

CSCI 4320 (Principles of Operating Systems), Fall 2006

Homework 5

Assigned: November 29, 2006.

Due: December 6, 2006, at 5pm. *Not accepted past 5pm December 7* (except for extra-credit problems).

Credit: 20 points.

1 Reading

Be sure you have read (or at least skimmed) Chapter 6.

2 Problems

Answer the following questions. You may write out your answers by hand or using a word processor or other program, but please submit hard copy, either in class or in my mailbox in the department office.

1. (5 points) Linux includes code to access several types of Windows filesystems, including FAT-32. So a system where one of the disk partitions holds a FAT-32 filesystem, one can configure Linux to access this filesystem through pathname `/windows/fat` for example. However, all the files in `/windows/fat` appear to be owned by user `root`, and attempts to change their ownership (with the `chown` command) fail with an error message “Operation not permitted”. What’s wrong?
2. (5 points) Consider a digital camera that records photographs in some non-volatile storage medium (e.g., flash memory). Photographs are recorded in sequence until the medium is full; at that point, the photographs are transferred to a hard disk and the camera’s storage is cleared. If you were implementing a file system for the camera’s storage, what strategy would you use for file allocation (contiguous, linked-list, etc.) and why? Notice that this camera does not have the ability to delete photographs from its storage one at a time, so your file system does not need to support that.
3. (5 points) The textbook describes more than one strategy for keeping track of free blocks in a file system (free blocks, bitmaps, and FATs). What would happen if whatever data structure is used to keep track of free blocks was lost because of a system crash — is there a way to recover, or do you have to just reformat the disk and hope you backed up any really important files? Answer separately for MS-DOS FAT-16 and UNIX V7 filesystems.
4. (5 points) Consider a UNIX filesystem (as described in section 6.4.5) in which each i-node contains 10 direct entries, one single-direct entry, one double-indirect entry, and one triple-indirect entry. If a block is 1KB (1024 bytes) and a disk addresses is 4 bytes, what is the maximum file size, in KB? (*Hint:* Use the blocksize and size of disk addresses to determine how many entries each indirect block contain.)

3 Programming Problems

Do as many of the following programming problems as you like (none, all, or some). You will end up with at least one code file per problem. Submit your program source (and any other needed files) by sending mail to `bmassing@cs.trinity.edu`, with each file as an attachment. Please use a subject line that mentions the course number and the assignment (e.g., “csci 4320 homework 5”). You can develop your programs on any system that provides the needed functionality (something UNIX-like), but I will test them on one of the department’s Fedora Core 5 Linux machines, so you should probably make sure they work in that environment before turning them in.

1. (Up to 10 extra-credit points) Write a program that given a directory D , blocksize B , and maximum number of blocks M as command-line arguments prints out how many files in D and its subdirectories are of size B or less, how many are of size between B and $2B$, etc., up to size MB . (This might be useful in getting an idea of what size files are typical, so if you had a choice of blocksize you would know what choice might make the most sense.) Include directories and symbolic links (but count the size of the link and not the file/directory it links to). Here is sample output for running the program with $D = /lib$, $B = 512$, and $M = 20$, on Xena00:

Results for directory `/lib` with blocksize 512:

```
157 files of size      1 blocks
 87 files of size      2 blocks
 69 files of size      3 blocks
113 files of size      4 blocks
 38 files of size      5 blocks
 26 files of size      6 blocks
 33 files of size      7 blocks
 24 files of size      8 blocks
 19 files of size      9 blocks
 21 files of size     10 blocks
 14 files of size     11 blocks
 27 files of size     12 blocks
 18 files of size     13 blocks
  8 files of size     14 blocks
 34 files of size     15 blocks
 92 files of size     16 blocks
102 files of size     17 blocks
 59 files of size     18 blocks
 61 files of size     19 blocks
 69 files of size     20 blocks
2342 files of size     21 blocks or more
```

(Of course, you won’t be able to examine files in directories you don’t have access to. It’s okay in to just print error messages for files/directories you can’t access.)

To get maximum points, your program should be in C or C++ and make no use of system commands such as `ls`. (You can use another language, or even write a shell script, but you will get fewer points.) Library functions `opendir`, `readdir`, and `lstat` will probably be helpful. You might also be interested in functions `chdir` and `strerror`. These functions are described by man pages. (Remember also that `man -a foo` gives all man pages for `foo`. This can be helpful if there is both a command `foo` and a function `foo`.)

2. (Up to 10 extra-credit points) Write a program that given a directory D as a command-line argument prints all the “broken” symbolic links in D or any of its subdirectories — that is, symbolic links that point to a file that doesn’t exist. Here is sample output for running the program with $D = /users/bmassing/Local/HTML-Documents/CS4320/Homeworks/HW05/Problems$:

```
Broken symbolic links in /users/bmassing/Local/HTML-Documents/CS4320/Homeworks/HW05/Problems:
/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/TestData/barfoo
/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/TestData/foobar
```

(Again, you won’t be able to examine files in directories you don’t have access to, so just print error messages. You should be able to access everything in the above directory, however. If you want to create some test data of your own, remember that to make a symbolic link called `sym` pointing to `foo`, you type `ln -s foo sym`.)

To get maximum points, your program should be in C or C++ and make no use of system commands such as `ls`. (You can use another language, or even write a shell script, but you will get fewer points.) The library routines mentioned for the previous problem may be helpful.

3. (Up to 10 extra-credit points) Write a program that given a directory D as a command-line argument finds all the files in D or any of its subdirectories to which there are two or more hard links and prints, for each of them, all the paths within D that point to that file. Here is sample output for running the program with $D = /users/bmassing/Local/HTML-Documents/CS4320/Homeworks/HW05/Problems$:

```
Files with multiple hard links in /users/bmassing/Local/HTML-Documents/CS4320/Homeworks/HW05/Problems:
/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/homework/index.html
/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/homework/homework.html

/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/TestData/bbbb
/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/TestData/bbb
/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/TestData/b
/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/TestData/bb

/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/TestData/dd
/users/bmassing/Local/HTML-Documents/Classes/CS4320_2006fall/Homeworks/HW05/Problems/TestData/d
```

This output means that the two pathnames in the first group reference the same file, the four pathnames in the second group reference the same file, etc. Output can be in any order as long as paths that reference the same file are grouped together. (Again, you won’t be able to examine files in directories you don’t have access to, so just print error messages. You should be able to access everything in the above directory, however. If you want to create some test data of your own, remember that to make a hard link called `sym` pointing to `foo`, you type `ln foo sym`.)

To get maximum points, your program should be in C or C++ and make no use of system commands such as `ls`. (You can use another language, or even write a shell script, but you will get fewer points.) The library routines mentioned for the previous problems may be helpful.