

Due: Monday, 18 September 2006, midnight

The main purpose of these exercises is to get you familiar with both the Pyth language that we'll be working on this semester, and the Subversion software. You *must* register a team to complete this homework, since that's the only way to get one of our repositories, and that's where you're handing things in. You will not have immediate access to a repository after registering; we only update our administrative information periodically. Sending mail to `cs164` to indicate when you've changed something (added or modified a team or registered a new SSH key) is a good idea.

It's *really important* that these tools work for you. Do *not* just accept error messages as if they were Acts of Nature beyond your control; *whatever it takes, make it work!*

1. Check out a copy of the content of your team's initial Subversion repository. The command-line version should look like this:

```
svn checkout svn+ssh://cs164-tb@host/team
```

where *host* is one of the instructional servers and *team* is your team's name. This should create a subdirectory in your current directory whose name is *team*, and which contains subdirectories named `staff` and `project`. `staff` contains various potentially useful files from us. The `project` directory is organized as suggested in the Subversion documentation: with three (initially empty) subdirectories `trunk`, `branches`, and `tags`.

2. Under your checked-out team directory (i.e., the directory that contains `staff` and `project`, create an empty subdirectory named with your class login (the one you registered with). We'll call this your "personal directory." Don't use another name, or your partner will be able to read and modify it. Within this new directory, create a directory called `hw2`. Arrange for your personal directory and its contents to be checked into your repository. Check any homework solution files from the problems below into the repository under this `hw2` subdirectory.

3. Either you or your partner should use `svn copy` to copy files from one of the staff directories `staff/java` or `staff/c++` into your `project/trunk` directory (and be sure this copy is committed). Whichever one of you did *not* do this should make sure that when you do an `svn update` on your copy of the `project` subdirectory, you get the files your partner checked in.

Modify some of the files you've copied in (for example, add your name to a source file). Commit your change and make sure your partner can receive the change (using `svn update`).

4. Have either you or your partner use `svn remove` to delete your entire `project/trunk` directory. Commit this change (if needed). Then recreate the directory you just deleted using `svn copy` (with a `-r` option to indicate the version you are copying from).

5. Write a Python program that reads text from the standard input, keeping track of the number of times each word in the text appears, and then prints out the ten most frequently occurring words that are at least 4 letters long, one per line, in order of decreasing frequency. A “word” here is a contiguous substring of letters; anything other than a letter delimits a word. Ignore case (and print words in lower case). Check the result in as the file `count.py` in the `hw2` subdirectory of your personal directory. This program should be doable in the intersection of Pyth and Python.