

UNIVERSITY OF CALIFORNIA
Department of Electrical Engineering
and Computer Sciences
Computer Science Division

CS 61B
Fall 2012

P. N. Hilfinger

Using Subversion

1 Introduction

A *version-control system* (also *source-code control system* or *revision-control system*) is a utility for keeping around multiple versions of files that are undergoing development. Basically, it maintains copies of as many old and current versions of the files as desired, allowing its users to retrieve old versions (for example, to “back out” of some change they come to regret), to compare versions of a file, to merge changes from independently changing versions of a file, to attach additional information (change logs, for example) to each version, and to define symbolic labels for versions to facilitate later reference. No serious professional programmer should work without version control of some kind.

SUBVERSION, which we’ll be using this semester, is an open-source version-control system that enjoys considerable use. The basic idea is simple: SUBVERSION maintains *repositories*, each of which is a set of *versions* (plus a little global data). A version is simply a snapshot of the contents of an entire directory structure—a hierarchical collection of directories (or “folders” in the Mac and PC worlds) that contain files and other directories. A repository’s administrator (in this class, that’s the staff), can control which parts of the repository any given person or group has access to. SUBVERSION is an example of a *client/server* system. That is, running it requires two programs, one of which (the server) accesses the repository and responds to requests from the other program (the client), which is what you run. The client and server need not run on the same machine, allowing remote access to the repository. We have one server program, which the staff will look after, and several alternatives for the client program, depending on how you are working. In this document, we’ll describe the standard command-line client (called `svn`), which you run in a shell.

The staff will maintain a SUBVERSION repository containing subdirectories for each of you on the instructional machines. The SUBVERSION server program manipulates this repository and runs on one of the instructional machines. You will be able to run the SUBVERSION client program (`svn`) either from your class account or remotely from any machine where `svn` is installed.

At any given time, you may have *checked out* a copy of all or part of one of the versions of your repository subdirectory into one of your own directories (known as a *working directory*). You can make whatever changes you want (without having any effect on the repository) and then create (*commit* or *check in*) a new version containing modifications, additions, or deletions you’ve made in your working directory.

If you are working in a team, another member might check out a completely independent copy of the same version you are working on, make modifications, and commit those. You may both, from time to time, *update* your working directory to contain the latest committed files, including changes from other team members. Should several of you have changed a particular file—you in a committed revision, let us say, and others in their working copies—then the others can update their working

copies with your changes, *merging* in any of your changes to files you've both modified. After these team members then commit the current state of their working directories, all changes they've made will be available to you on request. If you attempt to commit a file that someone else has modified and committed, then SUBVERSION will require you to update your file (merging these committed changes into it) before it will allow you to commit.

In this class, I may give you various code skeletons to be filled in with your solutions. Sometimes, these skeletons may contain errors I wish to fix, or I might take pity and add a useful method or two after handing out the assignment. By distributing these skeletons to you in the repository, I make it possible for you to merge my changes in the skeletons into your own work, without your having to find our changes and apply them by hand.

Should you mess something up, all the versions you created up to that point still exist unchanged in the repository. You can simply fall back to a previous version entirely, or replace selected files from a previous version. Of course, to make use of this facility, you must check your work in frequently.

Even though the abstraction that SUBVERSION presents is that of numbered snapshots of entire directory trees full of files, its representation is far more efficient. It actually stores *differences* between versions, so that storing a new version that is little changed from a previous one adds little data to the repository. In particular, the representation makes it particularly fast and cheap to create a new version that contains an additional copy of an entire subdirectory (but in a different place and with a different name). No matter how big the subdirectory is, the repository contains only the information of what subdirectory in what version it was copied from. As we'll see, you can use these cheap copies to get the effect of *branching* a project, and of *tagging* (naming) specific versions. These copies have other uses as well. In particular, they allow you to import a set of skeleton files for an assignment without significantly increasing the size of the repository.

What follows is specific to this course. SUBVERSION actually allows much more flexible use than I illustrate here.

2 Subversion from the command line

2.1 Setting up a working directory

With SUBVERSION, you are almost always working on a *copy* of something in the repository. Only when you *commit* your copy (usually with the command `svn commit`) does the repository get changed. It is impossible to emphasize this too much—failing to commit changes (either by not issuing the command at all or by failing to notice error messages when you do) is the root of most evil that befalls beginning SUBVERSION users.

In particular, a *working directory* is a copy of some directory in the repository. I suggest that you set up such a directory under your account. In what follows, I'll assume this directory's name is **work** (any name will do) and that it is in your home directory. I'll also assume that your class login is **cs61b-yu** (I need your login because that's the name of the part of the class repository that you're allowed to change). You can create this directory with the following sequence of commands:

```
% cd      # Go to your home directory
% svn checkout svn+ssh://cs61b-ta@torus.cs.berkeley.edu/cs61b-yu/trunk work
Checked out revision 101
```

The part that begins "`svn+ssh:`" is the *URL* (Universal Resource Locator) for a part of the repository. It identifies the protocol for communicating (`svn+ssh`), a host machine for those files, the directory containing the repository (`cs61b-yu`), and a subdirectory of the repository (`trunk`). SUBVERSION will create directory **work** if needed, and will add to it a hidden directory called `work/.svn`. Basically, you'll never have to worry about these `.svn` directories; they contain administrative information that

SUBVERSION leaves around for itself for use in later commands, rather like the “cookies” that remote websites are always storing on your disk so that they remember who you are when next you talk to them.

You would soon get heartily tired of typing “`svn+ssh://cs61b-ta...`,” but fortunately, you have two things working for you:

- Once you have set up a working directory, the hidden `.svn` contains the information about the repository’s URL (the `svn info` command will show it to you. Therefore, you will not have to use the full URL too often.
- Feeling your pain, we have used the marvelous facilities provided by modern operating system command processors to define shorthand notation. Specifically, on the instructional machines, we have defined `$MYREPOS` to be short for

```
svn+ssh://cs61b-ta@torus.cs.berkeley.edu/cs61b-yu,
```

so that the command to create your `work` directory is simply

```
svn co $MYREPOS/trunk work
```

The part of the repository that you own is the subdirectory `cs61b-yu`. We’ve set this up for you so that it contains one (initially empty) subdirectory called `tags`, which is the traditional name for the directory where you copy specific, named, versions of a project you are developing. In real life, you might use the `tags` directories to hold releases of your software. In this class, you’ll use it to put versions that you hand in. We can read the entire repository (we own it, after all), and so we will see everything you put in your `tags` directory and automatically treat it as something you’ve handed in.

The staff keeps a subdirectory of the repository containing its publicly available files (such as the skeleton we provide for homeworks and projects). For example, the skeleton for lab #2 is named `svn+ssh://cs61b-ta@torus.cs.berkeley.edu/staff/lab2`. Again, we have a shorthand: `$STAFFREPOS/lab2`.

2.2 Starting a project

Suppose that you’ve just started work on some files for Project #2, and have placed them in the subdirectory `work/proj2` (I’m assuming that you’ve created `work` as in §2.1. We’ll call the `proj2` directory in the repository the *trunk* version of `proj2`—the one you do development on. Your submissions will be snapshots of this directory. In other “real-life” uses of Subversion, you’ll see these trunk directories stored in subdirectories named `trunk`, but we won’t bother with that here.

At the moment, SUBVERSION knows nothing about these files. You can see this by running the `status` subcommand:

```
% cd ~/work
% ls
proj2
% svn status
?      proj2
```

which means “I see this directory called `proj2`, but it does not seem to taken from the repository that this working directory (`work`) came from.” As soon as possible, you should put all these files under version control with the `add` subcommand, which might produce the following output:

```
% cd ~/work
% svn add proj2
A      proj2
A      proj2/Main.java
A      proj2/Main.java~
A      proj2/Main.class
A      proj2/doc
A      proj2/doc/INTERNALS
```

The repository *has not changed*, but now SUBVERSION client knows that you intend for all these files to eventually be part of the repository, as you'll see if you check the status:

```
% svn status
A      proj2
A      proj2/Main.java
A      proj2/doc
A      proj2/doc/INTERNALS
```

Additionally, `svn add` has modified the working directories `proj2` and `proj2/doc` by adding `.svn` directories to them, turning them into official working directories.

If you'd been working on it before this step, your directory might contain files such as `Main.java~` (an Emacs back-up file) or `Main.class` (the result of compilation), which you probably do not need to archive (they are either useless or reconstructable). On the instructional machines, we've configured your version of the client to ignore these and several other files. For those of you who want to set up SUBVERSION clients at home, take a look at the contents of the file `~/subversion/config` to see what we did (look for 'global-ignores'). If unwanted files such as those do get added for some reason, you can remove them from version control:

```
% svn revert proj2/Main.java~ proj2/Main.class
```

which undoes any changes we've made to these files since last creating a version (in this case, it "unadds" them).

The procedure I've outlined so far assumes that you start from scratch with your own files. If we have provided a skeleton, you can use it by first copying the skeleton into a working directory, like this:

```
% cd ~/work
% svn copy $STAFFREPOS/proj2 proj2
A      proj2
```

(As described previously, `$STAFFREPOS` is a shorthand defined on the instructional machines. The full command is

```
% svn copy svn+ssh://cs61b-ta@torus.cs.berkeley.edu/staff/proj2 proj2
```

which should work for you on systems other than the instructional computers.) You now will have a directory named `~/work/proj2` with the same files that we ended up with originally.

So far, you've worked exclusively with your working files. The repository has not changed at all. If you were to erase your files at this point, everything you've done would be lost permanently. So now it's time to record all the changes you've made (specifically, these files and directories you've added) in the repository. As I said above, this is called *committing* the changes:

```
% cd ~/work/proj2
% svn commit -m 'Start Project 2'
... various messages
Committed revision 1294.
% svn update
At revision 1294.
```

Each version has a unique *revision number*. Because of how we've arranged the CS61B repository, the revision number will reflect the total number of versions created by all members of the class, even though their versions have no impact on yours. In this course, you'll probably always want to add that final `svn update` command, as we did above, so that your working directory and the repository are completely in sync. Let's not go into an explanation of why SUBVERSION's designers made this step optional just now.

SUBVERSION requires that each version have a *log message* attached to it, which in this case you've supplied with the `-m` option. For other changes, your log messages should be rather more substantial than this. You can also take them from a file:

```
% svn commit -F my-commit-notes
% svn update
```

If you don't supply either, SUBVERSION will try to invoke your favorite editor. If that is Emacs, which is our default, you'll find yourself looking at a buffer that contains a list of the files you've added. Add an additional message to the buffer, save it, and exit.

Suppose that at this point, some malicious gremlin executes

```
% rm -rf ~/work
```

You can get it all back with

```
% svn checkout $MYREPOS ~/work
```

2.3 The typical work cycle

At this point, you enter the usual cycle of adding and editing files, compiling, and debugging. SUBVERSION will notice any files you change the next time you commit your changes, without any other action by you. Each time you introduce a new source file, tell SUBVERSION about it:

```
% svn add Manager.java
A      Manager.java
```

Occasionally, you'll need to delete a file that you previously committed:

```
% svn delete Munger.java
D      Munger.java
```

or rename it:

```
% svn move Munger.java Mangler.java
A      Mangler.java
D      Munger.java
```

In these last two cases, SUBVERSION will remove or move (rename) your working files appropriately and make a note to itself of these actions for the next commit operation. Neither editing, adding, removing, nor moving a working file has any effect on the repository, which remains unchanged until you commit your changes:

```
% svn commit -m "Add a Manager and mung the Munger"
...
Committed revision 1301.
% svn update
```

In other words, revision 1301 now contains a snapshot of all the files in the current directory (or its subdirectories) that have changed since you last committed them, minus those you have removed, plus those you have added.

Feel free to commit *frequently*, whenever you think you've finished working on a file, or finished a minor change, or just feel like knocking off or going to lunch. You will not be wasting space because SUBVERSION stores only changes (or *deltas*) from one commit operation to the next.

You may find yourself wanting to work both from your home computer and your instructional account. SUBVERSION will help keep you synchronized. If, for example, you've been working from home, the command

```
% svn update
...
Updated to revision 1350.
```

issued from within a working directory, will bring your latest commit into that working directory. Be careful, though: if you haven't been careful to commit your changes at home, you won't get your most recent work. Likewise, if you did not commit the last changes you made in your instructional account (so that the working files have changes that are not reflected in the repository), you'll get messages about how those changes have been "merged" with the changes you committed from home. Your working copy will still be out of sync with the repository, and you may want to do a commit right away.

2.4 Comparing

One of the advantages to keeping around history is that you can see what you've done and recover what you've lost. After you've done some editing on your files, you can compare them with previous versions. To see what changes you've made to file `Main.java` since you committed your changes:

```
svn diff Main.java
```

You'll get a listing that looks something like this:

```
Index: Main.java
=====
--- f2 (revision 1350)
+++ f2 (working copy)
@@ -1,3 +1,5 @@
-import java.lang.ArrayList;
+/* Project #2: Main program. */
+
+import java.util.ArrayList;
+import java.util.Scanner;

@@ -6,4 +8,5 @@
+int N;
+String v;
+Scanner inp;

for (String arg : args) {
```

The ‘+’s indicate lines added in your working copy; the ‘-’s indicate lines removed in your working copy; and the other lines indicate unchanged lines of context. The remaining lines give the line numbers of these changes in the two different versions.

To see all changes to all files, just leave off any file names:

```
% svn diff
Index: Main.java
...
Index: Manager.java
...
```

You can also compare your working files to previous revisions by number, as in

```
% svn diff -r 1294
...
```

or compare two committed revisions of the current directory, as in

```
% svn diff -r 1294:1300
```

Of course, revision numbers are not the easiest things to deal with, so there are other options for specifying differences. To find the differences between the current working copy and what you had at 1PM, you can write

```
% svn diff -r {13:00}
```

or between the current working copy and noon on the 25th of October:

```
% svn diff -r "{2007-10-25 12:00}" # You need quotes because of the space
```

or between the last copy you checked in and the previous one (ignoring changes in the working directory):

```
% svn diff -r PREV:BASE Main.java
```

Very often, you’re simply interested in knowing what files you’ve changed in the working directory and not committed yet. Use the ‘status’ command for this purpose:

```
% svn status
M      Main.java
A      Utilities.java
A      doc/Manual.txt
D      Junk.java
```

This lists changes with a flag indicating modified files (M), added files (A), and deleted files (D).

2.5 Retrieving previous versions

Suppose you’ve modified `Main.java`, have not committed it, but have decided that you’ve really messed it up and should simply roll back to the version you committed. This is particularly easy:

```
% svn revert Main.java
Reverted 'Main.java'
```

You can get fancier. Suppose that after committing `Main.java`, you have second thoughts, and want to fall back to the previous version of `Main.java`. Here’s a simple way to do so:

```
% svn commit
...
Committed revision 1522.
% svn update
At revision 1522.
% svn status
% svn cat -r PREV Main.java > Main.java
% svn status
M      Main.java
% svn diff -r PREV Main.java
No output
```

The `cat` command, like the similarly named command in Unix, lists file contents from the repository. The `> Main.java` here is Unix notation for “send the standard output of this command to `Main.java`.” Thus, the command overwrites `Main.java` with revision `PREV`—the version before the last that was checked in. The repository is not changed yet; the `status` commands in this example show that `Main.java` has been modified from the last committed version. You will have to commit this change to make it permanent.

The problem with this particular technique (less important in this class, but you might as well get exposed to real-world considerations) is that it messes up the historical information that SUBVERSION keeps around. The `log` command in SUBVERSION lists all the changes that a given file or directory has undergone. With the short method above, all you’ll be told is that the file committed in 1522 is somehow *different* from that in later revisions, rather than being told that the later revision is the *same* as that in revision 1521. So I can get a little fancy and recover the previous version in two steps:

```
% svn delete Main.java
...
% svn commit -m 'Remove Main.java in order to restore from previous version.'
...
% svn copy -r 1521 $MYREPOS/proj2/Main.java .
```

(don’t forget the dot on the end, meaning “current directory”). Now when you next commit, you’ll have the version of `Main.java` as of 1521 and the log will say so.

2.6 URLs and paths

At this point, you’ve seen two different syntaxes for denoting a file or directory. Just plain Unix paths or Windows file names:

```
svn delete Main.java
```

and entire URLs (*Universal Resource Locators*, just as in your browsers):

```
svn checkout svn+ssh://cs61b-ta@torus.cs.berkeley.edu/cs61b-yu/trunk/proj2 ~/work/proj2
or
svn checkout $MYREPOS/trunk/proj2 ~/work/proj2
```

Plain file names generally refer to working files (files in *your* directories); URLs refer to files in the repository. As you can see, URLs are rather tedious to write, which is why SUBVERSION uses those `.svn` directories to keep track of administrative details, including the original URL, allowing commands to be much shorter.

Sometimes, however, you really need to reference a file or directory in the repository. The checkout command and copy commands illustrated previously are two examples. Another very important one is producing a copy of a directory in the repository itself. Suppose that you've completed work on Project #2 and are ready to submit it. Suppose, in fact, that you've just finished committing your final version, which you keep in working directory `~/work/proj2`. In this class, we've established the convention that all submissions go in a subdirectory called `tags` of your repository. The command to submit is

```
% svn copy $MYREPOS/proj2 $MYREPOS/tags/proj2-1 -m "Project 2, first submission"
Committed revision 1600.
```

This does not change your working directory, but it causes the entire directory tree `proj2` in the repository to get copied as a new subdirectory `tags/proj2-1` in the repository. Any future changes to `proj2` either in your working files or in the repository have no effect on this copy. If you have a reason later to resubmit, just use the same command, but with a slight modification to distinguish submissions:

```
% svn copy $MYREPOS/proj2 $MYREPOS/tags/proj2-2 -m "Project 2, second submission"
Committed revision 1701.
```

This particular application of `svn copy` is an example of *tagging*, and you might refer to the copy in directory `tags` as a *release* of your project.

2.7 Merging

By far the most complex part of the version-control process is the process of *merging*: combining independent changes to a file. The main problem is that there is no automatic way to do it; at some point, a human must intervene to resolve conflicting changes. SUBVERSION's facilities here are rather primitive as these things go, but for our limited purposes, they will probably suffice. In this course, the only time the problem is likely to come up is in cases where you have started from some skeleton files I give you and I later change the skeleton. Assuming you want to take advantage of my changes, you'll want to perform some kind of merge.

In general, each time I change files that I expect you to modify, I will tag the resulting directory, as described above. So, for example, I will keep the current version of the initial files for Project #2 in the repository directory `.../staff/proj2`, and snapshots of each "released" version of the files in directories `.../staff/proj2-0`, `.../staff/proj2-1`, etc., with the highest-numbered tag being a copy of `.../staff/proj2`.

Let's assume that there are two tags for Project #2, `proj2-0` and `proj2-1`, and that when you copied the public version, `proj2-0` was the version you copied. Now that `proj2-1` is out, you want to incorporate its changes. First, commit your current trunk version (this is a safety measure that gives you a convenient way to undo the merge if it gets fouled up somehow). Now merge in the changes between `proj2-0` and the current trunk version like this:

```
% cd ~/work/proj2
% svn merge $STAFFREPOS/proj2-0 $STAFFREPOS/proj2
U      Main.java
C      Manager.java
```

The merge operation modifies your working directory only; the repository is not changed. The message tells you that `Main.java` has been updated with changes that occurred between `proj2-0` and `proj2`, and that `Manager.java` has also been updated, but there were some *conflicts*—changes between `proj2-0` and the trunk that overlap changes you made since `proj2-0`. These conflicts are marked in `Manager.java` like this:

```

<<<<<<< .working
    System.out.println ("Welcome to my project");
    initialize ();
=====
        initialize (args);
>>>>>>> .merge-right.r1009

```

The lines between "`<<<<<<< .working`" and `=====` are from your version of the file `Manager.java`, and the ones below `=====` up to the `>>>>>>>...` line are from `proj2-1`. In this case, `proj2-0` probably contained `initialize ()`; and you added a `println` call, while `v1` added an argument to the call to `initialize`. SUBVERSION decided that was close enough to the new `proj2`'s change to suggest an overlap, and so has asked you to fix the problem. Simply edit the file appropriately. In this case you probably want

```

    System.out.println ("Welcome to my project");
    initialize (args);

```

Now tell Subversion that the problem is solved with the command

```

% svn resolved Manager.java
Resolved conflicted state of 'Manager.java'

```

Finally, when all conflicts are resolved, commit your merge just as you would any change to your files.

2.8 Quick guide

Here are some common version-control tasks and the SUBVERSION commands that perform them. In all of the following, I'll assume that

- Your class account is `cs61b-yu`.
- You having chosen to keep working copies of your directories in directory `~/work`.
- You use the naming scheme described above, with one subdirectory per assignment.

We'll use "change directory" (`cd`) commands in these examples just to make clear what directory we're in. Many of these will be redundant.

Set up your working directory `~/work`:

```
% svn checkout $MYREPOS/trunk ~/work
```

Create a directory for a new project:

```

% cd ~/work
% svn mkdir proj2
% svn commit -m "Set up Project 2"
% svn update

```

You'd put your work in `~/work/proj2`.

Create a directory for a new project from our template:

```

% cd ~/work
% svn copy $STAFFREPOS/proj2 proj2
% svn commit -m "Set up project 2"
% svn update

```

Tell SUBVERSION about a new file:

```
% cd ~/work/proj2
create Main.java
% svn add Main.java
```

This command does *not* change the repository. You still have to commit the change.

Tell SUBVERSION about a whole new subdirectory:

```
% cd ~/work/proj2
create directory util and files util/IO.java and util/Manager.java
% svn add util
```

This command does *not* change the repository. You still have to commit the change.

Delete a file or directory:

```
% cd ~/work/proj2
% svn del Main.java
D      Main.java
% svn del util
D      util/IO.java
D      util/Manager.java
D      util
```

These change only your working directory. You must commit in order to change the repository.

Get a brief summary of changes since the last commit:

```
% cd ~/work/proj2
% svn status
M      util/Manager.java
A      util/Tools.java
```

Commit changes:

```
% cd ~/work/proj2
% svn commit -m "Log message"
various messages
Committed revision 2000.
% svn update
```

This transmits all additions, deletions, renamings, and changes to version-controlled files in your working directory to the repository, and creates a new revision that contains them. Leave off the `-m "Log message"` to get SUBVERSION to use a text editor to compose the message.

Compare your current working files against the repository: To compare against the most recently committed version:

```
% cd ~/work/proj2
% svn diff
```

Against another version by revision number:

```
% svn diff -r 1756
```

Against another version in the repository by time:

```
% svn diff -r '{13:00}'
```

Against another version in the repository by time and date:

```
% svn diff -r '{2007-10-17 13:00}'
```

Compare a particular working file with the repository:

```
% svn diff Main.java
```

You can also use `-r` here to look at previous versions of the file in the repository.

Merge corrections in our template into your trunk: First commit the trunk directory. Let's assume that you started from the initial release of the skeleton, which will be called `proj2-0`, and want to merge in the changes we've made to `proj2-1`.

```
% cd ~/work/proj2
% svn merge $STAFFREPOS/proj2-0 $STAFFREPOS/proj2
U      Main.java
C      Manager.java
Edit any files marked C to resolve clashes
% svn resolved Manager.java      # Tell svn about resolution
...
% svn commit -m "Merged in skeleton changes between v0 and v1"
Committed revision 2009.
```

Submit a copy of your trunk: First, make sure your work is committed:

```
% cd work/proj2
% svn commit -m "Project 2, ready to submit"
Committed revision 2010.
% svn update
```

and now create a submission copy:

```
% svn copy $MYREPOS/trunk/proj2 $MYREPOS/tags/proj2-0 -m "Project 2, first submission"
Committed revision 2011.
```

Follow the same procedure for additional submissions of the same project, changing `proj2-0` to `proj2-1`, etc.

Find out what you have submitted: To find out what submissions are actually in the repository (as opposed to merely sitting, uncommitted, in your working directory), use

```
% svn ls $MYREPOS/tags
```

A slight variation of this command lists all the files (not just the top-level directories):

```
% svn ls -R $MYREPOS/tags
```

although you will probably want to restrict this to a single submission:

```
% svn ls -R $MYREPOS/tags/proj2-1
```

Of course, you can see what you actually submitted by just checking it out (I suggest doing so someplace *other* than `work`, to avoid confusion):

```
svn co $MYREPOS/tags/proj2-1 DIR
```