

This assignment includes both programming and pen-and-paper exercises. For the former, you should follow the electronic submission instructions on page 7. For the latter, you should submit a hardcopy of pages 8–12 (only). You are welcome (and encouraged) to use any resources (e.g., Web sites) to help you with your work. However, *all such help must be clearly noted* in your submissions. Further, no matter what you use, *you must be able to explain* how and why it works.

For this homework, questions marked with  $\star$  are required for COS 580 students only. A  $\star$  mark on a question in this homework does *not* mean a similar question is not applicable to COS 480 students in future homeworks or exams! Therefore, COS 480 students are encouraged to answer the  $\star$  questions for extra credit on this homework as well.

COS 480 students earn points as indicated for each question below. The main programming task described below is worth 160 points and the non- $\star$  pen-and-paper questions are worth 40 points, for a total of 200, the full score. Any points earned on  $\star$  questions (potentially another 100 points) are simply added to the score. If the resulting score is greater than 200, the remaining points are extra credit. COS 580 students must answer all questions to yield a raw score (maximum 300) that will be scaled by  $2/3$  to obtain the score (maximum 200) on the homework. (COS 580 students who wish to work for extra credit should contact me for a suitable assignment.)

**Goal:** This assignment asks you to build a small database application for managing information related to rooms, facilities, and reservations. The main goal of this assignment is to gain experience using one or more application programming interfaces (APIs) to PostgreSQL. Secondary goals include practice in writing SQL queries and a study of the impact of database design on ease of querying and updating data, and on maintaining database consistency. Hopefully, this homework will provide a concrete motivation for these topics, which we will study soon.

**The Programming Environment:** An important part of this assignment is learning the interface between a typical programming environment and the database system. You are free to use a programming language and database interface library of your choice. However, only C (with embedded or dynamic SQL) and Java (with JDBC) are supported. While we will try to help you with other languages and libraries, you should not assume any particulars without first checking with me. Figuring out the details of the database system interface and the necessary libraries usually takes people a lot longer than they expect, so please start working on at least this part early.

**Database Tables:** The application uses the two database tables `Rooms` and `Facilities` from the previous assignment, with the following changes:

1. The `Rooms` table should include a column named `directions` that holds brief textual directions to the corresponding room (e.g., “Fourth door on the left heading north from the south stairwell.”).
2. The `Facilities` table should include a column named `comments` that holds comments from users of the rooms and facilities, all concatenated in one field (e.g., “The room is noisy. The chairs are uncomfortable. Fix the lights please!”).
3. Each of `Rooms` and `Facilities` should include a column named `id` that uniquely identifies rows in the that table.

We also need a third table, `Reservations`, with columns `cnum`, `rstart`, `rend`, `building`, `room`, and `name`. Each row in this table represents a reservation, with the columns representing, in order, the confirmation number, start time, end time, building name, room number, and name of the person making the reservation. Choose appropriate types for the columns. Unless otherwise specified, you may assume that all string-valued attributes contain at most 100 characters. An exception is the `comments` attribute of `Facilities`, which may contain up to 100,000 characters.

**The Application Program:** As described further in the packaging instructions below, your submission should produce an executable file called `roomsdb`, that uses PostgreSQL to implement the application described below. You must implement your application program as a Unix command-line program that reads from standard input and writes to standard output. Your program will be tested and graded on Gandalf. If you use some other machine for development, please check very carefully that your code runs on Gandalf. This application must implement the user functions described below. When the work (both internal processing and output to user) for each function is done, your application should write (to standard output) five dashes (-----) followed by a single newline character. We will refer to this string of five dashes followed by a newline as the **function termination string**. The following description also refers to a **separator string**, which consists of the three-character sequence space-colon-space. We will assume that the separator string is not a substring of any valid string input to this application. Except the output described in this homework, your program should not produce any extra output, such as diagnostic messages.

These functions will be invoked from standard input by listing the function name followed by its arguments, one per line. For example, the `connect` function described below takes two arguments and may be invoked as follows (using example values for the arguments):

```
connect
bigmoose
xyzzzy
```

String arguments will be listed verbatim, with no quotes or other demarcation. You may assume that function arguments do not contain any newline characters. Integers will be listed in conventional format (e.g., 123, 74). You may assume that all numbers are in the range [0...100,000]. Date-time values are in UTC, with the format YYYY-MM-DD HH:MM:SS. For example, 2005-09-27 19:04:09 denotes nine seconds past 7:04 p.m. on the 27th of September, 2005, UTC (i.e., 3:04:09 p.m. EDT).

The input will contain, in general, several function calls in the above format, listed one after the other. Your program should ignore lines with # (pound sign) as the first character. It should also ignore blank lines, but blank lines separating function invocations are *not* required. Since you know the number of arguments each function takes, there is no need for such separation. (Note that the function termination string is used only for output, not in the input.) Your application should read and process the functions in the order in which they appear in the input and should terminate gracefully (e.g., by closing open database connections) when the end of input is reached. There is no special end-of-input marker. You do not need to provide any error-handling features; your program will only be tested on valid input.

**Functions:** The functions that your program should implement are described below. Note that the descriptions use a conventional functional notation of the form  $f(a, b)$ , but the input is presented in the form described above.

**connect( $u, p$ ):** This function will be the first one invoked in any test run, and it will be invoked exactly once per run. In response, your application should perform all necessary initialization and connect to the database server as user  $u$  with password  $p$ . (Strictly speaking, your program need not perform any of these actions, since its observable behavior for this function does not depend on them. However, it is probably a good idea.)

We will test your program using a temporary account  $u$  that is *not* your class account. You may assume that the database for account  $u$  initially contains no user tables. Make sure you do not assume anything specific to your own class account. For example, you cannot rely on any initialization you have in your `.login` or `.bashrc` files, since these files will not be the same for the test account. Please be sure to understand the implications of this requirement. Creating code that can be easily run by someone else is an important part of this homework. For testing, you should use your own account name and password in place of  $u$  and  $p$ . (You may wish to test your submission by temporarily replacing your customized account files, if any, with the default ones that came with your account.)

**createTables():** This function should result in the creation of the database tables described on page 1. This function will be called before any of the functions below. It does not return any results.

**destroyTables():** This function should cause the removal of the database tables created by `createTables`. After `destroyTables`, the database should be in its initial pristine state (with no user tables). You may assume that after this function is called, a call to `createTable` will precede a call to any of the functions described below. This function does not return any results.

**addRoom( $b, f, r, d$ ):** When this function is invoked, your application should add a row ( $b, f, r, d, i$ ) to the `Rooms` table, where  $b, f, r,$  and  $d$  denote, respectively, the building, floor,

room number, and directions, and where  $i$  is an identifier of your program's choosing. (See the note on identifiers below.) The output of this function is the identifier  $i$ .

**findRoom( $s$ ):** This function should search for rooms for which  $s$  occurs as a substring of either building name or directions (or both). This search, and **all searches on string-valued attributes**, should be case-insensitive unless specified otherwise. The matching room records should be printed one per line. Each line should contain the room's identifier (see above), building name, and room number, separated using the separator string described earlier (page 2). The output should be sorted in ascending lexicographic order of building name (case insensitive) and ascending order of room numbers (secondary sort order). Output lines here and elsewhere should be terminated by a single newline character.

**describeRoom( $i$ ):** This function should print descriptive information about the room identified by the given identifier  $i$  (*exact, case-sensitive* string match). If there is no room with identifier  $i$ , no output should be produced and this condition is not an error. If the room identified by  $i$  exists, the following information should be printed on a single line (in this order): building name, room number, floor, number of chairs, and number of screens.

For this and other functions, attribute values and other items printed on an output line should be separated using the separator string described earlier (page 2). Strings should be printed literally (with no quotes, padding, or other artifacts). Integers and dates should be printed in the format used for the input.

**addFacilities( $b, r, c, s$ ):** When this function is invoked, your application should add a row ( $f, b, r, c, s$ ) to the **Facilities** table, where  $b, r, c$ , and  $s$  denote, respectively, the building name, room number, number of chairs, and number of screens, and where  $f$  is an identifier similar to that used in **addRoom**. The output of this function is the identifier  $c$ .

**addReservation( $s, e, b, r, n$ ):** When this function is invoked, your application should add a row ( $s, e, b, r, n, c$ ) to the **Reservations** table, where  $s, e, b, r$ , and  $n$  denote, respectively, the start and end times of the reservation, the building name, the room number, and the name of the person making the reservation. The identifier  $c$  is a reservation confirmation number (identifier) that is also the output of this function.

**Note on Identifiers:** The identifiers generated by your program in response to the **addRoom**, **addFacilities**, and **addReservation** functions must uniquely identify the rows in the respective tables. Your application is responsible for generating and managing these identifiers. Once your application has exposed a room's identifier (by printing it as output), the identifier may be presented as an argument of the **describeRoom** function at any point in the future. These identifiers must persist between sessions. For example, if your program exposes a room identifier `xyzy182` during one session a **describeRoom** function call with `xyzy182` as the argument must produce details of the corresponding record. Unless this record has been deleted or otherwise modified in the interim, the output of this **describeRoom** function

invocation should be the same as if it had been invoked in the original session. All matching for identifiers should be exact. (If you use strings as identifiers, the match should be case-sensitive, exact string match, for example.) There is a similar constraint on reservation identifiers: Once exposed, they must permit lookup using the `describeReservation` function below.

**describeReservation(*c*):** This function should print descriptive information about the room reservation identified by the given identifier (confirmation number) *c* (*exact, case-sensitive* string match). If there is no reservation matching identifier *c*, no output should be produced and this condition is not an error. If the reservation by *c* exists, the following information should be printed on a single line (in this order): building name, room number, start time, end time, and name of reservation holder.

**makeReservations(*p, l, b, r, n*):** This function is similar to `addReservation`, differing only in how the begin and end times of the reservation are specified. Instead of these times being specified explicitly, they are specified using a *time-pattern* *p* and time-length *l* as described below. (The other arguments (*b, r, and n*) are treated as in `addReservation`.) Recall that all times are in UTC.

The begin times are based on interpreting the time-pattern *p* as a *crontab*<sup>1</sup> expression. The expression *p* consists of five fields (minute, hour, day of month, month, and day of week) separated by whitespace. The following excerpt (slightly modified) from the crontab manual describes the semantics.

[A timestamp matches *p*] when the minute, hour, and month of year fields match the current time, and when at least one of the two day fields (day of month, or day of week) match the current time. The time and date fields are:

field	allowed values
minute	0–59
hour	0–23
day of month	1–31
month	1–12 or names
day of week	0–7 (0 and 7 mean Sun.)
[year	integer year of Gregorian calendar]

A field may be an asterisk (\*), which always stands for “first–last.”

Ranges of numbers are allowed. Ranges are two numbers separated with a hyphen. The specified range is inclusive. For example, 8–11 for an “hours” entry specifies [a reservation] at hours 8, 9, 10 and 11. [Similarly, 2005–2007 specifies a reservation for years 2005, 2006, and 2007 A.D. We will assume a temporal granularity of one minute. Thus the time-pattern \* \* 1 1 \* 2006 specifies the

---

<sup>1</sup>Paul Vixie, *crontab*—tables for driving *cron*, Manual, 4th Berkeley Distribution January 1994.

time points marking each minute of each hour of January 1st, 2006 (a total of  $60 \times 60$  points in time.)

Lists are allowed. A list is a set of numbers (or ranges) separated by commas. Examples: “1,2,5,9”; “0-4,8-12.”

Step values can be used in conjunction with ranges. Following a range with “/<number>” specifies skips of the number’s value through the range. For example, “0-23/2” can be used in the hours field to specify command execution every other hour (the alternative in the V7 standard is “0,2,4,6,8,10,12,14,16,18,20,22”). Steps are also permitted after an asterisk, so if you want to say “every two hours,” just use “\*/2.”

Names can also be used for the “month” and “day of week” fields. Use the first three letters of the particular day or month (case doesn’t matter). Ranges or lists of names are not allowed.

A single invocation of `makeReservations` results in the creation of several reservations (several entries in the `Reservations` table). We will refer to this group of reservations created by a `makeReservations` invocation as a *reservation group*. The time-length argument specifies the duration of each reservation in the reservation group, in minutes. For example, the following invocation results in 24 reservations for January 15th, 2006 for the first 10 minutes of each hour of that day:

```
makeReservation(0 * 15 1 * 2006, 10, Neville, 210, Alice)
```

For this assignment, you may assume that reservations beyond 2007-12-31 23:59:59 may be safely ignored. For patterns that specify reservations both before and after this date, only the reservations after the date may be ignored; the earlier ones must be managed properly. This assumption may simplify your implementation of invocations such as the following:

```
makeReservation(0 * 15 1 * *, 10, Neville, 210, Alice)
```

**matchReservation(*p*, *b*, *r*):** This function finds reservations for room number *r* in building *b* such that there is at least one instant of time that matches the time-pattern *p* and that lies between the begin and end times of the reservation (including the end-points). The interpretation of *p* is identical to that used by `makeReservations`. The output consists of the confirmation numbers of the matching reservations, one per line.

**getFreeRoom(*p*, *l*, *c*, *s*):** This function finds rooms that are unreserved at all times matching the time-pattern *p* and *l*, which are interpreted as in `makeReservations`. In addition, the matching rooms are required to have at least *c* chairs and *s* screens. The output consists of the building and room-number for each matching room, one per line. A room is considered unreserved at a time matching *p* only if it is unreserved for the entire length of each interval specified by *p* and *l*. If the facilities (number of chairs and screens) of a room are not known, the room does not match.

**Packaging and Submission:** You should generate a gzip-compressed tar archive file called *M-hw02-N.tgz*, replacing *M* with your last-name and replacing *N* with an arbitrary 4-digit integer (e.g., *Doe-hw01-4242.tgz*). The execution of the following sequence of commands on Gandalf (replacing *Doe-hw02-4242.tgz* with the name of your file, which resides in */tmp*) should result in the creation of a directory */tmp/Doe-hw02*:

```
cd /tmp; gzip -dc Doe-hw02-4242.tgz | tar xf -
```

Typing `make` at the Unix shell prompt in the */tmp/Doe-hw02* directory should result in the complete compilation of your program, producing an executable file called `roomsdb`. You will need to include an appropriate Makefile for this procedure to work. You should also include a short `README` file describing the files in your submission, along with anything that may be helpful in fixing your submission if it does not work as above. You must make sure your program works when `stdin` and `stdout` are redirected. For example, we may run your program as follows, where `datafile` is a text file contains the input of the program: `roomsdb < datafile`. Please check carefully that your file satisfies these requirements. Submit the file using anonymous FTP as described in the previous assignment.





5. (10 pts.) Briefly explain your method for implementing `makeReservations`. Indicate exactly what rows are inserted into various tables. If your implementation uses tables other than those described earlier, indicate their schema. Hint: Use an illustrative example.

6. (10 pts.) ★ Indicate how the schema may be modified to fix the problems you identified in your answer to Question 4. Justify your answer.

7. (10 pts.) ★ Indicate two significant changes necessary to your answer to Question 6 to further improve your schema (in order to fix problems not noted in your answer to Question 4). Identify the problems and their solutions.

8. (10 pts.) ★ In processing reservation groups in `makeReservations`, we permitted reservations beyond 2007-12-31 to be ignored (page 6). Describe how this restriction may be removed, permitting management of reservations that extend indefinitely into the future.

9. (10 pts.) ★ Suppose we wish to permit the selective modification and cancellation of reservations, including canceling some, but not all, of those created by a single `makeReservations` invocation. Indicate how the method described in your answer to Question 8 can support this feature. You may modify your answer if needed.

10. (60 pts.) ★ Either (1) implement your answer to Question 9 in your application and write *implemented* below or (2) simply write *not implemented* to receive ten points. If you write implemented, you must also indicate how these features are exposed to the user by an interface similar to that used for other functions, and how they may be tested. Hint: Include examples.