
Problem Set 2: Crypto

This is CS50. Harvard University. Fall 2015.

Table of Contents

Objectives	1
Recommended Reading	1
Academic Honesty	2
Reasonable	2
Not Reasonable	3
Assessment	4
Getting Ready	4
Getting Started	5
Initializing	6
Hail, Caesar!	7
Parlez-vous français?	12
How to Submit	14
Step 1 of 2	14
Step 2 of 2	15

Questions? Head to [CS50 Discuss](https://cs50.net/discuss)¹ or join classmates at [office hours](https://cs50.net/hours)²!

Objectives

- Become better acquainted with functions and libraries.
- Dabble in cryptography.

Recommended Reading

- Pages 11 – 14 and 39 of <http://www.howstuffworks.com/c.htm>.
- Chapters 6, 7, 10, 17, 19, 21, 22, 30, and 32 of *Absolute Beginner's Guide to C*.
- Chapters 7, 8, and 10 of *Programming in C*.

¹ <https://cs50.net/discuss>

² <https://cs50.net/hours>

Academic Honesty

This course's philosophy on academic honesty is best stated as "be reasonable." The course recognizes that interactions with classmates and others can facilitate mastery of the course's material. However, there remains a line between enlisting the help of another and submitting the work of another. This policy characterizes both sides of that line.

The essence of all work that you submit to this course must be your own. Collaboration on problem sets is not permitted except to the extent that you may ask classmates and others for help so long as that help does not reduce to another doing your work for you. Generally speaking, when asking for help, you may show your code to others, but you may not view theirs, so long as you and they respect this policy's other constraints. Collaboration on quizzes is not permitted at all. Collaboration on the course's final project is permitted to the extent prescribed by its specification.

Below are rules of thumb that (inexhaustively) characterize acts that the course considers reasonable and not reasonable. If in doubt as to whether some act is reasonable, do not commit it until you solicit and receive approval in writing from the course's heads. Acts considered not reasonable by the course are handled harshly. If the course refers some matter for disciplinary action and the outcome is punitive, the course reserves the right to impose local sanctions on top of that outcome that may include an unsatisfactory or failing grade for work submitted or for the course itself.

If you commit some act that is not reasonable but bring it to the attention of the course's heads within 72 hours, the course may impose local sanctions that may include an unsatisfactory or failing grade for work submitted, but the course will not refer the matter for further disciplinary action except in cases of repeated acts.

Reasonable

- Communicating with classmates about problem sets' problems in English (or some other spoken language).
- Discussing the course's material with others in order to understand it better.
- Helping a classmate identify a bug in his or her code at office hours, elsewhere, or even online, as by viewing, compiling, or running his or her code, even on your own computer.

- Incorporating snippets of code that you find online or elsewhere into your own code, provided that those snippets are not themselves solutions to assigned problems and that you cite the snippets' origins.
- Reviewing past semesters' quizzes and solutions thereto.
- Sending or showing code that you've written to someone, possibly a classmate, so that he or she might help you identify and fix a bug.
- Sharing snippets of your own code online so that others might help you identify and fix a bug.
- Turning to the web or elsewhere for instruction beyond the course's own, for references, and for solutions to technical difficulties, but not for outright solutions to problem set's problems or your own final project.
- Whiteboarding solutions to problem sets with others using diagrams or pseudocode but not actual code.
- Working with (and even paying) a tutor to help you with the course, provided the tutor does not do your work for you.

Not Reasonable

- Accessing a solution to some problem prior to (re-)submitting your own.
- Asking a classmate to see his or her solution to a problem set's problem before (re-)submitting your own.
- Decompiling, deobfuscating, or disassembling the staff's solutions to problem sets.
- Failing to cite (as with comments) the origins of code or techniques that you discover outside of the course's own lessons and integrate into your own work, even while respecting this policy's other constraints.
- Giving or showing to a classmate a solution to a problem set's problem when it is he or she, and not you, who is struggling to solve it.
- Looking at another individual's work during a quiz.
- Paying or offering to pay an individual for work that you may submit as (part of) your own.
- Providing or making available solutions to problem sets to individuals who might take this course in the future.
- Searching for, soliciting, or viewing a quiz's questions or answers prior to taking the quiz.

- Searching for or soliciting outright solutions to problem sets online or elsewhere.
- Splitting a problem set's workload with another individual and combining your work.
- Submitting (after possibly modifying) the work of another individual beyond allowed snippets.
- Submitting the same or similar work to this course that you have submitted or will submit to another.
- Submitting work to this course that you intend to use outside of the course (e.g., for a job) without prior approval from the course's heads.
- Using resources during a quiz beyond those explicitly allowed in the quiz's instructions.
- Viewing another's solution to a problem set's problem and basing your own solution on it.

Assessment

Your work on this problem set will be evaluated along four axes primarily.

Scope

To what extent does your code implement the features required by our specification?

Correctness

To what extent is your code consistent with our specifications and free of bugs?

Design

To what extent is your code written well (i.e., clearly, efficiently, elegantly, and/or logically)?

Style

To what extent is your code readable (i.e., commented and indented with variables aptly named)?

All students, whether or not taking the course for a letter grade, must ordinarily submit this and all other problem sets to be eligible for a satisfactory grade unless granted an exception in writing by the course's heads.

Getting Ready

First, curl up with the below shorts on loops, functions, Caesar's cipher, and command-line arguments. If you happen to see (and are confused by!) `char *` in these and other shorts, know for now that `char *` simply means `string`. But more on that soon!

<https://www.youtube.com/watch?v=HHmiHx7GGLE>

Be sure you're reasonably comfortable answering the below when it comes time to submit this problem set's form!

- How does a while loop differ from a do-while loop? When is the latter particularly useful?
- What does `undeclared identifier` usually indicate if outputted by `make` (or, really, `clang`)?
- Why is Caesar's cipher not very secure?
- What's a function?
- Why bother writing functions when you can just copy and paste code as needed?

Next, take a self-paced tour through Week 2's examples, the source code for which can be found at <http://cdn.cs50.net/2015/fall/lectures/2/m/src2m/> and <http://cdn.cs50.net/2015/fall/lectures/2/w/src2w/>, fast-forwarding through any programs with which you're already comfortable:

<https://www.youtube.com/watch?v=9zoRoz8Pq4E&list=PLhQjrBD2T380sc-fXwl1sviA-twxFduVU>

You might also find these shorts on arrays, return values, and Vigenère's cipher of interest as well!

<https://www.youtube.com/watch?v=7mOJN1c1JEo>

Getting Started

Alright, here we go again!

Log into cs50.io³ and execute

```
.....  
update50  
.....
```

within a terminal window to make sure your workspace is up-to-date. If you somehow closed your terminal window (and can't find it!), make sure that **Console** is checked under the **View** menu, then click the green, circled plus (+) in CS50 IDE's bottom half, then select **New Terminal**. If you need a hand, do just ask via [CS50 Discuss](https://cs50.net/discuss)⁴!

³ <https://cs50.io/>

⁴ <https://cs50.net/discuss>

Next, execute

```
mkdir ~/workspace/pset2
```

at your prompt in order to make a directory called `pset2` in your `workspace` directory. Take care not to overlook the space between `mkdir` and `~/workspace/pset2` or any other character for that matter! Keep in mind that `~` denotes your home directory, `~/workspace` denotes a directory called `workspace` therein, and `~/workspace/pset2` denotes a directory called `pset2` within `~/workspace`.

Now execute

```
cd ~/workspace/pset2
```

to move yourself into (i.e., open) that directory. Your prompt should now resemble the below.

```
username@ide50:~/workspace/pset2 $
```

If not, retrace your steps and see if you can determine where you went wrong. You can actually execute

```
history
```

at the prompt to see your last several commands in chronological order if you'd like to do some sleuthing. You can also scroll through the same one line at a time by hitting your keyboard's up and down arrows; hit Enter to re-execute any command that you'd like. If still unsure how to fix, remember that [CS50 Discuss⁵](#) is your friend!

All of the work that you do for this problem set must ultimately reside in your `pset2` directory for submission.

Initializing

Alright, let's get more comfortable with `string`.

⁵ <https://cs50.net/discuss>

Write, in a file called `initials.c`, a program that prompts a user for their name (using `GetString` to obtain their name as a `string`) and then outputs their initials in uppercase with no spaces or periods, followed by a newline (`\n`) and nothing more. You may assume that the user's input will contain only letters (uppercase and/or lowercase) plus single spaces between words. Folks like `Joseph Gordon-Levitt`, `Conan O'Brien`, and `David J. Malan` won't be using your program. (If only!)

So that we can automate some tests of your code, your program must behave per the examples below. Assumed that the underlined text is what some user has typed.

```
.....
username@ide50:~/workspace/pset2 $ ./initials
Zamyla Chan
ZC
username@ide50:~/workspace/pset2 $ ./initials
robert thomas bowden
RTB
.....
```

If you'd like to check the correctness of your program with `check50`, you may execute the below.

```
.....
check50 2015.fall.pset2.initials initials.c
.....
```

And if you'd like to play with the staff's own implementation of `initials`, you may execute the below.

```
.....
~cs50/pset2/initials
.....
```

Hail, Caesar!

Recall from David DiCiurcio's short that Caesar's cipher encrypts (i.e., scrambles in a reversible way) messages by "rotating" each letter by k positions, wrapping around from `Z` to `A` as needed (cf. http://en.wikipedia.org/wiki/Caesar_cipher). In other words, if p is some plaintext (i.e., an unencrypted message), p_i is the i^{th} character in p , and k is a secret key (i.e., a non-negative integer), then each letter, c_i , in the ciphertext, c , is computed as:

$$c_i = (p_i + k) \% 26$$

This formula perhaps makes the cipher seem more complicated than it is, but it's really just a nice way of expressing the algorithm precisely and concisely. And computer scientists love precision and, er, concision.

For example, suppose that the secret key, k , is 13 and that the plaintext, p , is "Be sure to drink your Ovaltine!" Let's encrypt that p with that k in order to get the ciphertext, c , by rotating each of the letters in p by 13 places, whereby:

.....
Be sure to drink your Ovaltine!
.....

becomes:

.....
Or fher gb qevax lbhe Binygvar!
.....

We've deliberately printed the above in a monospaced font so that all of the letters line up nicely. Notice how `O` (the first letter in the ciphertext) is 13 letters away from `B` (the first letter in the plaintext). Similarly is `r` (the second letter in the ciphertext) 13 letters away from `e` (the second letter in the plaintext). Meanwhile, `f` (the third letter in the ciphertext) is 13 letters away from `s` (the third letter in the plaintext), though we had to wrap around from `z` to `a` to get there. And so on. Not the most secure cipher, to be sure, but fun to implement!

Incidentally, a Caesar cipher with a key of 13 is generally called ROT13 (cf. <http://en.wikipedia.org/wiki/ROT13>). In the real world, though, it's probably best to use ROT26, which is believed to be twice as secure⁶.

Anyhow, your next goal is to write, in `caesar.c`, a program that encrypts messages using Caesar's cipher. Your program must accept a single command-line argument: a non-negative integer. Let's call it k for the sake of discussion. If your program is executed without any command-line arguments or with more than one command-line argument, your program should yell at the user and return a value of `1` (which tends to signify an error) immediately as via the statement below:

.....
`return 1;`
.....

⁶ <http://www.urbandictionary.com/define.php?term=ROT26>

Otherwise, your program must proceed to prompt the user for a string of plaintext and then output that text with each alphabetical character "rotated" by k positions; non-alphabetical characters should be outputted unchanged. After outputting this ciphertext, your program should exit, with `main` returning `0`, as via the statement below:

```
return 0;
```

If you don't explicitly return an `int` from within `main`, `0` is actually returned for you automatically. (Indeed, per its "return type," `main` does need to return an `int`. But more on that another time.) Now that you're returning `1` explicitly to signify errors, it's best to return `0` (by convention) explicitly to signify success. Whereas `0` generally represents success, any non-`0` `int` generally represents an error. That way, you can represent (gasp) upwards of four billion errors (since an `int` is generally 32 bits)!

Anyhow, even though there exist only 26 letters in the English alphabet, you may not assume that k will be less than or equal to 26; your program should work for all non-negative integral values of k less than $2^{31} - 26$. (In other words, you don't need to worry if your program eventually breaks if the user chooses a value for k that's too big or almost too big to fit in an `int`. Now, even if k is greater than 26, alphabetical characters in your program's input should remain alphabetical characters in your program's output. For instance, if k is 27, `A` should not become `[` even though `[` is 27 positions away from `A` in ASCII; `A` should become `B`, since 27 modulo 26 is 1, as a computer scientists might say. In other words, values like $k = 1$ and $k = 27$ are effectively equivalent.

Your program must preserve case: capitalized letters, though rotated, must remain capitalized letters; lowercase letters, though rotated, must remain lowercase letters.

Where to begin? Well, this program needs to accept a command-line argument, k , so this time you'll want to declare `main` with:

```
int main(int argc, string argv[])
```

Recall that `argv` is an "array" of strings. You can think of an array as row of gym lockers, inside each of which is some value (and maybe some socks). In this case, inside each such locker is a `string`. To open (i.e., "index into") the first locker, you use syntax like `argv[0]`, since arrays are "zero-indexed." To open the next locker, you use syntax like `argv[1]`. And so on. Of course, if there are n lockers, you'd better stop opening lockers

once you get to `argv[n - 1]`, since `argv[n]` doesn't exist! (That or it belongs to someone else, in which case you still shouldn't open it.)

And so you can access `k` with code like

```
string k = argv[1];
```

assuming it's actually there! Recall that `argc` is an `int` that equals the number of strings that are in `argv`, so you'd best check the value of `argc` before opening a locker that might not exist! Ideally, `argc` will be `2`. Why? Well, recall that inside of `argv[0]`, by default, is a program's own name. So `argc` will always be at least `1`. But for this program you want the user to provide a command-line argument, `k`, in which case `argc` should be `2`. Of course, if the user provides more than one command-line argument at the prompt, `argc` could be greater than `2`, in which case it's time for some yelling.

Now, just because the user types an integer at the prompt, that doesn't mean their input will be automatically stored in an `int`. Au contraire, it will be stored as a `string` that just so happens to look like an `int`! And so you'll need to convert that `string` to an actual `int`. As luck would have it, a function, `atoi`⁷, exists for exactly that purposes. Here's how you might use it:

```
int k = atoi(argv[1]);
```

Notice, this time, we've declared `k` as an actual `int` so that you can actually do some arithmetic with it. Ah, much better. Incidentally, you can assume that the user will only type integers at the prompt. You don't have to worry about them typing, say, `foo`, just to be difficult (even though the staff's solution does catch such); `atoi` will just return `0` in such cases.

Because `atoi` is declared in `stdlib.h`, you'll want to `#include` that header file atop your own code. (Technically, your code will compile without it there, since we already `#include` it in `cs50.h`. But best not to trust another library to `#include` header files you know you need.)

Okay, so once you've got `k` stored as an `int`, you'll need to ask the user for some plaintext. Odds are CS50's own `GetString` can help you with that.

⁷ <https://reference.cs50.net/stdlib.h/atoi>

Once you have both `k` and some plaintext, it's time to encrypt the latter with the former. Recall that you can iterate over the characters in a string, printing each one at a time, with code like the below:

```
for (int i = 0, n = strlen(p); i < n; i++)
{
    printf("%c", p[i]);
}
```

In other words, just as `argv` is an array of strings, so is a `string` an array of chars. And so you can use square brackets to access individual characters in strings just as you can individual strings in `argv`. Neat, eh? Of course, printing each of the characters in a string one at a time isn't exactly cryptography. Well, maybe technically if k is 0. But the above should help you help Caesar implement his cipher! For Caesar!

Incidentally, you'll need to `#include` yet another header file in order to use `strlen`⁸.

And Zamyala has some tips for you as well:

<https://www.youtube.com/watch?v=V6IDxI-3WAA>

So that we can automate some tests of your code, your program must behave per the below. Assumed that the underlined text is what some user has typed.

```
username@ide50:~/workspace/pset2 $ ./caesar 13
Be sure to drink your Ovaltine!
Or fher gb qevax lbhe Binygvar!
```

Besides `atoi`, you might find some handy functions documented at <https://reference.cs50.net/> under `ctype.h` and `stdlib.h`. For instance, `isdigit` sounds interesting. And, with regard to wrapping around from `Z` to `A` (or `z` to `a`), don't forget about `%`, C's modulo operator. You might also want to check out <http://asciitable.com/>, which reveals the ASCII codes for more than just alphabetical characters, just in case you find yourself printing some characters accidentally.

If you'd like to check the correctness of your program with `check50`, you may execute the below.

⁸ <https://reference.cs50.net/string.h/strlen>

check50 2015.fall.pset2.caesar caesar.c

And if you'd like to play with the staff's own implementation of `caesar`, you may execute the below.

~cs50/pset2/caesar

BTW, `uggc://jjj.lbhghor.pbz/jngpu?i=bUt5FWLEUN0`.

Parlez-vous français?

Well that last cipher was hardly secure. Fortunately, [Nate's short on Vigenère's cipher⁹](#), there's a more sophisticated algorithm out there. Suffice it to say it's French, per http://en.wikipedia.org/wiki/Vigen%C3%A8re_cipher. Though do not be misled by the article's discussion of a tabula recta. Each c_i can be computed with relatively simple arithmetic! You do not need a two-dimensional array.

Vigenère's cipher improves upon Caesar's by encrypting messages using a sequence of keys (or, put another way, a keyword). In other words, if p is some plaintext and k is a keyword (i.e., an alphabetical string, whereby `A` and `a` represent 0, while `Z` and `z` represent 25), then each letter, c_i , in the ciphertext, c , is computed as:

$$c_i = (p_i + k_j) \% 26$$

Note this cipher's use of k_j as opposed to just k . And recall that, if k is shorter than p , then the letters in k must be reused cyclically as many times as it takes to encrypt p .

Your final challenge this week is to write, in `vigenere.c`, a program that encrypts messages using Vigenère's cipher. This program must accept a single command-line argument: a keyword, k , composed entirely of alphabetical characters. If your program is executed without any command-line arguments, with more than one command-line argument, or with one command-line argument that contains any non-alphabetical character, your program should complain and exit immediately, with `main` returning `1` (thereby signifying an error that our own tests can detect). Otherwise, your program must proceed to prompt the user for a string of plaintext, p , which it must then encrypt according

⁹ <https://youtu.be/9zASwVoshiM>

to Vigenère's cipher with k , ultimately printing the result and exiting, with `main` returning `0`.

As for the characters in k , you must treat `A` and `a` as 0, `B` and `b` as 1, ..., and `Z` and `z` as 25. In addition, your program must only apply Vigenère's cipher to a character in p if that character is a letter. All other characters (numbers, symbols, spaces, punctuation marks, etc.) must be outputted unchanged. Moreover, if your code is about to apply the j^{th} character of k to the i^{th} character of p , but the latter proves to be a non-alphabetical character, you must wait to apply that j^{th} character of k to the next alphabetical character in p ; you must not yet advance to the next character in k . Finally, your program must preserve the case of each letter in p .

Not sure where to begin? As luck would have it, this program's pretty similar to `caesar`! Only this time, you need to decide which character in k to use as you iterate from character to character in p .

And here's Zamyra again with some tips:

<https://www.youtube.com/watch?v=Uma2HZMPm2M>

So that we can automate some tests of your code, your program must behave per the below; highlighted in bold are some sample inputs.

```
.....  
username@ide50:~/workspace/pset2 ./vigenere bacon  
Meet me at the park at eleven am  
Negh zf av huf pcfx bt gzrwep oz  
.....
```

How to test your program, besides predicting what it should output, given some input? Well, recall that we're nice people. And so we've written a program called `devigenere` that also takes one and only one command-line argument (a keyword) but whose job is to take ciphertext as input and produce plaintext as output.

To use our program, execute

```
.....  
~cs50/pset2/devigenere k  
.....
```

at your prompt, where k is some keyword. Presumably you'll want to paste your program's output as input to our program; be sure, of course, to use the same key. Note that you do not need to implement `devigenere` yourself, only `vigenere`.

If you'd like to check the correctness of your program with `check50`, you may execute the below.

```
check50 2015.fall.pset2.vigenere vigenere.c
```

And if you'd like to play with the staff's own implementation of `vigenere`, you may execute the below.

```
~cs50/pset2/vigenere
```

How to Submit

Step 1 of 2

1. When ready to submit, log into [CS50 IDE¹⁰](#).
2. Toward CS50 IDE's top-left corner, within its "file browser" (not within a terminal window), control-click or right-click your `initials.c` file (that's within your `pset2` directory) and then select **Download**. You should find that your browser has downloaded `initials.c`.
3. Repeat for `caesar.c`.
4. Repeat for `vigenere.c`.
5. In a separate tab or window, log into [CS50 Submit¹¹](#), logging in if prompted.
6. Click **Submit** toward the window's top-left corner.
7. Under **Problem Set 2** on the screen that appears, click **Upload New Submission**.
8. On the screen that appears, click **Add files....** A window entitled **Open Files** should appear.
9. Navigate your way to `initials.c`. Odds are it's in your **Downloads** folder or wherever your browser downloads files by default. Once you find `initials.c`, click it once to select it, then click **Open** (or the like).
10. Click **Add files...** again, and a window entitled **Open Files** should appear again.

¹⁰ <https://cs50.io/>

¹¹ <https://cs50.harvard.edu/submit>

- 11 Navigate your way to `caesar.c` as before. Click it once to select it, then click **Open** (or the like).
- 12 Navigate your way to `vigenere.c` as before. Click it once to select it, then click **Open** (or the like).
- 13 Click **Start upload** to upload all of your files at once to CS50's servers.
- 14 On the screen that appears, you should see a window with **No File Selected**. If you move your mouse toward the window's lefthand side, you should see a list of the files you uploaded. Click each to confirm the contents of each. (No need to click any other buttons or icons.) If confident that you submitted the files you intended, consider your source code submitted! If you'd like to re-submit different (or modified) files, simply return to [CS50 Submit](#)¹² and repeat these steps. You may re-submit as many times as you'd like; we'll grade your most recent submission, so long as it's before the deadline.

Step 2 of 2

Head to <https://forms.cs50.net/2015/fall/psets/2/> where a short form awaits. Once you have submitted that form (as well as your source code), you are done! If you end up resubmitting your files (per step 1 of 1), no need to resubmit the form.

This was Problem Set 2.

¹² <https://cs50.harvard.edu/submit>