Unix Tools	Lecture 7
CS114 Fall 2003	Wednesday, October 15, 2003

More text processing

Recall from last time that when we execute a command, we can redirect `stdout` to `file` by appending `> file`. If you want to *add* the output to the end of an existing file, you can use the `>> file` redirection operator.

(What happens if you attempt to redirect to a file that already exists? In fact, it depends. You may either get an error saying that the file already exists, or you may have the file deleted and replaced by the output of the command. The behavior depends on the shell that you are using the access Unix. We will see shells in the next lecture.)

Here are some interesting programs that can be used with redirection and pipelines:

- `cat` *`file1 file2 ...`* outputs all the files to `stdout`, concatenating them together.
- `head` *`file1 file2 ...`* lists the first 10 lines of each file to `stdout` (you can specify how many lines to list by using a `-`*`number`* option).
- `tail` *`file1 file2 ...`* lists the last 10 lines of each file to `stdout` (you can specify how many lines to list by using a `-`*`number`* option).
- `less` *`file`* displays *`file`* a screenful at a time.

For example, assume that you have a file `f`, for which you want the first ten lines in alphabetical order containing a match for `rp65`. The following pipeline does this.

```
fgrep 'rp65' f | sort | head
```

Since sorting can be quite time consuming, what you should do with the above is send the output for a file `out`, and compute the whole thing in the background. (Recall last lecture.):

```
fgrep 'rp65' f | sort | head > out  &
```

Stream processing with sed

A very convenient utility to use in a pipeline is `sed`, which is a program that can perform subsitutions based on regular expression. There are two ways of invoking `sed`:

- `sed -e 'script' file` which applies the instructions in *script* to the content of *file* and sends the result to `stdout`, or
- `sed -f scriptfile file` which applies the instructions in the file *scriptfile* to the content of *file* and sends the result to `stdout`.

If you do not specify a file from which to get input, `sed` will takes its input from `stdin`.

A script is a set of instructions that you can use to indicate to `sed` what actions you want performed on the file. Refer to the man pages for `sed` for a complete description of possible instructions. Here, I will only describe one, namely a global form of regular expression substitution. The instruction:

`s/`*regexp*`/`*substexp*`/g`

will perform the substitution of any string matching *regexp* with *substexp*. So, if you want to replace every instance of, say, `rp56` in a file `f2` by `<ric netid>`, you could use:

```
sed -e 's/rp65/<ric netid>/g' f2
```

and if file `f2` contains:

```
this is rp65 a line
this rp65 is also a line
```

we get the following:

```
babbage% sed -e 's/rp65/<ric netid>/g' f
this is <ric netid> a line
this <ric netid> is also a line
```

More complex substitutions can occur, and I'll refer you again to the man pages for `sed` for more information.