

CS2112—Fall 2014

Assignment 2

Ciphers and Encryption

Due: Tuesday, September 16, 11:59PM

Design Overview due: Monday, September 8, 11:59PM

In this assignment, you will build multiple systems for the encryption and decryption of text. The first part focuses on simple ciphers and cipher cracking, while the second part explores the widely used RSA public-key encryption algorithm, which you use multiple times a day. You are tasked with creating a command-line application that can generate, save, and use the ciphers you build. Your implementation of the system should use inheritance to share code between different ciphers.

0 Updates

- Clarified requirements on handling large input. (9/5)
- Updated Main.java by inserting a break after line 151. (9/7)
- Modifications to Main.java are allowed, but not necessary (9/11)

1 Instructions

1.1 Grading

Solutions will be graded on both correctness and style. A correct program compiles without errors or warnings, and behaves according the requirements given here. A program with good style is clear, concise, and easy to read.

A few suggestions regarding good style may be helpful. You should use brief but mnemonic variables names and proper indentation. Your code should include comments as necessary to explain how it works, but without explaining things that are obvious.

1.2 Partners

You *must* work alone for this assignment. But remember that the course staff is happy to help with problems you run into. Use Piazza for questions, attend office hours, or set up meetings with any course staff member for help.

1.3 Documentation

For this assignment, we ask that you document all of your methods with Javadoc-style comments. How to write Javadocs will be covered in lab, and the course staff will be able to answer any questions during office hours.

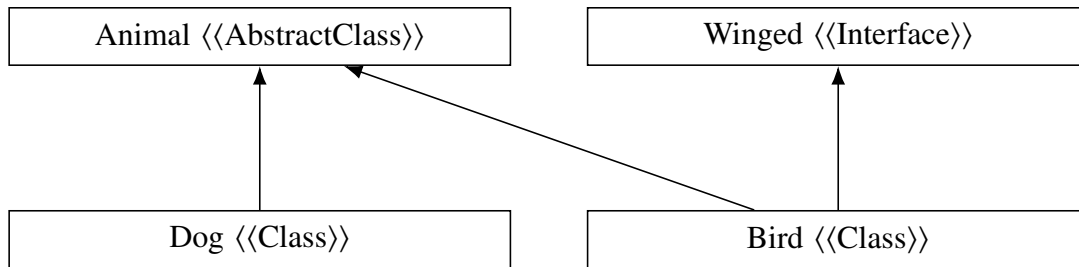


Figure 1: An example of a class diagram. Dog, which is a class, extends Animal, which is an abstract class. Bird, which is a class, extends Animal and implements Winged, which is an interface.

1.4 Provided interfaces

You must implement all methods provided, even if you do not use them. You may not change the signature and return type of any provided method. Formal parameters may be renamed, however, but this is discouraged. You may also add `throws` declarations to methods if you believe they improve your design. You may (and are encouraged to) add as many additional methods as you need.

2 Design overview document

Many components of this assignment share common functionality. For example, many ciphers need to read input both from the command line and from files. A program that implements input reading once in a common place and allows each cipher to use this functionality is better than one having such code duplicated for each cipher. A clean and readable program exploits inheritance that collects common code among individual components into a single place.

To ensure that your design is reasonable for this assignment, we require that you submit a design overview document before the assignment is due. A design document should contain a diagram of the class hierarchy you are planning to implement, along with a short paragraph briefly justifying your design decisions. When designing the class hierarchy, consider factoring out shared functionality between the implementation of multiple ciphers to eliminate duplicate code. Figure 1 shows an example of a class diagram that you can use as a guide.

Submit your design overview document as a PDF file named `a2_designOverview.pdf`. Scans of handwritten diagrams are acceptable.

3 Monoalphabetic and transposition ciphers

3.1 Overview

A cipher is a way to protect a message by changing its letters or characters so that only the desired recipient(s) can read it. The original message is called the *plaintext*, and the transformed message is called the *ciphertext*. Monoalphabetic ciphers provide what is perhaps the most rudimentary

encryption, as they create a one-to-one correspondence between letters in the plaintext and letters in the ciphertext.

3.1.1 The Caesar cipher

The Caesar cipher is a monoalphabetic cipher that functions by mapping an alphabet to a ‘shifted’ version of itself. For example, a cipher where each letter is shifted to the right by one ($A \rightarrow B, B \rightarrow C, \dots, Z \rightarrow A$) would encrypt the message CAT as DBU. We say that the above Caesar cipher has a *shift parameter* of 1. A shift parameter of 0 results in the original alphabet. Shift parameters are not limited to values 1–26; they can be any integer, including negative numbers. A shift of -1 maps $A \rightarrow Z, B \rightarrow A$, and so on.

3.1.2 Random substitution cipher

The Caesar cipher is a special case of a *monoalphabetic substitution cipher*, a cipher that consistently replaces individual plaintext symbols with corresponding ciphertext symbols. With a Caesar cipher, the shift parameter can be learned from knowing how a single symbol was encrypted, determining the entire mapping. A substitution cipher is harder to break if the mapping between plaintext and ciphertext symbols is random. In this case, knowing how any one symbol is mapped by the cipher gives very little information about how other symbols are mapped.

Note that random number generators used in computing are almost never truly random, but rather *pseudorandom*. They are produced by an algorithm called a *pseudorandom number generator* (PRNG) whose output is difficult to distinguish from true random numbers. A *cryptographic* random number generator is a random number generator for which practical algorithms are not known that can distinguish their pseudorandomness from true randomness.

3.1.3 The Vigenère cipher

Another way to strengthen the Caesar cipher is to use different substitutions on different letters. The Vigenère cipher¹ was once considered to be unbreakable. Rather than using a uniform shift for all letters, a repeating pattern of shifts is applied. Traditionally, the key is represented as a word, with A representing a shift of 1, B a shift of 2, and so on. Encrypting CATALOG with key ABC yields DCWBNRH, because the shifts ABCABCA are applied to the plaintext. The Vigenère cipher makes frequency-based cryptanalysis more difficult, especially if the key is long, because even if the same letter appears many times in the plaintext, it may appear in the ciphertext as many different letters.

3.1.4 The route cipher

Another family of ciphers are transposition ciphers, in which the letters of a message are rearranged in a systematic manner to form an anagram. One such cipher is called a *route cipher*. First, the

¹The cipher was actually invented by an Italian cryptologist [Giovann Battista Bellaso](#) in 1553. [Blaise de Vigenère](#), a French cryptographer, created a different, stronger [autokey cipher](#) in 1586.

plaintext is written in a grid of predetermined dimensions. The ciphertext is then read off according to a certain pattern. For this assignment, the pattern is to read down the columns going from left to right. For example, the plaintext ABCDEFGHIJKLMNOP encrypted in a route cipher with a set width of four characters is arranged as

```
ABCD
EFGH
IJKL
MNOP
```

And so the ciphertext reads AEIMBFJNCGKODHLP. Whitespace must be added to the end of the block when the plaintext does not fit nicely in the prescribed dimensions. For example,

```
ABCD
EFGH
IJKL
MN_ _
```

has two spaces added at the end of the fourth line to make everything fit nicely, and the ciphertext reads AEIMBFJNCGK_ _DHL_ _.

3.2 Implementation

Your task is to implement Caesar cipher, random substitution cipher, Vigenère cipher, and route cipher. Your cipher implementations will be accessed through class `CipherFactory`. Two interfaces `EncryptionCipher` and `DecryptionCipher` define necessary methods for encryption and decryption. In addition to implementing these interfaces, your ciphers should extend abstract class `AbstractCipher`.

We will be looking for elegant program design that *minimizes the repetition of common code*. The best program is rarely one with the most lines of code, but rather one that accomplishes the task most simply and with the *least* code. You will find inheritance a valuable language feature for avoiding repetitions. Additional abstract classes may be useful to make your code shorter.

3.2.1 Letter encoding

A consistent standard is important for representing characters in both the encrypter and decrypter. All letters should be converted to their uppercase equivalents. Whitespace—specifically, spaces, tabs, and newlines—should be maintained. All other characters should be discarded. For example, the sentence I really like Cornell, don't you? would become the plaintext I REALLY LIKE CORNELL DONT YOU. You may assume that when decrypting, the program will encounter only uppercase letters and whitespace characters.

3.2.2 Saving the cipher

To save a Caesar or random substitution cipher to a file, simply print the encrypted alphabet to the file, followed by a newline. For example, saving a Caesar cipher with shift parameter 1 creates a

file whose content is as follows:

```
1 BCDEFGHIJKLMNOPQRSTUVWXYZA
2
```

To save a Vignère cipher, print the key to the file, followed by a newline. For example, a Vignère cipher with the key KEY saves to a file that looks like

```
1 KEY
2
```

To save a route cipher, print the width of the grid to a file, followed by a newline. A route cipher with grid width 4 saves to a file with content

```
1 4
2
```

4 Cipher cracking using frequency analysis

Monoalphabetic ciphers are easiest to break using frequency analysis. A cryptanalyst analyzes the frequency of letters in the target language and in the encoded message. This information can be used to reconstruct the cipher and decrypt the message.

4.1 Implementation

You should implement a tool to analyze the frequency of letters over multiple unencrypted texts in the target language, and then use this analysis to crack messages encrypted with a Caesar cipher. You should do so by completing the methods provided in class `FrequencyAnalyzer`. Like the encrypter, `FrequencyAnalyzer` should keep track of only uppercase English letters and handle other characters appropriately (convert or ignore). How you do this is up to you, but here is a hint: there are only 26 possible Caesar ciphers. Find the one that best explains the frequencies of the symbols seen in the ciphertext, under the assumption that the sample text provided contains letters with frequencies typical of the plaintext (i.e., the frequencies found in English-language text). If you are looking for large chunks of English text for testing you may find [Project Gutenberg](#) useful.

5 RSA encryption

RSA is probably the most widely used encryption schema in the world today. RSA is a *public-key cipher*: anyone can encrypt messages using the public key; however, knowledge of the private key is required in order to decrypt messages, and knowing the public key does not help crack the private key. Public-key cryptography makes the secure Internet possible. Before public-key cryptography, keys had to be carefully exchanged between people who wanted to communicate, often by non-electronic means. Now RSA is routinely used to exchange keys without allowing anyone snooping on the channel to understand what has been communicated.

RSA is believed to be very secure, based on the widely held assumption that no one has an efficient algorithm for factoring large numbers; deriving the private key from the public key appears to be as hard as factoring.

5.1 The algorithm

5.1.1 Key generation

1. Choose two random and distinct prime numbers p and q . These must be kept secret. The larger p and q are, the stronger the encryption will be.
2. Compute $n = p \cdot q$. This is the *modulus* used for encryption.
3. Compute $\phi(n)$, the *totient* of n . It is the number of positive integers less than n that are relatively prime to it. For a product of two primes p and q , the totient is easy to compute: $\phi(n) = (p - 1)(q - 1)$. Notice that computing $\phi(n)$ requires knowledge of p and q .
4. Choose an integer e such that $1 < e < \phi(n)$ and e is relatively prime to $\phi(n)$. That is, the greatest common divisor of e and $\phi(n)$ is 1.
5. Compute the decryption key d as the multiplicative inverse of e modulo the totient, written as $e^{-1} \bmod \phi(n)$. This is a value d such that $1 \equiv e \cdot d \bmod \phi(n)$. Such a value d can be found using [Euclid's extended greatest-common-divisor algorithm](#).

The public key is the pair (n, e) and the private key is the pair (n, d) .

5.1.2 Encryption

A plaintext message s is encrypted as ciphertext c via the following formula: $c = s^e \bmod n$. Note that it can be done using only the publicly known n and e .

5.1.3 Decryption

An encrypted message c is decrypted as plaintext s via the following formula: $s = c^d \bmod n$. Note that this *cannot* be done with just the public key.

5.1.4 Why does this work?

If we encrypt and then decrypt a plaintext message s , we obtain a new message $s' = (s^e \bmod n)^d \bmod n$. For the cryptosystem to work, we must have $s = s'$. It is not too hard to see that this is true if we lean on one well-known result from number theory.

By the properties of modular arithmetic, we can pull the $\bmod n$ to the outside: $s' = (s^e)^d \bmod n = s^{ed} \bmod n$. Recall e and d are chosen to be multiplicative inverses modulo $\phi(n)$, so $s^{ed} = s^{1+k\phi(n)} = s \cdot s^{k\phi(n)}$ for some integer k . It turns out that the sequence $s^0, s^1, s^2, s^3, \dots$, taken modulo n , always repeats with period $\phi(n)$. Therefore $s^{\phi(n)} \equiv s^0 \equiv 1 \bmod n$, so $s \cdot s^{k\phi(n)} \equiv s \cdot (s^{\phi(n)})^k \equiv s \cdot 1^k \equiv s \bmod n$. Therefore $s^{ed} \equiv s \bmod n$, as desired.

For example, $\phi(10) = \phi(2 \cdot 5) = 1 \cdot 4 = 4$, and $3^4 = 81 \equiv 1 \equiv 2401 = 7^4 \bmod 10$. In fact, we can use 3 and 7 as e and d , since $3 \cdot 7 = 21 \equiv 1 \bmod \phi(10)$. Let's try it out on the message 8. We encrypt it as $8^3 = 512 \equiv 2 \bmod 10$. Going back the other way, $2^7 = 128 \equiv 8 \bmod 10$. It works!

5.2 Implementation

5.2.1 Dealing with large numbers

RSA involves large numbers, so you should use class `java.math.BigInteger` for all arithmetic. To generate large prime numbers, you should use the appropriate `BigInteger` constructor with `certainty = 20`. The numbers generated by this constructor are only ‘probably’ prime, but given a high enough `certainty`, this is good enough. You’ll want to choose a bit length (the `bitlength` parameter) for p and q such that their product contains the right number of bits. Recall that the product of two n -digit numbers contains at most $2n$ digits.

5.2.2 Message format and padding

Encryption The most challenging part of implementing RSA is not the arithmetic (at least with the help of a class such as `BigInteger`). Rather, it is formatting the message so that it can be correctly encrypted. The plaintext should be broken down into chunks of size 117 bytes by implementing interface `ChunkReader`. The bits from the sequence of plaintext bytes should then be used to construct the `BigInteger` object to which the RSA algorithm is applied, with the least significant bits of the number as the first byte of each chunk. Note that we no longer need to convert case or special characters, as we are using their underlying representations.

Since the input length is unlikely to be an even multiple of 117, it is necessary for each chunk to keep track of the actual number of bytes of data contained in the chunk. This is done by extending the chunk with a 118th byte containing the number of actual data bytes in the chunk (from 1 to 117). There are always 118 total bytes in the data that is converted to a `BigInteger`: up to 117 data bytes, plus one extra byte to keep track of the chunk size. If fewer than 117 data bytes are available, padding bytes are inserted so that the size byte is still the 118th.

In general, encryption using the RSA algorithm can make the number larger. For this reason, the numbers generated by encryption are converted back into 128-byte arrays that are written to the output file, and the length of encrypted output is always a multiple of 128. The difference between 118 and 128 leaves plenty of room for the number to grow when it is encrypted, so encryption never overflows the available space.

For stronger encryption, the padding bytes added to short chunks would be random bytes, protecting against a dictionary attack. However, your program will be easier to debug if you set all such bytes to zero, and we will accept such a solution.

When writing out an encryption result to a file, exactly the encrypted 128-byte chunks should be written, not the textual representation of the number. Hint: Use `OutputStream` directly rather than converting to a string and back.

It is **strongly** recommended that you test converting text and files to and from chunks without any encryption before you try implementing the RSA algorithm itself. Doing this will help you catch bugs while the number of possible causes is still small.

Example: Consider converting the plaintext string CS2112 into an appropriately formatted byte array. A character is really a signed 16-bit integer, and takes up two bytes. Specifically, the characters have the following codes (note that 0x indicates a base-16 number):

char	Value
C	0x0043
S	0x0053
2	0x0032
1	0x0031

How these character codes are translated into an array of bytes depends on the encoding chosen. With the default ISO-8859-1 encoding, only character codes 0x00–0xFF are supported, and the byte array will be { 0x43, 0x53, 0x32, 0x31, 0x31, 0x32 }. With the UTF-8 encoding, you’ll get the same sequence for this string, but characters above 0x7F will translate to multi-byte sequences. The UCS-2 encoding will preserve give you 12 bytes of data and will not lose any information. For this assignment, we recommend using the ISO-8859-1 encoding.

With 6 bytes of plaintext data, there will be $117 - 6 = 111$ bytes of padding and the final byte containing 6 to signify the size of the actual “payload.”

When interacting with data read from a file, the data should be handled directly as bytes, since the data may not make sense as characters under any character encoding. You should be able to encrypt .class files, for exapmle.

Appendix A contains an example showing expected results when encrypting and decrypting with RSA.

Decryption Decryption is simply the inverse of encryption. The input is read in chunks of size 128, which are converted to `BigIntegers` and run backward through the transformation. The 117th byte in the decrypted result specifies how many bytes of data to extract from the array as the decrypted output. Encrypting a file and then decrypting the result should give back exactly the original file.

5.2.3 Saving the keys

The keys, like the cipher schema, can be saved to files.

Public Key When a public key is stored to a file, the file should contain the decimal representation of n , followed by a newline, followed by the decimal representation of e , and end with a newline. That is, the file should look like this:

```
1 <n>
2 <e>
3
```

Private Key The storage of the private key is the same as the storage of the public key, except for the addition of a third line containing d . More precisely, the file for a private key looks as follows:

```
1 <n>
2 <e>
3 <d>
4
```

5.2.4 Additional requirement

Your implementation of the various ciphers should not crash when attempting to encrypt or decrypt input larger than the Java heap. For large input files, you will not be able to bring the entire input into memory at once.

5.3 Useful resources

- class [BigInteger](#)
- class [String](#)
- class [Byte](#)
- abstract class [InputStream](#)
- Wikipedia article on [RSA](#)
- Wikipedia article on [Factory design pattern](#)

For this assignment it may be helpful to know what is really stored in a file, especially if you are having trouble with reading or writing to a file. Viewing a file using a text editor is not reliable when you're dealing with binary data. Suppose you have a file `output.txt` and you want to see what's in it.

On Linux and Mac, either of these commands will show you what's there very clearly:

```
od -Ad -txC -tc output.txt
xxd output.txt
```

On Windows there are various “hex editors” that you can use. People seem to like HxD; hexedit is also okay. Hex editors also exist for Linux/Mac, of course. There is also a plug-in called EHEP for Eclipse that purports to provide hex editor support.

6 Command line invocation

We have provided a command-line interface which allows users to interact with the code that you write in this assignment. It is included in the file `Main.java`. You are encouraged to read through and understand this code. You may also modify this code, but modifications are not required for this assignment. The user may provide any of the following commands via the console:

```
java -jar <YOUR_JAR> <CIPHER_TYPE> <CIPHER_FUNCTION> <OUTPUT_OPTIONS>
```

or

```
java -cp <CLASS_FILE_DIR> Main.java
    <CIPHER_TYPE> <CIPHER_FUNCTION> <OUTPUT_OPTIONS>
```

Cipher type There are three different cipher types that we have asked you to implement, and the flags for each of them are as follows.

- `--monosub <cipher_file>`: A monoalphabetic substitution cipher is loaded from the file specified.
- `--caesar <shift_param>`: A Caesar cipher with the given shift parameter is used for these operations.
- `--random`: A monoalphabetic substitution cipher is randomly generated and used by this program.
- `--crackedCaesar [-t <examples> | -c <encrypted>]`: A Caesar cipher is constructed using frequency analysis with the files flagged `-t` listed in examples as the unencrypted language and files tagged `-c` as the encrypted language. A user may provide any number of `-t` and `-c` flags, and in any order.
- `--vigenere <key>`: Creates a Vigenère cipher using the given keyword (given as a string, max length 128 characters)
- `--vigenereL <cipher_file>`: Loads a Vigenère cipher from the given file.
- `--route <width>`: Creates a route cipher with the width given as an integer.
- `--rsa`: Creates a new RSA cipher.
- `--rsaPr <file>`: Creates an RSA encrypter/decrypter from the private key stored in the specified file.
- `--rsaPu <file>`: Creates an RSA (encrypter) from the public key stored in the specified file.

Cipher functions Next, at most one of the following options may also be specified by the user.

- `--em <message>`: Encrypts the given message.
- `--ef <file>`: Encrypts the provided file using the specified cipher scheme.
- `--dm <message>`: Decrypts the given message.
- `--df <file>`: Decrypts the provided file using the specified cipher scheme.

Output options Finally, the user may add as many output flags as they wish.

- `--print`: Prints the result of applying the cipher (if any) to the console.
- `--out <file>`: Prints the result of applying the cipher (if any) to the specified file.
- `--save <file>`: Saves the current cipher to the provided file (if the current cipher is RSA, this saves the private key).
- `--savePu <file>`: If the current cipher is RSA, this saves the public key to the given file.

6.1 Examples

- Make a new Caesar cipher with shift parameter 15, apply it to the provided message, output the result to file `encr.txt`, and save the cipher to file `ca15`:

```
java -jar <your_jar> --caesar 15 --em 'ENCrypt Me!'
--out encr.txt --save ca15
```

- Load the cipher from `ca15`, decrypt the message in `encr.txt`, and print the result to the console:

```
java -jar <your_jar> --monosub ca15 --df encr.txt --print
```

- Create a frequency analyzer using 3 English texts and 1 encrypted text. Use the resulting cipher to decrypt the encrypted text, print the result, and save the cipher:

```
java -jar <your_jar> --crackedCaesar -t moby-dick.txt -c mystery.txt
-t frankenstein.txt -t macbeth.txt --df mystery.txt
--save brokenCiph --print
```

- Create an RSA encrypter, encrypt the given message, save the ciphertext to a file, and save the two keys to different files:

```
java -jar <your_jar> --rsa --em 'rsa is alright, i guess'
--out encr.txt --save priv.pr --savePu pub.pu
```

- Load an RSA private key, decrypt a message, print the resulting plaintext, and also save it to a file:

```
java -jar <your_jar> --rsaPr --df encr.txt --out decr.txt --print
```

6.2 Errors

Should anything go wrong during execution, including user-error (malformed requests, missing files), your program should not simply 'die'. For example, if a user attempts to execute two incompatible actions such as

```
java -jar <your_jar> --random --savePu outfile.pu
```

a reasonable warning should be printed to the console (System.out). Also, no Java exception should ever be shown to the user. Instead, your program should detect the error and find a sensible way to resolve or communicate the problem.

7 Advice

This assignment involves writing much more code than Assignment 1 required, and demands more careful design. Par for this assignment is about 900 lines of code, assuming you design a reasonable class hierarchy that allows effective code reuse.

It can also be a challenging debugging exercise if you make mistakes, especially for the RSA cipher. Think carefully about the code you are writing and to convince yourself that it is correct. You will also need to test your code carefully. Try not only “normal” inputs but also corner cases.

7.1 Start early

You should start on this assignment as early as possible. It will require careful thoughts about your design to use object-oriented programming methodology in the most effective way. Trying to do it all at the last minute is nearly certain to result in code that is both messy and broken.

7.2 Use assertions

Getting RSA to work can be the most challenging part of the assignment if you are not careful of design, testing, and debugging. Much of the problem comes from small issues when chunking bytes, padding, and converting to and from `BigInteger` objects. If a bug is propagated through the encryption process, it becomes infeasible to tell where the problem is.

We strongly recommend using assertions at each step to help pinpoint bugs. You can use assertions to confirm things like the length of your chunks, the fact the values do not change when converted to and from a `BigInteger` object without encryption, and other properties that you expect to be true. If there are complicated things you would like to check, write extra methods to check them and call those methods from an assertion.

Enabling assertions By default, assertions are disabled. This is done so that programmers can use computationally expensive assertions without hurting the performance of production code. To enable assertion checking, the program must be run with the `-ea` flag. This flag can be passed as a VM argument in the Run Configuration. Check out [this stackoverflow post](#) for help. The [Oracle guide on how to use assertions](#) may also be a helpful resource.

7.3 Build and test incrementally

When building large programs, it is helpful to test your code as you go. Continuous testing increases the chance that any new bugs that show up are the result of code you wrote recently. Ideally, at every point during development, you have some incomplete (but correct) code that offers a firm foundation for further work.

Think about how to develop your code in an order that allows you to test it as you go. Assertions that check preconditions and class invariants are helpful ways to test your code as you develop it.

It is also helpful to design your test cases ahead of time. A good set of test cases will make incremental testing much more effective. A good set of test cases will actually help you pinpoint the key issues your code has to deal with before you even write the code.

7.4 Read specs carefully

The specification for `BigInteger` has some subtle issues that may complicate your task. Read the specifications carefully, keeping the following issues in mind:

Endian The RSA algorithm specified here calls for the first byte from the file to represent the least significant bits of the number passed into the algorithm. Such a numeric representation is said to be *little-endian*. However, the relevant constructors for the `BigInteger` class expect a *big-endian* byte array in which the most significant byte comes *first*. These are opposite but equally valid conventions. Your code will need to deal with the difference correctly to earn full credit. Switching endianness can be achieved by reversing the byte array representing a number. In your program, this switch will need to occur whenever a `BigInteger` is created or when its bytes are retrieved during RSA encryption and decryption.

Sign The `BigInteger()` constructor expects a byte array in the *two's-complement representation*. This means that the most significant bit of the first byte represents the sign of the number. In fact, bytes themselves are in two's complement: using 8 bits, they represent numbers between -128 and 127 . A positive `BigInteger` may therefore need an extra zero byte in the most significant position to avoid having the number interpreted as negative. For example, the array `{-128}` represents -128 . To represent positive 128 , we need the longer byte array `{0, -128}`. You are likely to encounter this issue both when constructing `BigInteger`s and when converting them back into byte arrays.

Character encoding Plaintext bytes outside the ASCII range of $[0, 127]$ will show up as negative values in the range $[-128, -1]$. For the purpose of representing the number to be encrypted, however, we are interpreting bytes as unsigned integers in the range $[0, 255]$, so -128 becomes 128 and -1 becomes 255 . The `BigInteger` constructor should take care of most of this behind the scenes, except for the most significant byte of the given array.

The ISO-8859-1 character set should be helpful for this; see `java.nio.charset.Charset`. You can create a `Charset` using the `forName` method and then use that character set to do encoding and decoding of strings to and from byte arrays. Equivalently, you can also pass the string "ISO-8859-1" as an argument to certain `String` methods and constructors to specify that you want to use this encoding into bytes.

7.5 Hex editors

For this assignment, you may find it useful to be able to view and edit bytes in hexadecimal form. This can be done using a hex editor such as HexFiend for Mac or HexEdit for Windows. You can also view an integer n in base 16 using `Integer.toString(n, 16)`.

8 🌀 HARMMA

For full credit, you are not required to do anything more than what is specified so far, but for good **HARMMA** you may add additional features. Possible extensions include but are not limited to the following:

- Randomized RSA padding.
HARMMA: 🌀 🌀 🌀 🌀 🌀
- Cipher block chaining for stronger RSA (instead of the current “electronic codebook” encryption mode).
HARMMA: 🌀 🌀 🌀 🌀 🌀
- Cryptanalysis for simple substitution ciphers, or other ciphers such as Vigenère, and a command to decrypt such messages.
HARMMA: 🌀 🌀 🌀 🌀 🌀
- Digraph (two-character sequence) frequency analysis to more accurate automatic decryption than single-letter frequency analysis.
HARMMA: 🌀 🌀 🌀 🌀 🌀

- A route cipher that uses a different pattern, e.g., going diagonally or spiraling inwards (FYI, diagonal patterns only work on grids with more columns than rows). Be sure to specify what the pattern is through documentation and your `README.txt` file.

HARMA: 🍹 🍹 🍹 🍹 🍹

- More secure storage of ciphers in files.

HARMA: 🍹 🍹 🍹 🍹 🍹

- Additional ciphers of your choice.

HARMA: 🍹 🍹 🍹 🍹 🍹

Make sure to document anything you do that goes beyond what is requested, and be especially sure that any extensions you make do not break the required functionality of your program.

9 Submission

You should compress exactly these files into a zip file that you will then submit on CMS:

- *Source code*: Because this assignment is more open than the last, you should include all source code required to compile and run your project.
- `README.txt`: This file should contain your name, your NetID, all known issues you have with your submitted code, and the names of anyone you have discussed the assignment with. It should also include descriptions of any extensions you implemented.

Do not include any files ending in `.class`.

All `.java` files should compile and conform to the prototypes we gave you. We write our own classes that use your classes' public methods to test your code. *Even if you do not use a method we require, you should still implement it for our use.*

A Full RSA Example

Consider encrypting and decrypting the plaintext message I love CS2112 using an RSA cipher with

```
n = 1093733205710804998199890334068767166455920362828704765160762785
    7499483958630357076021298990643044443078830573207441251383385554
    0595306533636002574137640210725086799272942906615619662751138490
    7081299171319768555238149258992942399208388665959233718761302432
    26847351820673092753217938229621222197931912046784793
e = 157307166947463324293191
d = 5051344722702808401081282712266153060988904798305842766107188088
    5908654445646155458691942425898014706411932503960880426195998687
    5931918963365832849565604072385256625773752537810929752797947864
    1740341834159319848666223848546326674700490685819268998953869293
    3170796325885997807827620143533305258355518889959511
```

After encrypting the plaintext using the above RSA public key, the following is the expected ciphertext (in hex), which might not be readable otherwise.

```

d6db 6e83 4048 c291 593f 6ea9 69ba 30c2
c6a1 d152 37f4 51e9 684a 32e2 ea07 a9e6
7d35 27bc 41a7 f58f ed4c 0ac4 afa5 f06b
3270 6c3b c730 0772 3c23 c795 a3a2 bfae
1ac8 5188 9cc8 319c fb37 b9d9 75b3 4e82
719e 6892 8171 fac8 2885 ce54 88dd 4d83
e663 6ee9 9dc5 618e 4728 aa69 ea0c ac1d
e89e 4d81 5b9c 9443 f399 42aa 9b8e d996

```

If your RSA cipher is working correctly, but does *not* switch endianness during encryption, you would see the following output after encryption (in hex). If this is your output, make sure you are switching endianness correctly before creating a `BigInteger` object.

```

fd77 0c6e dde7 aace c2d9 1d82 8b0f 7175
6f55 4721 e3ad b241 fc43 7b4c 9d07 c330
9727 4bc5 ae96 5361 a869 9c60 3e56 8012
dfdc f01a 76da 29e7 2d2c 6170 b7db 00b2
596f 5c98 ebbf 49d0 beeb a633 9804 45a3
03d3 a4b6 f1dd 489c 7838 4d46 2965 0e58
bb95 bdef 70d4 2691 ae88 76dd 1f78 95db
524c 1a4b 27d6 ba53 284d 584f 6fc6 ec8b

```

Decrypting the correct ciphertext should yield the original message, I love CS2112.

If you had the correct encrypted message, but you do *not* switch endianness during decryption, you would see this output after decryption: 2112SC evol I, the original message backwards. If this is your output, make sure you are switching endianness correctly after reading bytes from a `BigInteger` object.

If your output differs from above after either encryption or decryption, something else is likely going wrong. Ensure the correctness of your `ChunkReader` independent of the encryption process and pay special attention to the order in which you are performing operations on your data bytes. The order should be symmetric for the encryption and decryption steps.