

CMSC/Math 456: Cryptography (Fall 2022)

Lecture 18

Daniel Gottesman

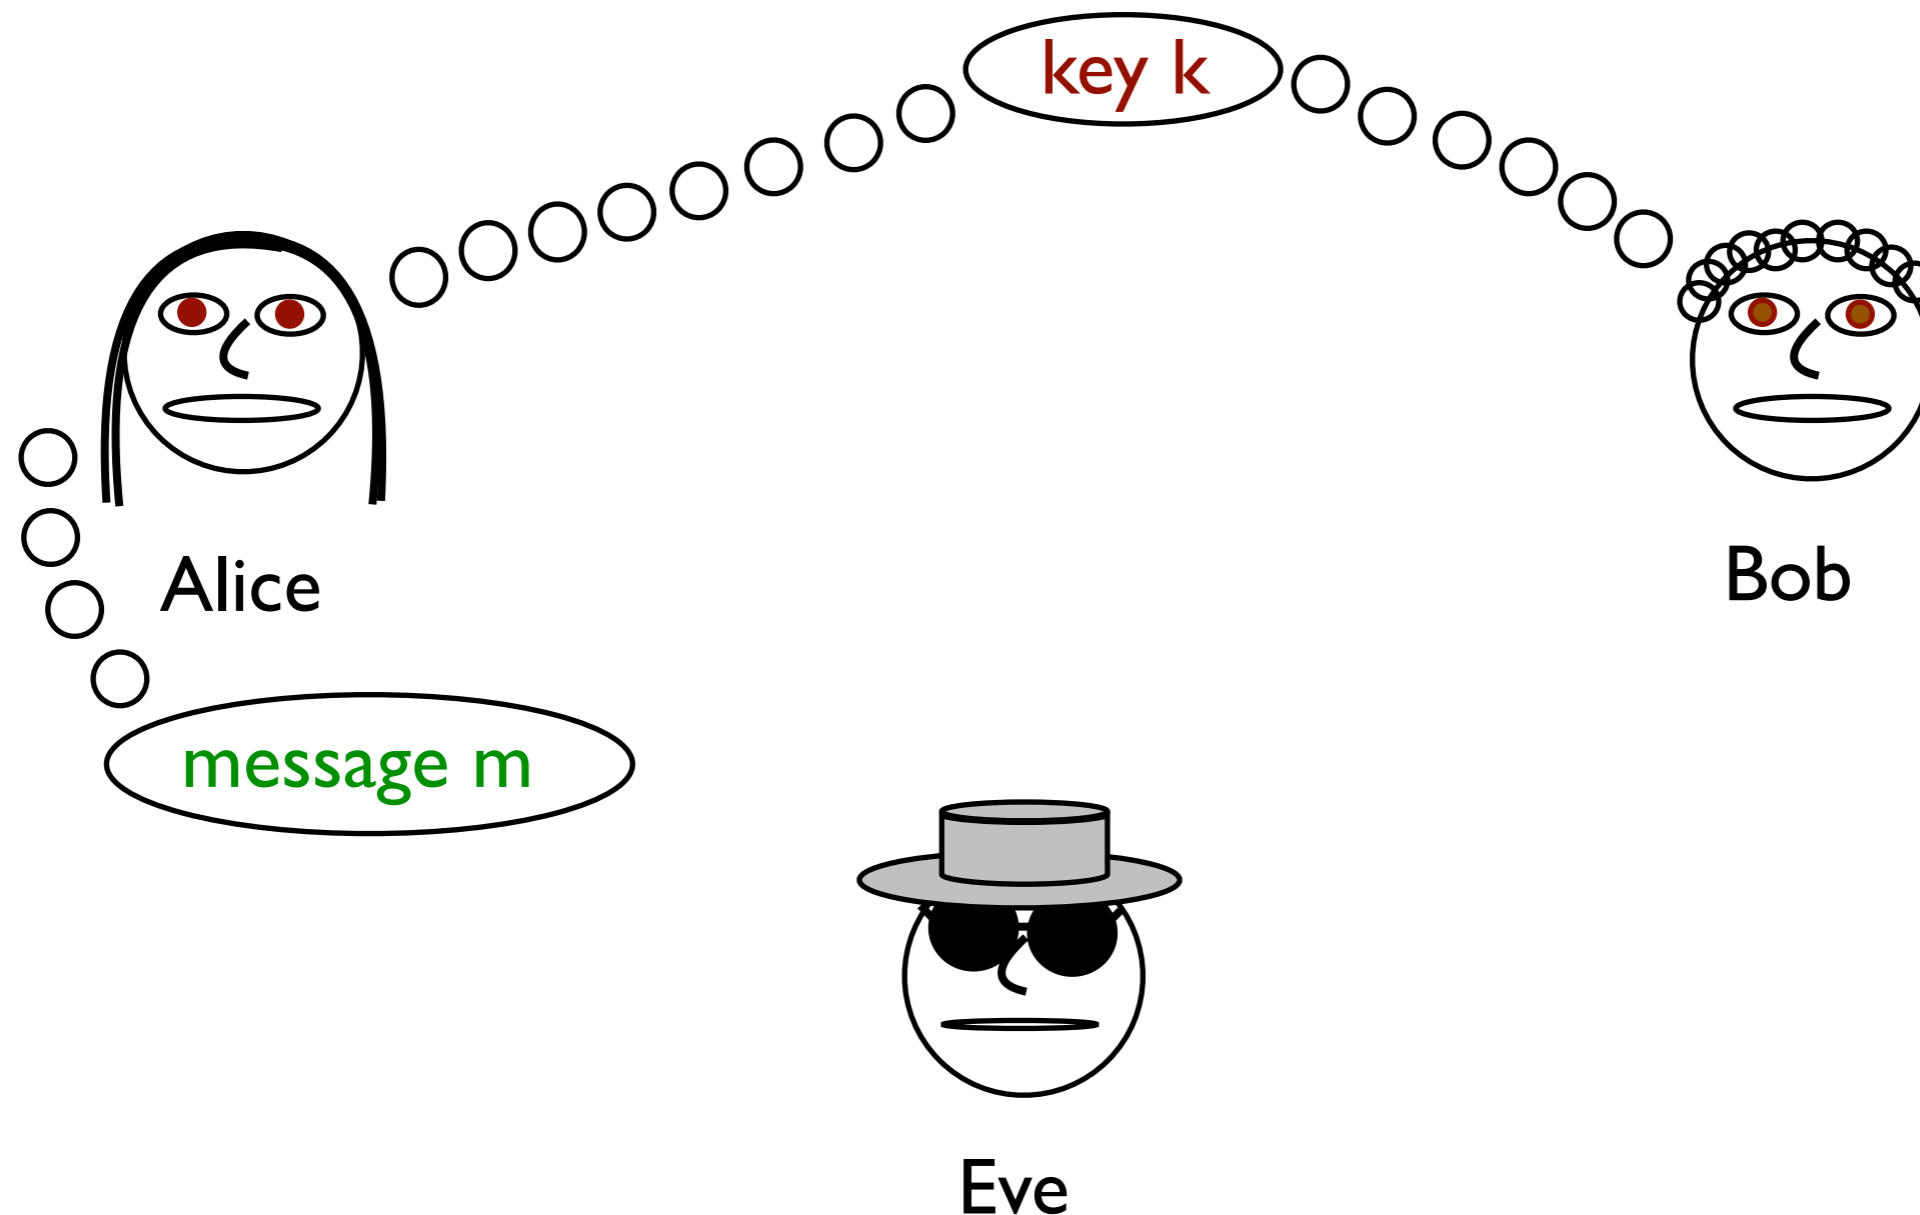
Administrative

Grades for Problem Set 5 are out.

Mid-term grading is not quite done, but is close.

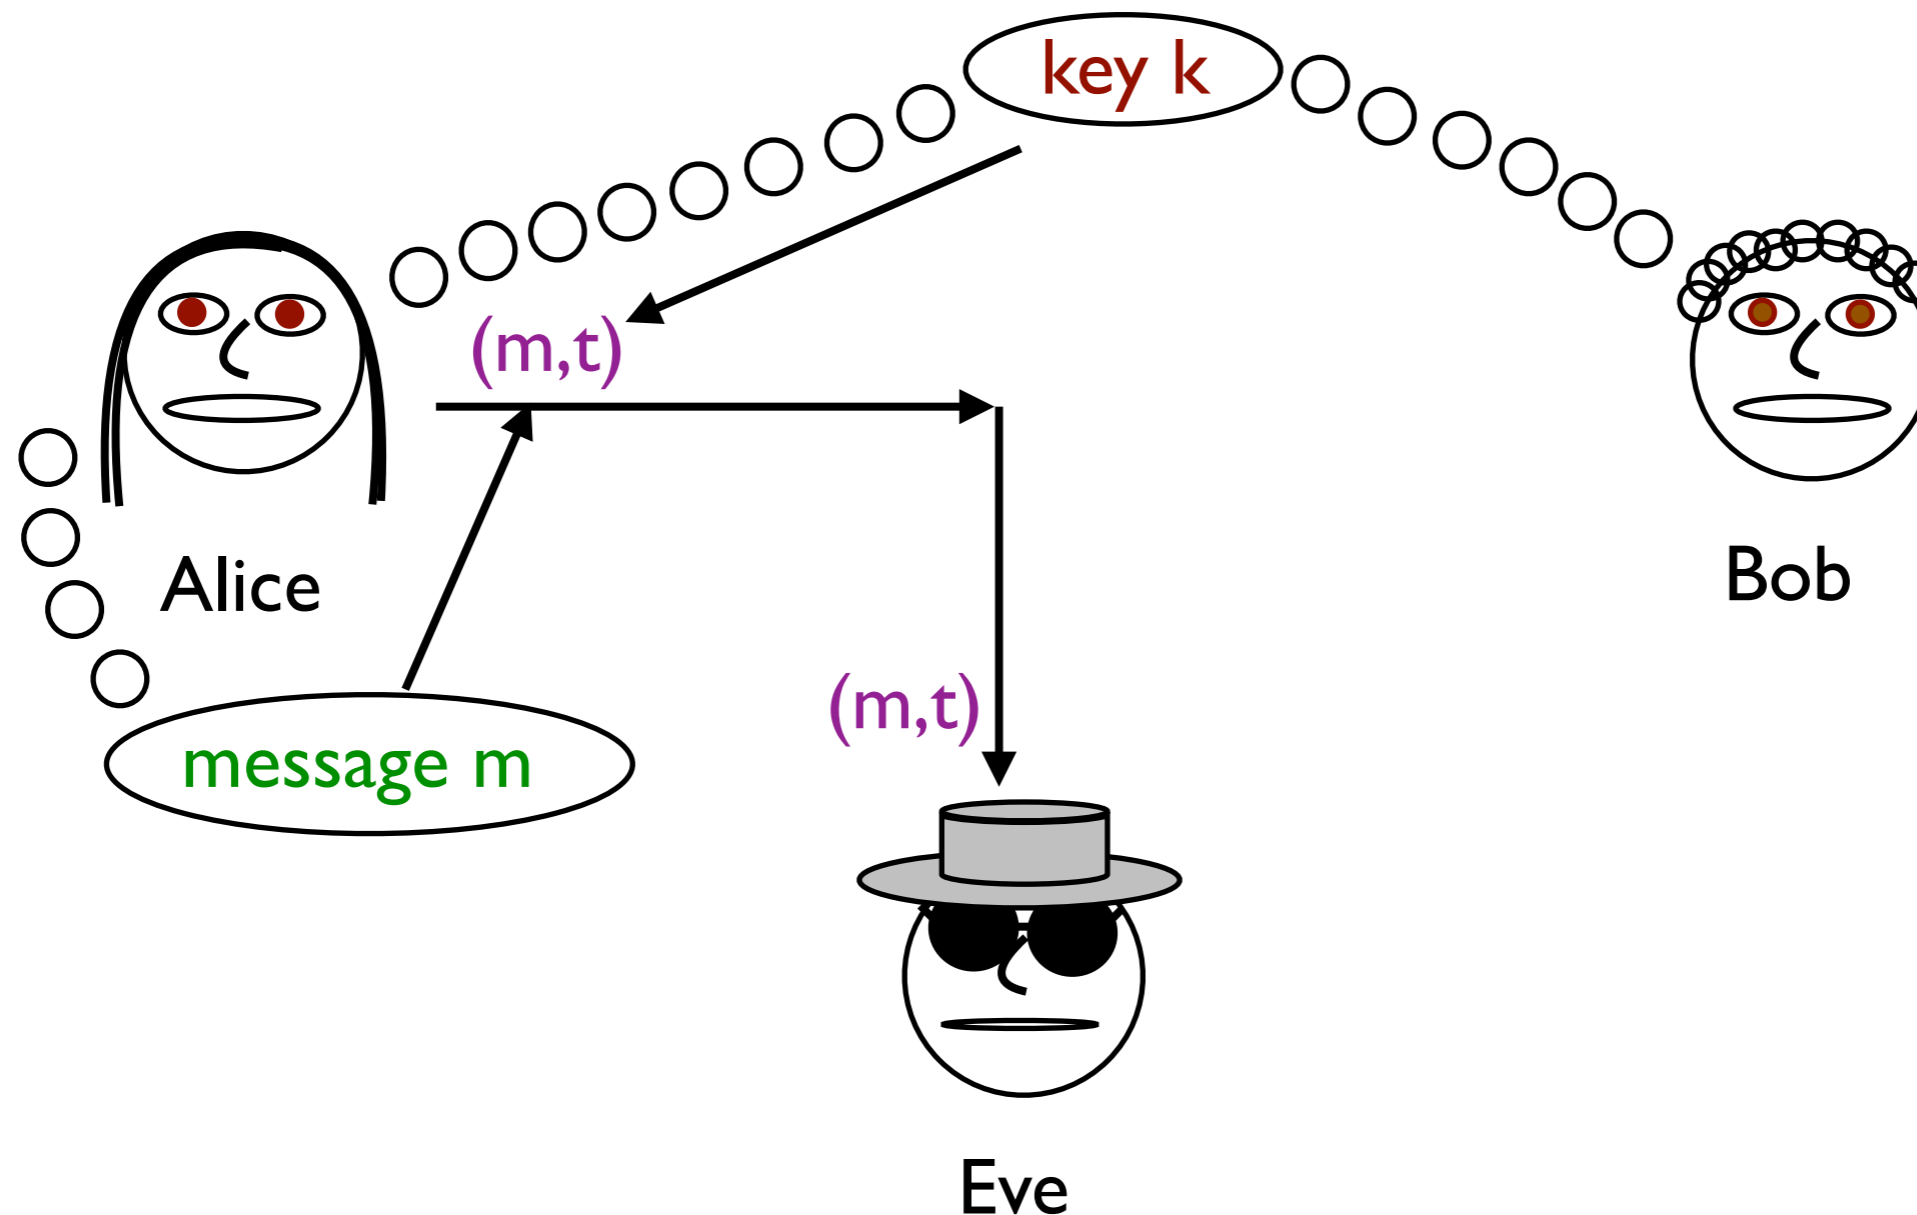
Message Authentication

We discussed message authentication, whereby Alice can send a message to Bob without encryption and Bob can be sure it came from Alice unchanged.



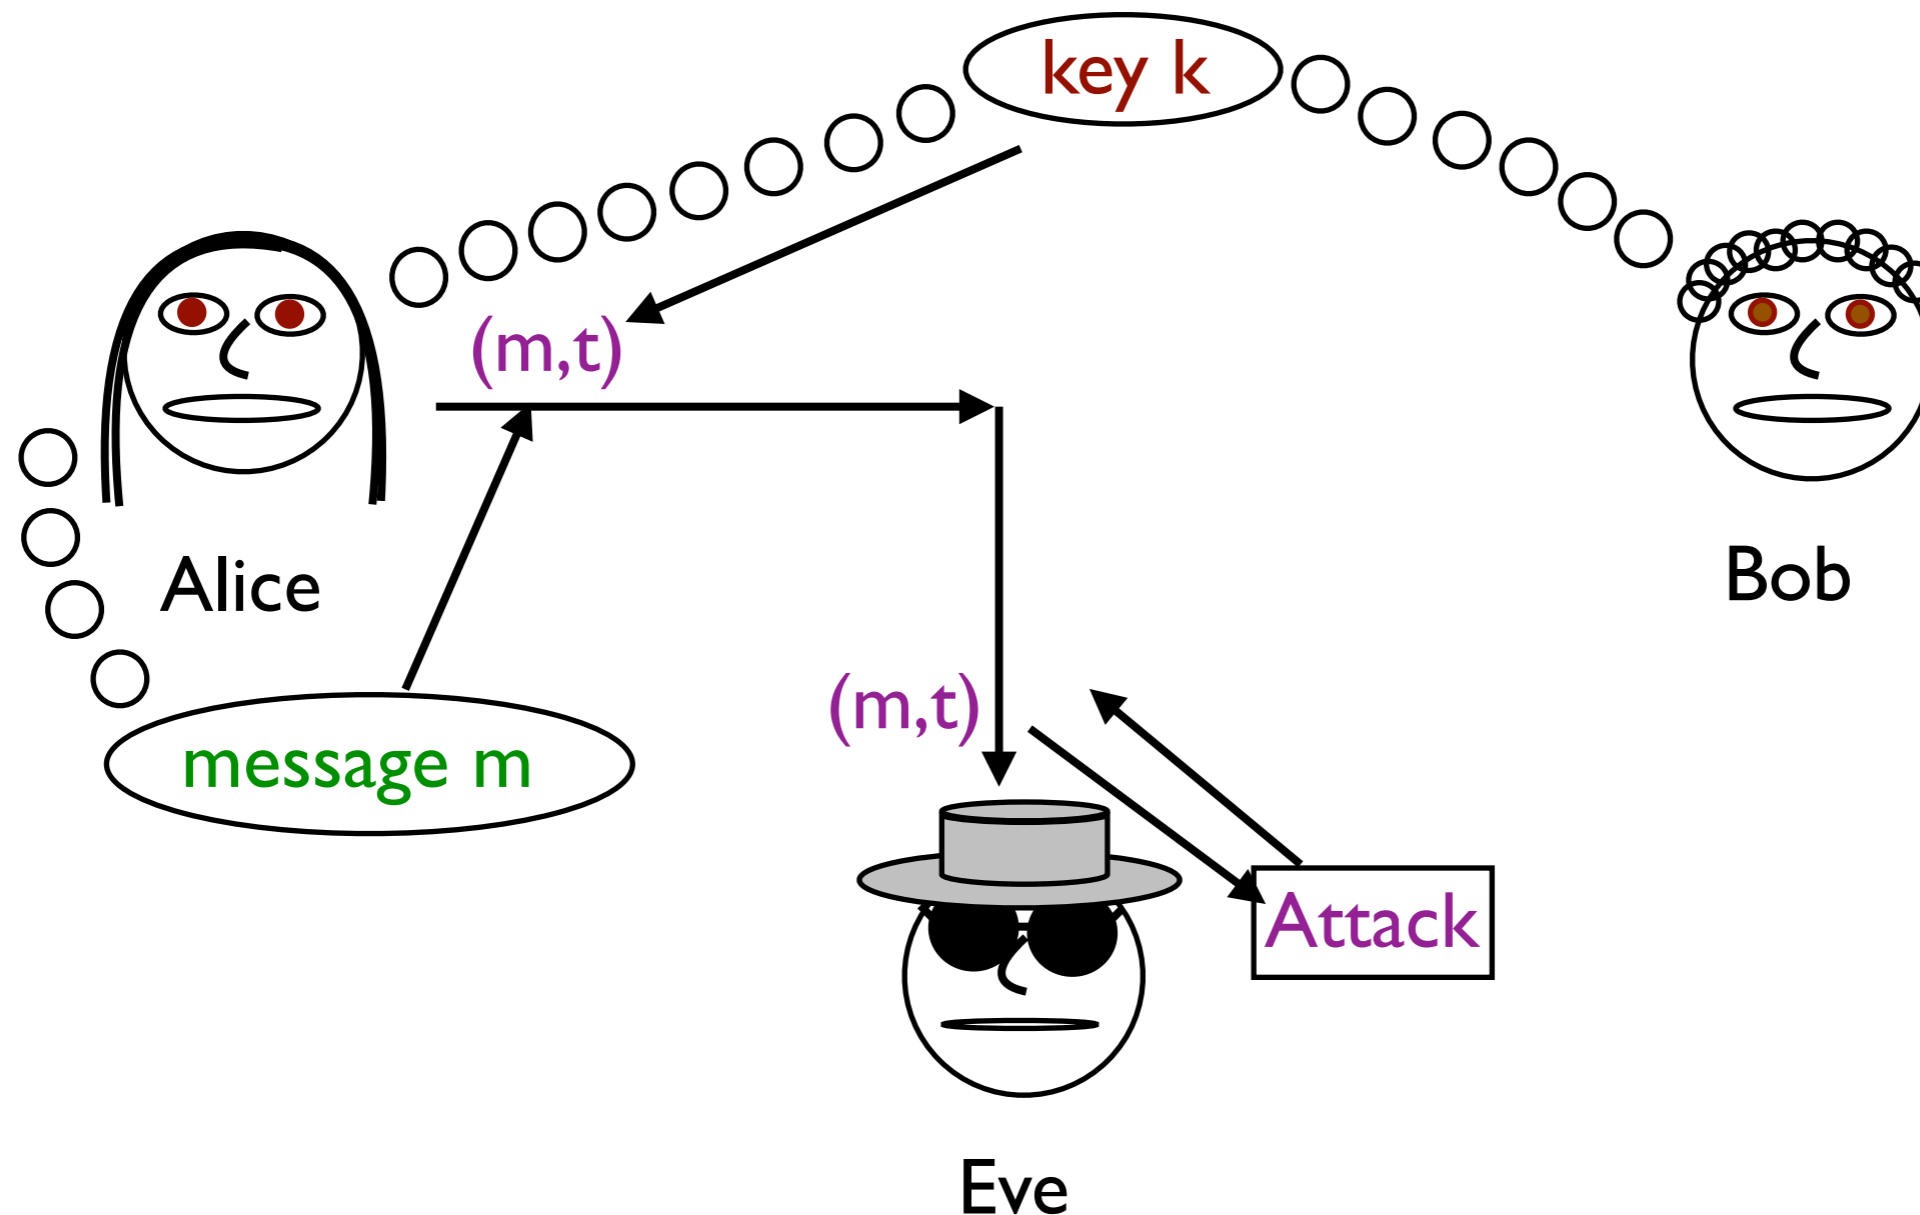
Message Authentication

We discussed message authentication, whereby Alice can send a message to Bob without encryption and Bob can be sure it came from Alice unchanged.



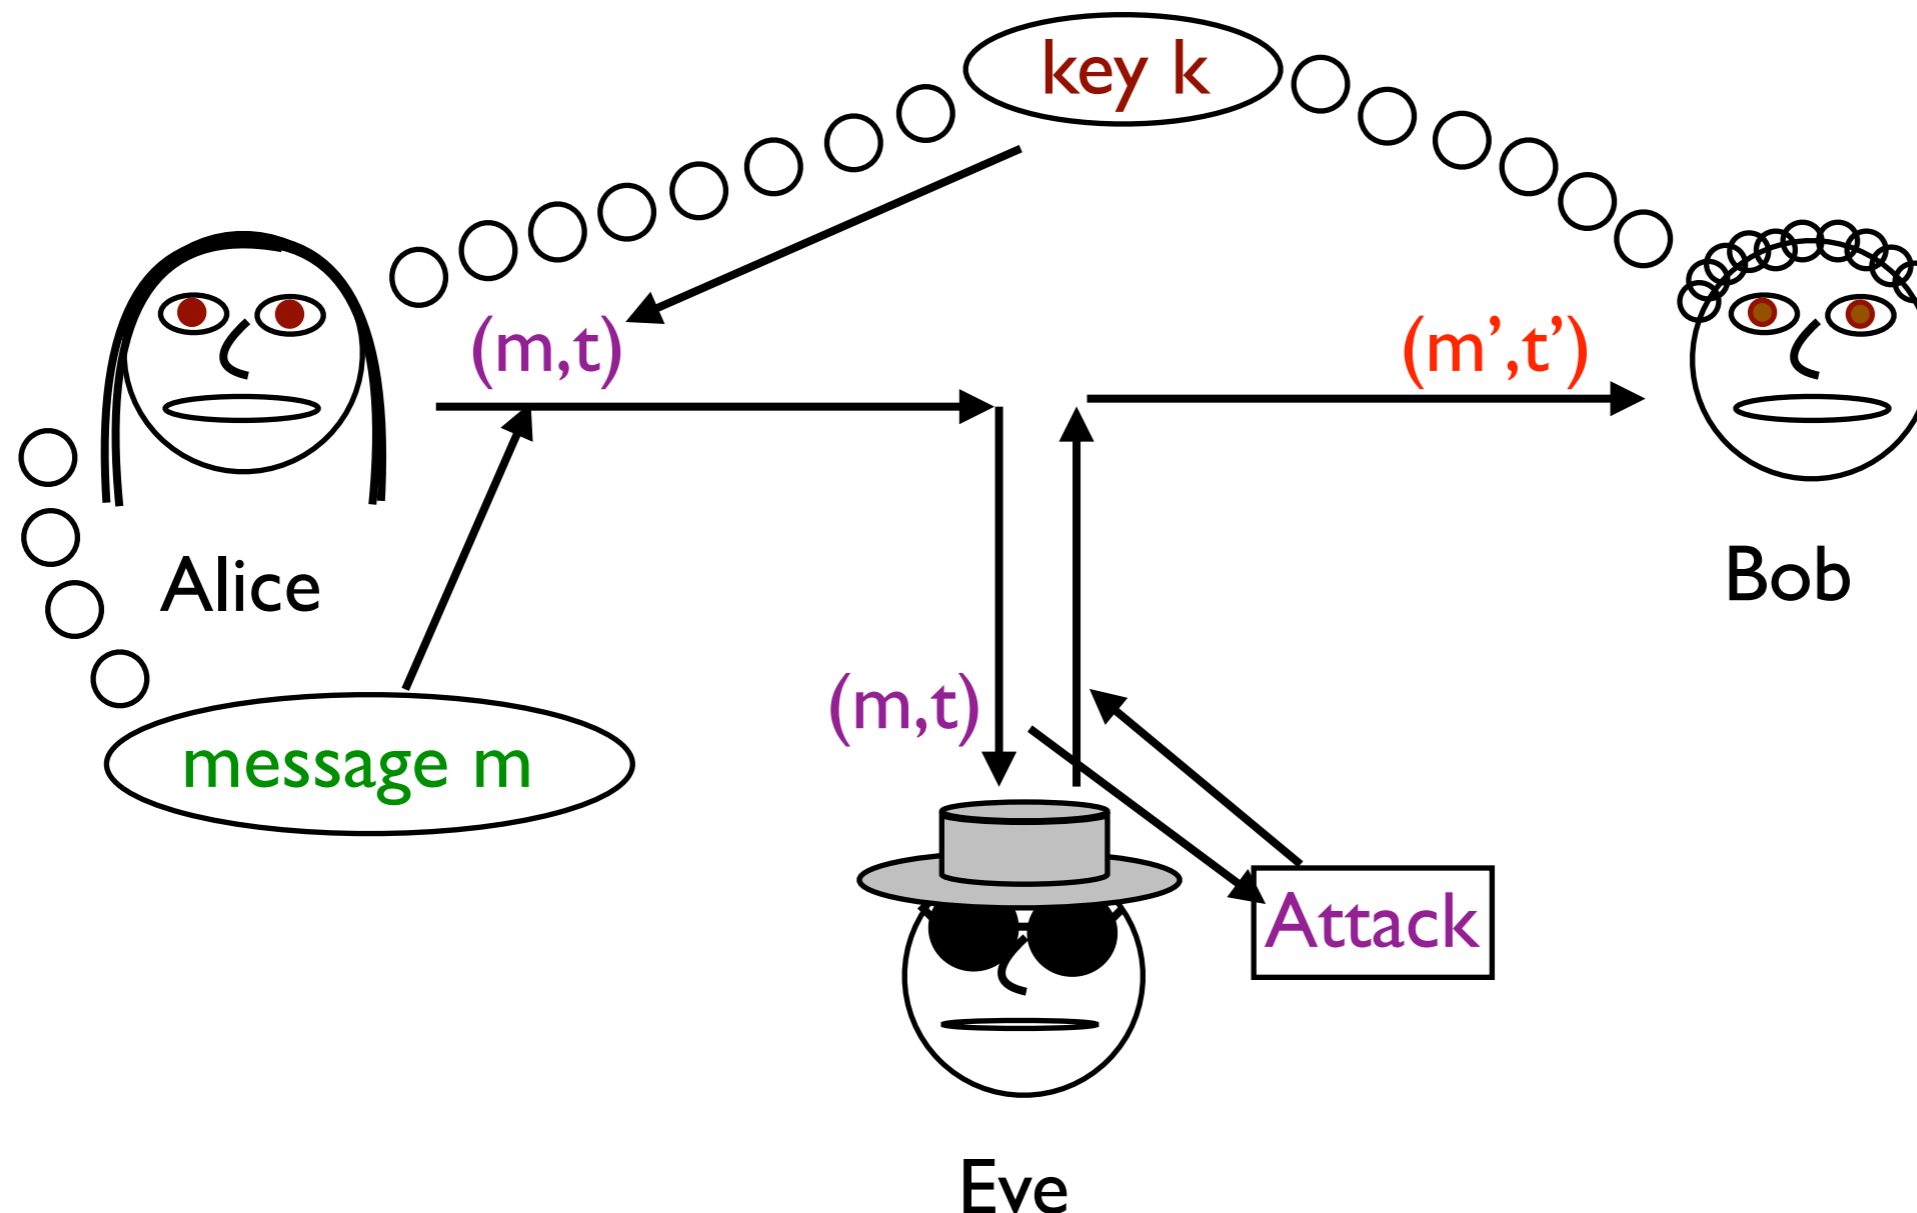
Message Authentication

We discussed message authentication, whereby Alice can send a message to Bob without encryption and Bob can be sure it came from Alice unchanged.



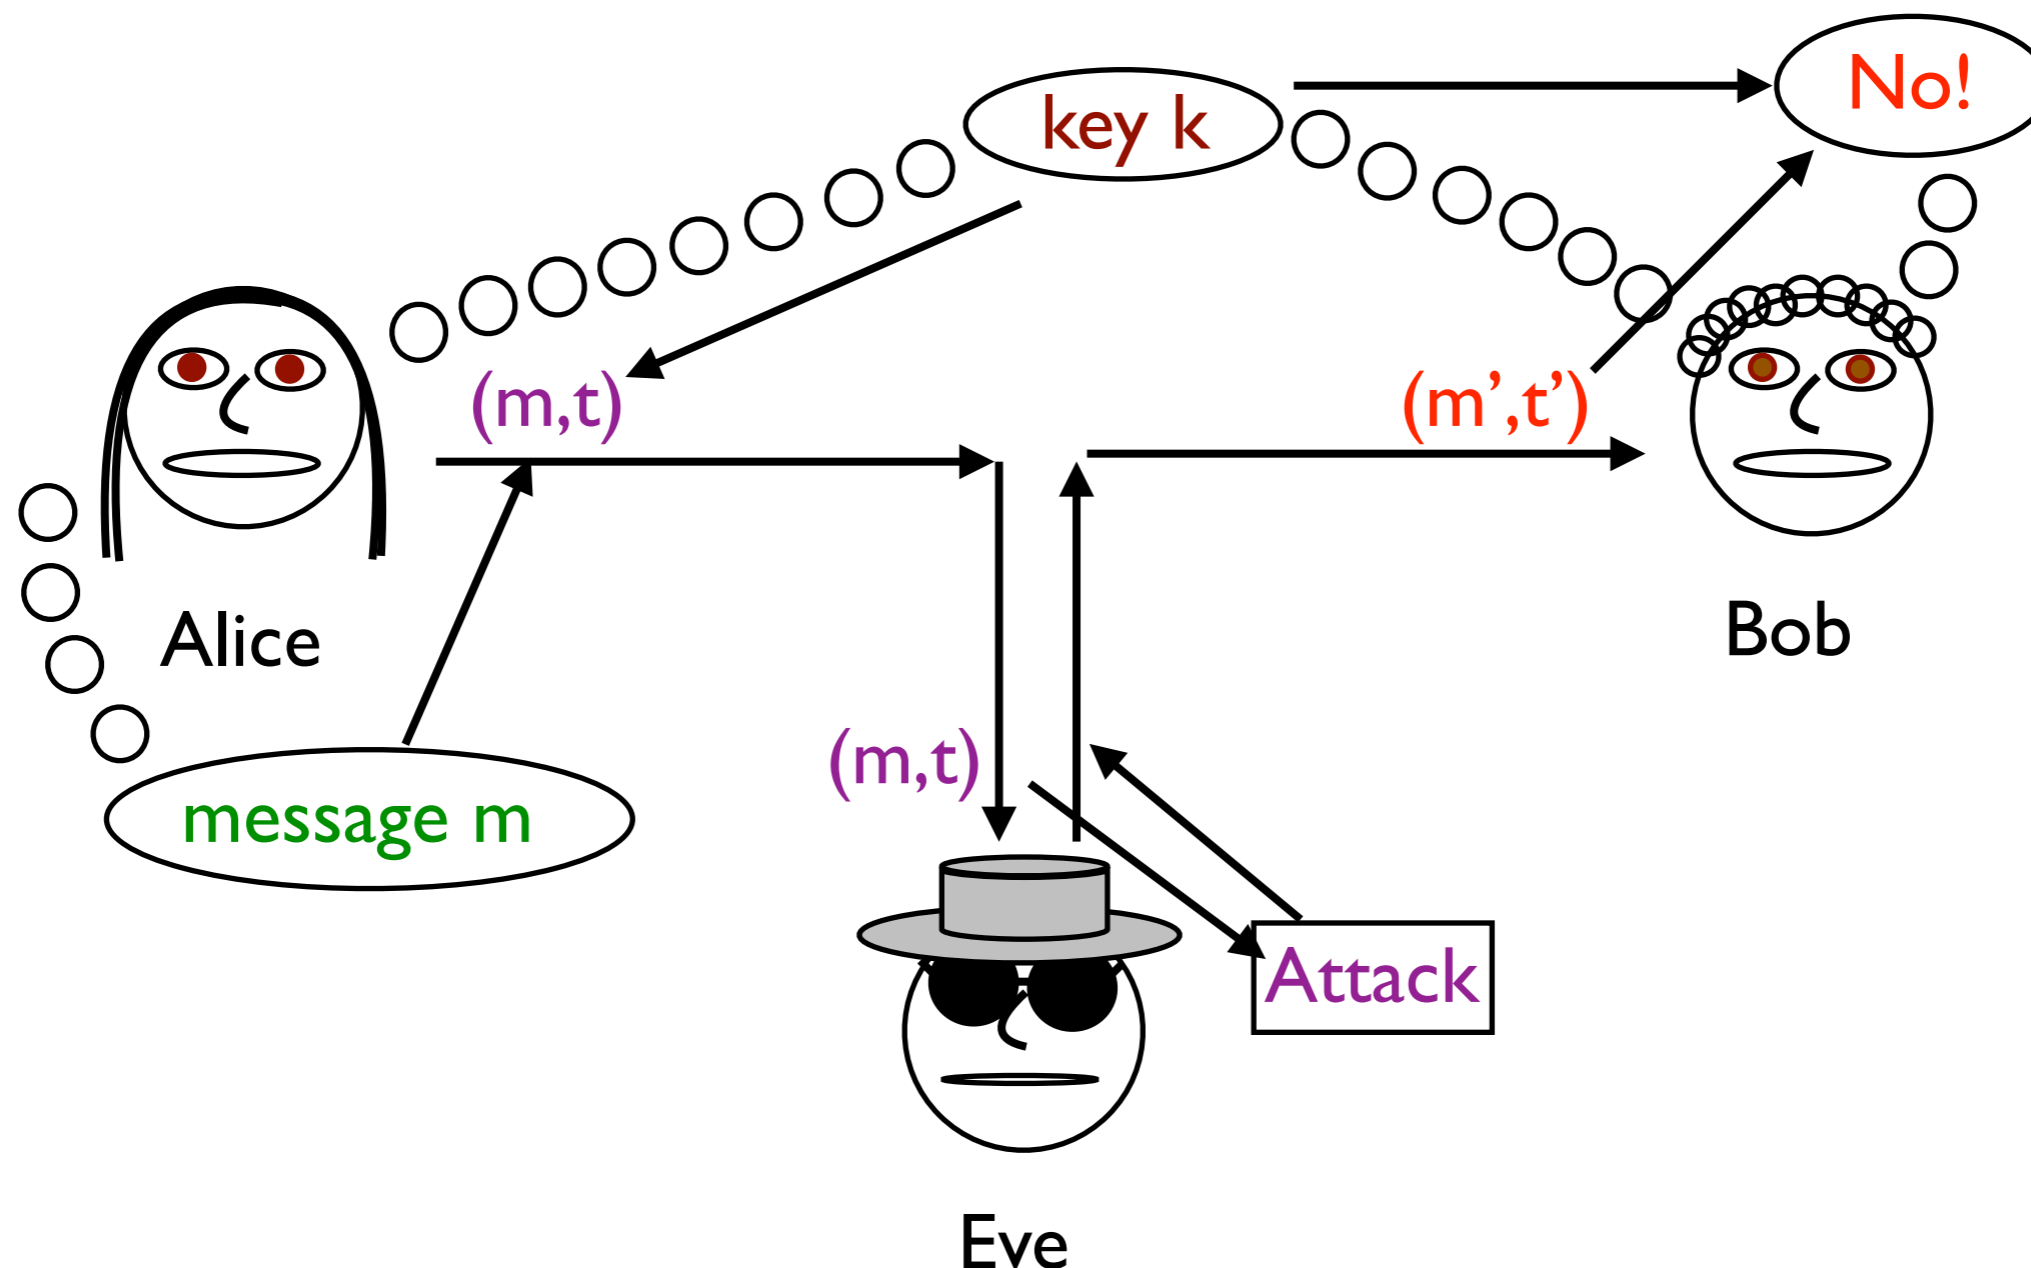
Message Authentication

We discussed message authentication, whereby Alice can send a message to Bob without encryption and Bob can be sure it came from Alice unchanged.



Message Authentication

We discussed message authentication, whereby Alice can send a message to Bob without encryption and Bob can be sure it came from Alice unchanged.

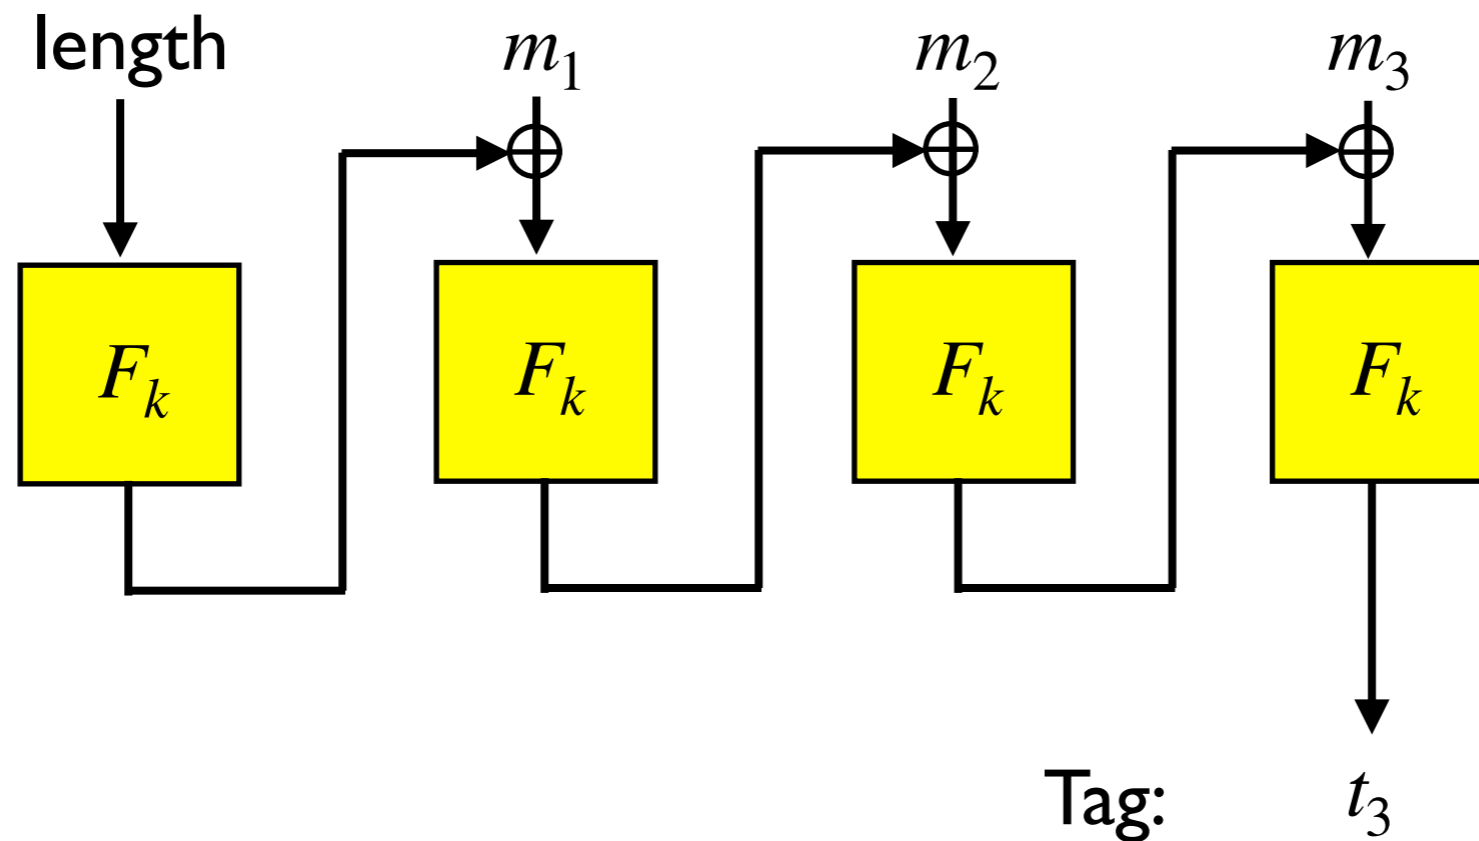


MACs from Pseudorandom Functions

We can make a secure MAC from a pseudorandom function $F_k(m)$. I.e.,


$$\text{Mac}(k, m) = F_k(m).$$

We can also make MACs with a tag much shorter than the message using a CBC-MAC structure:



Pigeonhole Principle


When the tag is shorter than the message, there are more possible messages than tags, which means the **pigeonhole principle** applies:



Suppose we have n pigeons and m holes.

Pigeonhole Principle

When the tag is shorter than the message, there are more possible messages than tags, which means the **pigeonhole principle** applies:




Suppose we have n pigeons and m holes.

If $n \leq m$, then each pigeon can be in its own hole.

Pigeonhole Principle

When the tag is shorter than the message, there are more possible messages than tags, which means the **pigeonhole principle** applies:

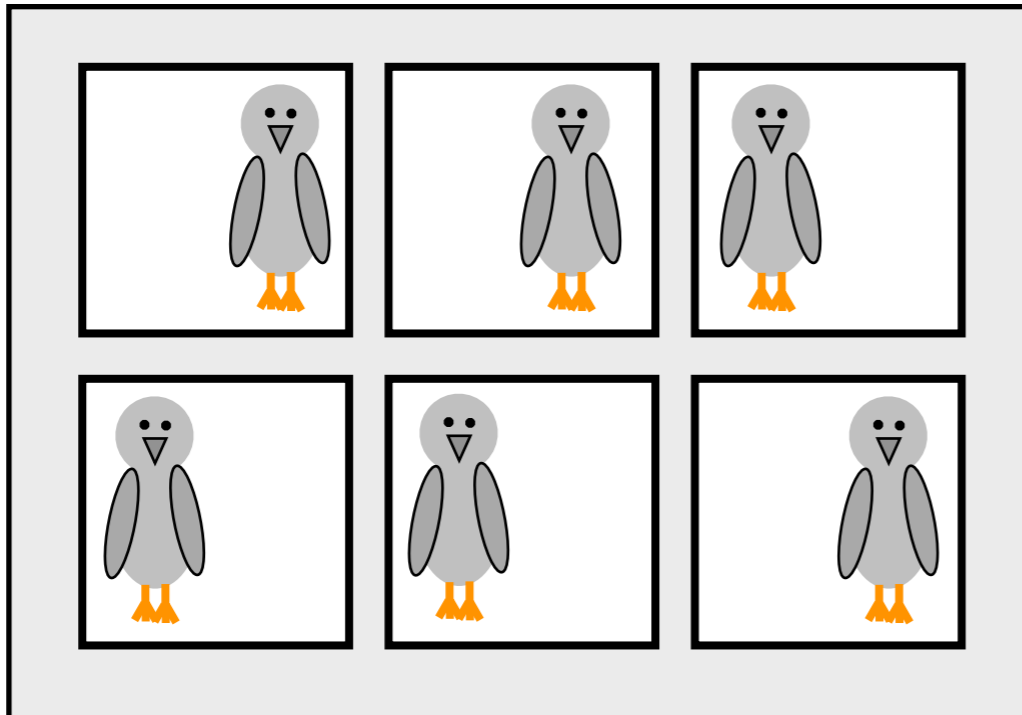


Suppose we have n pigeons and m holes.

If $n \leq m$, then each pigeon can be in its own hole.

Pigeonhole Principle

When the tag is shorter than the message, there are more possible messages than tags, which means the **pigeonhole principle** applies:




Suppose we have n pigeons and m holes.

If $n \leq m$, then each pigeon can be in its own hole.

Pigeonhole Principle

When the tag is shorter than the message, there are more possible messages than tags, which means the **pigeonhole principle** applies:



Suppose we have n pigeons and m holes.

If $n \leq m$, then each pigeon can be in its own hole.

But if $n > m$, then more than one pigeon has to be in the same hole. This is the **pigeonhole principle**.

Collisions and MACs

If we want to have long messages and short tags, we are necessarily going to have more than one message that has the same tag. For a MAC to be secure, it should be hard for Eve to find two such messages:

E.g.: Suppose Eve knows that messages with all bits flipped have the same tag.

Then she sees message **00110011** with tag **1001**.

Now she can forge **11001100** also with tag **1001**.

Thus, we need such a MAC to be **collision-resistant**: It is hard for Eve to find two messages **m, m'** such that **Mac(k,m) = Mac(k,m')**.

Hash Functions

Perhaps surprisingly, there are functions that have collision resistance (as far as we can tell) even *without a key* or if the key is known.

A **collision-resistant hash function** is a function $H(x)$ such that it is **hard to find** two values x, x' such that $H(x) = H(x')$.

“Hard” means polynomial-time as usual for us, so to make this a rigorous definition, you actually need to look at a family of hash functions with larger output lengths.

Informally, we can still consider a hash function of fixed size to be collision-resistant if it is hard in practice to find a collision.

This notion only really makes sense if $|H(x)| < |x|$, so the output is shorter than the input.

Cryptographic Hash Functions

Hash functions are also used in non-cryptographic settings. (E.g. as hash tables.)

In a **non-cryptographic hash function**, collisions between typical elements from the domain should be rare.

In a **cryptographic hash function**, **all** collisions should be hard to find.

E.g.: A simple non-cryptographic hash function is mod:

$$H(x) = x \bmod p$$

But it is very easy to find collisions in this function if you are actively trying to.

Try it: $p = 563$, $x = 822$, $H(x) = 259$

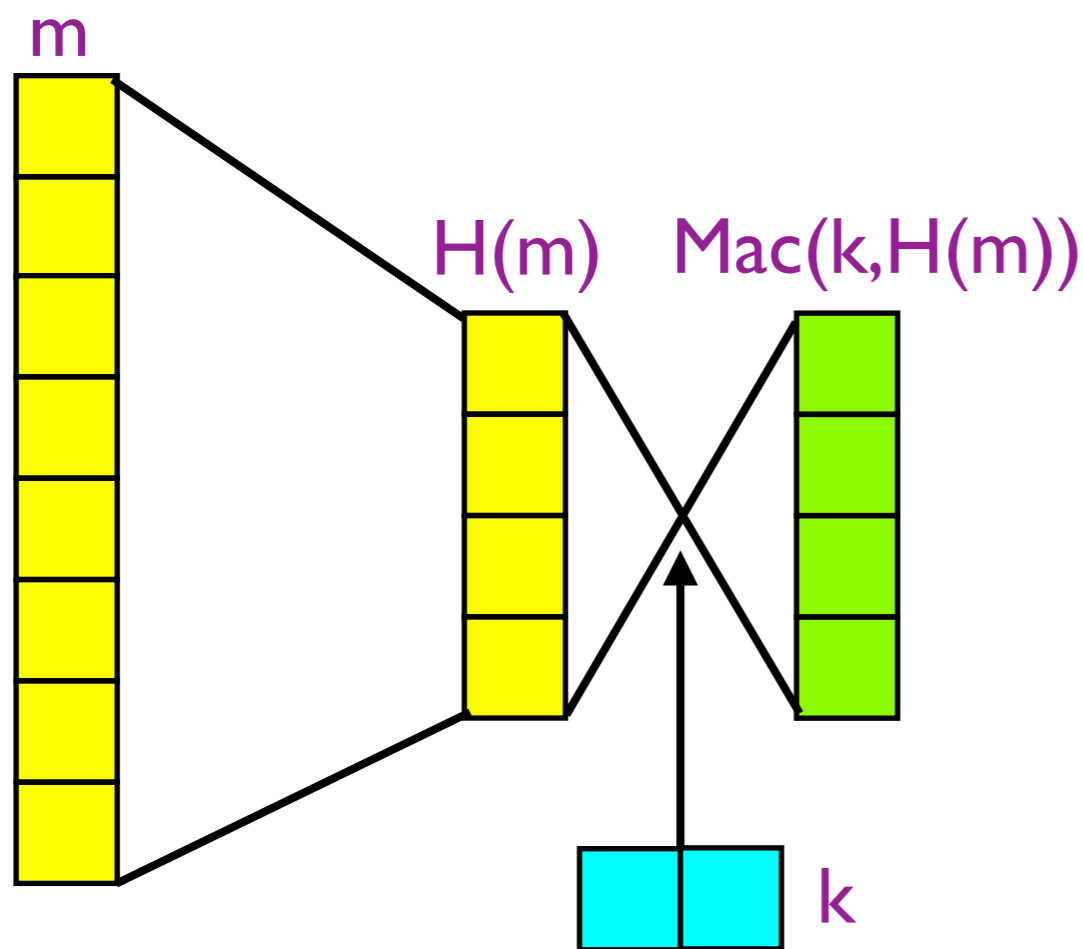
Can you find x' (not 822) with $H(x') = 259$?

Hash Functions for MACs

Given a hash function $H(x)$ and a MAC $\text{Mac}(k,m)$, we can make a new MAC:

$$\text{Mac}'(k, m) = (H(m), \text{Mac}(k, H(m)))$$

If $H(x)$ and $\text{Mac}(k,m)$ are both secure, then Mac' is also secure.



This allows us to authenticate even long messages using short tags and a small key.

Generally this “**hash-and-MAC**” protocol is faster than a CBC-MAC because one hash function is easier than many block ciphers.

HMAC is an even better variant of this idea.

Security of Hash-and-MAC

Suppose Eve is trying to forge a message authenticated with hash-and-MAC.

She has two options to forge $(m, H(m), \text{MAC}(k, H(m)))$:

- Forge a message with $H(m) \neq H(m')$ for all m' for which she has seen a tag. But that would require breaking the original **MAC** protocol.
- Forge a message with $H(m) = H(m')$ for some $m' \neq m$. But that would require breaking the collision-resistance of the hash function.

Note that **both** elements are needed.

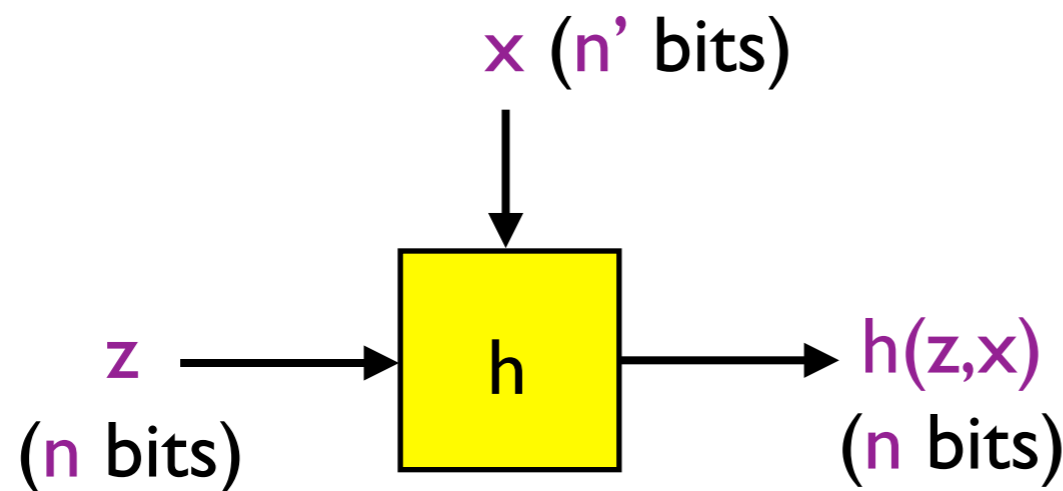
Try it: Forge a message if the tag doesn't include $\text{MAC}(k, H(m))$ but only $H(m)$.

Also: Forge a message if $H(m)$ is not collision-resistant.

Constructing Hash Functions


Strategy: Construct a hash function for fixed-size input blocks. This is called a **compression function**. Then for larger inputs, break them up into chunks of an appropriate size and apply the compression function repeatedly, adding one chunk each time.

Specifically, we want a compression function $h(z,x)$ that takes **two** inputs of size n and n' , and has **one** output of size n .



$h(z,x)$ should be collision-resistant.

Merkle-Damgard Construction



The input x is broken up into (x_1, x_2, x_3, \dots) ; each x_i is n' bits long.

z_0 is set to some IV. (It is fixed based on the specific construction, not chosen randomly.)

But ... since x might not be a multiple of n' , we need to pad it to fill out the blocks. A bad choice of padding makes this construction insecure.

Padding for Merkle-Damgard


For instance, suppose we pad with 0s. Then an input x that is not a multiple of n' bits long will hash to the same value as $(x||0)$.

Padding Construction:


Construct \tilde{x} by following x by $|$ and then a string of 0s, followed by the length of x . Choose a number of 0s so that the length of \tilde{x} is a multiple of n' .

Then break \tilde{x} up into t pieces \tilde{x}_i of n' bits each.

Then $H(x) = z_t$, where $z_i = h(z_{i-1}, \tilde{x}_i)$.



Merkle-Damgard security



Theorem: If $h(z,x)$ is collision-resistant, then $H(x)$ given by the Merkle-Damgard construction is collision-resistant as well.

Proof Idea:

Suppose we have two inputs x, x' to H such that $z_i = z'_i$.

Recall that $z_i = h(z_{i-1}, \tilde{x}_i) = z'_i = h(z'_{i-1}, \tilde{x}'_i)$.

- If $(z_{i-1}, \tilde{x}_i) \neq (z'_{i-1}, \tilde{x}'_i)$, then we have a collision in h .
- Otherwise $z_{i-1} = z'_{i-1}$ and we can apply induction to either find a collision in h or show that $x = x'$.

Making a Compression Function

A common strategy: the **Davies-Meyer construction**

Given block cipher $F_k(x)$, let

$$h(z, x) = F_z(x) \oplus x$$

(This doesn't always work. The block cipher **F** needs to be sufficiently similar to an ideal cipher, which essentially has no exploitable internal structure.)

Practical Ciphers

Examples of hash functions constructed via Davies-Meyer and Merkle-Damgard:

- MD5 (no longer secure)
- SHA-1 (no longer secure)
- SHA-2 (still OK)

These hash functions use their own block ciphers which involve multiple rounds, in each of which the input is shifted cyclically with some minor changes to most of it and major change to one segment.

- SHA-3 is the newest standard and it is based on different principles.

Birthday Attacks

Recall the **birthday paradox**: With N possible dates for birthdays, with high probability, a group of M people has two people with the same birthday if $M \approx \sqrt{N}$.

We can take advantage of this to look for collisions in a hash function: If the hash function $H(x)$ has s -bit output, we will find a collision after computing H on about $2^{s/2}$ random input values.

Consequence: A hash function needs to have output size twice as long as a block cipher to be secure in practice!

In particular, SHA-2 and SHA-3 have 224-bit and higher output sizes, compared to 128-bit block size and minimum key length for AES.

Key Derivation

Recall in the RSA-based KEM, we needed to pick a key derivation function to convert the key to bits. We also needed it with Diffie-Hellman.

RSA KEM:

Gen: Pick a (public) key derivation function $H(x)$, then as usual for RSA, i.e., generate two random primes p and q which are s bits long. Let $N = pq$. Choose $e, d \in \mathbb{Z}_N^*$ s.t. $ed = 1 \pmod{\varphi(N)}$. The public key is (N, e) and the private key is (N, d) .

Encaps: Choose random x . The ciphertext is $c = x^e \pmod{N}$ and the key is $H(x)$.

Decaps: Given c and d , compute $x' = c^d \pmod{N}$. Then the key is $H(x')$.

Hash Functions for Key Derivation

One solution (although not the only one) is to use a hash function for key derivation.



We don't care so much about collisions here. Instead what we want is that $H(k)$ should be roughly uniformly distributed and that any partial information Eve might have about k should get squeezed out, so she has little information about $H(k)$.

These come from a stronger property of the hash function, namely that it can be well modeled by a **random oracle**.

Hash Functions as Random Oracles

A **random oracle** is a random function which is implemented as a black box. That is, we count complexity by how many times the function is evaluated.

If $H(x)$ is a random oracle, knowing the value of $H(x_i)$ for any set $\{x_i\}$ doesn't help you predict $H(y)$ if $y \notin \{x_i\}$.

A random oracle is **automatically collision resistant**:

For any pair $H(x), H(y)$ ($x \neq y$), the probability that $H(x) = H(y)$ is exactly 2^{-s} , where H has an s -bit output. This is precisely the limit on collisions set by birthday attacks.

But collision resistance doesn't necessarily imply that the hash function is a good random oracle.

Random Oracle vs. Pseudorandom

A random oracle can be used to make a pseudorandom generator or function, but **they are not the same**.

Random Oracle

- No key
- Can be evaluated by everyone
- Outputs always look random

Pseudorandom Function

- Keyed
- Evaluated only using the key
- Outputs look random only if key is unknown.

Both are things that *look like* random functions, but in different contexts.

What Does it Mean to be RO?

When we say a hash function is “well-modeled as a random oracle,” we mean that it has no *useful* structure for an attacker to exploit.

It means

- The output strings have no special properties and no discernible correlations.
- Birthday attacks are the best or only way to find collisions.
- Computing intermediate values (i.e., stopping the evaluation partway through) is not useful for speeding up searches.
- In general, brute force attacks are the best or only attacks that will work.

Any real function must fail to match a random oracle to some degree, but some functions seem to be close enough in practice for security proofs based on random oracles to be useful.

