

EEN118 LAB TEN

In this lab, you will be performing some important data-processing operations, specifically sorting a large database file. Sorting data is a very important operation in computing for many reasons. One of those reasons is that it makes the data more accessible to humans once it is printed (imagine trying to use a telephone directory in which the names do not appear in any particular order). Another reason is that it makes the data more quickly searchable by the computer.

There are many large data files to use for this lab, but you will not need to down-load them, the files can all be opened directly because you will be working with unix on rabbit again.

Important Note

This lab is to be run under a Unix system, not windows. You must also use only the standard C++ and Unix library files. Do not `#include library.h`. Try to remember what the standard includes are, but if you can't remember, the lab guys will remind you.

Look at the file `"/home/118/people1.txt"` with a text editor. You will see that it contains data for a number of people. Each line contains exactly seven items: a person's social security number, their date of birth (year-month-day), their first name, their last name, and their zip code. The seven items are separated by spaces, but no item will ever contain a space. Here is a sample from near the beginning of the file:

```
114680858 19670607 Matilda Vincent 71734
114930037 19471024 Desdemona Hanover 69743
115550206 19790110 Xanadu Perlman 12193
116520629 19630921 Alexander Hall 94976
117050976 19301016 David Lamprey 50895
119610646 19650202 Thomas Porlock 83582
120330928 19621126 Cary Cartman 90387
122460462 19620411 Bella Oldman 94495
123040628 19220213 Elizabeth Watson 90369
123580905 19230308 Gustav Hornswaggle 37568
125040813 19840613 Godfrey Tumor 55645
125610677 19580903 Gustav Trentham 73590
126470499 19521219 Justin Oddly 33458
127700250 19300616 Ursula Farnes 74163
129540334 19791114 Betty Eaton 25807
130020412 19361114 Maggie McIntosh 16628
132680826 19631118 Raul Kringle 53633
135040001 19490427 Arthropod Gravedigger 68521
135590854 19561012 Aloysius Pornman 27649
136870683 19521224 Alexandra Hanover 74067
138350783 19240518 Tammy Iverson 18015
139160698 19730717 Norbert Pringle 25257
139310201 19731224 Wotan Limburger 85408
139480872 19611205 Bella Fimburg 37692
139620582 19290612 Lillian Morlock 56064
```

This real file contains exactly 1000 lines of data.

1. *Read the Data*

Write a program that creates an array large enough to hold all the data, then reads all the data from the file into that array. Of course, it will have to be an array of `structs` that you will also need to define. Make your program close the file, then print out the first 10 items of data from the array, so that you can make sure everything was read correctly.

2. *Basic Search*

Make your program ask the user to enter a name. It should then search through the data in the array (don't read the file again), finding any entry with a matching name. Correct matches with either first or last name should be accepted. For every matching entry that is found, print out all six data items, everything that is known about each matching person.

Remember that if you use the `==` operator to compare strings, the test is case-sensitive. The user (i.e. you) will have to type the name exactly correctly, with capital letters in the right places.

Important: Good clean design will make this lab much easier. Write a separate function that searches the array, do not put all the work in `main`.

3. *Find the Oldest*

Modify your program so that after closing the file, instead of just printing the first ten items of data, it searches through 1000 of them all to find the oldest person represented. It should print all six items for the oldest person. In each file, the oldest person has a fairly distinctive name, so you'll have a good clue that you've got it right.

Important: As for part two, good clean design will make this lab much easier. Write a separate function that searches the array to find the oldest person, do not put all the work in `main`.

4. *Promote the Youngest*

For some unfathomable reason, the management wants the youngest person to occupy the first position in the array. Give yourself a function, just like the previous one, that finds the youngest person in the array. Once you know where the youngest person is, you can swap his or her data with the data already occupying the first position in the array. That way nothing is lost. Remember that the first position in an array is numbered zero, not one.

5. *Now Promote the Second Youngest.*

The management has now decided not only that the youngest person must occupy the first position in the array, but also that the second-youngest person must occupy the second position in the array. So, after searching for the youngest and moving their data to the front of the array, now search the remainder of the array (all except the first element), and move the youngest person you find (which must be the second youngest of all) into the second position of the array. Make sure you swap data, so that whoever was originally in the second position is not lost.

6. *More of the Same.*

The management are going to keep on adding requirements like this, next putting the third-youngest in the third position, then the fourth, then the fifth. There is no knowing when they will grow out of this petty obsession, so make things easier for yourself. Modify your search function so that it can be told how much of the array to search. That is, give it two `int` parameters (let's call them `a` and `b`); its job is now to search only the portion of the array between position `a` and position `b`, to find the youngest person therein. This makes it very easy to search the remainder of the array to find the second and third youngest.

7. *The Ultimate Demand.*

Now the management make their final demand. You are to repeat the process of moving the `n`th-youngest person into the `n`th position 1000 times. (remember, 1000 is the number of data records in the whole file).

This will result in the array being completely sorted. Do it, and check that it worked. Make your program prints the contents of the array after it has finished. Look at the output to make sure that everyone is printed in order of increasing age.

8. *Sorting the File.*

Once you have sorted the contents of the array, it might be a good idea to save the sorted data in a file. Make your program create a new file, and write all the contents of the array into that file in a sensible format. Use a text editor to look at the file and verify that it has the same format as the original file, and all the data is properly sorted.

9. *How Fast Is It?*

It is important to know how long computer operations are going to take when they have to work on a large amount of data. The standard Unix functions that give accurate timing are a little mysterious, so here is a little function that you can just copy and paste into your program. It requires two extra library files to be included, they are:

```
#include <time.h>
#include <sys/resource.h>
```

Here is the function

```
double get_cpu_time()
{ struct rusage ruse;
  getrusage(RUSAGE_SELF, &ruse);
  return ruse.ru_utime.tv_sec+ruse.ru_utime.tv_usec/1000000.0 +
         ruse.ru_stime.tv_sec+ruse.ru_stime.tv_usec/1000000.0; }
```

It returns the time as a double, and is accurate to a couple of milliseconds.

Use this function (twice) to time how long it takes the computer to sort the array of 1000 data items. Do not include the time it takes to read the file or the time it takes to write the new file, just the pure sorting time. Note the time that you observe.

Now you know how long it takes to sort a database of 1,000 items. How long do you think it would take to sort a database of 2,000 names? 3,000 names? 10,000 names?

Think about those questions, and work out what you believe the answer is. Then find out what the real answer is. The other files have exactly the same format as `people1.txt`, but are longer. `people(N).txt` contains N thousand data records. If your program was nicely written, it will be a few seconds' work to change the array size and make it read a different file. These are the files that are available:

```
/home/118/people1.txt  
/home/118/people2.txt  
/home/118/people3.txt  
/home/118/people5.txt  
/home/118/people10.txt  
/home/118/people20.txt  
/home/118/people30.txt  
/home/118/people50.txt  
/home/118/people100.txt
```

See how long it takes to sort these larger files, and compare the results to your predictions. If your predictions weren't substantially correct, make sure you understand why. You have just demonstrated a very important phenomenon of computing.