

Further comments on RSA and ElGamal

Theorem 182. Determining the secret private key d in RSA is as difficult as factoring N .

Proof. Let us show how to factor $N = pq$ if we know e and d such that $d \equiv e^{-1} \pmod{(p-1)(q-1)}$.

- Write $ed - 1 = 2^t m$, where t is chosen as large as possible such that 2^t divides $ed - 1$. Since $ed - 1 \equiv 0 \pmod{(p-1)(q-1)}$ and 2^2 divides $(p-1)(q-1)$, we have $t \geq 2$.
- Pick a random invertible residue x . Observe that $x^{ed-1} \equiv 1 \pmod{N}$. In other words, $(x^m)^{2^t} \equiv 1$. Hence, the multiplicative order of x^m must divide 2^t .
- Suppose that x^m has different order modulo p than modulo q .

Note. One can show that this works for at least half of the (invertible) residues x . As such, if we are unlucky, we just select another x .

Since both orders must divide 2^t , we may suppose x^m has order 2^s modulo p , and larger order modulo q . Then, $x^{2^s m} \equiv 1 \pmod{p}$ but $x^{2^s m} \not\equiv 1 \pmod{q}$.

Consequently, $\gcd(x^{2^s m} - 1, N) = p$ so that we have found the factor p of N .

Note. Of course, we don't know s (because we don't know p and q), but we can just go through all $s = 0, 1, 2, \dots, t - 1$. One of these has to reveal the factor p . \square

However. It is not known whether knowing d is actually necessary for Eve to decrypt a given ciphertext c . This remains an important open problem.

Advanced comment. Recall that we don't necessarily need to have $d \equiv e^{-1} \pmod{(p-1)(q-1)}$ but that it suffices to have the same with $(p-1)(q-1)$ replaced by $\text{lcm}(p-1, q-1)$. This means that, in the above argument, we only get $t \geq 1$ but the rest of the argument still applies.

Example 183. (homework) Bob's public RSA key is $N = 323$, $e = 101$. Knowing $d = 77$, factor N using the approach of the previous theorem.

Solution. Here, $de - 1 = 7776 = 2^5 \cdot 243$ so that $t = 5$ and $m = 243$.

- Let's pick $a = 2$. $a^m = 2^{243} \equiv 246 \pmod{323}$ must have order dividing 2^5 . $\gcd(246^2 - 1, 323) = 19$ (so we don't even need to check $\gcd(246^{2^s} - 1, 323)$ for $s = 2, 3, 4$)
Hence, we have factored $N = 17 \cdot 19$.

Comment. Among the $\phi(323) = 16 \cdot 18 = 288$ invertible residues a , only 36 would not lead to a factorization. The remaining 252 residues all reveal the factor 19.

Another project idea. Run some numerical experiments to get a feeling for the number of residues that result in a factorization.

Definition 184. Bob’s public key cryptosystem is **semantically secure** if Eve cannot do better than guessing in the following challenge:

- Bob determines a random public and private key. The public key is given to Eve.
- Eve selects two plaintexts m_1 and m_2 .
- Alice flips a fair coin and, accordingly, using the public key encrypts m_1 or m_2 as c .
- Eve now needs to decide whether c is the encryption of m_1 or m_2 .

For this definition to make precise mathematical sense, we need to assume that Eve’s computing power is somehow limited (typically, she is limited to polynomial-time algorithms).

Comment. Also, many variations exist of what semantic security exactly is. All of these try to capture the idea that an attacker does not learn anything about m from knowing c . The one above is often referred to as IND-CPA (Indistinguishability under Chosen Plaintext Attack).

Important comment. Realize that semantic security is a very strong property to ask for! In particular, this is much stronger than what we usually think about in terms of security: you might call a cipher secure if it is “impossible” for an attacker to get m from c . Semantic security is requiring that an attacker gets so little information from c that she cannot even tell whether it came from (her own choices) m_1 or m_2 .

Example 185. Is vanilla RSA semantically secure?

Solution. No. Eve can just encrypt both m_1 and m_2 herself, and compare with c . She then knows for sure which of the two was encrypted.

Comment. As mentioned before, in practice, RSA is never used in its vanilla (or “textbook”) version (unless random plaintexts are encrypted). Instead, it is randomized (like ElGamal is by design) by padding the plaintext with random stuff.

Check out OAEP: https://en.wikipedia.org/wiki/Optimal_asymmetric_encryption_padding

The resulting RSA-OAEP has been proven semantically secure (under the “RSA assumption” that finding m from c is hard).

Example 186. Is ElGamal semantically secure?

Solution. Essentially, yes.

Recall that the public key is $(p, g, h) = (p, g, g^x)$.

The ciphertext is $(c_1, c_2) = (g^y, h^y m) = (g^y, g^{xy} m)$. Eve needs to decide whether the m in there is m_1 or m_2 .

Equivalently, she needs to decide whether $r = c_2 / m_1$ (or $r = c_2 / m_2$) equals g^{xy} or not.

This is essentially the DDH problem.

Strictly speaking. Because of the issue with quadratic residues mentioned when we introduced the DDH problem, ElGamal is not semantically secure in the sense we defined things. However, if we wanted (this is more of a theoretical point), this issue could be fixed by not computing with all invertible residues modulo p , but only with quadratic residues. We could further select p to be a **safe prime**, meaning that $(p - 1) / 2$ is prime again, in which case all quadratic residues (except 1) have order $(p - 1) / 2$ (so that no similar games can be played using orders of elements).

Practical implications. Indeed, Diffie–Hellman and ElGamal in practice often use safe primes p . In that case, as we observed in Example 180, there are no elements of small order (besides 1 and -1). Since generating such primes can be a bit expensive, it is common to use preselected ones. For instance, RFC 3526 lists six such primes (together with a generator g) with 1536, 2048, ..., 8192 bits.

<https://www.ietf.org/rfc/rfc3526.txt>

Important. It is perfectly fine that p and g are not random in Diffie–Hellman or ElGamal. However, it is absolutely crucial that x (and y) are random (generated using a cryptographically secure PRG).

Example 187. What is your feeling? Can we make RSA even more secure by allowing N to factor into more than 2, say, 3 primes?

Solution. That doesn't seem like a good idea. Namely, observe that the security of RSA relies on adversaries being unable to factor N . Allowing more factors of N (while keeping the size of N fixed) makes that task easier, because more factors means that the factors are necessarily smaller.

Example 188. RSA has proven to be secure so far. However, it is easy to implement RSA in such a way that it is insecure. One important but occasionally messed up part of RSA is that **p and q must be unpredictable**, and the only way to achieve that is to choose p, q completely randomly in some huge interval $[M_1, M_2]$.

- For instance, if $N = pq$ has m digits and we know the first (or last) $m/4$ digits of p , then we can efficiently factor N .

An adversary might know many digits of p if, for instance, we make the mistake of generating the random prime p by considering candidates of the form $2^{1023} + k$ for small (random) values of k (2^{1023} was chosen so that the resulting number has 1024 bits).

- Also, we must use a cryptographically secure PRG to generate p and q .

If using a “bad” PRG or choosing seeds with too little entropy, then (especially among a large number of public keys generated this way) it becomes likely that (different) public keys N and N' share a prime factor p . In that case, everybody can determine $p = \gcd(N, N')$ and break both public keys.

Indeed. For instance, in a study of Lenstra et. al., millions of public keys were collected and compared. Among the RSA moduli, about 0.2% shared a common prime factor with another one. That's terrible: if (different) public keys N and N' share a prime factor p , then everybody can determine $p = \gcd(N, N')$ and break both public keys.

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

- In that direction, is the security of public key cryptosystems like RSA in any way compromised when used by tens of millions of users?

As noted above, millions of people using “bad” PRGs for generating RSA public keys make it likely that this weakness can be practically exploited.

Similarly, for Diffie–Hellman and ElGamal, it is common to use fixed primes p . While fine in principle, this may be an issue if used by millions of users faced against an adversary Eve with vast resources. See, for instance: <https://threatpost.com/prime-diffie-hellman-weakness-may-be-key-to-breaking-crypto/>

Example 189. (side-channel attacks) For instance, by measuring the time it takes to decrypt messages as $m = c^d \pmod{N}$ in RSA, Eve might be able to reconstruct the secret key d .

This **timing attack**, first developed by Paul Kocher (1997), is particularly unsettling because it illustrates that the security of a system can be compromised even if mathematically everything is sound. This sort of attack is called a **side-channel attack**. It attacks the implementation (software and/or hardware) rather than the cryptographic algorithm.

See Section 6.2.3 in our book for more details on how d can be obtained in this attack.

In a similar spirit, there exist power attacks (measuring power instead of time during decryption) or fault attacks (for instance, injecting errors during computations):

https://en.wikipedia.org/wiki/Side-channel_attack

How to prevent? Implement RSA in such a way that no inferences can be drawn from the time and power consumption.

Lesson. Do not implement crypto algorithms yourself!! Instead, use one of the well-tested open implementations.

It's kind of sad, isn't it? Don't come up with your own ciphers. Don't implement ciphers yourself...

But it is important to realize just how easy it is to implement these algorithms in such a way that security is compromised (even if the idea, intentions and algorithms are all sound and secure).

After advertising open implementations, let us end this discussion with a cautionary example in that regard.

Example 190. The following story made lots of headlines in 2016:

<https://threatpost.com/socat-warns-weak-prime-number-could-mean-its-backdoored/116104/>

After a year, it was noticed that, in the open-source tool Socat ("Netcat++"), the Diffie-Hellman key exchange was implemented using a hard-coded 1024 bit prime p (nothing wrong with that), which wasn't prime! Explain how this could be used as a backdoor.

Solution. The security of the Diffie-Hellman key exchange relies on the difficulty of taking discrete logarithms modulo p . If we can compute x in $h = g^x \pmod{p}$, then we can break the key exchange.

Now, if $p = p_1 p_2$, then we can use the CRT to find x by solving the two (much easier!) discrete logarithm problems

$$h = g^x \pmod{p_1}, \quad h = g^x \pmod{p_2}.$$

This is an example of a **NOBUS backdoor** ("nobody but us"), because the backdoor can only be used by the person who knows the (secret) factorization of p .

Comment. In the present case, the Socat "prime" p actually has the two small factors 271 and 13597, and $p/(271 \cdot 13597)$ is still not a prime (but nobody has been able to factor it). This might hint more at a foolish accident than a malicious act.

Important follow-up question. Of course, the issue has been fixed and the composite number has been replaced by the developers with a large prime. However, should we trust that it really is a prime?

We don't need to trust anyone because primality checking is simple! We can just run the Miller–Rabin test N times. If the number was composite, there is only a 4^{-N} chance of us not detecting it. (In OpenSSL, for instance, $N = 40$ and the chance for an error, 2^{-80} , is astronomically low.) Both Fermat and Miller–Rabin instantly detect the number here to be composite (for certain).

Comment. This illustrates both what's good and what's potentially problematic about open source projects. The potentially problematic part for crypto is that Eve might be among the people working on the project. The good part is that (hopefully!*) many experts are working on or looking into the code. Thus, hopefully, any malicious acts on Eve's part should be spotted soon (in fact, with proper code review, should never make it into any production version). Of course, this "hope" requires ongoing effort on the parts of everyone involved, and the willingness to fund such projects.

*However, sometimes very few people are involved in a project, despite it being used by millions of users. For instance, see: <https://en.wikipedia.org/wiki/Heartbleed>

Example 191. (short plaintext attack on RSA) Suppose a 56bit DES key (or any other short plaintext) is written as a number $m \approx 2^{56} \approx 10^{16.9}$ and encrypted as $c = m^e \pmod{N}$.

Eve makes two lists:

- $cx^{-e} \pmod{N}$ for $x = 1, 2, \dots, 10^9$
- $y^e \pmod{N}$ for $y = 1, 2, \dots, 10^9$

If there is a match between the lists, that is $cx^{-e} = y^e \pmod{N}$, then $c = (xy)^e \pmod{N}$ and Eve has learned that the plaintext is $m = xy$.

This attack will succeed if m is the product of two integers x, y (up to 10^9). This is the case for many integers m .

Another project idea. Quantify how many integers factor into two small factors.

How to prevent? To prevent this attack, the plaintext can be padded with random bits before being encrypted. Recall that we should actually never use vanilla RSA (unless with random plaintexts) and always use a securely padded version instead!

Example 192. For RSA, does double (or triple) encryption improve security?

- Say, if Bob asks people to send him messages first encrypted with a first public key (N, e_1) and then encrypted with a second public key (N, e_2) .
- Or, what if Bob asks people to send him messages first encrypted with a first public key (N_1, e_1) and then encrypted with a second public key (N_2, e_2) .

Solution.

- No, this does not result in any additional security.

After one encryption, $c_1 = m^{e_1} \pmod{N}$ and the final ciphertext is $c_2 = c_1^{e_2} \pmod{N}$. However, note that $c_2 = m^{e_1 e_2} \pmod{N}$, which is the same as encryption with the single public key $(N, e_1 e_2)$.

- This adds only a negligible bit of security and hence is a bad idea as well. The reason is that an attacker able to determine the secret key for (N_1, e_1) is likely just as able to determine the secret key for (N_2, e_2) , meaning that the attack would only take twice as long (or two computers). That's only a tiny bit of security gained, somewhat comparable to increasing N from 1024 to 1025 bits. If heightened security is wanted, it is better to increase the size of N in the first place.

[Make sure you see how the situation here is different from the situation for 3DES.]

Example 193. (common modulus attack on RSA) Alice encrypts m using each of the RSA public keys (N, e_1) and (N, e_2) so that the ciphertexts are $c_1 = m^{e_1} \pmod{N}$ and $c_2 = m^{e_2} \pmod{N}$. Eve might be able to figure out m from c_1 and c_2 !! How and when?

Solution. The crucial observation is that $c_1^x c_2^y \equiv m^{e_1 x + e_2 y} = m^{e_1 x + e_2 y} \pmod{N}$. Eve can choose x and y .

She knows m if she can arrange x and y such that $e_1 x + e_2 y = 1$. This is possible if $\gcd(e_1, e_2) = 1$, in which case Eve would use the extended Euclidean algorithm to determine appropriate x and y .

A scenario. Bob's public RSA key is (N, e) . However, when Alice requests this public key from Bob, her message gets intercepted by Eve who instead sends (N, e_2) back to Alice, where e_2 differs from e in only one bit. Alice uses (N, e_2) to encrypt her message and sends c_2 to Bob. Of course, Bob fails to decrypt Alice's message and so resends his public key to Alice (this time, Eve doesn't intervene). Alice now uses (N, e) to encrypt her message and sends c to Bob.

Since $e - e_2 = \pm 2^r$, we have $\gcd(e, e_2) = 1$ (why?!), so that Eve can determine m as explained above.

Comment on that scenario. From a practical point of view, we can argue that, if Eve can trick Alice into using a modified version of Bob's public key, then she might as well give a completely new public key (that Eve created) to Alice, in which case she can immediately decipher c_2 . That's certainly true. However, that way, Eve's malicious intervention would be plainly visible as such.

Example 194. (chosen ciphertext attack on RSA) Show that RSA is not secure under a chosen ciphertext attack.

First of all, let us recall that in a chosen ciphertext attack, Eve has some access to a decryption device. In the present case, we mean the following: Eve is trying to determine m from c . Clearly, we cannot allow her to use the decryption device on c (because then she has m and nothing remains to be said). However, Eve is allowed to decrypt some other ciphertext c' of her choosing (hence, “chosen ciphertext”).

You may rightfully say that this is a strange attacker who can decrypt messages except the one of particular interest. This model is not meant to be realistic—instead, it is important for theoretical security considerations: if our cryptosystem is secure against this (adaptive) version of chosen ciphertext attacks, then it is also secure against any other reasonable chosen ciphertext attacks.

Solution. RSA is not secure under a chosen ciphertext attack:

Suppose $c = m^e \pmod{N}$ is the ciphertext for m .

Then, Eve can ask for the decryption m' of $c' = 2^e c \pmod{N}$. Since $c' = (2m)^e \pmod{N}$, Eve obtains $m' \equiv 2m$, from which she readily determines $m = 2^{-1}m' \pmod{N}$.

Comment. On the other hand, RSA-OAEP is provably secure against chosen ciphertext attacks. Recall that, in this case, m is padded prior to encryption. As a result, $2m$ or, more generally am , is not going to be a valid plaintext.

Example 195. What we just exploited is that RSA is **multiplicatively homomorphic**.

Multiplicatively homomorphic means the following: suppose m_1 and m_2 are two plaintexts with ciphertexts c_1 and c_2 . Then, (the residue) m_1m_2 has ciphertext c_1c_2 .

[That is, multiplication of plaintexts translates to multiplication of ciphertexts, and vice versa. Mathematically, this means that the map $m \rightarrow c$ is a homomorphism (with respect to multiplication).]

Indeed, for RSA, $c_1 = m_1^e$ and $c_2 = m_2^e$, so that $c_1c_2 = m_1^e m_2^e = (m_1m_2)^e \pmod{N}$ is the ciphertext for m_1m_2 .

Why care? In our previous example, being multiplicatively homomorphic was a weakness of RSA (which is “cured” by RSA-OAEP). However, there are situations where homomorphic ciphers are of practical interest. With a homomorphic cipher, we can do calculations using just the ciphertexts without knowing the plaintexts (for instance, the ciphertexts could be encrypted (secret) votes, which could be publicly posted; then anyone could add up (in an additively homomorphic system) these votes into a ciphertext of the final vote count; the advantage being that we don’t need to trust an authority for that count). The search for a fully **homomorphic encryption** scheme is a hot topic. For a nice initial read, you can find more at:

<https://blog.cryptographyengineering.com/2012/01/02/very-casual-introduction-to-fully/>

Example 196. (chosen ciphertext attack on ElGamal) Show that ElGamal is not secure under a chosen ciphertext attack.

Solution. Recall, again, that in a chosen ciphertext attack, Eve is trying to determine m from c and Eve has access to a decryption device, which she can use, except not to the ciphertext c in question.

Suppose $c = (c_1, c_2) = (g^y, g^{xy}m)$ is the ciphertext for m . Then $(c_1, 2c_2) = (g^y, g^{xy}2m)$ is a ciphertext for $2m$. Hence, Eve can ask for the decryption of $c' = (c_1, 2c_2)$, which gives her $m' = 2m$, from which she determines $m = 2^{-1}m' \pmod{p}$.

In fact, again, the reason that ElGamal is not secure under a chosen ciphertext attack is that it is multiplicatively homomorphic.

Example 197. Show that ElGamal is multiplicatively homomorphic.

Solution. Let $(g^{y_1}, g^{xy_1}m_1)$ be a ciphertext for m_1 , and $(g^{y_2}, g^{xy_2}m_2)$ a ciphertext for m_2 .

The product (component-wise) of the ciphertexts is $(g^{y_1+y_2}, g^{x(y_1+y_2)}m_1m_2)$, which is a ciphertext for m_1m_2 . So, again, the product of ciphertexts corresponds to the product of plaintexts.

A quick summary of some aspects of RSA and ElGamal.

- As long as appropriate key sizes are used, both RSA and ElGamal appear secure.
About the same key size needed for both: at least 1024 bits. By now, better 2048 bits.
- The security of both RSA and ElGamal can be compromised by using a cryptographically insecure PRG to generate the secret pieces p, q (for RSA) or x (for ElGamal).
- It is important to have different ciphers, especially ones that rely on the difficulty of different mathematical problems.
Comment. Factoring $N = pq$ and computing discrete logarithms modulo p are the two different problems for RSA and ElGamal, respectively. It is not known whether the ability to solve one of them would make it significantly easier to also solve the other one. However, historically, advances in factorization methods (like the number field sieve) have subsequently lead to similar advances in computing discrete logarithms. Both problems seem of comparable difficulty.
- Both are multiplicatively homomorphic, but RSA loses this property when padded.