

P

preparedSQLStatement	private
----------------------	---------

R

readCommitted	repeatableRead	rollBack
readUncommitted	resultSet	

S

scanReverse	fetchLast	scrollingResultSetNotification
scanTables	fetchNext	selectdb
current	fetchPrior	selecthost
SchemaName	fetchRelative	serializable
scrollingCursorsSupported	scanReverse	SQLState
scrollingResultSet	scrollingResultSetException	SQLStatement
fetchAbsolute	scrollingResultSetNotification	
fetchFirst	scrollingResultSetException	

T

table	TableName	level
CatalogName	TableType	transactionsSupported
SchemaName	transactionLevel	
TableName	level	isSupported

TableType transactionLevelSupported