

Module 10.7: Exploring the RSA Cipher

Gregory V. Bard

February 6, 2020

- This is a practice workbook for the RSA cipher, but it assumes that you have already read and mastered the previous module.
- There is a with-answers version, and a without-answers version.
- In the with-answers version of this workbook, the black ink represents the question, and the blue ink represents the answer. The red ink is additional information that would normally only be calculated with technology (e.g. the computer-algebra system called Sage, or perhaps Maple).

Problem 10-7-1

A friend of yours is generating an RSA public-private key pair. He has followed all the steps of the key-generation process, but has forgotten what to do with each of the six things he has computed.

Of the six items computed during RSA key generation, p , q , N , φ , e , d , label each one as “public,” “private,” or “destroyed.”

- When finished, d should be:
- When finished, N should be:
- When finished, q should be:
- When finished, φ should be:
- When finished, e should be:
- When finished, p should be:

Problem 10-7-2

There is a distrustful user, Vincent. He doesn't fully understand the RSA protocol and wants to verify why each particular specification is important. Let's help Vincent verify these things, one at a time.

First, Vincent doesn't understand why φ needs to be destroyed. Suppose Vincent's φ becomes public, but nothing else (besides Vincent's public key) becomes public. (Let's also assume that no one publishes a new algorithm to make it easy to factor huge numbers.) What would happen then?

Problem 10-7-3

Continuing with the previous problem, Vincent now understands that φ should be destroyed. We learn that RSA is very dependent on the hardness of factoring. Vincent wants to know why it would be so bad if factoring huge numbers became feasible. In particular, what would happen if nothing other than Vincent's public key was public, but factoring huge numbers became easy because of some newly published algorithm?

Problem 10-7-4

Suppose that we know that the ambassador to Kuwait has the following RSA public key:

$$N = 34,189 \text{ and } e \equiv 6393$$

Now suppose that the Russian government builds a quantum computer, and factors

$$N = 34,189 = (191)(179)$$

Tell me how this makes it easy to find the ambassador's private key.

Problem 10-7-5

We return to Vincent, who doesn't trust all the details of the RSA algorithm's specification. Let's assume that no one publishes an algorithm to make it easier to factor huge numbers. Suppose Vincent's q becomes public, but nothing else (besides Vincent's public key) becomes public. What could happen then?

Problem 10-7-6

Background Information

Whenever you use your web browser in `https` mode, instead of `http` mode, and you therefore see that key icon or lock icon that indicates an encrypted communication, then you're using

either RSA, or a related algorithm called Diffie-Hellman, for a step that is called “establishing a shared secret” or “key exchange.” The mechanism for doing this is a protocol called SSL (the secure sockets layer) and its successor, TLS (transport layer security).

In 2012, some researchers (see bibliographic references below) took a bunch of RSA public keys (11,400,000 of them) from the Electronic Frontier Foundation’s SSL Observatory, and started taking gcds of pairs of N s. It turns out that this exceptionally simple process—downloading lots of public keys, picking random pairs of N s, and taking gcds —factored 0.2% of the N s. That’s a lot larger than it sounds like. There were 26,965 public keys with N that were thusly factored, which is alarming since most programmers who use cryptography think that RSA is unbreakable.

Citations:

A. Lenstra, J. Hughes, M. Augier, J. Bos, T. Kleinjung, and C. Wachter. “Ron was wrong, Whit is right.” Posted to eprint.iacr.org, 17 February, 2012.

<https://eprint.iacr.org/2012/064.pdf>

There is also a NY Times article about this: J. Markoff. “Flaw Found in an Online Encryption Method.” *The New York Times*. February 15th, 2012.

<http://www.nytimes.com/2012/02/15/technology/researchers-find-flaw-in-an-online-encryption-method.html>

Let’s see if we can do it ourselves. . .

Suppose $e_1 \equiv 151, 153$ and $e_2 \equiv 176, 303$. Further suppose $N_1 = 302, 303$ and $N_2 = 352, 603$. First, you compute $\text{gcd}(N_1, N_2) = 503$ using the Extended Euclidean Algorithm, or Sage.

- What is the factorization of N_1 ?
- What is φ_1 ?
- How can we obtain the private key, d_1 ?
- What is the factorization of N_2 ?
- What is φ_2 ?
- How can we obtain the private key, d_2 ?
- What is the impact of d_1 and d_2 becoming known?

Problem 10-7-7

Flip back to “the short proof of correctness of RSA,” which was given in the previous module, between Problems 10-6-11 and 10-6-12. Carefully examine the proof. What makes it incorrect?

Problem 10-7-8

We return to Vincent, who doesn't trust all the details of the RSA algorithm's specification. Suppose Vincent's p becomes public, but nothing else (besides Vincent's public key) becomes public. Let's further suppose that no one publishes an algorithm to make it easier to factor huge numbers. What could happen then?

Problem 10-7-9

We return to Vincent, who doesn't trust all the details of the RSA algorithm's specification. When we require that $2 < e < \varphi - 2$, we are (of course) explicitly ruling out $e \equiv 1$. What would go wrong if $e \equiv 1$?

Problem 10-7-10

We continue with Vincent, who doesn't trust all the details of the RSA algorithm's specification. When we require that $2 < e < \varphi - 2$, we are (of course) explicitly ruling out $e \equiv 0$. What would go wrong if $e \equiv 0$?

Problem 10-7-11

We are still talking with Vincent, who doesn't trust all the details of the RSA algorithm's specification. When we require that $2 < e < \varphi - 2$, we are (of course) explicitly ruling out $e \equiv \varphi$. What would go wrong if $e \equiv \varphi$?

Problem 10-7-12

We have yet another inquiry from Vincent, who doesn't trust all the details of the RSA algorithm specification. When we require that $2 < e < \varphi - 2$, we are (of course) explicitly ruling out $e \equiv 2$. What would go wrong if $e \equiv 2$?

Problem 10-7-13

We have still another inquiry from Vincent, who doesn't trust all the details of the RSA algorithm specification. When we require that $2 < e < \varphi - 2$, we are (of course) explicitly ruling out $e \equiv \varphi - 1$. What would go wrong if $e \equiv \varphi - 1$?

Problem 10-7-14

Let's return to Alice and Bob at the beginning of their romantic relationship. This was discussed in Problem 10-6-9 to Problem 10-6-11 of the previous module. Suppose Bob had chosen $e \equiv 32,831$ instead of $e \equiv 19$, but his N had remained unchanged, namely 4,056,187.

For this problem, you'll definitely want to carry out the computations in Sage to get final numerical answers.

- Suppose that Bob wants to pre-encrypt some times for dates with Alice, in case he might need to respond to Alice while he's away from his laptop. While Sage will work on a smartphone, it can be a bit clumsy, so he's going to encrypt 1830, 1900, and 1930, to represent 6:30 PM, 7:00 PM, and 7:30 PM. He'll keep the resulting ciphertexts handy for use when needed. What are those ciphertexts?
- Now suppose that he decides to encrypt those ciphertexts again, effectively "double encrypting" them. What are the new ciphertexts?
- What is $e^2 \bmod \varphi$? To save you from flipping back to the other module, permit me to share that Bob's $\varphi = 4,052,160$.
- What can you tell me about Bob's d ?
- Stepping away from this particular problem, if a set of RSA credentials has been properly generated by carefully following the key-generation procedure, but $e^2 \equiv 1 \pmod{\varphi}$, then what can you tell me about d ?

Here is the Sage code that I used, in case you are curious about how I found 32,831.

```
for k in range(1, 1000):
    temp = 1 + 4_052_160*k
    if (sqrt(temp) in ZZ):
        print(temp, "= 1 + 4,056,052 *", k, "=", sqrt(temp), "^ 2.")
```

That code produced the output:

```
1077874561 = 1 + 4,056,052 * 266 = 32831 ^ 2.
3351136321 = 1 + 4,056,052 * 827 = 57889 ^ 2.
```

This means that I could also have chosen $e \equiv 57,889 \equiv d$ instead of $e \equiv 32,831 \equiv d$.

The following code will encrypt (and encrypt again) every possible time from 1800 to 1858.

```
for time in range(1800, 1859):
    ciphertext1 = pow( time, 32_831, 4_056_187 )
    ciphertext2 = pow( ciphertext1, 32_831, 4_056_187 )

    print(time, "--->", ciphertext1, "--->", ciphertext2)
```

Problem 10-7-15

We continue with Vincent, who doesn't trust all the details of the algorithm specification. When we require that $2 < e < \varphi - 2$, we are (of course) explicitly ruling out $e \equiv \varphi - 2$. What would go wrong if $e \equiv \varphi - 2$?

Problem 10-7-16

This problem was inspired by Problem 6-8-8 in *Cryptography with Coding Theory*, 2nd edition, written by my former dissertation supervisor, Prof. Lawrence Washington, and his student Prof. Wade Trappe.

Let's now consider the possibility that Porter has some highly confidential contacts, Boris and Natasha. To protect their highly confidential reports, he generates two public keys, with the same N but with two different encryption exponents, e_1 and e_2 . Of course, he also generates the two matching private keys d_1 and d_2 . When either Boris or Natasha wants to send him something, they are instructed to first encrypt with e_1 , and then encrypt that ciphertext again with e_2 .

- What must Porter do when receiving a message that is encrypted in this way? Boris thinks that Porter must first decrypt with d_1 and then decrypt with d_2 . However, Natasha thinks that Porter must first decrypt with d_2 and then decrypt with d_1 . Who is correct? Or are both contacts wrong?
- As it turns out, this two-stage encryption process is identical to encrypting once with a different public key. What is that public key? What is the corresponding private key?

The Real Lesson of the Previous Problem:

Tell me, what is the real lesson of the problem that we just finished?

Problem 10-7-17

Now Vincent is getting bolder, and wants to question why RSA uses a product of two prime numbers. What if there were only one prime number? In particular, suppose that $N = 54,323$ and $e \equiv 2003$, both of which are prime numbers. What goes wrong in this case?

Problem 10-7-18

Let's say that an investment bank gets an international call from a wealthy customer, Rich, who is on vacation overseas. The person on the phone is claiming that he has lost his wallet, and with that, his credit cards and his cash. He wants to withdraw a few thousand dollars,

and has wiring instructions so that the bank can send the money directly to him overseas. Surely the bank needs to know that it is really Rich that they are speaking to, and not some criminal who is impersonating Rich.

Suppose that Rich has the following credentials:

$$d = 192,191$$

$$\varphi = 461,032$$

$$e = 179,607$$

$$N = 462,391$$

$$p = 683$$

$$q = 677$$

Further suppose that the bank already has Rich's public key on file. The bank is going to ask Rich two questions.

- The bank asks Rich, "What number, encrypted with your public key, will yield 50,669?" What does Rich have to do, in order to compute that? Your answer to this subproblem, and the three after it, should be a formula with numbers and operations, but no variables.
- The bank asks Rich, "What number, encrypted with your public key, will yield 46,375?"
- How will they check Rich's answer to the first question? What will they compute?
- How will they check Rich's answer to the second question? What will they compute?
- If both answers are correct, what does this prove?
- Suppose it really was a criminal on the phone and not Rich. Suppose that when the bank posed the questions to the criminal on the phone, that the criminal merely guessed. What is the probability that the criminal would have guessed correctly both times, if the criminal responded with random numbers strictly between 1 and N ?
- Suppose the bank had instead asked the person on the phone to provide their four-digit PIN (personal identification number), and the person was a criminal impersonating Rich. If the criminal does not know Rich's PIN, and has to guess, then what is the probability of success? You can assume that the PIN is a 4-digit number, possibly starting with one or several zeroes.
- Using RSA in this way, as a method of authentication, is called "the RSA-challenge-response method." It is an alternative method of authentication compared to using passwords. Suppose a less technically proficient bank were to give each customer a password to memorize. In situations like this, a customer would call on the phone, and dictate the password. What advantages does the challenge-response method have over the password method?

Note, conceptually we could think of 399,842 as the number 50,669 having been encrypted “negative one” times, and similarly, 214,689 is the number 46,375 having been encrypted “negative one” times. By using Rich’s public key to encrypt 399,842 and 214,689 once (positive one times), these actions “cancel out” leaving the original integers (50,669 and 46,375) as a result.

Authentication is as vast a topic as encryption. The above is only the simplest version of the RSA challenge-response authentication method. Additional safeguards can be added.

For example, in this problem, if Rich already knew the phone number of his bank, had it in his smartphone, or if he found it on the web, then he can be confident that he’s communicating with his bank and not an entity that is pretending to be his bank. However, if two email servers communicate with each other across the internet, until they finish their mutual authentication process, they should not trust each other. Rich and his bank have a one-way authentication problem. On the internet, we typically encounter two-way authentication problems, which are more complicated.

To make matters even more complex, there could be an adversary in the middle of the communications process. This gives rise to a family of cryptographic attacks called man-in-the-middle attacks.

Problem 10-7-19

Near the beginning of the previous module, I made the claim that if I were to replace someone’s e with $e+z\varphi$, for any z in the positive integers, then all encryption and decryption behavior would remain unchanged. Vincent does not understand why this must be true. Therefore, write a proof (using mathematical induction) that if e is replaced with $e+z\varphi$ for any positive integer z , that the value of $m^e \bmod N$ remains unchanged.

Hint: for all $x \bmod pq$, it is true that $x^{\varphi+1} \equiv x \bmod pq$, where $\varphi = (p-1)(q-1)$, $p \neq q$, and both p and q are primes.

Because $e, e+\varphi, e+2\varphi, e+3\varphi, e+4\varphi, \dots$, would all encrypt identically, I consider e and d to be living in the world of “mod φ ” arithmetic. As I mentioned near the start of the previous module, this is why I use the \equiv sign with e and d , and not the $=$ sign.

For small moduli, you can test the hint given above with the following code. Please don’t let the modulus be more than 1000, however. Otherwise you will put too much strain on the `sagecell.sagemath.org` servers, which are being provided to the world for free, but which nonetheless have only finite computing power.

```
modulus = 13*17

#####
##### Don't change anything that follows...
##### ...unless you know what you are doing

assert modulus in ZZ
```



```

assert modulus > 1
assert modulus < 1500

exponent = euler_phi(modulus) + 1

found = 0

for k in range(0, 2*modulus):
    test = pow( k, exponent, modulus )
    if (test != k):
        print(k, "^", exponent, "mod", modulus, "=", test)
        found = found + 1

if (found == 0):
    print("No exceptions found.")
else:
    print(found, "exceptions were found.")

```

Interestingly, the more general statement of the hint is false. If I replace pq with N , and I allow N to be something other than the product of two distinct primes, then the hint becomes “for all $x \bmod N$, it is true that $x^{\varphi+1} \equiv x \bmod N$, where $\varphi = \Phi(N)$.” As a counterexample, consider $N = 13^2 = 169$, which is not the product of two distinct primes. As it turns out, $\Phi(169) = 156$. We might expect that raising 26 to the 157th power would yield $26 \bmod 169$, but this is false. You can test that by typing `pow(26, 157, 169)` into Sage, and you will see that the answer is zero. It is also fun to change the line `modulus = 13*17` into `modulus = 13*13` in the Sage code that was given above.

Problem 10-7-20

Consider the standard “alphabet soup” of RSA. Give me a formula for $N - \varphi$, in terms of p and q . (This might seem pointless, but trust me, this will be useful momentarily.)

A Taste of Computational Hardness

Suppose there existed a device that could take N and return $\varphi = \Phi(N)$, for any N that is a product of two large distinct primes. Clearly, such a device would be a catastrophe for RSA, because we saw in Problem 10-7-2 (from Page 2 of this module) that when φ becomes known, it is easy to compute the private key that matches the corresponding public key. We will now prove that this device can be used, in essentially the same amount of time, to factor N . (That might seem to be an irrelevant objective right now, but it will turn out that this will allow us to get a small glimpse of an otherwise extremely advanced topic: the connections between complexity theory and cryptography.)

We will actually be making use of an old trick from algebra classes in middle schools (and high schools). Stated as a theorem:

Two real numbers r and t satisfy $r + t = S$ and $rt = P$ if and only if r and t are the roots of this polynomial:

$$x^2 - Sx + P = 0$$

In the previous problem, we learned that $N - \varphi = p + q - 1$, but you probably were wondering what the use of that might be. We will vary this formula only slightly to write $N - \varphi + 1 = p + q$, by adding one to each side.

Now imagine that Alice has Bob's public key, and this device that takes N as an input, and which outputs φ . First, she will take the N from the public key and use the device to get φ . While she does not yet know what p and q are, she does know that $pq = N$ and that $p + q = N - \varphi + 1$. After computing $S = N - \varphi + 1$, and $P = N$, all Alice only needs to solve the quadratic equation

$$x^2 - Sx + P = 0$$

and the two distinct roots will be $x = p$ and $x = q$. The quadratic formula from middle school (or high school) will do that work for Alice.

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} = \frac{S \pm \sqrt{S^2 - (4)(1)(P)}}{2}$$

We can model the amount of computation to do this either with time, with the number of arithmetic operations, or with the number of assembly-language instructions that a microprocessor will require. Personally, I prefer the last choice, but it is actually the least popular among cryptologists. Most cryptologists count operations, so we'll do that instead.

- We need two operations to find $S = p + q = N - \varphi + 1$.
- We need two more to compute S^2 and $4P$, and a third to subtract them.
- The sixth operation would be the square root of $S^2 - 4P$.
- Then $(S - \sqrt{S^2 - 4P})/2$ and $(S + \sqrt{S^2 - 4P})/2$ each require two more operations.
- In total, we need only 10 additional operations. (This would be a matter of nanoseconds on a modern computer.)

We have now proven that “if a device can (in time T) take N and return $\varphi = \Phi(N)$, for any N that is a product of two large distinct primes, then one can factor N in only slightly more time.” Alternatively, “if a device can (with K operations) take N and return $\varphi = \Phi(N)$, for any N that is a product of two large distinct primes, then one can factor N using only $K + 10$ operations.”

Before we continue, I think it is useful to perform this operation with two specific cases, using actual numbers.

Problem 10-7-21

Consider Boris, who has a public key of $e \equiv 262,795$ and $N = 588,983$, and Natasha, who has a public key of $e \equiv 192,073$ and $N = 687,811$. Alice has the device that we were discussing moments ago, that can take N and return $\varphi = \Phi(N)$, for any N that is a product of two large distinct primes.

- First, Alice puts Boris's $N = 588,983$ into her device and obtains $\Phi(588,983) = 587,448$. What does this tell us about the sum of Boris's primes?
- What will the product of Boris's primes be?
- What quadratic equation must Alice now solve?
- What are the roots of that quadratic equation?
- Please note that Alice would probably never solve the quadratic equations above. That's because once she has Boris's φ in her hands, she can just compute $d \equiv e^{-1} \pmod{\varphi}$ and that will let her see Boris's d (as we saw in Problem 10-7-2, from Page 2 of this module). There is no additional advantage for her to see Boris's two primes as well, unless she wants to see them.
- Second, Alice puts Natasha's $N = 687,811$ into her device and obtains $\Phi(687,811) = 686,152$. What does this tell us about the sum of Natasha's primes?
- What will the product of Natasha's primes be?
- What quadratic equation must Alice now solve?
- What are the roots of that quadratic equation?

It is worth noting that I generated these public keys for Boris and Natasha randomly. It is unfortunate luck that Boris and Natasha had the prime 797 in common. This would make their public keys easily breakable, using the gcd method explained earlier in Problem 10-7-6 (on Page 3 of this module). I did not cause that to happen deliberately—the shared prime occurred by coincidence.

The Real Lesson of the Previous Problem

With considerable effort over the last few pages, we proved that “if a device can (using K operations) take N and return $\varphi = \Phi(N)$, for any N that is a product of two large distinct primes, then one can factor N using only $K + 10$ operations.” I'm now going to state a sequence of sentences that are logically equivalent, so that you can get an idea of what complexity theory statements sound like. It would take a large number of pages to rigorously define all the concepts laid out here, so please don't be frustrated if some of what

follows seems to be vague. Entire books¹ have been written on the connections between complexity theory and cryptography, so there is no need for this textbook to dwell too long on this topic.

We can phrase our conclusions more mathematically by saying “if there exists a device that can (in polynomial time) take N and return $\varphi = \Phi(N)$, for any N that is a product of two large distinct primes, then there exists a device that can factor N in only slightly more time.”

The contrapositive is much more exciting. The contrapositive is “if there does not exist a device that can factor N (in polynomial time), for any N that is a product of two large distinct primes, then there does not exist a device that can take N and return $\varphi = \Phi(N)$.”

Computer scientists will abbreviate this by saying “if factoring the product of two distinct large primes is computationally infeasible, then computing the Euler Totient Function of the product of two distinct large primes is computationally infeasible.”

We can compress this even further by saying “if factoring the product of two distinct large primes is hard, then computing the Euler Totient Function of the product of two distinct large primes is hard.” (Personally, this is how I prefer to say such things when I am thinking about this topic in my research.)

Some would even say that “computing the Euler Totient Function of the product of two distinct large primes is no easier than factoring the product of two distinct large primes.”

Of course, if I have a device that can factor the product of two distinct large primes, in addition to being able to break RSA (and read everyone’s mail plus sign important documents, such as checks, as if I were any person whose public key I have), I could also compute $\Phi(N)$ for those numbers. Once I’ve factored $N = pq$, then it is trivial to compute $\varphi = (p - 1)(q - 1) = \Phi(N)$. Therefore, I can confidently say “computing the Euler Totient Function of the product of two distinct large primes is no harder than factoring the product of two distinct large primes.”

In conclusion, I should summarize this discussion by saying we’ve proven that “computing the Euler Totient Function of the product of two distinct large primes is computationally equivalent to factoring the product of two distinct large primes.”

A Happy Note to End this Module

Believe it or not, this insight that knowledge of $\Phi(N)$ can be used to factor N launched the undergraduate research project of a student under my supervision, when I was teaching at Fordham University in The Bronx. His name is Kyle Kloster, and his Honor’s Bachelor’s thesis was a new method of factoring integers, aimed at integers that are the product of two large distinct primes. You can read that Honor’s Bachelor’s thesis, entitled “Factoring a semiprime n by estimating $\phi(n)$,” which he presented on May 7th, 2010, and I’m sure you’ll agree that it could have been a Master’s thesis. Kyle later went on to earn a PhD in Mathematics from Purdue University in 2016. That Honor’s Bachelor’s thesis is located at <http://www.gregory-bard.com/publications.html>

¹A notable example is *Complexity Theory and Cryptology*, by Jörg Rothe, published by Springer in 2005.