Description of the pgUtils package

Johannes Rainer*

April 20, 2005

Tyrolean Cancer Research Institute
Innrain 66, 6020 Innsbruck, Austria, `http://www.tcri.at`
and Institute for Genomics and Bioinformatics, Graz University of Technology,
8010 Graz, Austria, `http://www.genome.tugraz.at`

Contents

1 Introduction

The package *pgUtils* depends on the *RdbiPgSQL* package and provides some utility functions for databases in special for a PostgreSQL database. The package provides functions for creating database tables with autoincrementing primary keys and the posibility to referencing to other tables using foreign keys to allow a referential integrity of the data. Other functions can be used to insert or update the information in database tables. All functions run within transactions, so if any error occurs during a database call, the original state of the database before the call can be restored. Additionally the packages provides a simple logging mechanism, which writes log messages to a file (which can be specified).

---

*johannes.rainer@tugraz.at

2 Creating a database table

First of all a connection to the database has to be established. This can be done using the `dbConnect`. The code below connects to the database *template1* which is installed by default in every PostgreSQL database. As we do not want to store any information into this datbase we create a new database by sending the appropriate SQL call and connect then to this new database.

```
> library(pgUtils)

Loading required package: Rdbi
Loading required package: RdbiPgSQL

> con.su <- dbConnect(PgSQL(), host = "localhost", user = "postgres",
+     dbname = "template1")
> if (sum(dbListDatabases(con.su)[, "datname"] == "pgutils") >
+     0) {
+     dbSendQuery(con.su, "DROP DATABASE pgutils")
+ }

status = 1
status.string = PGRES_COMMAND_OK
rows = 0
columns = 0
is.binary = FALSE
command.response = DROP DATABASE

> dbSendQuery(con.su, "CREATE DATABASE pgutils")

status = 1
status.string = PGRES_COMMAND_OK
rows = 0
columns = 0
is.binary = FALSE
command.response = CREATE DATABASE

> dbDisconnect(con.su)
> con <- dbConnect(PgSQL(), host = "localhost", user = "postgres",
+     dbname = "pgutils")
> log.file <- "test.log"
```

The first database table that we will create is a simple table with 3 columns, two of them contain numbers and one character strings. The `log.file <- "test.log"` specifies the name for the log file. By default this file is called *pgUtils.log* and resides in the current working directory.

```
> createDBTable(con, name = "firsttest", attributes = c("a", "b",
+     "c"), data.types = c("TEXT", "REAL", "REAL"))

[1] FALSE

> dbColnames(con, "firsttest")

[1] "firsttest_pk" "a"            "b"            "c"
```

2

This call created a table called *firsttest* with three columns. The function automatically created a additional column with the name `firsttest_pk` that is used as primary key column and which is automatically incremented upon data insertion. The `data.types` attribute allows to specify the data types for the columns (if not submitted all data types will automatically set to *TEXT*). To insert data into the database the function `insertIntoTable` can be used.

```
> MyTable <- data.frame(a = c("some", "text"), b = c(2, 3), c = c(1.3,
+     3.5))
> MyTable

     a b   c
1 some 2 1.3
2 text 3 3.5

> insertIntoTable(con, name = "firsttest", data = MyTable)
```

To read the data from the database the functions `dbSendQuery` and `dbGetResult` from the *RdbiPgSQL* can be used (in combination with SQL).

```
> result <- dbSendQuery(con, "SELECT * FROM firsttest")
> dbGetResult(result)

  firsttest_pk    a b   c
1            1 some 2 1.3
2            2 text 3 3.5
```

3 Handling table references

Relational databases allow to combine informations between database tables and to concatenate the informations. In the next example we will also create two database tables that are related to each other (to be exact we create a *1 to n* relation, that means that *n* entries of the one table are related (belong) to one entry in the other table).

```
> createDBTable(con, name = "species", attributes = c("speciesname",
+     "value"), data.types = c("TEXT"))

Setting the data types for all attributes to   TEXT

[1] FALSE

> insertIntoTable(con, name = "species", data = data.frame(speciesname = c("hobbit",
+     "human", "orc"), value = c("good", "bad", "very bad")))
> createDBTable(con, name = "individual", attributes = c("name",
+     "age"), data.types = c("TEXT"), references = "species")

Setting the data types for all attributes to   TEXT

[1] FALSE

> ref <- new("AutoReference", source.table = "individual", ref.table = "species",
+     source.table.column = "speciesref", ref.table.column = "speciesname")
> insertIntoTable(con, name = "individual", data = data.frame(name = c("frodo",
+     "fred", "bilbo"), speciesref = c("hobbit", "human", "hobbit"),
+     age = c(23, 32, 111)), references = ref)
```

The main points about the calls above are, that it is important to define which table relates to which and which column of the submitted data table can be used to establish the relation between the tables (in the example above this is done with the `AutoReference` object and the column *speciesref* in the second data table). The function `insertIntoTable` uses this information to insert the primary keys of the entry to which the rows of the second table link into theyr foreign key field. The database table *species* has the following data stored

```
> dbGetResult(dbSendQuery(con, "SELECT * FROM species"))

  species_pk speciesname    value
1          1      hobbit     good
2          2       human      bad
3          3         orc very bad
```

and the table *individual*

```
> dbGetResult(dbSendQuery(con, "SELECT * FROM individual"))

  individual_pk species_fk  name age
1             1          1 frodo  23
2             2          2  fred  32
3             3          1 bilbo 111
```

Both tables have a primary key column (usually the name of the table ending in _pk) and the *individual* table has also a foreign key column (usually the name of the referenced table ending in _fk) that is used to link to the *species* table. So n rows of the *individual* table relate to one entry in the *species* table.

The information between the two tables can now be joined like

```
> dbGetResult(dbSendQuery(con, "SELECT * FROM individual JOIN species ON (species_fk=species_pk) WHERE speciesname='hobbit'"))

  individual_pk species_fk  name age species_pk speciesname value
1             1          1 frodo  23          1      hobbit  good
2             3          1 bilbo 111          1      hobbit  good
```

Another nice feature of this foreign keys concept is, that it is not possible to delete entries from the database that are related by another entry. In our case it means that we cannot delete the *hobbit* entry as long as there are *hobbits* in our *individual* table.

```
> dbSendQuery(con, "DELETE FROM species WHERE speciesname='hobbit'")

status = 7
status.string = PGRES_FATAL_ERROR
error.message = ERROR:  update or delete on "species" violates foreign key constraint "$1" on "individual"DETAIL:  Key (specie
rows = 0
columns = 0
is.binary = FALSE

> dbDisconnect(con)
```

In this way the referential integrity will be mantained and the data in the database will not get inconsistent.

4   The logging mechanism

The package *pgUtils* provides also a simple logging mechanism, that logs debug pr error messages into a file by default called *pgUtils.log*. We can also read the logging messages that have been written to the file until now (we called the log file in the first example *"test.log"*).

```
> log.level

[1] "DEBUG"

> readLines("test.log")

 [1] "> Wed Apr 20 12:05:01 2005 : createDBTable: INFO: creating table... CREATE TABLE \"firsttest\" ( firsttest_pk BIGSERIAL
 [2] "> Wed Apr 20 12:05:01 2005 : insertIntoTable:  INSERT INTO \"firsttest\" (a, b, c) VALUES (  'some', '2', '1.3' ); : PGR
 [3] "> Wed Apr 20 12:05:01 2005 : insertIntoTable:  INSERT INTO \"firsttest\" (a, b, c) VALUES (  'text', '3', '3.5' ); : PGR
 [4] "> Wed Apr 20 12:05:01 2005 : createDBTable: INFO: creating table... CREATE TABLE \"species\" ( species_pk BIGSERIAL PRIM
 [5] "> Wed Apr 20 12:05:01 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'hobbit', 'good' )
 [6] "> Wed Apr 20 12:05:01 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'human', 'bad' );
 [7] "> Wed Apr 20 12:05:01 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'orc', 'very bad'
 [8] "> Wed Apr 20 12:05:01 2005 : createDBTable: INFO: creating table... CREATE TABLE \"individual\" ( individual_pk BIGSERIA
 [9] "> Wed Apr 20 12:05:01 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[10] "> Wed Apr 20 12:05:01 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[11] "> Wed Apr 20 12:05:01 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[12] "> Wed Apr 20 12:07:05 2005 : createDBTable: INFO: creating table... CREATE TABLE \"firsttest\" ( firsttest_pk BIGSERIAL
[13] "> Wed Apr 20 12:07:05 2005 : insertIntoTable:  INSERT INTO \"firsttest\" (a, b, c) VALUES (  'some', '2', '1.3' ); : PGR
[14] "> Wed Apr 20 12:07:05 2005 : insertIntoTable:  INSERT INTO \"firsttest\" (a, b, c) VALUES (  'text', '3', '3.5' ); : PGR
[15] "> Wed Apr 20 12:07:05 2005 : createDBTable: INFO: creating table... CREATE TABLE \"species\" ( species_pk BIGSERIAL PRIM
[16] "> Wed Apr 20 12:07:05 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'hobbit', 'good' )
[17] "> Wed Apr 20 12:07:05 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'human', 'bad' );
[18] "> Wed Apr 20 12:07:05 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'orc', 'very bad'
[19] "> Wed Apr 20 12:07:05 2005 : createDBTable: INFO: creating table... CREATE TABLE \"individual\" ( individual_pk BIGSERIA
[20] "> Wed Apr 20 12:07:06 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[21] "> Wed Apr 20 12:07:06 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[22] "> Wed Apr 20 12:07:06 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[23] "> Wed Apr 20 12:08:40 2005 : createDBTable: INFO: creating table... CREATE TABLE \"firsttest\" ( firsttest_pk BIGSERIAL
[24] "> Wed Apr 20 12:08:40 2005 : insertIntoTable:  INSERT INTO \"firsttest\" (a, b, c) VALUES (  'some', '2', '1.3' ); : PGR
[25] "> Wed Apr 20 12:08:40 2005 : insertIntoTable:  INSERT INTO \"firsttest\" (a, b, c) VALUES (  'text', '3', '3.5' ); : PGR
[26] "> Wed Apr 20 12:08:40 2005 : createDBTable: INFO: creating table... CREATE TABLE \"species\" ( species_pk BIGSERIAL PRIM
[27] "> Wed Apr 20 12:08:40 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'hobbit', 'good' )
[28] "> Wed Apr 20 12:08:40 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'human', 'bad' );
[29] "> Wed Apr 20 12:08:40 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'orc', 'very bad'
[30] "> Wed Apr 20 12:08:40 2005 : createDBTable: INFO: creating table... CREATE TABLE \"individual\" ( individual_pk BIGSERIA
[31] "> Wed Apr 20 12:08:41 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[32] "> Wed Apr 20 12:08:41 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[33] "> Wed Apr 20 12:08:41 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[34] "> Wed Apr 20 12:10:46 2005 : createDBTable: INFO: creating table... CREATE TABLE \"firsttest\" ( firsttest_pk BIGSERIAL
[35] "> Wed Apr 20 12:10:47 2005 : insertIntoTable:  INSERT INTO \"firsttest\" (a, b, c) VALUES (  'some', '2', '1.3' ); : PGR
[36] "> Wed Apr 20 12:10:47 2005 : insertIntoTable:  INSERT INTO \"firsttest\" (a, b, c) VALUES (  'text', '3', '3.5' ); : PGR
[37] "> Wed Apr 20 12:10:47 2005 : createDBTable: INFO: creating table... CREATE TABLE \"species\" ( species_pk BIGSERIAL PRIM
[38] "> Wed Apr 20 12:10:47 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'hobbit', 'good' )
[39] "> Wed Apr 20 12:10:47 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'human', 'bad' );
[40] "> Wed Apr 20 12:10:47 2005 : insertIntoTable:  INSERT INTO \"species\" (speciesname, value) VALUES (  'orc', 'very bad'
[41] "> Wed Apr 20 12:10:47 2005 : createDBTable: INFO: creating table... CREATE TABLE \"individual\" ( individual_pk BIGSERIA
[42] "> Wed Apr 20 12:10:47 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[43] "> Wed Apr 20 12:10:47 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
[44] "> Wed Apr 20 12:10:47 2005 : insertIntoTable:  INSERT INTO \"individual\" (species_fk, name, age) VALUES (  (SELECT spec
```