

This programming assignment should be submitted electronically by following the instructions on page 7. You are welcome, and encouraged, to use any resources, such as 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, functions marked with \star are required for COS 580 students only. The main programming task described below is worth 200 points. For COS 480 students, any points earned on \star functions (potentially another 50 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 implement all functions to yield a raw score (maximum 250) that will be scaled by $4/5$ 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 ferns and observations. 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. This homework should provide a concrete motivation for these topics.

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 database tables `Ferns`, `Habitats`, and `FruitDates` from the previous assignment, with the following changes:

1. The `Ferns` table should include a column named `notes` that holds textual descriptions of ferns, such as the following description for the Ostrich Fern:¹ “Fronds two to eight feet high, growing in a crown; broadly lanceolate, pinnate, the numerous pinnae deeply pinnatifid, narrowed toward the channeled stipe. Fertile fronds shorter, pinnate with margins of the pinn revolute into a necklace form containing the sori.”

¹George Henry Tilton, *The Fern Lover's Companion: A Guide for the Northeastern States and Canada* (Project Gutenberg, <http://www.gutenberg.org/ebooks/11365>, 2004).

2. The `name` column in both `Habitats` and `FruitDates` tables should each be replaced by columns `genus` and `species` that hold the corresponding data for each fern, instead of its common name.

We also use the following additional tables:

1. `Places(name, directions, habitat)`. A row (n, d, h) in this table indicates that the location with name n has habitat h , with d being textual directions to that location. Each location is identified by a unique name, but may have multiple habitat designations and, similarly, multiple directions.
2. `Individuals(id, genus, species, location, height, fruit, odate)`. Each row in this table represents an individual fern, with the columns representing, in order, an attribute that uniquely identifies the individual, the genus and species of the individual, the name of its location (as in the name attribute in `Places`), its height in inches, an indicator if it was fruiting (boolean), and the date of the observation. Although each individual has a unique `id` attribute, there are, in general, multiple observations per individual.
3. `OSchedule(id, obs, indid, odatetime)`. Each row in this table represents a scheduled observation, with the columns representing, in order, a unique identifier for the scheduled observation, the name of the observer (person), the individual fern to be observed (identified using the `id` attribute of `Individuals`), and the date and time of the scheduled observation.

Choose appropriate types for the columns of these tables. Unless otherwise specified, you may assume that all string-valued attributes contain at most 100 characters. Exceptions are the `notes` attribute of `Ferns` and the `directions` attribute of `Places`, which may each contain up to 100,000 characters.

The Application Program: As described further in the packaging instructions below, your submission should *produce* (not contain) an executable file called `ferndb`, 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:

```
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 $[0 \dots 10^5]$. Date-time values are in UTC, with the format `YYYY-MM-DD HH:MM:SS`. For example, `2007-10-31 23:01:06` denotes six seconds past 11:01 p.m. on the 31st of October, 2007, UTC.

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 are encouraged, but not required, 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 all the database tables required by this application. 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.

addFern(n,g,s,l,u,d): When this function is invoked, your application should add a row (n, g, s, l, u, d) to the `Ferns` table, where $n, g, s, l, u,$ and d denote, respectively, the common name, genus, species, lower and upper limits of the height range, and notes on the fern. This function does not return any results.

findFern(s): This function should search for ferns for which s occurs as a substring of one or more of the common name, genus, and species. This search, and **all searches on string-valued attributes**, should be case-insensitive unless specified otherwise. The matching records should be printed one per line. Each line should contain the common name, genus, and species, separated using the separator string described earlier (page 2). The output should be sorted in ascending lexicographic order of the common name (case insensitive). Output lines here and elsewhere should be terminated by a single newline character.

describeFern(n): This function should print descriptive information about all ferns (there may be several, or none) with common name n . If there is no match no output should be produced and this condition is not an error. For every matching fern record, the following information should be printed on a single line (in this order): common name, genus, species, lower range of its height, upper range of its height, and notes.

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.

addPlace(n,d,h): When this function is invoked, your application should add a row (n, d, h) to the `Places` table, where $n, d,$ and $h,$ denote, respectively, the name, directions, and habitat of the place. This function does not return any results.

describePlace(j): This function should print descriptive information about the place with name n (*exact, case-sensitive* string match). The output is the set of matching (n, h, d) tuples, where h and d denote habitat and directions. (There may be multiple matches or no matches; these conditions are not errors.)

addIndividual(*n,l,h,f,d*): When this function is invoked, your application should add a row (or rows, as noted below) (*i, g, s, l, h, f, d*) to the `Individuals` table, where *i* is an identifier generated by your program to uniquely identify the individual; *g* and *s* (genus and species) are the botanical name corresponding to the common name *n*; and *l, h, f, and d* are values for the location, height, fruit, and odate columns. Botanical names (*g* and *s* values above) corresponding to given common names should be determined using the `Ferns` table. If there is no fern in that table with the required common name, the genus and species fields of the newly inserted row in `Individuals` should be null. If there are multiple matches, there should be one row for each (genus, species) pair matching the common name (with the same identifier value *i*). The output of this function is the identifier *i*. (See the note on identifiers below.)

addOSchedule(*o,i,d*): When this function is invoked, your application should add a row (*l, o, i, d*) to the `OSchedule` table, where *l* is an identifier generated by your program to uniquely identify this scheduled observation, and *o, i, and d* denote values for the remaining columns of the `OSchedule` table in the order they were introduced earlier. The output of this function is the identifier *l*.

addOSchedulePlace(*o,p,d*): This function is similar to `addOSchedule`, except that instead of adding a scheduled-observation record for a single individual (fern), it adds a set of records, one for each individual known to be located at the given place *p* (a place name). That is, the net effect is as though `addOSchedule(o,i,d)` were invoked for each individual *i* that is located at *p*. If there are no individuals at the given location, nothing is inserted; this condition is not an error. The output of this function is the set of identifiers of the newly inserted observation-schedule records (one per line).

describeIndividual(*k*): This function is to the `Individuals` table what `describePlace` is to the `Places` table, with the individual identifier *i* replacing the place name. The output, if any, consists of the values of the columns of the `Individuals` table, in the order they were introduced earlier.

describeSchedule(*i*): This function is to the `OSchedule` table what `describeIndividual` is to the `Individuals` table. The output, if any, consists of the values of the columns of the `OSchedule` table, in the order they were introduced earlier.

Note on Identifiers: Every identifier generated by your program in response to an invocation of the `addIndividual`, `addOSchedule`, or `addOSchedulePlace` functions must uniquely identify an individual or scheduled-observation. Your application is responsible for generating and managing these identifiers. Once your application has exposed an individual's identifier (by printing it as output), the identifier may be presented as an argument of the `describeIndividual` function at any point in the future. These identifiers must persist between sessions. For example, if your program exposes a tree identifier `xyzy182` during one

session then a `describeIndividual` function call in a subsequent session with `xyzz182` as the argument must produce details of the corresponding individual. 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 are similar constraints on identifiers of scheduled observations.

addOScheduleX(*o,i,p*): This function is similar to `addOSchedule`, differing only in how the time (and date) of the scheduled observation is specified. Instead of it being specified explicitly, it is specified using a *time-pattern* *p* as described below. As noted below, a single invocation of this function may result in the creation of several scheduled observations, unlike the `addOSchedule` function, which always creates one scheduled observation. The other arguments, *o* and *i*, are treated as in `addOSchedule`. Recall that all times are in UTC.

The scheduled observation times are based on interpreting the time-pattern *p* as a *crontab*² 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 the common era]

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, `2007–2009` specifies a reservation for years 2007, 2008, and 2009. We will assume a temporal granularity of one minute. Thus the time-pattern `* * 1 1 * 2008` specifies the time points marking each minute of each hour of January 1st, 2008 (a total of 60×60 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.”`

²Paul Vixie, *crontab*—tables for driving *cron*, Manual, 4th Berkeley Distribution, 1994.

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 `addOScheduleX` results in the creation of several scheduled observations (several entries in the `OSchedule` table). We will refer to this group of scheduled observations created by a `addOScheduleX` invocation as a *observation group*. For example, the following invocation results in 24 scheduled observations for January 15th, 2008 for the first 10 minutes of each hour of that day:

```
addOScheduleX(Alice, t101ab, 0 * 15 1 * 2008)
```

The output of this function is a list of observation-schedule identifiers, listed one per line.

For this assignment, you may assume that reservations beyond 2010-12-31 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:

```
addOScheduleX(Alice, t101xy, 0 * 15 1 * *)
```

matchObservations(p, l, o): Here p is a time-pattern as described above, while l is a floating-point number. This function finds all scheduled observations for observer (person) o in the time interval $[i, i + l]$, where i is any instant of time that matches the pattern p and l is interpreted as a time duration, in hours. The output consists of the identifiers of the matching scheduled-observations, one per line.

* **getReqdObs(p, l):** This function finds individual ferns that are in need of observations in the interval specified by time-pattern p and duration l , interpreted as in `matchObservations`. We say an individual f needs an observation in an interval I if there is at least one time instant $i \in I$ such that there is no scheduled observation of f in the interval $[i - M, i]$, where M denotes the time duration of one month. The output of this function is a list of records, one per line, indicating the identifier of a matching individual and the earliest instant of time i satisfying the condition described above.

Packaging and Submission: You should follow the packaging and submission procedure suggested by the previous assignment, replacing `hw01` with `hw02` as appropriate. Unpacking your submission should create a directory `M-hw02` (as a subdirectory of the working

directory), where M is your last name. Typing `make` at the Unix shell prompt in this directory should result in the complete compilation of your program, producing an executable file called `ferndb`. 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: `treedb < datafile`. Please check carefully that your file satisfies these requirements. Proper packaging is an important part of this assignment.