



Programming Project 1: Processing New York City Open Data

Due date: Sept. 29, 11:59PM EST.

You may discuss any of the assignments with your classmates and tutors (or anyone else) but **all work for all assignments must be entirely your own**. Any sharing or copying of assignments will be considered academic dishonesty. If you get significant help from anyone, you should acknowledge that help in your submitted documents, in which case your grade will adjusted in proportion to the part that you completed on your own. You are responsible for every line in your program: you need to know what it does and why. You should not use any data structures and features of C++ that have not been covered in this class or any of its prerequisite classes. If you have doubts about whether or not you are allowed to use certain language features, inquire first.

Summary

This assignment is designed to expose you to the use of open data. Wikipedia has a good description of open data: “open data is the idea that some data should be freely available to everyone to use and republish as they wish, without restrictions from copyright, patents or other mechanisms of control.” New York City (NYC) has a treasure trove of open datasets that are available to download. You can explore them at <https://nycopendata.socrata.com/>.

For this assignment, you will work with the dataset consisting of the most popular baby names given in New York City in the four years from 2011 through 2014. The data is stored in a comma-separated-values (CSV) file which has been downloaded and stored on our server for your convenience. You can find it in the directory

```
/data/biocs/b/student.accounts/cs235_sw/data/
```

You will be writing parts of a program that will process this dataset to answer various questions about the data, such as which name was the most popular name given to a baby girl in 2012 among Hispanic women, or what name was the most popular among all women in 2014. The exact types of questions will be discussed below.

This assignment gives you experience in working with existing code. You are given the interfaces to two different classes. In one case you have to write an implementation for that interface, and in the other, you are given the implementation as a binary file and you just need to write client code that uses the interface.

Objectives

The objectives of this first programming project are for you to master the following tasks:

- working with multi-file programs using separate compilation
- reading, understanding, and working with existing source code
- reading data from input files
- writing data to output files
- using and understanding command line arguments
- working with large data sets



- using arrays
- writing class interfaces and their implementations

Most, if not all, of the skills that you need to complete this project are based on the material covered in CSci 135, but it is possible that your instructor there did not cover that skill enough, resulting in your having to spend more time improving it for this assignment. It is even possible that there might be certain topics that your section did not cover at all, or that you did not get much practice on, in which case you are at a greater disadvantage, but you should view this as an opportunity to get up to speed now.

Program Input

Your program will read the data from a file whose name is given as a command line argument. It will also read the commands it must carry out from a file given as a second command line argument. There is no other input provided. In particular, there is no user interaction.

Input Files

The program is given the pathnames of two input text files as its command line arguments: The first is the pathname of a file containing the data set in CSV (comma separated values) format. The second is the pathname of a file containing a list of commands that your program needs to perform. For example, if my program is named `proj1_sweiss`, and the data is stored in a file in my working directory named `babynames.csv` and the commands are in a file named `test_commands`, then I would type

```
proj1_sweiss babynames.csv test_commands
```

to run my program on the `babynames.csv` dataset using the commands in the set of commands in `test_commands`. It is an error if someone attempts to run the program without two file names. The program must check that two arguments were given on the command line and that each file can be opened for reading by the program. If there are any errors, such as missing arguments, arguments that do not exist, or that cannot be opened, the program must display a short but descriptive error message that indicates the nature of the error (such as “the file `mycommands` cannot be opened.”) The message should be as specific as possible. It is not enough to write something like “could not open file.” The error message should be written to the standard error stream (`cerr` in C++, `stderr` in C).

Data File

The data file was downloaded from the webpage

<https://data.cityofnewyork.us/Health/Most-Popular-Baby-Names-by-Sex-and-Mother-s-Ethnic/25th-nujf> as a CSV file. The first line in the file contains the names of the six columns in the file:

```
BRTH_YR,GNDR,ETHCTY,NM,CNT,RNK
```

respectively, the birth year, gender, mother’s ethnicity, baby’s name, count, and rank. Each line after the first represents a single data record consisting of six, comma-separated values with the following meaning:

BRTH_YR The year in which babies with the given name were born

GNDR The gender of the babies with the given name

ETHCTY The ethnicity of the mothers of the babies with the given name

NM The actual baby name

CNT The number of babies with the given name, born in the given year, with the given gender and mother’s ethnicity

RNK The rank of that name of all baby names in the given year, given mother’s ethnicity, and baby gender, with 1 being the highest rank, i.e., the most frequently occurring name.



As an example, the line

```
2013,MALE,ASIAN AND PACIFIC ISLANDER,Dylan,45,24
```

represents the fact that 45 male babies born in 2013 to women who self-identified as “Asian and Pacific Islander” were named Dylan, and that of all male babies born in 2013 to such women, Dylan was the 24th most popular.

The downloaded data file had duplicate lines, although I do not know why they are there. For your convenience, I removed the duplicate lines, so that each line represents a unique (year,gender, ethnicity, name) record. The only space characters in the lines are those found between the words of the ethnicity field.

Commands File

The file that contains commands has a single command per line. Each line is in one of the three forms described below. The exact meaning of each command is specified in the section below entitled *Program Behavior*:

```
print tofilename
```

where *tofilename* is replaced by the name of a file to be created to hold the output.

```
sort how
```

where *how* is replaced by one of the words `byname` or `byrank`, whose meanings are explained below

```
getmax year gender ethnicity
```

where *year* is either a number between 2011 and 2014 or the digit 0,
gender is either M, F, or the digit 0,
ethnicity is either W, B, H, A, or the digit 0.

Notice that the above descriptions do not describe *what the program is supposed to do* when it reads a line, but merely whether a line is in a legal form. If it is not syntactically correct, the program must report the error on the standard error stream and advance to the next line, and if it is correct, the program must process the given command.

Output Files

The program writes all correct output to one or more output files. All error messages generated by your code should be written to the standard error stream (`cerr` in C++, `stderr` in C.)

Program Behavior

If the data file and commands file are opened successfully, the program must then process each command one after the other. Each command must be processed according to the description presented below. When all commands have been processed, any open files should be closed and the program terminated.

```
print
```

Syntax: `print tofilename`

This command causes the data to be written to the given file in whatever order it is stored in when the command is issued, one record per line. If the file already exists, it is overwritten. The records should be written one per line. For each record, the following information should be displayed in a field of exactly the width specified:

- `birth_year`, field width: 4 characters, left justified



- gender, field width: 6 characters, as either “MALE” or “FEMALE”, left justified
- ethnicity, field width: 26 characters, exactly as contained in original data file, left justified
- name, field width: 16 characters, left justified, in UPPER CASE
- count, field width: 6 characters, right justified
- rank, field width: 6 characters, right justified

The entries should be separated by tab characters (a total of 5 tabs). The field widths must be exactly as specified, no more, no less.

sort

Syntax: `sort byname | sort byrank`

If `byname`, the records are sorted with four keys: first by birth year in ascending order, then by gender with FEMALE < MALE, then by ethnicity in ascending alphabetical order, then by name, in ascending alphabetical order.

If `byrank`, the records are sorted with four keys: first by birth year in ascending order, then by gender with FEMALE < MALE, then by ethnicity in ascending alphabetical order, then by rank, in descending numeric order.

In both cases, the records are sorted in place in the object storing the records.

getmax

Syntax: `getmax year gender ethnicity`

where *year* is either a number between 2011 and 2014 or the digit 0,

gender is either M, F, or the digit 0,

ethnicity is either W, B, H, A, or the digit 0,

The `getmax` command is given three mandatory arguments. If no argument is the digit 0, then the `getmax` command finds the name whose count is the greatest for the given year, gender, and ethnicity and ***appends the entire record to a file named results in the current working directory.*** using the same output format as is used in the `print` command described above.

The digit 0 is a ***wild-card*** for the given field. This means that `getmax` ignores that field when looking for the record with greatest count. If there are multiple records with a maximal count, then it writes all of them to the file, using the same format as the `print` command. For example, if the *year* is 2011 but *gender* is 0 and *ethnicity* is 0, then `getmax` writes the record of the baby born in 2011 with greatest count value, and if there are ties, then all records that are tied at that count are written. If the *year* is 2011 and the *gender* is F, but *ethnicity* is 0, then `getmax` writes the record of female babies born in 2011 with the greatest count.

Implementation Constraints

You are provided a class named `Command` that represents a single command. That class provides the functionality to obtain the next command from the input command file and parse it. It stores the command's parameters in private data members but it also provides methods to get the type of command and the command's arguments. If it encounters any bad commands, it returns appropriate error values. The `Command` class interface and its implementation are stored on the server in two separate files named `command.h` and `command.o`. In particular you are not given the source code from which `command.o` was compiled. You will not have access to the source code of `command.o`, because it is not necessary. Your main program must use this `Command` class as its means of obtaining the commands. You need to read the `command.h` file to see what public methods it provides and design your program to use these methods.

You are also given the interface to a class named `Babynome`. The `Babynome` class represents a single Babynome record, i.e., the data contained in a single row of the CSV file. The interface includes a few public methods



and four friend functions. If friend functions are a new concept for you, do not fear - they will be explained in class. Your job is to write an implementation for this interface. You are not allowed to alter the interface in any way. The interface is included at the end of this assignment but it is also contained in the header file `babyname.h` on the server, the exact location to be given in class. Your implementation must be named `babyname.cpp` or `babyname.cxx`.

Your main program is responsible for opening input files, reading from and writing to files, and closing files. It is also responsible for checking the correctness of the command line arguments and issuing any error messages related to the command line. If you are not experienced in using command line arguments, please read my notes on the topic here:

<http://www.compsci.hunter.cuny.edu/~sweiss/resources/cmmdlineargs.pdf>.

Your program should create a class that represents the entire data set. This class should store the data in an array (or a vector), and provide methods for populating the array from an input stream, sorting the array, writing the array data to an output stream and getting the various maximal values needed by the `getmax` command. It should hide all details from the client (main) and should not read from or write to a file. I repeat, *it does not perform any file I/O, instead writing to and from streams*.

In order to sort with multiple keys, you will need to use a stable sorting algorithm. The simple *insertion sort* algorithm is a stable sorting algorithm and I strongly suggest that you use it, otherwise you will make things more complicated than they need to be. If you did not learn insertion sort in the prerequisite class, then look at *Listing 11-3* in *Walls and Mirrors* or look it up in another textbook. You were supposed to learn that a sorting algorithm is *stable* if it preserves the relative order of equal keys. If you never did, then read more from your textbook.

All class interfaces and implementations are to be in separate files from each other and from the main program. Thus, your project will have at least the following files: a main file, an interface for the `Command` class, an object file for the `Command` class, an interface and implementation for the `Babyname` class, an interface for the dataset class (or whatever you want to call it), and an implementation file for that class.

Advice

You should start right away! This program does not require you to write much code (more than you are used to perhaps, but less than future assignments will require), but it will take some time.

You should modularize your design so that you can test it regularly. You can implement methods that perform one task at a time. This way, if you run out of time, at least parts of your program will be functioning properly.

You should backup your code after you have spent some time working on it. Save it to a flash drive, email it to yourself, upload it to your Google drive, anything that gives you a second (or maybe third copy). Murphy's Law is always in force - things can go wrong at the last minute, so do not take a chance on not having something to submit.

Programming Rules

Your program must conform to the programming rules described in the *Programming Rules* document on the course website. It is to be your own work alone.

Grading Rubric

The program will be graded based on the following rubric.

- Compiling on a `cs1lab` machine: 30%
- Correctness of the implementation 40%
- Performance 5%
- Design (modularity and organization) 10%
- Documentation: 10%



- Style and proper naming: 5%

This implies that a program that does not compile on a lab computer cannot receive more than 70 points. Some implementations might be more efficient in terms of running time than others. Although you will not have learned enough to write this program in a very efficient way, there are certain programming decisions that can lead to very poor performance or to acceptable performance. The performance component of 5% will be given if the program is reasonably efficient.

Submitting the Assignment

This assignment is due by the end of the day (i.e. 11:59PM, EST) on September 29, 2016. Hopefully your experience submitting the first assignment has educated you so that submitting the second will be easy. There is a directory in the CSci Department network whose full path name is `/data/biocs/b/student.accounts/cs235_sw/projects/project1`.

That is where your submission will go. In order to put it there you must follow these steps. First you will create a zip file containing your source code, including the header file `command.h` and the object file `command.o`. To do this, create a directory named `proj1_username` where `username` is replaced by your login name, and put all of your files into that directory. Do not place anything else into this directory. You will lose 1 point for each file that does not belong there. With all files in your directory, change directory so that `proj1_username` is in the current directory and run the command

```
zip -r proj1_username.zip ./proj1_username
```

This will compress all of your files into the file named `proj1_username.zip`. Then you will use the program `submit235project` to deposit the zip file into the directory. The program requires two arguments: the number of the assignment and the pathname of your zip file. For example, if your username on our system is `Bugs.Bunny` and your zip file is named `proj1_Bugs.Bunny.zip` and it is in your current working directory (say your home directory) then you would type

```
/data/biocs/b/student.accounts/cs235_sw/bin/submit235project 1 proj1_Bugs.Bunny.zip
```

The program will create the file

```
/data/biocs/b/student.accounts/cs235_sw/projects/project1/proj1_Bugs.Bunny.zip.
```

You will not be able to read this file, nor will anyone else except for me. But you can verify that the command succeeded by typing the command

```
ls -l /data/biocs/b/student.accounts/cs235_sw/projects/project1
```

and making sure you see your file and that its size is the same as the size of the original `proj1_Bugs.Bunny.zip`.

If you decide to make any changes and resubmit, just run the command again and it will replace the old file with the new one. Do not try to put your file into this directory in any other way - you will be unable to do this.

If you put a solution there and then decide to change it before the deadline, just replace it by the new version. Once the deadline has passed, you cannot do this. I will grade whatever version is there at the end of the day on the due date. You cannot resubmit the program after the due date.



The Babyname Class Interface

```
#include "command.h"

class Babyname
{
public:
    /** Babyname() – default constructor */
    Babyname();

    /** set() – set method for the class
     * assigns given values to the member variables, while validating values
     * method ensures that all parameters are within the allowed ranges
     * as defined in the specification
     * @param int year [in] birth year
     * @param Gender gender [in] baby's gender
     * @param Ethnicity ethnicity [in] mother's ethnicity
     * @param string name [in] name of baby
     * @param int count [in] frequency (see comments)
     * @param int rank [in] rank (see comments)
     * @pre none
     * @post object is assigned all valid values
     */
    void set( int year,
             Gender gender,
             Ethnicity ethnicity,
             string name,
             int count,
             int rank
             );

    /** get() – get method for the class
     * retrieves values of the member variables, and passes to corresponding
     * parameters.
     * @param int year [out] birth year
     * @param Gender gender [out] baby's gender
     * @param Ethnicity ethnicity [out] mother's ethnicity
     * @param string name [out] name of baby
     * @param int count [out] frequency (see comments)
     * @param int rank [out] rank (see comments)
     * @pre none
     * @post object is assigned all valid values
     */
    void get( int & year,
             Gender & gender,
             Ethnicity & ethnicity,
             string & name,
             int & count,
             int & rank
             );

    /** print() – prints the baby name data onto the given ostream
     * assuming it is already open for writing, in the format specified
     * in the assignment.
     * @param ostream s [inout] ostream opened for writing
```



```
* @pre  the object has valid data
* @post  if the object has valid data, then it is written to ostream
*        in the specified format and the ostream is updated
*/
void print ( ostream & s );

/** Several friend functions that perform similar tasks;
 * Each is given two Babyname objects by constant reference and
 * returns true or false depending on whether one is less than or equal to
 * the other.
 * is_less_by_rank(lt , rt) is true if and only if:
 *   birth year of lt < birth year of rt
 *   or they are equal and gender of lt < gender of rt ,
 *   or they are equal and ethnicity of lt < ethnicity of rt ,
 *   or they are equal and rank of lt < rank of rt.
 * is_less_by_name( lt , rt) has the same logic except in the last
 *   case when name replaces rank.
 * The two is_equal functions return true if and only if year , gender ,
 *   and ethnicity are equal and in the case of is_equal_name , the names
 *   are identical , and in the case of rank , the ranks are identical.
 */
friend  bool is_less_by_rank(const Babyname & lt , const Babyname & rt );
friend  bool is_equal_rank  (const Babyname & lt , const Babyname & rt );
friend  bool is_less_by_name(const Babyname & lt , const Babyname & rt );
friend  bool is_equal_name  (const Babyname & lt , const Babyname & rt );

private:
    int      birth_year;           // 2011 through 2014
    Gender   baby_gender;         // male or female
    Ethnicity mother_ethnicity;   // asian , black , hispanic , or white
    string   baby_name;          // any valid string
    int      name_count;         // >= 0
    int      name_rank;          // > 0
};
```