

Application: hash functions

A hash function H is a function, which takes an input x of arbitrary length, and produces an output $H(x)$ of fixed length, say, b bit.

Example 193. (error checking) When Alice sends a long message m to Bob over a potentially noisy channel, she also sends the hash $H(m)$. Bob, who receives m' (which, he hopes is m) and h , can check whether $H(m') = h$.

Comment. This only protects against accidental errors in m (much like the check digits in credit card numbers we discussed earlier). If Eve intercepts the message $(m, H(m))$, she can just replace it with $(m', H(m'))$ so that Bob receives the message m' .

Eve's job can be made much more difficult by sending m and $H(m)$ via two different channels. For instance, in software development, it is common to post hashes of files on websites (or announce them otherwise), separately from the actual downloads. For that use case, we should use a one-way hash function (see next example).

- The hash function $H(x)$ is called **one-way** if, given y , it is computationally infeasible to compute m such that $H(m) = y$. [Also called **preimage-resistant**.]

Similarly, a hash function is called (weakly) **collision-resistant** if, given a message m , it is difficult to find a second message m' such that $H(m) = H(m')$. [Also called **second preimage-resistant**.]

- It is called **(strongly) collision-resistant** if it is computationally infeasible to find two messages m_1, m_2 such that $H(m_1) = H(m_2)$.

Comment. Every hash function must have many collisions. On the other hand, the above requirement says that finding even one must be exceedingly difficult.

Example 194. (error checking, cont'd) Alice wants to send a message m to Bob. She wants to make sure that nobody can tamper with the message (maliciously or otherwise). How can she achieve that?

Solution. She can use a one-way hash function H , send m to Bob, and publish (or send via some second route) $y = H(m)$. Because H is one-way, Eve cannot find a value m' such that $H(m') = y$.

Some applications of hash functions include:

- **error-checking:** send m and $H(m)$ instead of just m
- **tamper-protection:** send m and $H(m)$ via different channels (H must be one-way!)
If H is one-way, then Eve cannot find m' such that $H(m') = H(m)$, so she cannot tamper with m without it being detected.
- **password storage:** discussed later (there are some tricky bits)
- **digital signatures:** more later
- **blockchains:** used, for instance, for cryptocurrencies such as Bitcoin

Some popular hash functions:

	published	output bits	comment
CRC32	1975	32	not secure but common for checksums
MD5	1992	128	common; used to be secure (now broken)
SHA-1	1995	160	common; used to be secure (collision found in 2017)
SHA-2	2001	256/512	considered secure
SHA-3	2015	arbitrary	considered secure

- CRC is short for **Cyclic Redundancy Check**. It was designed for protection against common transmission errors, not as a cryptographic hash function (for instance, CRC is a linear function).
- SHA is short for **Secure Hash Algorithm** and (like DES and AES) is a federal standard selected by NIST. SHA-2 is a family of 6 functions, including SHA-256 and SHA-512 as well as truncations of these. SHA-3 is not meant to replace SHA-2 but to provide a different alternative (especially following successful attacks on MD5, SHA-1 and other hash functions, NIST initiated an open competition for SHA-3 in 2007). SHA-3 is based on Keccak (like AES is based on Rijndael; Joan Daemen involved in both). Although the output of SHA-3 can be of arbitrary length, the number of security bits is as for SHA-2.
https://en.wikipedia.org/wiki/NIST_hash_function_competition
- MD is short for **Message Digest**. These hash functions are due to Ron Rivest (MIT), the “R” in RSA. Collision attacks on MD5 can now produce collisions within seconds. For a practical exploit, see: [https://en.wikipedia.org/wiki/Flame_\(malware\)](https://en.wikipedia.org/wiki/Flame_(malware))
 MD6 was submitted as a candidate for SHA-3, but later withdrawn.

Constructions of hash functions

Recall that a hash function H is a function, which takes an input x of arbitrary length, and produces an output $H(x)$ of fixed length, say, b bit.

Example 195. (Merkle–Damgård construction) Similarly, a **compression function** \tilde{H} takes input x of length $b + c$ bits, and produces output $\tilde{H}(x)$ of length b bits. From such a function, we can easily create a hash function H . How?

Importantly, it can be proved that, if \tilde{H} is collision-resistant, then so is the hash function H .

Solution. Let x be an arbitrary input of any length. Let’s write $x = x_1x_2x_3\dots x_n$, where each x_i is c bits (if necessary, pad the last block of x so that it can be broken into c bit pieces).

Set $h_1 = 0$ (or any other initial value), and define $h_{i+1} = \tilde{H}(h_i, x_i)$ for $i \geq 1$. Then, $H(x) = h_{n+1}$ (b bits).

[In $\tilde{H}(h_i, x_i)$, we mean that the b bits for h_i are concatenated with the c bits for x_i , for a total of $b + c$ bits.]

Comment. This construction is known as a Merkle–Damgård construction and is used in the design of many hash functions, including MD5 and SHA-1/2.

Careful padding. Some care needs to be applied to the padding. Just padding with zeroes would result in easy collisions (why?), which we would like to avoid. For more details:

https://en.wikipedia.org/wiki/Merkle-Damgård_construction

Example 196. Consider the compression function $\tilde{H}: \{3 \text{ bits}\} \rightarrow \{2 \text{ bits}\}$ defined by

x	000	001	010	011	100	101	110	111
$\tilde{H}(x)$	00	10	11	01	10	00	01	11

[This was not chosen randomly: the first output bit is the sum of the digits, and the second output bit is just the second input bit.]

- Find a collision of \tilde{H} .
- Let H be the hash function obtained from \tilde{H} using the Merkle–Damgård construction (using initial value $h_1 = 0$). Compute $H(1101)$.
- Find a collision with $H(1101)$.

Solution.

- For instance, $\tilde{H}(001) = \tilde{H}(100)$.
- Here, $b = 2$ and $c = 1$, so that each x_i is 1 bit: $x_1x_2x_3x_4 = 1101$.
 $h_1 = 00$
 $h_2 = \tilde{H}(h_1, x_1) = \tilde{H}(001) = 10$
 $h_3 = \tilde{H}(h_2, x_2) = \tilde{H}(101) = 00$
 $h_4 = \tilde{H}(h_3, x_3) = \tilde{H}(000) = 00$
 $h_5 = \tilde{H}(h_4, x_4) = \tilde{H}(001) = 10$
Hence, $H(1101) = h_5 = 10$.
- Our computation above shows that, for instance, $H(1) = 10 = H(1101)$.

The construction of good hash functions is linked to the construction of good ciphers. Below, we indicate how to use a block cipher to construct a hash function.

Why linked? The ciphertext produced by a good cipher should be indistinguishable from random bits. Similarly, the output of a cryptographic hash function should look random, because the presence of patterns would likely allow us to compute preimages or collisions.

However. The design goals for a hash function are somewhat different than for a cipher. It is therefore usually advisable to not crossbreed these constructions and, instead, to use a specially designed hash function like SHA-2 when a hash function is needed for cryptographic purposes.

First, however, a cautionary example.

Example 197. (careful!) Let E_k be encryption using a block cipher (like AES). Is the compression function \tilde{H} defined by

$$\tilde{H}(x, k) = E_k(x)$$

one-way?

Solution. No, it is not one-way.

Indeed, given y , we can produce many different (x, k) such that $\tilde{H}(x, k) = y$ or, equivalently, $E_k(x) = y$. Namely, pick any k , and then choose $x = D_k(y)$.

Example 198. Let E_k be encryption using a block cipher (like AES). Then the compression function \tilde{H} defined by

$$\tilde{H}(x, k) = E_k(x) \oplus x$$

is usually expected to be collision-resistant.

Let us only briefly think about whether \tilde{H} might have the weaker property of being one-way (as opposed to the previous example). For that, given y , we try to find (x, k) such that $\tilde{H}(x, k) = y$ or, equivalently, $E_k(x) \oplus x = y$. This seems difficult.

Just getting a feeling. We could try to find such (x, k) with $x = 0$. In that case, we need to arrange k such that $E_k(0) = y$. For a block cipher like AES, this seems difficult. In fact, we are trying a known-plaintext attack on the cipher here: assuming that $m = 0$ and $c = y$, we are trying to determine the key k . A good cipher is designed to resist such an attack, so that this approach is infeasible.

Comment. Combined with the Merkle–Damgård construction, you can therefore use AES to construct a hash function with 128 bits output size. However, as indicated before, it is advisable to use a hash function designed specifically for the purpose of hashing.

For several other (more careful) constructions of hash functions from block ciphers, you can check out Chapter 9.4.1 in the *Handbook of Applied Cryptography* (Menezes, van Oorschot and Vanstone, 2001), freely available at: <http://cacr.uwaterloo.ca/hac/>

We have seen how hash functions can be constructed from block ciphers (though this is usually not advisable). Similarly, hash functions can be used to build PRGs (and hence, stream ciphers).

Example 199. A hash function $H(x)$, producing b bits of output, can be used to build a PRG as follows. Let x_0 be our b bit seed. Set $x_n = H(x_{n-1})$, for $n \geq 1$, and $y_n = x_n \pmod{2}$. Then, the output of the PRG are the bits $y_1y_2y_3\dots$

Comment. As for the B-B-S PRG, if b is large, it might be OK to extract more than one bit from each x_n .

Comment. Technically speaking, we should extract a “hardcore bit” y_n from x_n .

Comment. It might be a little bit better to replace the simple rule $x_n = H(x_{n-1})$ with $x_n = H(x_0, x_{n-1})$. Otherwise, collisions would decrease the range during each iteration. However, if b is large, this should not be a practical issue. (Also, think about how this alleviates the issue in the next example.)

Comment. Of course, one might then use this PRG as a stream cipher (though this is probably not a great idea, since the design goals for hash functions and secure PRGs are not quite the same). Our book lists a similar construction in Section 8.7: starting with a seed $x_0 = k$, bytes x_n are created as follows $x_1 = H(x_0)$ and $x_n = L_8(H(x_0, x_{n-1}))$, where L_8 extracts the leftmost 8 bits. The output of the PRG is $x_1x_2x_3\dots$. However, can you see the flaw in this construction? (Hint: it repeats very soon!)

Example 200. Suppose, with the same setup as in the previous example, we let our PRG output $x_1x_2x_3\dots$, where each x_n is b bits. What is your verdict?

Solution. This PRG is not unpredictable (at all). After b bits have been output, x_1 is known and $x_2 = H(x_1)$ can be predicted perfectly. Likewise, for all the following output.

Comment. While completely unacceptable for cryptographic purposes, this might be a fine PRG for other purposes that do not need unpredictability.

Here is a compression function, which is provably strongly collision-resistant.

However, it is rather slow and so is not practical for hashing larger data. On the other hand, its slowness could be beneficial for applications like password hashing.

Example 201. (the discrete log hash) Let p be a large safe prime (that is, $q = (p - 1) / 2$ is also prime). Let g_1, g_2 be two primitive roots modulo p . Define the compression function h as:

$$h: \{0, 1, \dots, q^2 - 1\} \rightarrow \{1, 2, \dots, p - 1\}, \quad h(m_1 + m_2q) = g_1^{m_1} g_2^{m_2} \pmod{p}.$$

[Note that, although not working with inputs and outputs of certain size in bits, this is a compression function, because the input space is much larger than the output space.]

Show that finding a collision of h is as difficult as determining the discrete logarithm x in $g_1^x = g_2 \pmod{p}$.

Solution. Suppose we have a collision: $g_1^{m_1} g_2^{m_2} \equiv g_1^{m'_1} g_2^{m'_2} \pmod{p}$

Hence, $g_1^{(m_1 - m'_1) + (m_2 - m'_2)x} \equiv 1 \pmod{p}$ or, equivalently, $(m_1 - m'_1) + (m_2 - m'_2)x \equiv 0 \pmod{p - 1}$ (because g_1 is a primitive root and so has order $p - 1$).

This final congruence can now be solved for x .

More precisely, if $d = \gcd(m_2 - m'_2, p - 1)$, there are actually d solutions for x . Since we chose p to be safe, the only factors of $p - 1$ are $1, 2, (p - 1) / 2, p - 1$.

Since $|m_2 - m'_2| < q$, the only possibilities are $d = 1, 2$ (unless $m_2 = m'_2$; however, this cannot be the case since then also $m_1 = m'_1$, so that we wouldn't have a collision in the first place).

Passwords

Let's say you design a system that users access using personal passwords. Somehow, you need to store the password information.

- The worst thing you can do is to actually store the passwords m .

This is an absolutely atrocious choice, even if you take severe measures to protect (e.g. encrypt) the collection of passwords.

Comment. Sadly, there are still systems out there doing that. An indication that this might* be happening is systems that require you to update passwords and then complain that your new password is too close to the original one. Any reasonably designed system should never learn about your actual password in the first place!

*: On the other hand, think about how you could check for (certain kinds of) closeness of passwords without having to store the actual password.

- Better, but still terrible, is to instead store hashes $H(m)$ of the passwords m .

Good. An attacker getting hold of the password file, only learns about the hash of a user's password. Assuming the hash function is one-way, it is infeasible for the attacker to determine the corresponding password (if the password was random!!).

Still bad. However, passwords are (usually) not random. Hence, an attacker can go through a list of common passwords (dictionary attack), compute the hashes and compare with the hashes of users (similarly, a brute-force attack can simply go through all possible passwords).

Even worse, it is immediately obvious if two users are using the same password (or, if the same user is using the same password for different services using the same hash function).

Comment. So, storing password hashes is not OK unless all passwords are completely random.

- Better, a random value s is generated for each user, and then s and $H(m, s)$ are stored. The value s is referred to as **salt**.

In other words, instead of storing the hash of the password m , we are storing the hash of the salted password, as well as the salt.

Why? Two users using the same password would have different salt and hence different hashes stored. As a consequence, an attacker can (of course) still mount a dictionary or brute-force attack but only against a single user, not all users at once.

Comment. Note how the concept of salt is similar to a nonce.

Comment. To be future-proof, the hash+salt is often stored in a single field in a format like (hash-algo, salt, salted hash).

Comment. There's also the concept of **pepper** (usually, sort of a secret salt). This provides extra security if the pepper is stored separately. [Sometimes pepper is used as a sort of small random salt, which is discarded; this only slows a brute-force attack down and should instead be addressed using the item below.]

[https://en.wikipedia.org/wiki/Pepper_\(cryptography\)](https://en.wikipedia.org/wiki/Pepper_(cryptography))

- Finally, we should not use the usual (fast!) hash functions like SHA-2.

Why? One of the things that makes SHA-2 a good hash function in practice is its speed. However, that actually makes SHA-2 a poor choice in this context of password hashing. An attacker can compute billions of hashes per second, which makes a dictionary or brute-force attack very efficient.

To make a dictionary or brute-force attack impractical, the hashing needs to be slowed down. See Example 202 for some scary numbers.

Hashing functions like SHA-2 are not secure password hashing algorithms.

Instead, options that are considered secure include: PBKDF2, bcrypt, scrypt.

Comment. For instance, WPA2 uses PBKDF2 based on SHA-1 with 4096 (fairly small!) iterations.

Comment. Only increasing the number of iterations increases computation time but not memory usage. scrypt is designed to also consume an arbitrarily specified amount of memory.

For a nice discussion about password hashing:

<https://security.stackexchange.com/questions/211/how-to-securely-hash-passwords>

Example 202. (the power of brute-force) In April 2021, the Bitcoin hashrate is about $150E=1.5 \cdot 10^{20}$ hashes per second. How long would it take to brute-force a (completely random!) 8 character password, using all 94 printable ASCII characters (excluding the space)?

Solution. There are $94^8 \approx 6.1 \cdot 10^{15}$ possible passwords. Hence, it would take about 0.00004 seconds!

Comment. Even using 10 random characters (almost no human password has that kind of entropy), there are $94^{10} \approx 5.4 \cdot 10^{19}$ possible passwords. It would take less than 0.36 seconds to go through all of these!

Comment. <https://bitinfocharts.com/comparison/bitcoin-hashrate.html>

Example 203. Your king's webserver contains the following code to check whether the king is accessing the server.

[As is far too common, his password derives from his girlfriend's name and year of birth.]

```
def check_is_king(password):
    return password=="Ludmilla1310"
```

Obviously, anyone who might be able to see the code (including its binary version) learns about your king's password. With minimal change, how can this be fixed?

Solution. The password should be hashed. For instance, in Python, using SHA-2 (why is that actually not a good choice here?) with 256 output bits:

```
from hashlib import sha256
def check_is_king(password):
    phash = sha256(password).hexdigest()
    return phash == "9e4b4fe180e22bc6cdf01fe9711cf2558507e5c3ae1c3c1f6607a25741941c66"
```

Comment. 256 bits are 64 digits in hexadecimal.

Python comment. Of course, a real implementation would use `digest()` instead of `hexdigest()`.

[For Python 3, if operating with strings (instead of bytes), `sha256(password)` needs to be replaced with something like `sha256(password.encode('utf-8'))`.]

Why is SHA not good here? Too fast to discourage brute-force attacks.

Example 204. Suppose you don't like the idea of creating random salt.

- (a) How about using the same salt for all your users?
- (b) Is it a good idea to use the username as salt?

Solution.

- (a) This is a terrible idea and defeats the purpose of a salt. (For instance, again an attacker can immediately see if users have the same password.)

Comment. Essentially, this is a form of pepper (if the value is kept secret, i.e. stored elsewhere).

- (b) That is a reasonable idea. One reason against it is that, ideally, the salt should be unique (globally). However, this could be easily achieved by using the username combined with something identifying your service (like your hostname).

Comment. A possible practical reason against choosing the username for salt is that the username might change.

Example 205. You need to hash (salted) passwords for storage. Unfortunately, you only have SHA-2 available. What can you do?

Solution. Iterate many times! (In order to slow down the computation of the hash.) The naive way would be to simply set $h_0 = H(m)$ and $h_{n+1} = H(h_n)$. Then use as hash the value h_N for large N .

In current applications, it is typical to choose N on the order of 100,000 or higher (depending on how long is reasonable to have your user wait each time she logs in and needs her password hashed for verification).

Application: digital signatures

Goal: Using a private key, known only to her, Alice can attach her **digital signature** s to any message m . Anyone knowing her public key can then verify that it could only have been Alice, who signed the message.

- Consequently, in contrast to usual signatures, digital signatures must depend on the message to be signed so that they cannot be simply reproduced by an adversary.
- This should sound a lot like public-key cryptography!

Cryptographically speaking, a digitally signed message (m, s) from Alice to Bob achieves:

- **integrity**: the message has not been accidentally or intentionally modified
- **authenticity**: Bob can be confident the message came from Alice

In fact, we gain even more: not only is Bob assured that the message is from Alice, but the evidence can be verified by anyone. We have “proof” that Alice signed the message. This is referred to as **non-repudiation**. We refer to a technical not a legal term here: if you are curious about legal aspects of digital signatures, see, e.g.:

<https://security.stackexchange.com/questions/1786/how-to-achieve-non-repudiation/6108>

For comparison, sending a message with its hash $(m, H(m))$ only achieves integrity.

Example 206. (authentication using digital signatures) Last time, we saw that using human generated and memorized passwords is problematic even if done “right” (Example 202). Among other things, digital signatures provide an alternative approach to authentication.

Authentication. If Alice wants to authenticate herself with a server, the server sends her a (random) message. Using her private key, Alice signs this message and sends it back to the server. The server then verifies her (digital) signature using Alice’s public key.

Obvious advantage. The server (like everyone else) doesn’t know Alice’s secret, so it cannot be stolen from the server (of course, Alice still needs to protect her secret from it being stolen).

(RSA signatures) Let H be a collision-resistant hash function.

- Alice creates a public RSA key (N, e) . Her (secret) private key is d .
- Her signature of m is $s = H(m)^d \pmod{N}$.
- To verify the signed message (m, s) , Bob checks that $H(m) = s^e \pmod{N}$.

This is secure if RSA is secure against known plaintext attacks.

Example 207. We use the silly hash function $H(x) = x \pmod{10}$. Alice’s public RSA key is $(N, e) = (33, 3)$, her private key is $d = 7$.

- How does Alice sign the message $m = 12345$?
- How does Bob verify her message?
- Was the message $(m, s) = (314, 2)$ signed by Alice?

Solution.

- (a) $H(m) = 5$. The signature therefore is $s = 5^7 \pmod{33}$ (note how computing that signature requires Alice's private key). Computing that, we find $s = 14$.
- (b) Bob receives the signed message $(m, s) = (12345, 14)$.
He computes $H(m) = 5$ and then checks whether $H(m) \equiv s^3 \pmod{33}$ (for which he only needs Alice's public key). Indeed, $14^3 \equiv 5 \pmod{33}$, so the signature checked out.
- (c) We compute $H(m) = 4$ and then need to check whether $H(m) \equiv s^3 \pmod{33}$. Since $2^3 \equiv 8 \not\equiv 4 \pmod{33}$, the signature does not check out. Alice didn't sign the message.

Just to make sure. What's a collision of our hash function? Why is it totally not one-way?

Example 208. Why should Alice sign the hash $H(m)$ and not the message m ?

Solution. A practical reason is that signing $H(m)$ is simpler/faster. The message m could be long, in which case we would have to do something like chop it into blocks and sign each block (but then Eve could rearrange these, so we would have to do something more clever, like for block ciphers). In any case, we shouldn't just sign $m \pmod{N}$ because then Eve can just replace m with any m' as long as $m \equiv m' \pmod{N}$.

There is another issue though. Namely, Eve can do the following **no message attack**: she starts with any signature s , then computes $m = s^e \pmod{N}$. Everyone will then believe that (m, s) is a message signed by Alice. This does not work if H is a one-way function: Eve now needs to find m such that $H(m) = s^e \pmod{N}$, but she fails to find such m if H is one-way.

Example 209. Is it enough if the hash function for signing is one-way but not collision-resistant?

Solution. No, that is not enough. If there is a collision $H(m) = H(m')$, then Eve can ask Alice to sign m to get (m, s) and later replace m with m' , because (m', s) is another valid signed message. (See also the comments after the discussion of birthday attacks.)

Comment. This question is of considerable practical relevance, since hash functions like MD5 and SHA-1 have been shown to not be collision-resistant (but are still considered essentially preimage-resistant, that is, one-way). In the case of MD5, this has been exploited in practice:

https://en.wikipedia.org/wiki/MD5#Collision_vulnerabilities

Example 210. Alice uses an RSA signature scheme and the (silly) hash function $H(x) = x_1 + x_2$, where $x_1 = x^{-1} \pmod{11}$ and $x_2 = x^{-1} \pmod{7}$ [with 0^{-1} interpreted as 0] to produce the signed message $(100, 13)$. Forge a second signed message.

Solution. Since we have no other information, in order to forge a signed message, we need to find another message with the same hash value as $m = 100$. From our experience with the Chinese remainder theorem, we realize that changing x by $7 \cdot 11$ does not change $H(x)$. Hence, a second signed message is $(177, 13)$.

Comment. The hash $H(m)$ for $m = 100$ is $H(100) = (100^{-1})_{\text{mod}11} + (100^{-1})_{\text{mod}7} = 1 + 4 = 5$.

Similar to what we did with RSA signatures, one can use ElGamal as the basis for digital signatures. A variation of that is the **DSA** (digital signature algorithm), another federal standard.

https://en.wikipedia.org/wiki/ElGamal_signature_scheme

https://en.wikipedia.org/wiki/Digital_Signature_Algorithm

Not surprisingly, the hashes mandated for DSA are from the SHA family.

Birthday paradox and birthday attacks

Example 211. (birthday paradox) Among $n = 35$ people (a typical class size; no leaplings), how likely is it that two have the same birthday?

Solution.

$$1 - \left(1 - \frac{1}{365}\right)\left(1 - \frac{2}{365}\right)\left(1 - \frac{3}{365}\right)\dots\left(1 - \frac{34}{365}\right) \approx 0.814$$

If the formula doesn't speak to you, see Section 8.4 in our book for more details or checkout:

https://en.wikipedia.org/wiki/Birthday_problem

Comment. For $n = 50$, we get a 97.0% chance. For $n = 70$, it is 99.9%.

Comment. In reality, birthdays are not distributed quite uniformly, which further increases these probabilities. Also note that, for simplicity, we excluded "leaplings", people born on February 29.

How is this relevant to crypto, and hashes in particular?

Think about people as messages and birthdays as the hash of a person. The birthday paradox is saying that "collisions" occur more frequently than one might expect.

Example 212. Suppose M is large ($M = 365$ in the birthday problem) compared to n . The probability that, among n selections from M choices, there is no collision is

$$\left(1 - \frac{1}{M}\right)\left(1 - \frac{2}{M}\right)\dots\left(1 - \frac{n-1}{M}\right) \approx e^{-\frac{n^2}{2M}}$$

Why? For small x , we have $e^x \approx 1 + x$ (the tangent line of e^x at $x=0$). Hence, $\left(1 - \frac{k}{M}\right) \approx e^{-k/M}$ and

$$\left(1 - \frac{1}{M}\right)\left(1 - \frac{2}{M}\right)\dots\left(1 - \frac{n-1}{M}\right) \approx e^{-1/M}e^{-2/M}\dots e^{-(n-1)/M} = e^{-(1+2+\dots+(n-1))/M} = e^{-\frac{n(n-1)}{2M}}$$

In the last step, we used that $1 + 2 + \dots + (n-1) = \frac{n(n-1)}{2}$. Finally, $e^{-\frac{n(n-1)}{2M}} \approx e^{-\frac{n^2}{2M}}$.

Decent approximation? For instance, for $M = 365$ and $n = 35$, we get 0.813 for the chance of a collision (instead of the true 0.814).

Important observation. If $n \approx \sqrt{M}$, then the probability of no collision is about $e^{-1/2} \approx 0.607$. In other words, a collision is quite likely!

In the context of hash functions, this means the following: if the output size is b bits, then there are $M = 2^b$ many possible hashes. If we make a list of $\sqrt{M} = 2^{b/2}$ many hashes, we are likely to observe a collision.

For collision-resistance, the output size of a hash function needs to have twice the number of bits that would be necessary to prevent a brute-force attack.

Practical relevance. This is very important for every application of hashes which relies on collision-resistance, such as digital signatures.

For instance, think about Eve trying to trick Alice into signing a fraudulent contract m . Instead of m , she prepares a different contract m' that Alice would be happy to sign. Eve now creates many slight variations of m and m' (for instance, by varying irrelevant things like commas or spaces) with the hope of finding \tilde{m} and \tilde{m}' such that $H(\tilde{m}) = H(\tilde{m}')$. (Why?!!)

Example 213. (chip based credit cards) Modern chip based credit cards use digital signatures to authenticate a payment.

How? The card carries a public key, which is signed by the bank, so that a merchant can verify the public key. The card then signs a challenge from the merchant for authentication. The private key used for that is not even known to the bank.

Note that all of this can be done offline, without needing to contact the bank during the transaction.

<https://en.wikipedia.org/wiki/EMV>

There's an interesting and curious story made possible by the fact that, around 2000, banks in France used 320 bit RSA (chosen in the 80s and then not fixed despite expert advice) for signing the card's public key:

https://en.wikipedia.org/wiki/Serge_Humpich

Comment. For contrast, the magnetic stripe just contains the card information such as card number.

Comment. This also leads to interesting questions like: can we embed a private key in a chip (or code) in such a way that an adversary, with full access to the circuit (or code), still cannot extract the key?

https://en.wikipedia.org/wiki/Obfuscation#White_box_cryptography

A digital signature is like a hash, which can only be created by a single entity (using a private key) but which can be verified by anyone (using a public key).

As one might expect, a symmetric version of this idea is also common:

Example 214. (MAC) A **message authentication code**, also known as a **keyed hash**, uses a private key k to compute a hash for a message.

Like a hash, a MAC provides integrity. Further, like a digital signature, it provides authenticity because only parties knowing the private key are able to compute the correct hash.

Comment. On the other hand, a MAC does not offer non-repudiation because several parties know the private key (whether non-repudiation is desirable or undesirable depends on the application). Hence, it cannot be proven to a third party who among those computed the MAC (and, in any case, such a discussion would make it necessary to reveal the private key, which is usually unacceptable).

From hash to MAC. If you have a cryptographic hash function H , you can simply produce a MAC $M_k(x)$ (usually referred to as a HMAC) as follows:

$$M_k(x) = H(k, x)$$

This seems to work fine for instance for SHA-3. On the other hand, this does not appear quite as secure for certain other common hash functions. Instead, it is common to use $M_k(x) = H(k, H(k, x))$. For more details, see:

https://en.wikipedia.org/wiki/Hash-based_message_authentication_code

From ciphers to MAC. Similarly, we can also use ciphers to create a MAC. See, for instance:

<https://en.wikipedia.org/wiki/CBC-MAC>

Elliptic curve cryptography

The idea of Diffie–Hellman (used, for instance, in DH key exchange, ElGamal or DSA) can be carried over to algebraic structures different from multiplication modulo p .

Recall that the key idea is, starting from individual secrets x, y , to share g^x, g^y modulo p in order to arrive at the joint secret $g^{xy} \pmod{p}$. That's using multiplication modulo p .

One important example of other such algebraic structures, for which the analog of the discrete logarithm problem is believed to be difficult, are elliptic curves.

https://en.wikipedia.org/wiki/Elliptic_curve_cryptography

Comment. The main reason (apart from, say, diversification) is that this leads to a significant saving in key size and speed. Whereas, in practice, about 2048bit primes are needed for Diffie–Hellman, comparable security using elliptic curves is believed to only require about 256bits.

For a beautiful introduction by Dan Boneh, check out the presentation:

https://www.youtube.com/watch?v=4M8_0o71piA

Points on elliptic curves

An **elliptic curve** is a (nice) cubic curve that can (typically) be written in the form

$$y^2 = x^3 + ax + b.$$

A point (x, y) is on the elliptic curve if it satisfies this equation. Each elliptic curve also contains the special point O (“the point at ∞ ”). [O will act as the neutral element when “adding points”]

Advanced comment. Sometimes it is useful (or necessary) to consider elliptic curves defined by more general cubic equations such as $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$ (however, in most cases, a linear change of variables can transform this equation into the simpler form $y^2 = x^3 + ax + b$ mentioned above).

Example 215. Determine some points (x, y) on the elliptic curve E , described by

$$y^2 = x^3 - x + 9.$$

Solution. We can try some small values for x (say, $x = 0, x = 1, x = 2, \dots$) and see what y needs to be in order to get a point on the elliptic curve. For instance, for $x = 1$, we get $x^3 - x + 9 = 9$ which implies that $(1, \pm 3)$ are points on the elliptic curve.

Doing so, we find the integral points $(0, \pm 3), (\pm 1, \pm 3)$.

On the other hand, for $x = 2$, we get $x^3 - x + 9 = 15$ which implies that $(2, \pm\sqrt{15})$ are points on the elliptic curve. However, for cryptographic purposes, we are usually not interested in such irrational points.


Much less obvious rational points include $(35, 207)$ or $(\frac{1}{36}, \frac{647}{216})$ (see Example 217).

Comment. In general, it is a very difficult problem to determine all rational points on an elliptic curve, and lots of challenges remain open in that arena.

Example 216. Plot the elliptic curve E , described by $y^2 = x^3 - x + 9$ and mark the point $(1, 3)$.

Solution. We let Sage do the work for us:

```
>>> E = EllipticCurve([-1,9])
>>> E.plot() + E(1,3).plot(pointsize=50, rgbcolor=(1,0,0))
```



Adding points on elliptic curves

Note. Simply adding the coordinates of two points P and Q on an elliptic curve will (almost always) not result in a third point on the elliptic curve. However, we will define a more fancy “addition” of points, which we will denote $P \boxplus Q$, such that the $P \boxplus Q$ is on the elliptic curve as well.

Given a point $P = (x, y)$ on E , we define $-P = (x, -y)$ which is another point on E .

Let us introduce an operation \boxplus in the following geometric fashion: given two points P, Q , the line through these two points intersects the curve in a third point R .

We then define $P \boxplus Q = -R$.

We remark that $P \boxplus (-P)$ is the point O “at ∞ ”. That’s the neutral (zero) element for \boxplus .

How does one define $P \boxplus P$? (Tangent line!)

Remarkably, the “addition” $P \boxplus Q$ is associative. (This is not obvious from the definition.)

Using \boxplus , we can construct new points: for instance, $(0, 3) \boxplus (1, -3) = (35, 207)$ as we will verify in the next example.

Easier to verify (but not producing anything new) is $(0, 3) \boxplus (1, 3) = (-1, -3)$.

Example 217. Consider again the elliptic curve E , described by $y^2 = x^3 - x + 9$.

- Determine $(0, 3) \boxplus (1, 3)$.
- Determine $(0, 3) \boxplus (1, -3)$.
- Determine $4(0, 3)$, which is short for $(0, 3) \boxplus (0, 3) \boxplus (0, 3) \boxplus (0, 3)$.

Solution. We let Sage do the work for us:

```
>>> E = EllipticCurve([-1,9])
>>> E(0,3) + E(1,3)
(-1:-3:1)
>>> E(0,3) + E(1,-3)
(35:207:1)
>>> 4*E(0,3)
(-1677023/60279696, 1406201395535/468011559744, 1)
```

We conclude that $(0, 3) \boxplus (1, 3) = (-1, -3)$ and $(0, 3) \boxplus (1, -3) = (35, 207)$ (one of the points mentioned in Example 215), while

$$4(0, 3) = \left(-\frac{1677023}{60279696}, \frac{1406201395535}{468011559744} \right).$$

Comment. Note how Sage represents the point (x, y) as $(x: y: 1)$. These are **projective coordinates** which make it easier to incorporate the special point O which is represented by $(0: 1: 0)$.

https://en.wikipedia.org/wiki/Projective_coordinates

The following computation demonstrates that adding O doesn't do anything:

```
>>> E(O)
(0:1:0)
>>> E(0,3) + E(O)
(0:3:1)
```

Comment. Note that, starting with a single point such as $(0, -3)$, we can generate other points such as $2(0, -3) = \left(\frac{1}{36}, \frac{647}{216} \right)$ (one of the points mentioned in Example 215). If the initial point is rational then so are the points generated from it.

Advanced comment. If you want to dig deeper, you can try to translate the geometric description of the addition $P \boxplus Q$ into algebra by deriving equations for the coordinates of $P \boxplus Q = (x_r, y_r)$ in terms of the coordinates of $P = (x_p, y_p)$ and $Q = (x_q, y_q)$. For instance, for the elliptic curve $y^2 = x^3 + ax + b$, one finds that

$$\begin{aligned}x_r &= \lambda^2 - x_p - x_q, \\y_r &= \lambda(x_p - x_r) - y_p,\end{aligned}$$

where $\lambda = (y_q - y_p)/(x_q - x_p)$ is the slope of the line connecting P and Q . If P and Q are the same point, then this line becomes the tangent line and the slope becomes $\lambda = (3x_p^2 + a)/(2y_p)$ instead. For more details:

https://en.wikipedia.org/wiki/Elliptic_curve

From these formulas, can you reproduce the computations we did in Sage?

Review. Addition $P \boxplus Q$ of points P, Q on an elliptic curve.

For cryptographic purposes, elliptic curves are usually considered modulo a (large) prime p .

Example 218. Let us consider $y^2 = x^3 - x + 9$ (the elliptic curve from the previous examples) modulo 7. List all points on that curve.

Solution. Note that, because we are working modulo 7, there are only 7 possible values for each of x and y . Hence, we can just go through all $7^2 = 49$ possible points (x, y) to find all points on the curve.

Or, better, we go through all possibilities for x (such as $x = 2$) and determine the corresponding possible values for y (if $x = 2$, then $y^2 = 2^3 - 2 + 9 = 15 \equiv 1 \pmod{7}$ which has solutions $y \equiv \pm 1 \pmod{7}$).

Doing so, we find 9 points: $O, (0, \pm 3), (\pm 1, \pm 3), (2, \pm 1)$.

[Recall that O is the special point “at ∞ ” which serves as the neutral element with respect to \boxplus .]

Comment. A theorem of Hasse–Weil says that the number of points on an elliptic curve modulo p is always close to p (this is indeed what we expect because, for each of the p choices for x , we get an equation of the form $y^2 \equiv a \pmod{p}$ which has 2 solutions if a is a nonzero quadratic residue [and for a random a the odds are about 50% that it is quadratic]). Moreover, we can compute the exact number of points very efficiently.

By taking everything modulo 7, we still have the previously introduced addition rule \boxplus .

For instance. $(0, 3) \boxplus (1, -3) = (35, 207) \equiv (0, -3) \pmod{7}$

Here is how we can use Sage to list all points, as well as add any two of them:

```
>>> E7 = EllipticCurve(GF(7), [-1,9])
>>> E7.points()
[(0: 1: 0), (0: 3: 1), (0: 4: 1), (1: 3: 1), (1: 4: 1), (2: 1: 1), (2: 6: 1), (6: 3: 1), (6: 4: 1)]
>>> E7(0,3) + E7(1,-3)
(0: 4: 1)
>>> E7(1,-3) + E7(0,-3)
(6: 3: 1)
```

Multiples of a point are simply denoted with nP . For instance, $3P = P \boxplus P \boxplus P$.

We then have a version of the **discrete logarithm** problem for elliptic curves:

(discrete logarithm) Given P, xP on an elliptic curve, determine x .
(computational Diffie–Hellman) Given P, xP, yP on an elliptic curve, determine $(xy)P$.

Comment. Interestingly, it appears that the computational Diffie–Hellman problem (CDH) is more difficult for elliptic curves modulo p than for regular multiplication modulo p . Indeed, suppose that p is an n -digit prime. Then the best known algorithms for regular CDH modulo p has runtime $2^{O(\sqrt[3]{n})}$, whereas the best algorithm for the elliptic curve CDH modulo p has runtime $\sqrt{p} \approx 2^{n/2} = 2^{O(n)}$.

As a consequence, it is believed that a smaller prime p can be used to achieve the same level of security when using elliptic curve Diffie–Hellman (ECDH). In practice 256bit primes are used, which is believed to provide security comparable to 2048bit regular Diffie–Hellman (DH); this makes ECDH about ten times faster in practice than DH.

Comment. On the other hand, due to that reduced bit size, quantum computing attacks on elliptic curve cryptography, if they become available, would be more feasible compared to attacks on ElGamal/RSA.

Comment. Apparently, more than 90% of webservers use one specific, NIST specified, elliptic curve: P-256:

$$y^2 = x^3 - 3x + 41058363725152142129326129780047268409114441015993725554835256314039467401291,$$

taken modulo $p = 2^{256} - 2^{224} + 2^{192} + 2^{96} - 1$ (the fact that $p \approx 2^{256}$ makes the computations on the elliptic curve much faster in practice). The initial point $P = (x, y)$ on the curve has huge coordinates as well.

Using this single curve is sometimes considered to be problematic, especially following the concerns that the NSA may have implemented a backdoor into Dual_EC_DRBG, which was a NIST standard 2006–2014.

https://en.wikipedia.org/wiki/Dual_EC_DRBG

A popular alternative is the curve Curve25519. In addition to some desirable theoretical advantages, its parameters are small and therefore not of similarly mysterious origin as the ones for P-256:

$$y^2 = x^3 + 486662x^2 + x, \quad p = 2^{255} - 19, \quad x = 9.$$

[Instead of points with (x, y) coordinates, one can actually work with just the x -coordinates for an additional speed-up.]

<https://en.wikipedia.org/wiki/Curve25519>

```
>>> E = EllipticCurve(GF(2^255-19), [0,486662,0,1,0])
>>> E
      y2 = x3 + 486662x2 + x
>>> E.order()
57896044618658097711785492504343953926856930875039260848015607506283634007912
>>> log(E.order(),2).n()
255.0000000000000
>>> P = E.lift_x(9)
>>> P
(9: 43114425171068552920764898935933967039370386198203806730763910166200978582548: 1)
>>> 100*P
(44032819295671302737126221960004779200206561247519912509082330344845040669336: 8626006\
392447572371634278060016659575750781271666323173891504901961672743344: 1)
>>> P.order()
7237005577332262213973186563042994240857116359379907606001950938285454250989
>>> log(P.order(),2).n()
252.0000000000000
>>> E.order() / P.order()
8
>>> 5*(20*P) == 20*(5*P)
1
```

On the other hand, it should be pointed out that it is not an easy task to “randomly generate” cryptographically secure elliptic curves plus suitable base point. That’s the reason pre-selected elliptic curves are of importance.

<http://blog.bjrn.se/2015/07/lets-construct-elliptic-curve.html>

P versus NP: A Millennium Prize Problem

The Clay Mathematics Institute has offered 10^6 dollars each for the first correct solution to seven **Millennium Prize Problems**. Six of the seven problems remain open.

https://en.wikipedia.org/wiki/Millennium_Prize_Problems

Comment. Grigori Perelman solved the Poincaré conjecture in 2003 (but refused the prize money in 2010).

https://en.wikipedia.org/wiki/Poincaré_conjecture

Example 219. (P vs NP) P versus NP is one of the Millennium Prize Problems that is of particular importance to cryptography.

“If the solution to a problem is easy to check for correctness, is the problem easy to solve?”

https://en.wikipedia.org/wiki/P_versus_NP_problem

Roughly speaking, consider decision problems which have an answer of yes or no. **P** is the class of such problems, which can be solved efficiently. **NP** are those problems, for which we can quickly verify that the answer is yes if presented with suitable evidence.

For instance.

- It is unknown whether factoring (in the sense of: does N have a factor $\leq M$?) belongs to **P** or not. The problem is definitely in **NP** because, if presented with a factor $\leq M$, we can easily check that.
- Deciding primality is in **P** (maybe not so shocking since there are very efficient nondeterministic algorithms for checking primality; not so for factoring).
- In the (decisional) travelling salesman problem, given a list of cities, their distances and d , the task is to decide whether a route of length at most d exists, which visits each city exactly once. The decisional TSP is clearly in **NP** (take as evidence the route of length $\leq d$). In fact, the problem is known to be NP-complete, meaning that it is in **NP** and as “hard” as possible (in the sense that if it actually is in **P**, then $P=NP$; that is, we can solve any other problem in **NP** efficiently).
- Other NP-complete problems include:
 - Sudoku: Does a partially filled grid have a legal solution?
 - Subset sum problem: Given a finite set of integers, is there a non-empty subset that sums to 0?

Comment. “Efficiently” means that the problem can be solved in time polynomial in the input size.

Take for instance computing $2^n \pmod n$, where n is the input (it has size $\log_2(n)$). This can be done in polynomial time if we use binary exponentiation (whereas the naive approach takes time exponential in $\log_2(n)$).

Comment. This is one of the few prominent mathematical problems which doesn’t have a definite consensus. For instance, in a 2012 poll of 151 researchers, about 85% believed $P \neq NP$ while about 10% believed $P=NP$.

Comment. **NP** are problems that can be verified efficiently if the answer is “yes”. Similarly, **co-NP** are problems that can be verified efficiently if the answer is “no”. It is an open problem whether $NP \neq co-NP$.

- Factoring is in both **NP** and **co-NP** (it is in **co-NP** because primality testing is in **P**).
- For all NP-complete problems it is unknown whether they are in **co-NP**. (If one of them is, then we would, unexpectedly, have $NP=co-NP$.)

The Riemann hypothesis: Another Millennium Prize Problem

The Riemann hypothesis is another one of the seven Millennium Prize Problems that is of importance to the underpinnings of cryptography. It is concerned with the distribution of primes.

Recall that we discussed the prime number theorem, which states that, up to x , there are about $x/\ln(x)$ many primes. The Riemann hypothesis gives very precise error estimates for an improved prime number theorem (using a function more complicated than the logarithm).

Example 220. (Riemann hypothesis) Consider the **Riemann zeta function** $\zeta(s) = \sum_{n \geq 1} \frac{1}{n^s}$. This series converges (for real s) if and only if $s > 1$.

The divergent series $\zeta(1)$ is the harmonic series, and $\zeta(p)$ is often called a p -series in Calculus II.

Comment. Euler achieved worldwide fame in 1734 by discovering and proving that $\zeta(2) = \frac{\pi^2}{6}$ (and similar formulas for $\zeta(4), \zeta(6), \dots$).

For complex values of $s \neq 1$, there is a unique way to “analytically continue” this function. It is then “easy” to see that $\zeta(-2) = 0, \zeta(-4) = 0, \dots$. The **Riemann hypothesis** claims that all other zeroes of $\zeta(s)$ lie on the line $s = \frac{1}{2} + a\sqrt{-1}$ ($a \in \mathbb{R}$). A proof of this conjecture (checked for the first 10,000,000,000 zeroes) is worth \$1,000,000.

<http://www.claymath.org/millennium-problems/riemann-hypothesis>

The connection to primes. Here’s a vague indication that $\zeta(s)$ is intimately connected to prime numbers:

$$\begin{aligned} \zeta(s) &= \left(1 + \frac{1}{2^s} + \frac{1}{2^{2s}} + \dots\right) \left(1 + \frac{1}{3^s} + \frac{1}{3^{2s}} + \dots\right) \left(1 + \frac{1}{5^s} + \frac{1}{5^{2s}} + \dots\right) \dots \\ &= \frac{1}{1 - 2^{-s}} \frac{1}{1 - 3^{-s}} \frac{1}{1 - 5^{-s}} \dots \\ &= \prod_{p \text{ prime}} \frac{1}{1 - p^{-s}} \end{aligned}$$

This infinite product is called the Euler product for the zeta function. If the Riemann hypothesis was true, then we would be better able to estimate the number $\pi(x)$ of primes $p \leq x$.

More generally, certain statements about the zeta function can be translated to statements about primes. For instance, the (non-obvious!) fact that $\zeta(s)$ has no zeros for $\text{Re } s = 1$ implies the prime number theorem.

<http://www-users.math.umn.edu/~garrett/m/v/pnt.pdf>