CMSC/Math 456: Cryptography (Fall 2022)

Lecture 8

Daniel Gottesman

Administrative

Problem Set #3 is out. This is a Python programming assignment, and is due in 1 week. The autograder will give you your score within a few minutes, but in case we discover a bug in the autograder, it is possible we will need to rerun it later, in which case your score could change.

Extensions: Recall that the course policy is that extensions must be approved in advance. This includes if you have difficulty uploading your assignment, so make sure you are able to upload and don't wait until the last minute.

This class is being recorded

Goals of Block Cipher Design

- Must be invertible to use with CBC mode (i.e., pseudorandom permutation rather than pseudorandom function).
- Even when the inputs are related, the outputs should be very different.

In particular, the change of even a single bit of the input should result in a totally different output. This is known as the "avalanche effect." It is often achieved by having multiple rounds, each of which magnifies small changes.

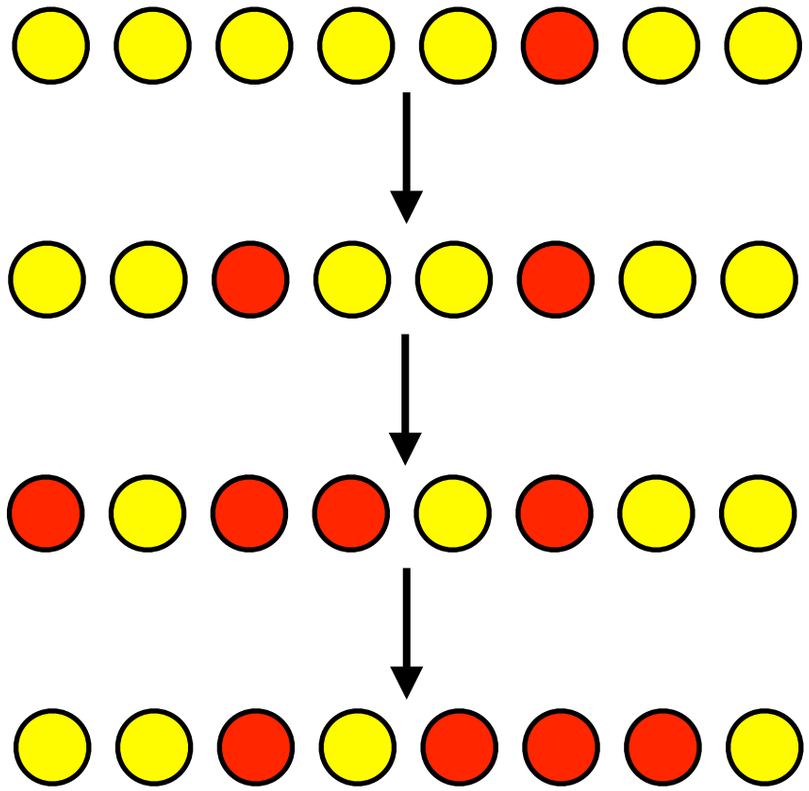

This class is being recorded

Avalanche Effect vs. Breaking Cipher

Why does the avalanche effect contribute to making a cipher secure against chosen plaintext attacks?

Because we are focusing on permutations, Eve can try to trace a ciphertext backwards to the input. Because she doesn't know the key, she has limited ability to do this.

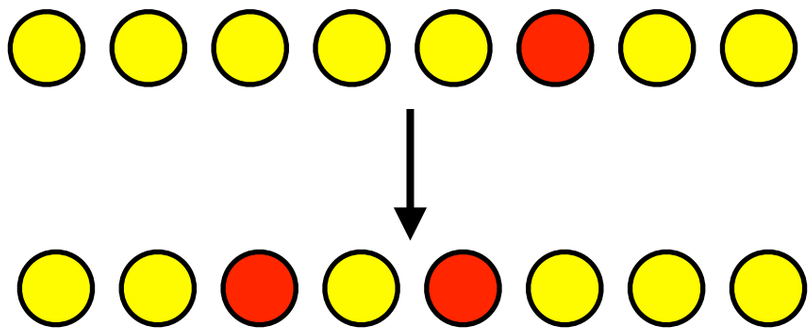

However, if there is no avalanche effect and Eve gets two (plaintext, ciphertext) pairs that differ in a small fraction of bits of the plaintext, the ciphertexts will also only differ in a limited set of places.

She can then focus on deducing the segment of the key that is relevant to those locations. This is a much smaller search than searching over all possible key values.

This class is being recorded

Substitution-Permutation Networks

The DES mangler function is a variant of a substitution-permutation network, a design paradigm for pseudorandom permutations.

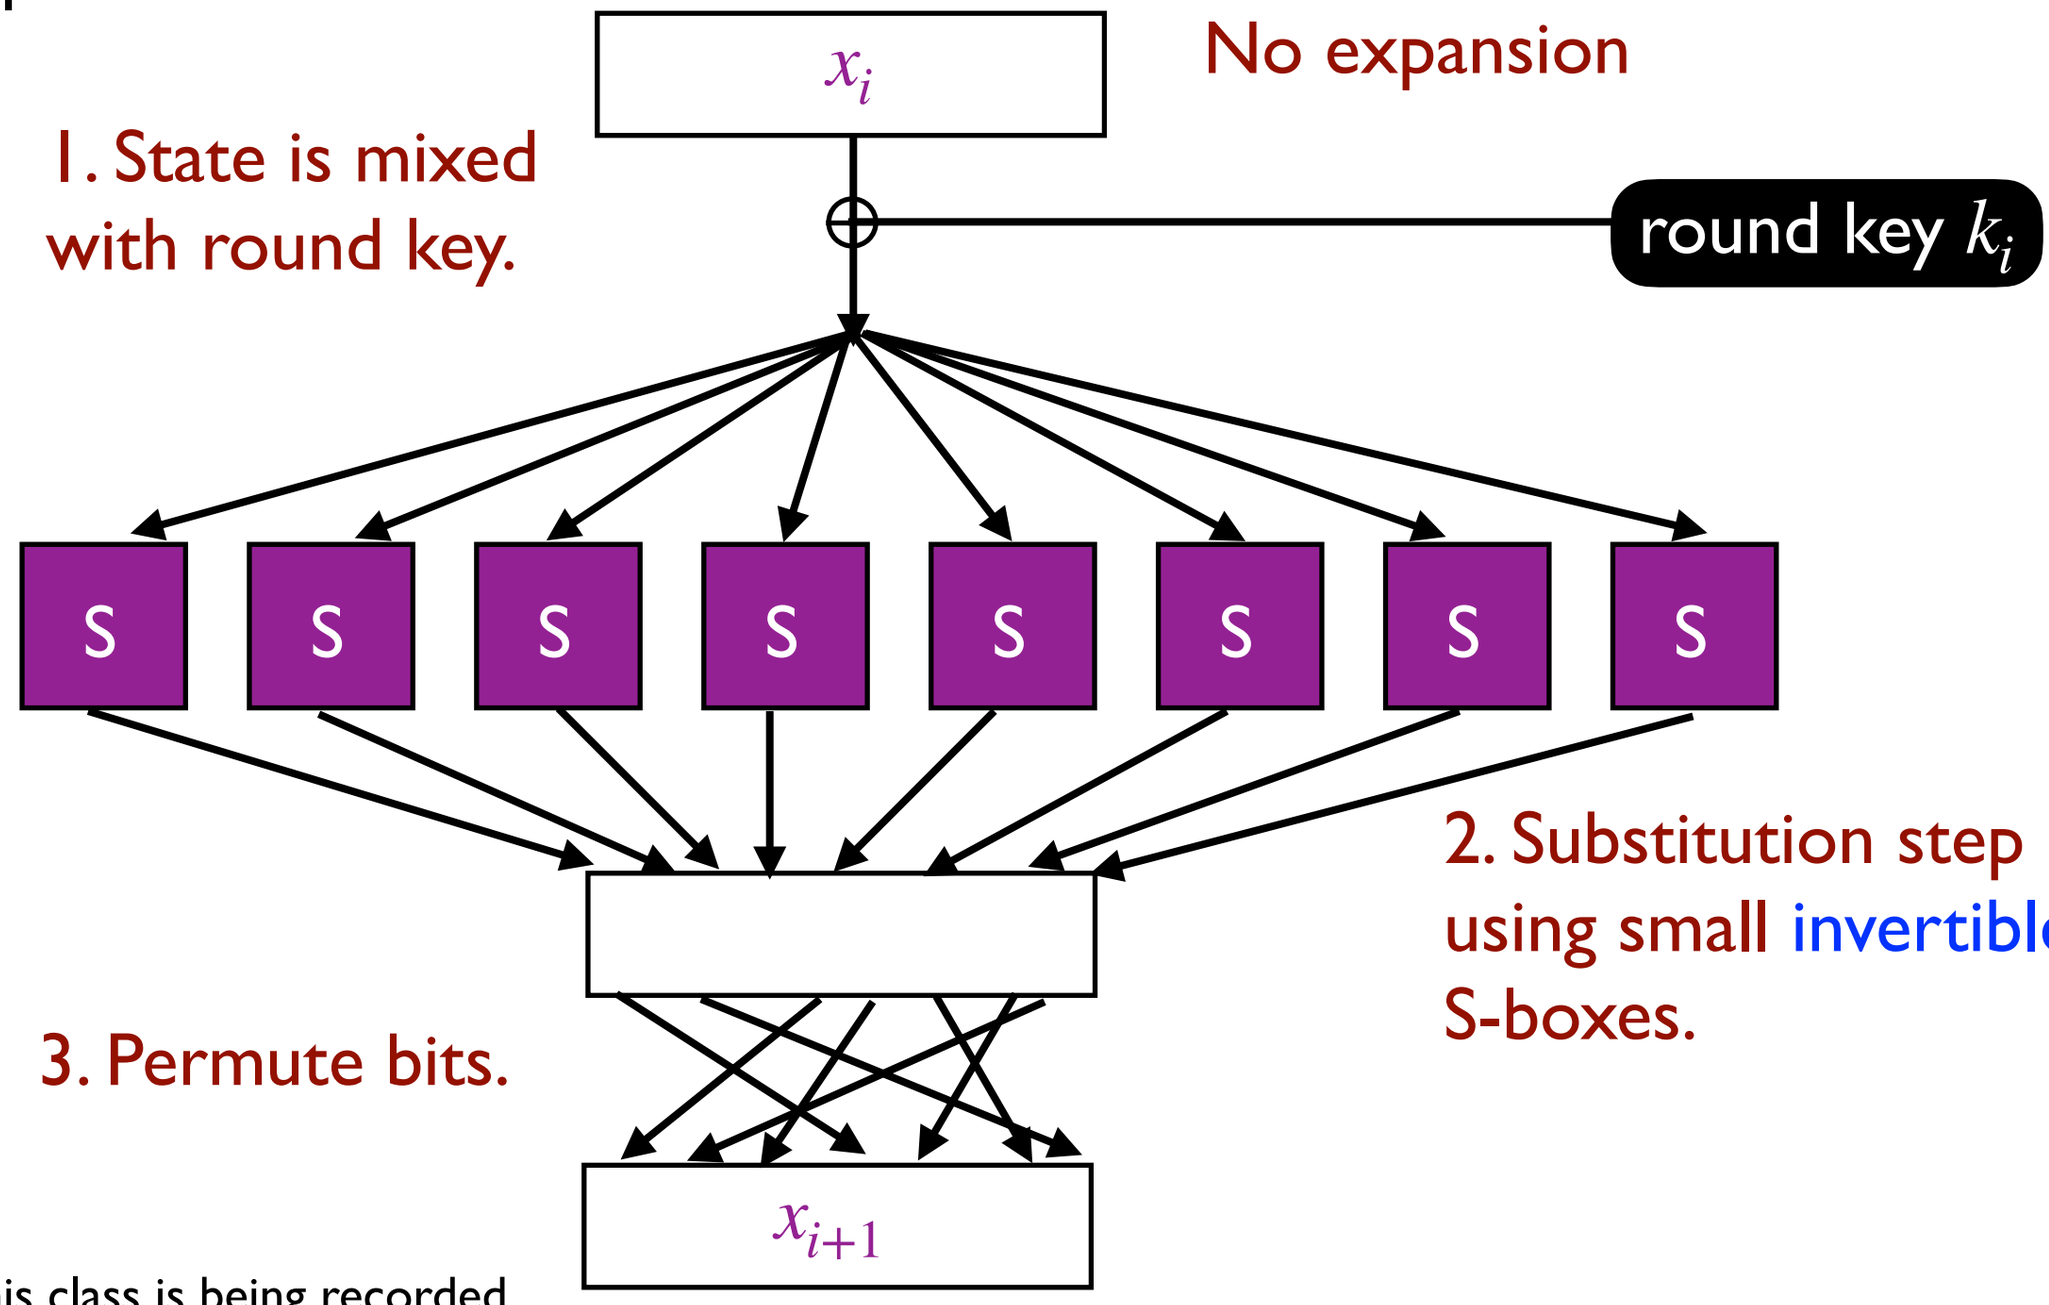

This class is being recorded

Confusion-Diffusion

The S-boxes introduce confusion: They change their inputs into totally different strings and magnify single-bit changes. However, the S-box is small and acts on only a few bits, so the confusion is only local.

Then the permutation step causes diffusion: whatever local confusion was introduced by the S-boxes spreads out to many different locations.

Multiple rounds of substitution and permutation cause the confusion to be magnified further and continue to spread around.

We need both to get an avalanche effect.

You also need key mixing: This is a permutation, and without the key, Eve can just trace the permutation backwards to get the input.

This class is being recorded

Substitution Only

Suppose we have S-boxes but no permutation.

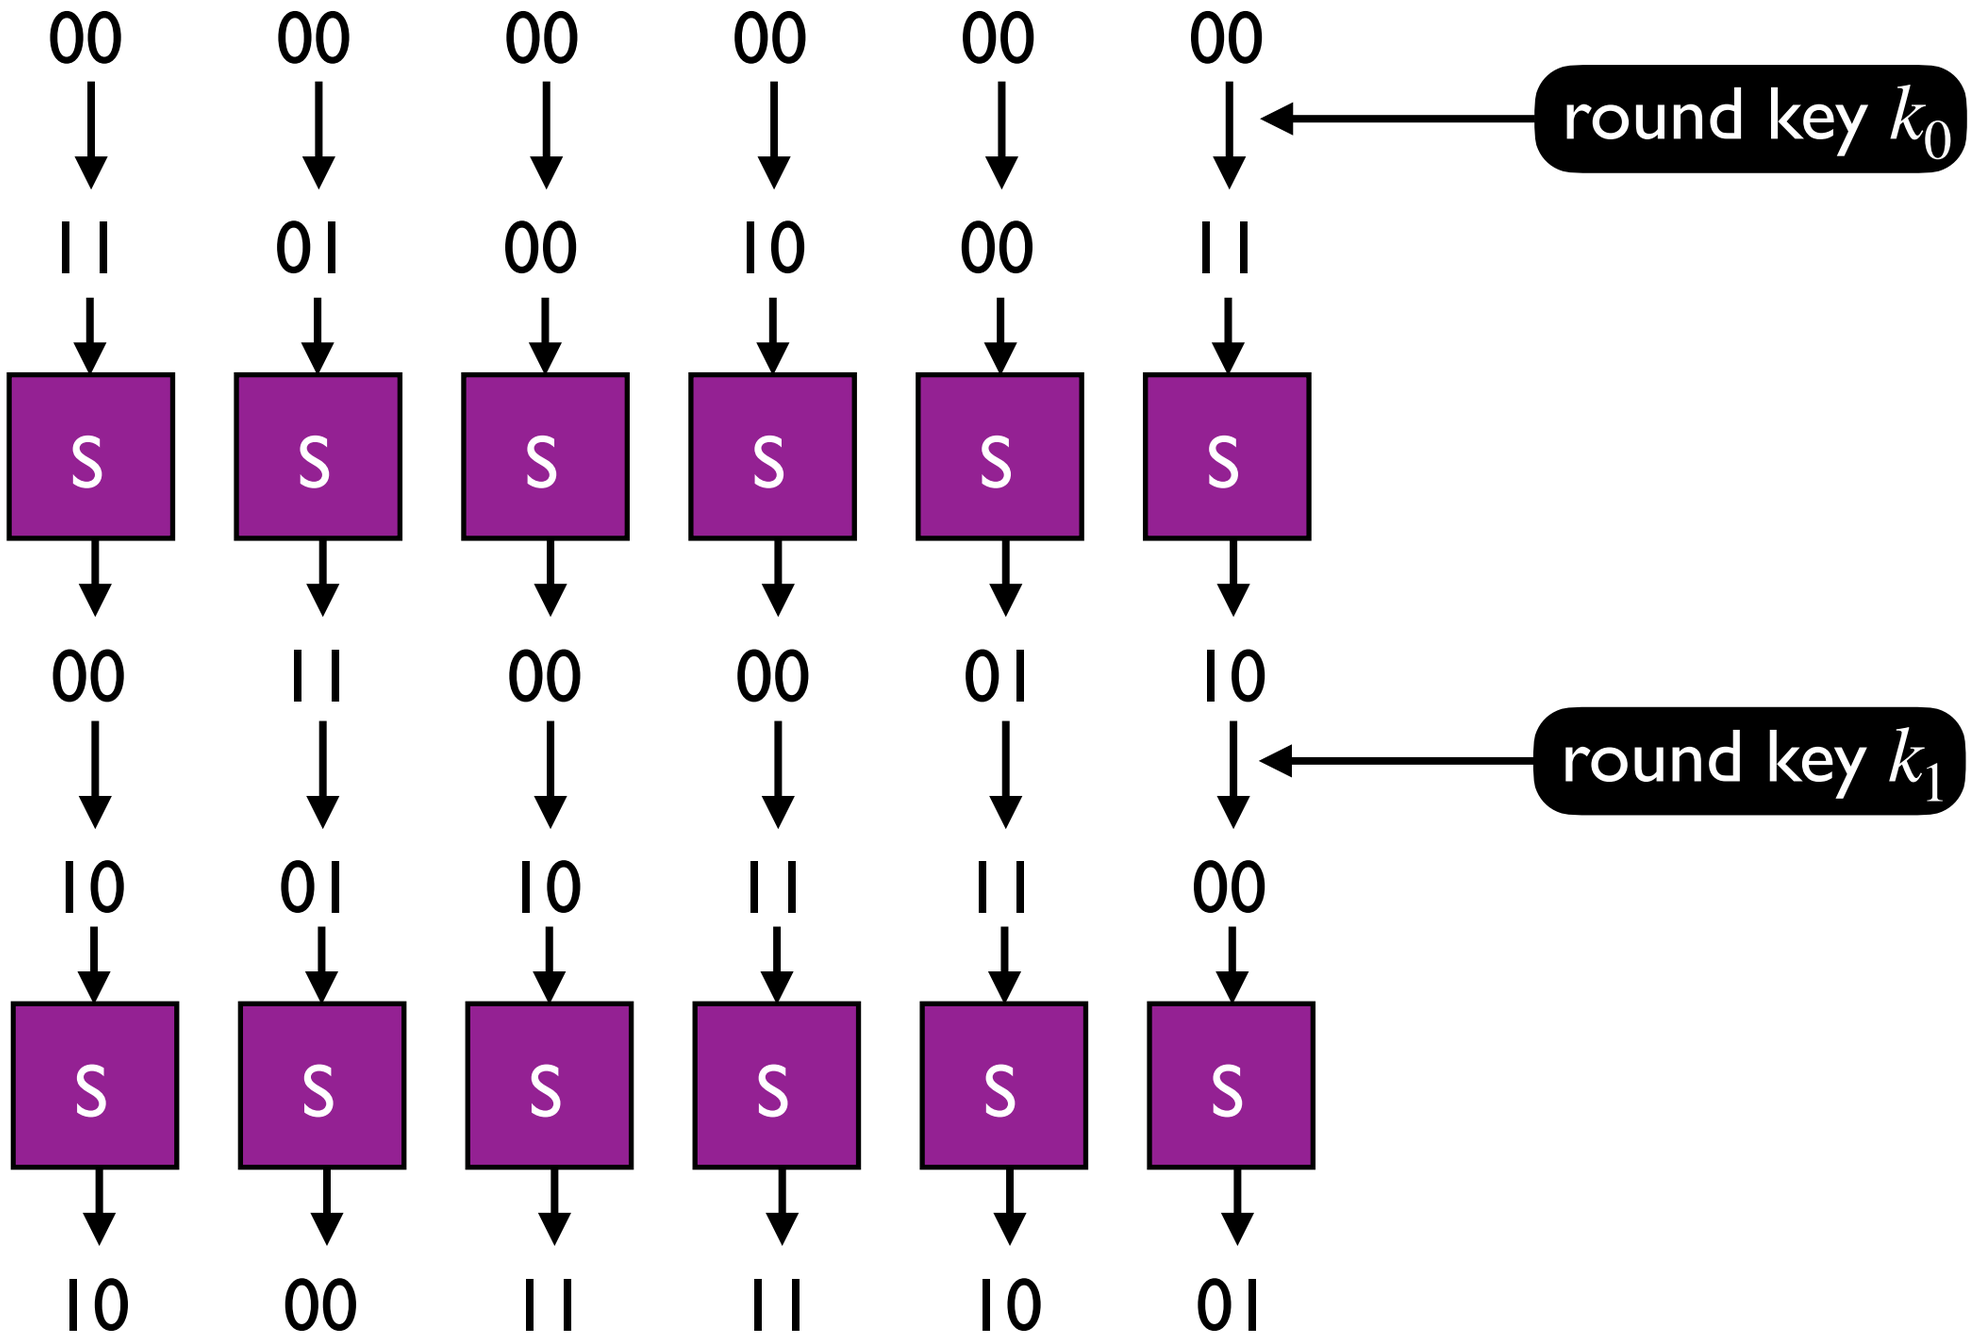

This class is being recorded

Substitution Only

Suppose we have S-boxes but no permutation.

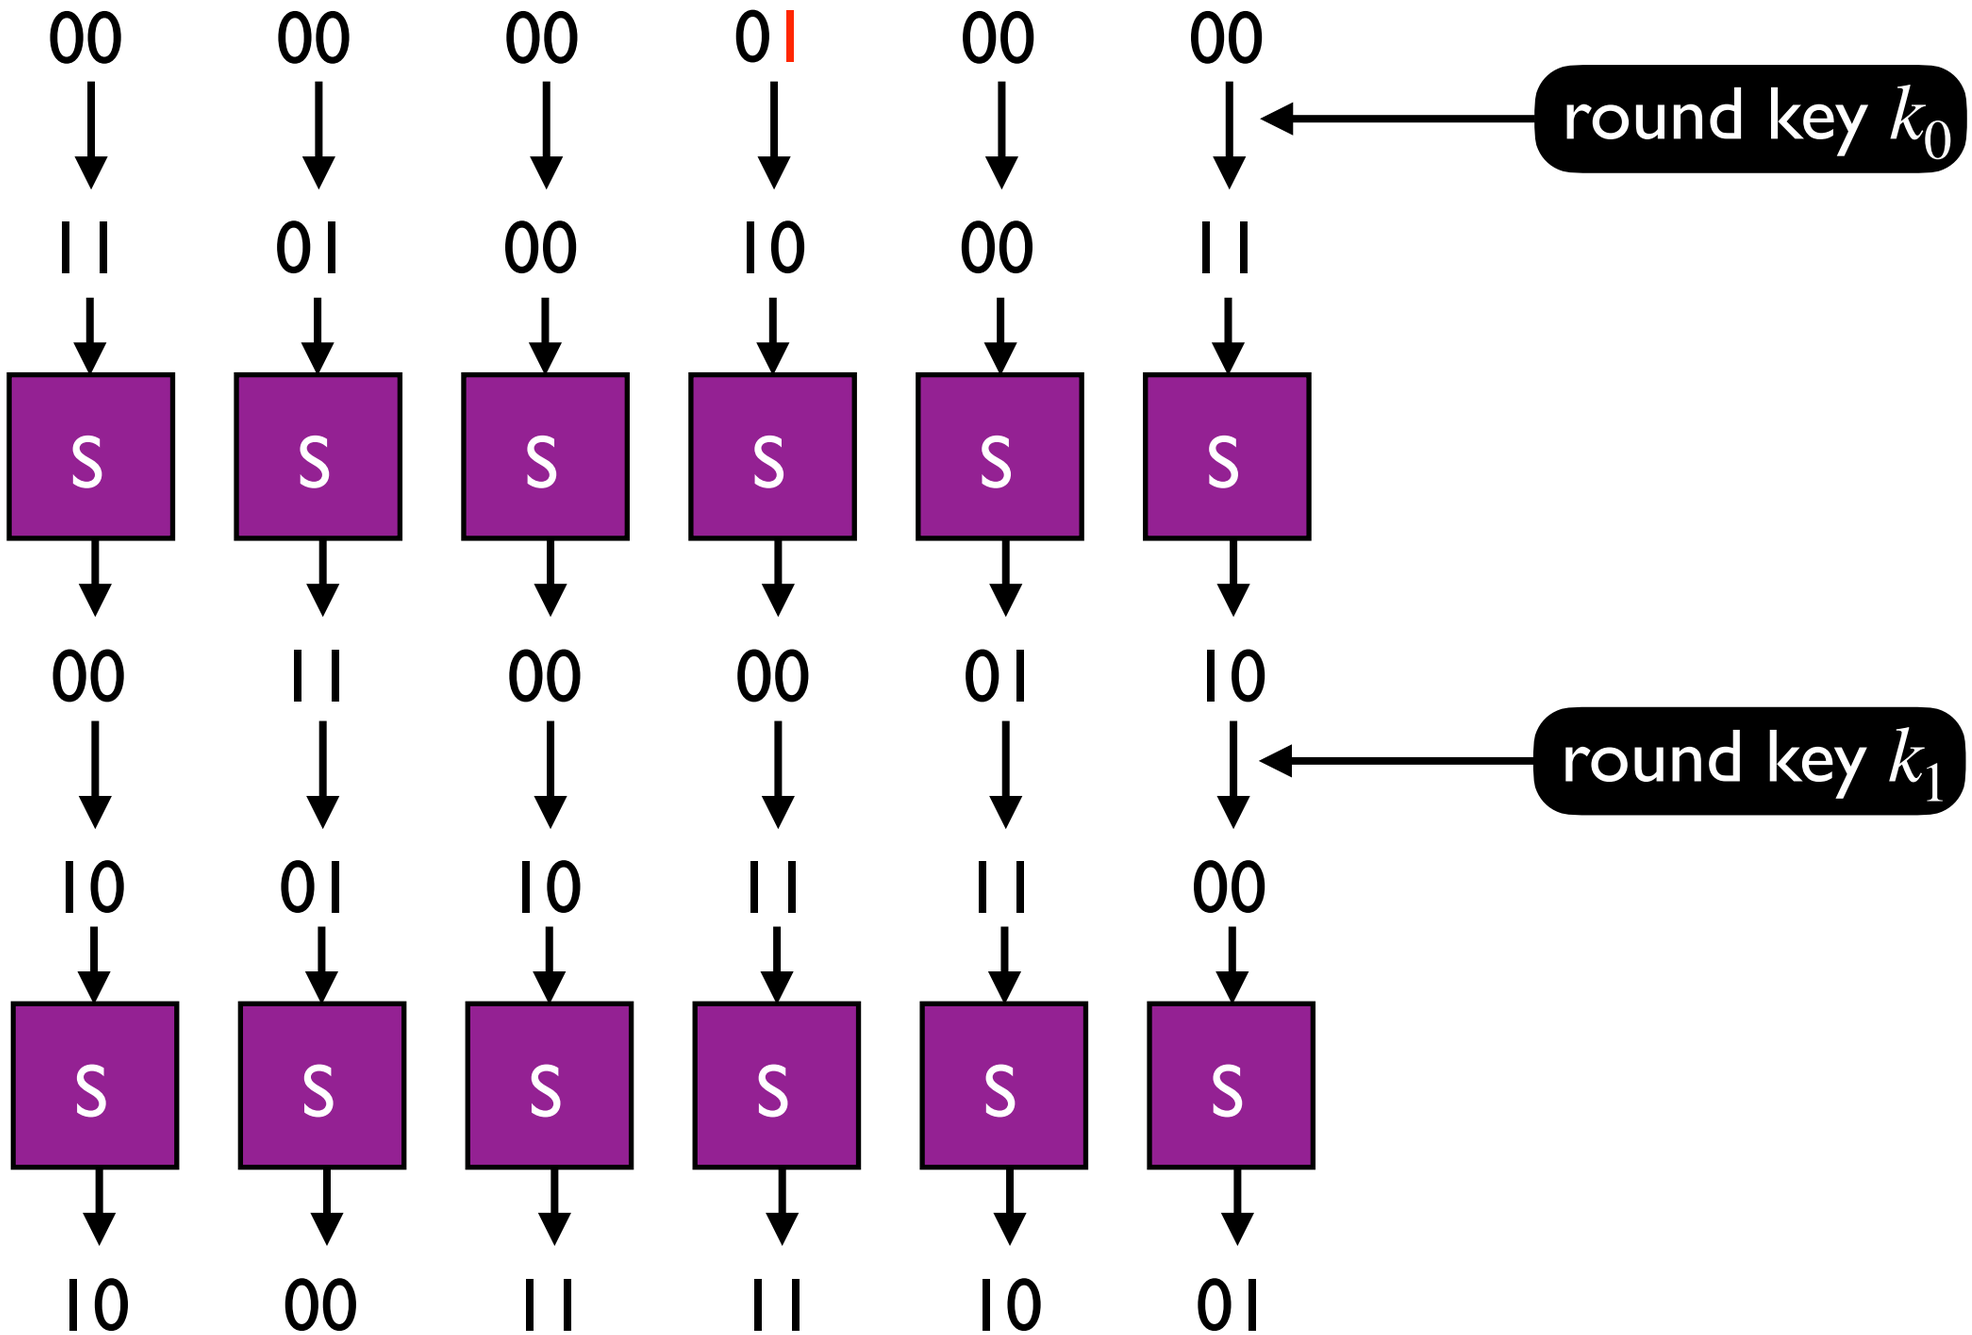

This class is being recorded

Substitution Only

Suppose we have S-boxes but no permutation.

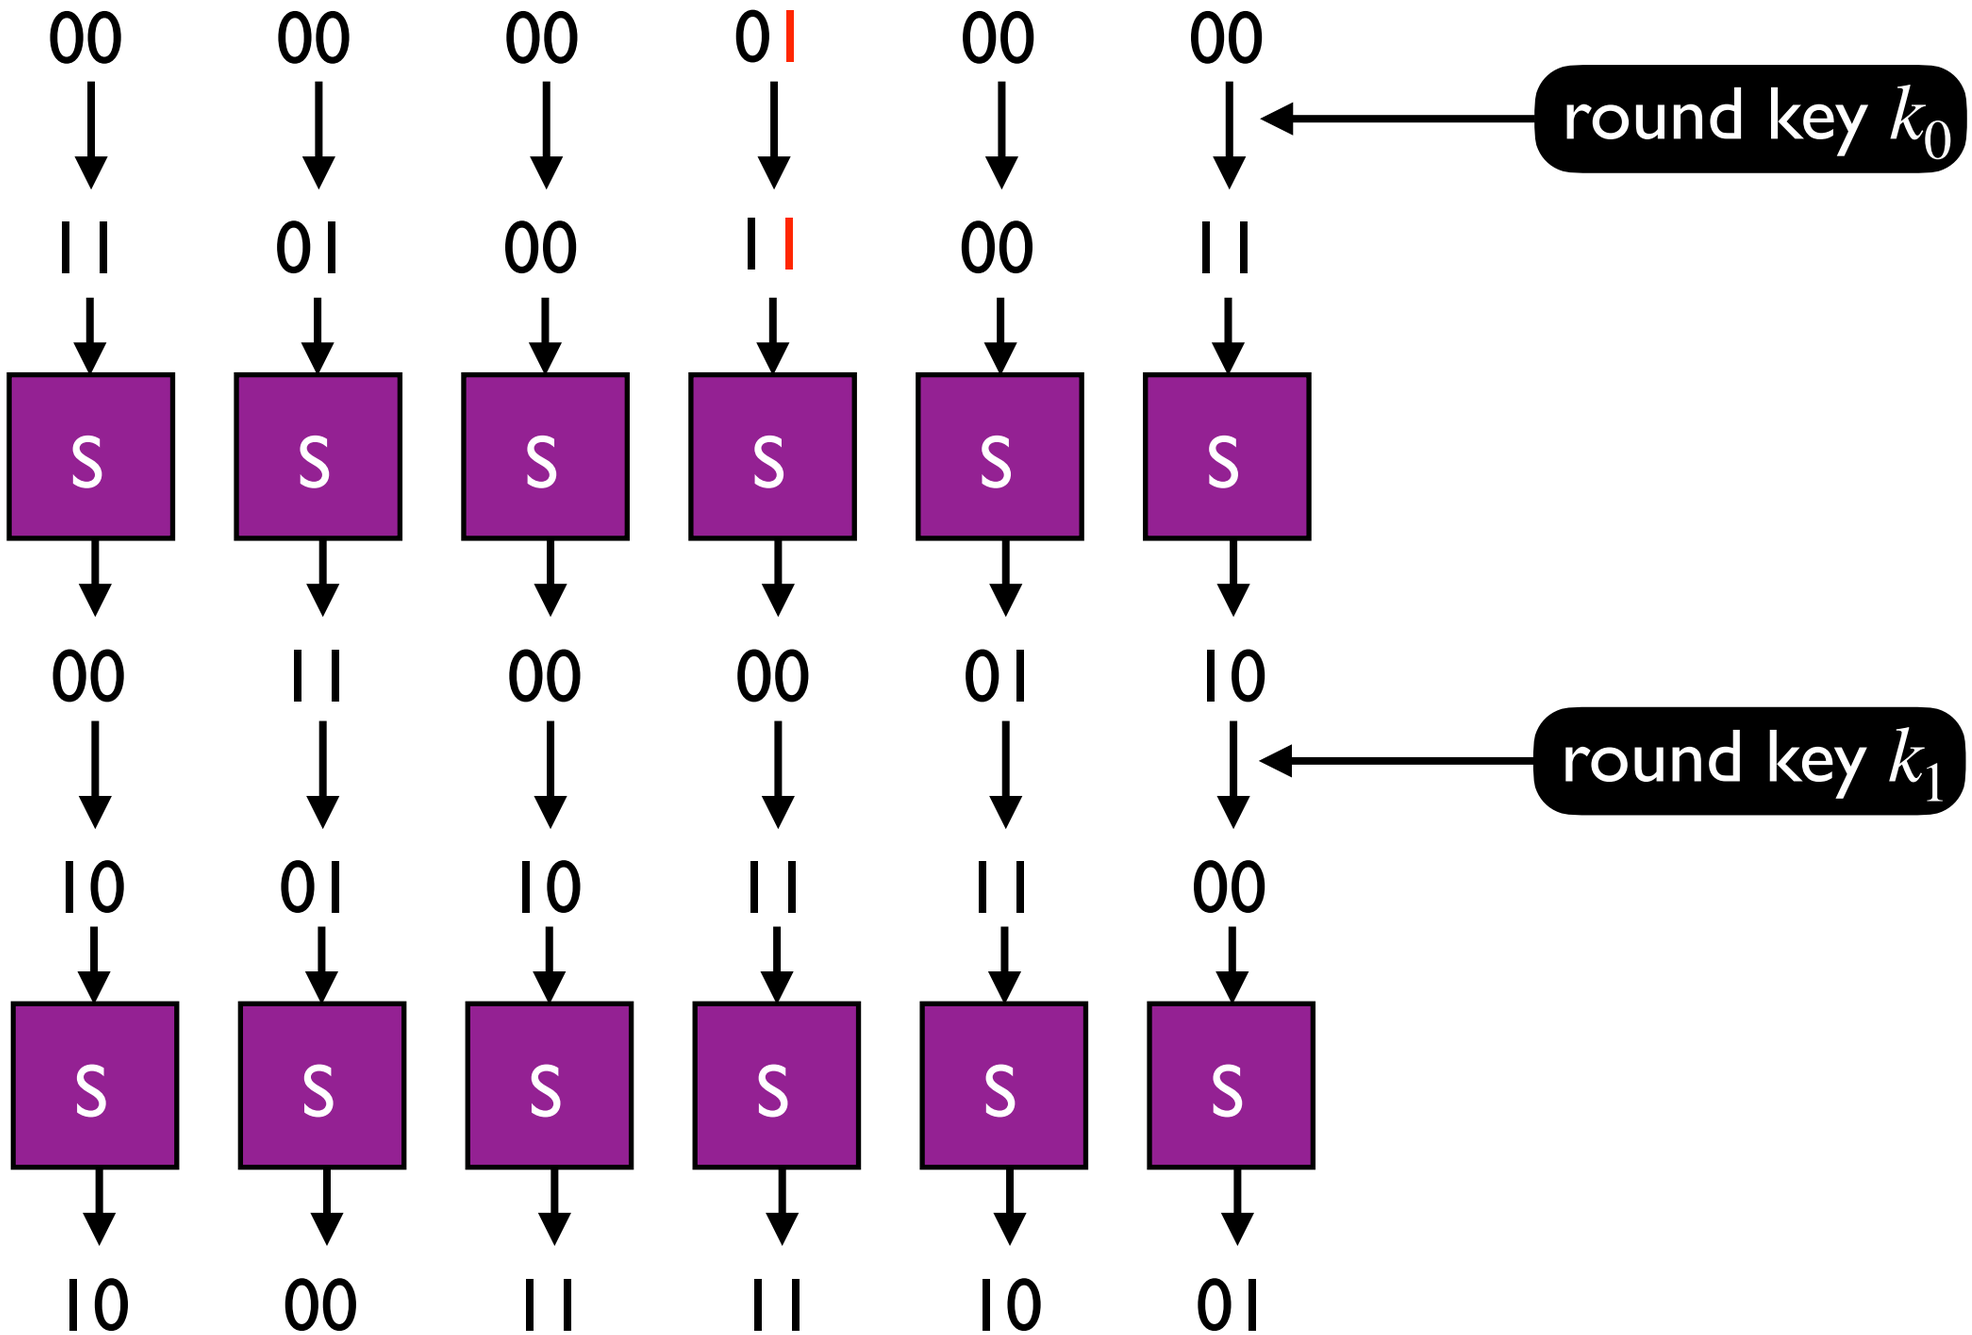

This class is being recorded

Substitution Only

Suppose we have S-boxes but no permutation.

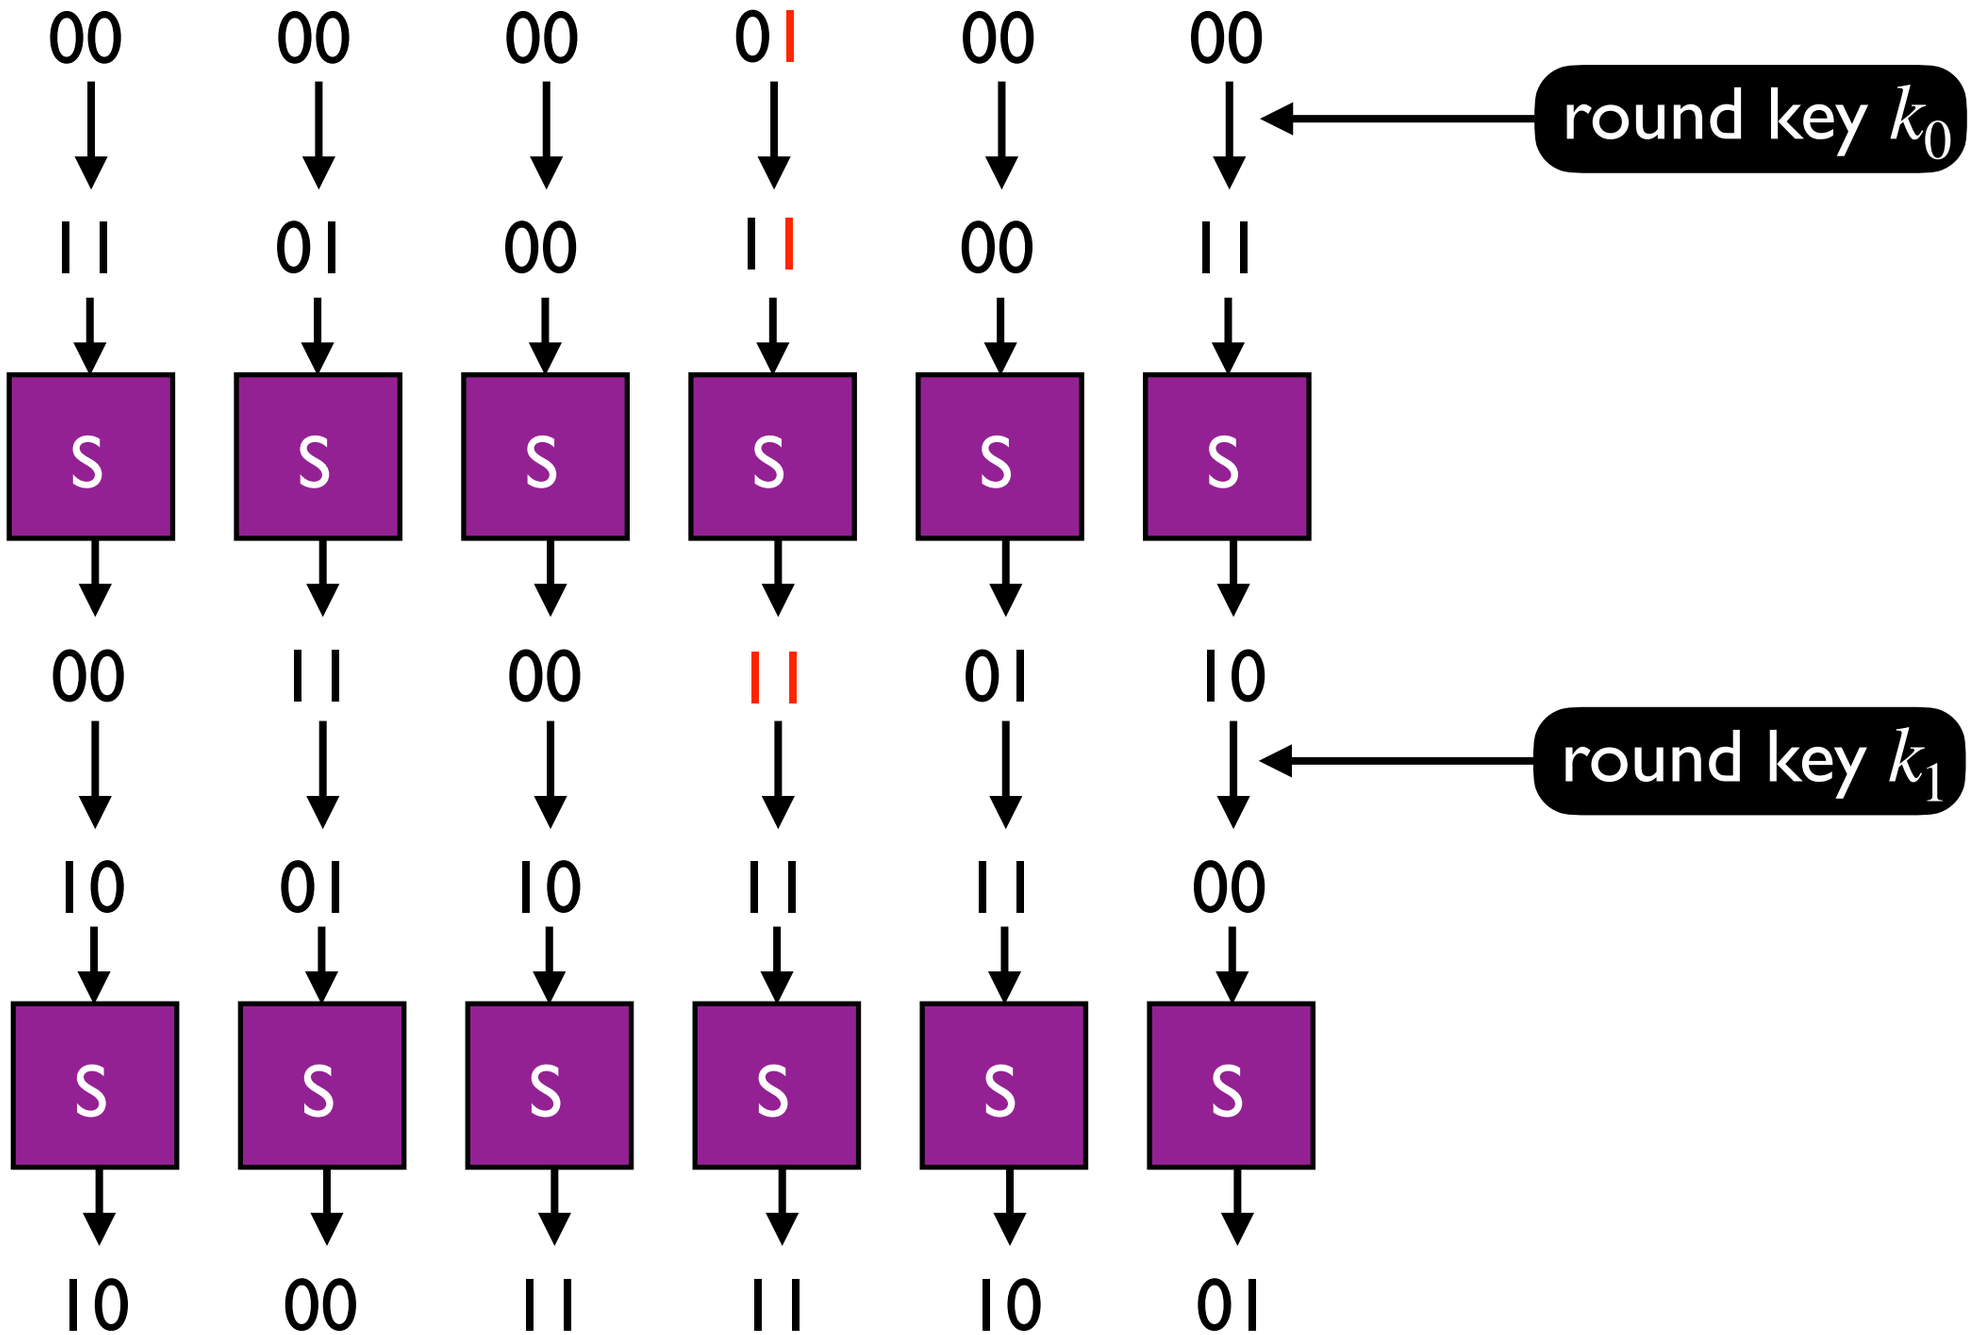

This class is being recorded

Substitution Only

Suppose we have S-boxes but no permutation.

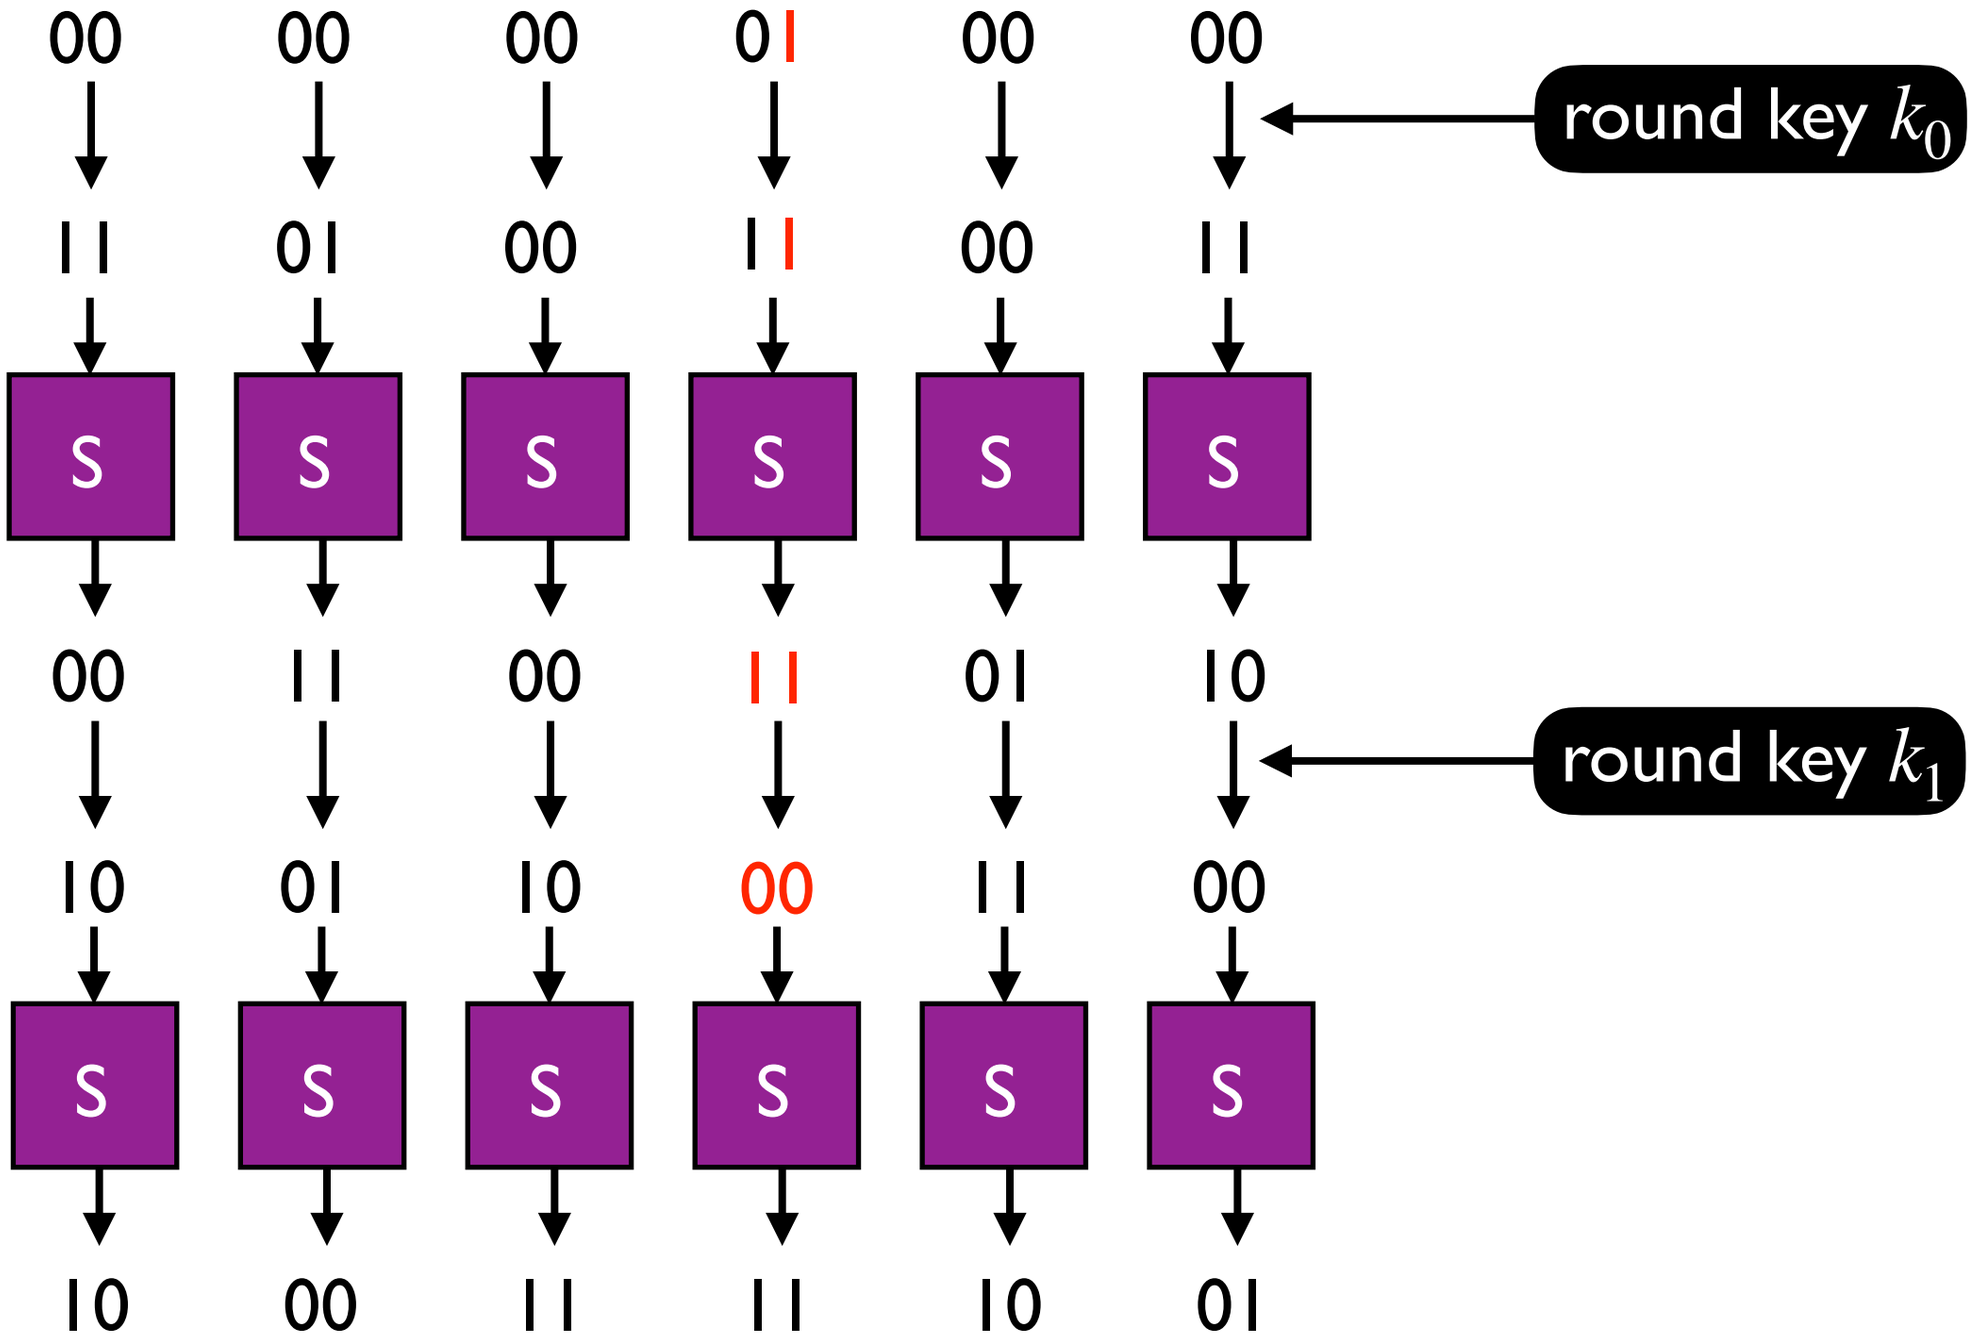

This class is being recorded

Substitution Only

Suppose we have S-boxes but no permutation.

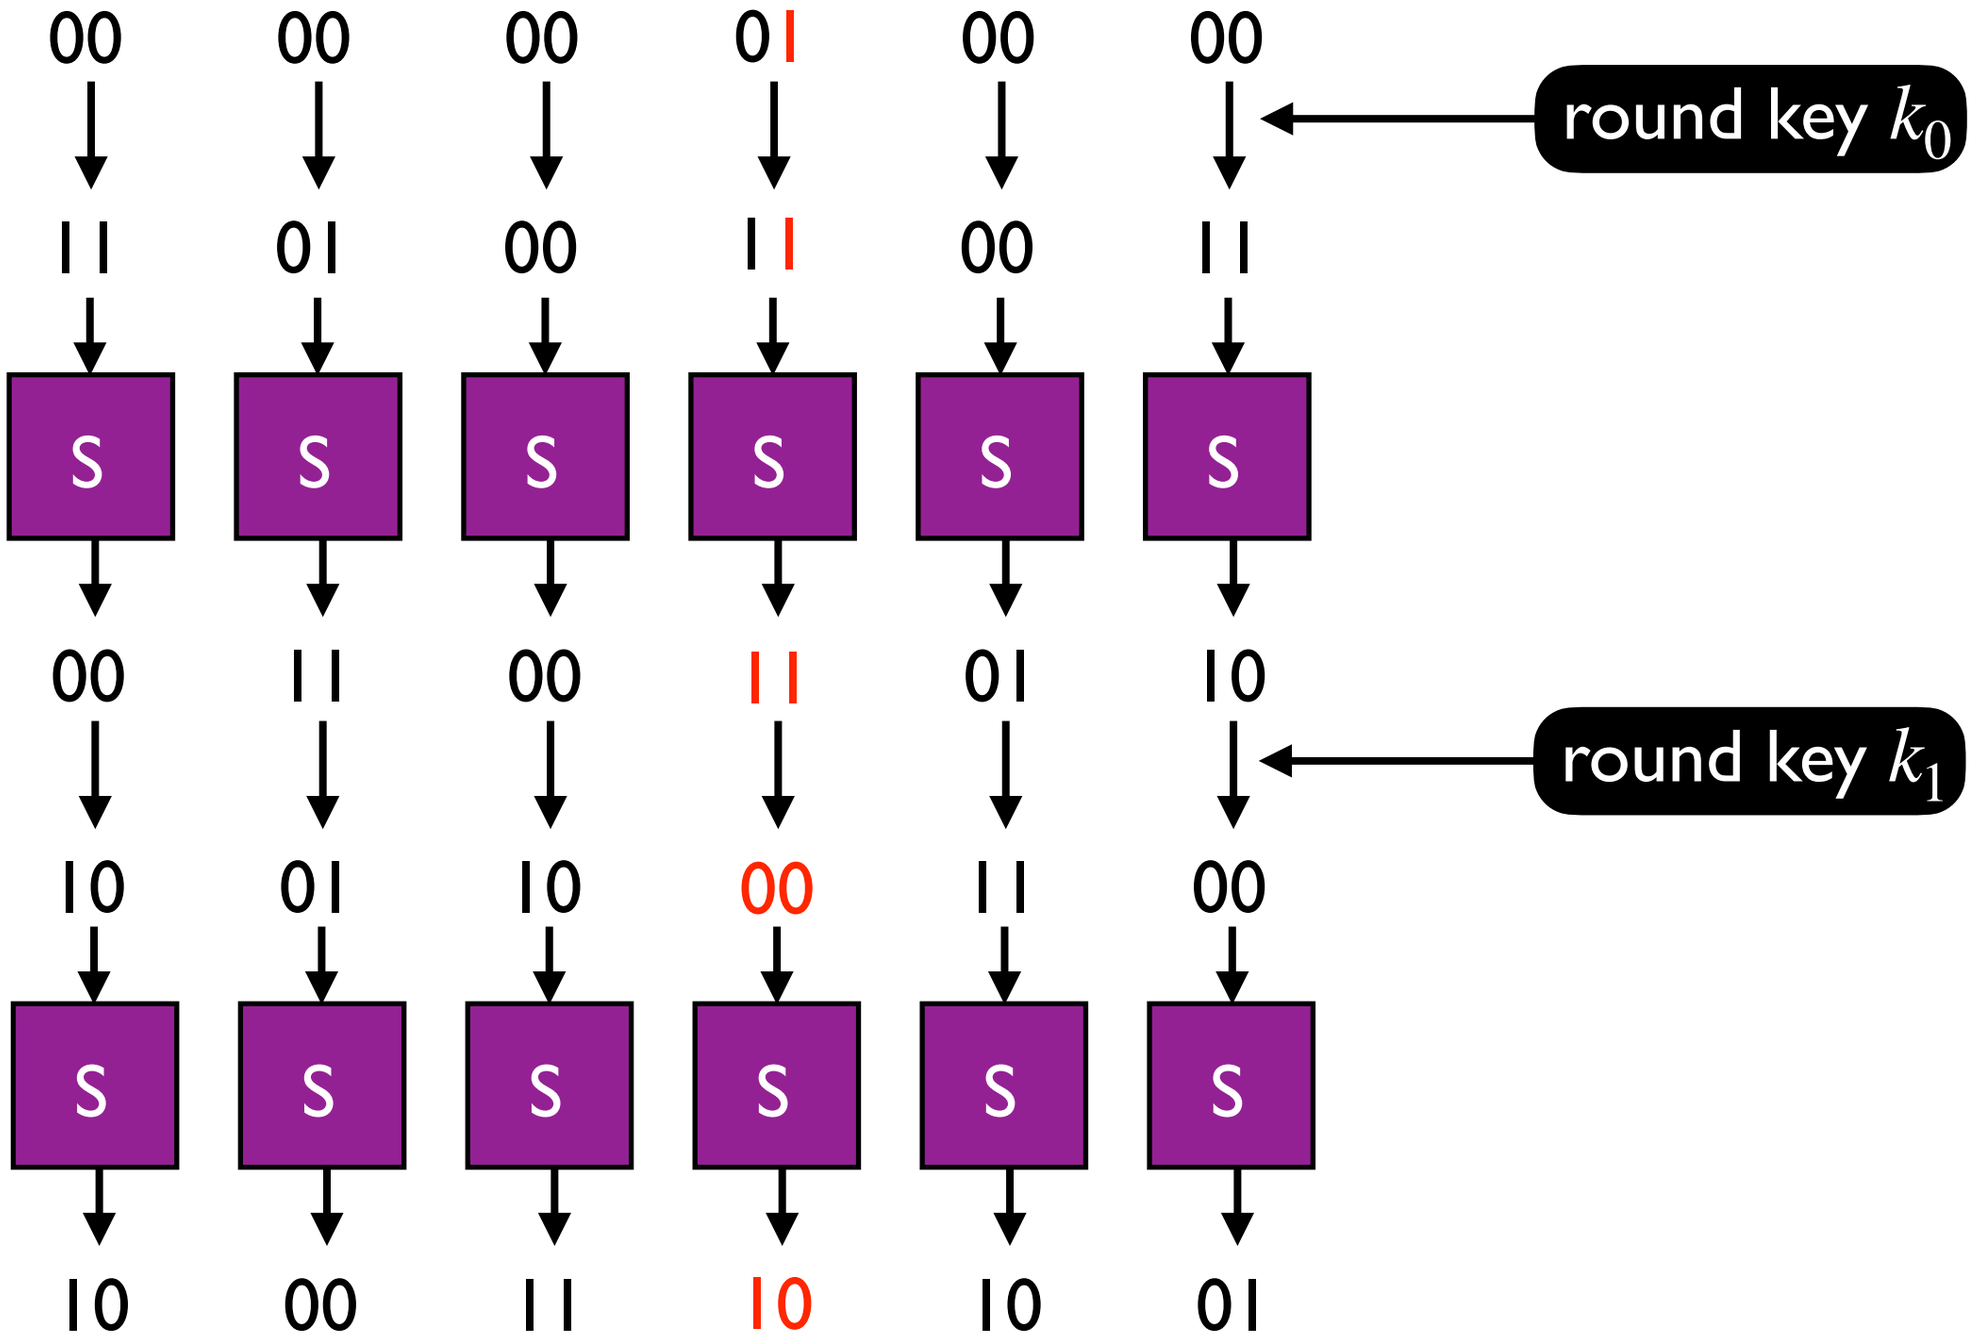

This class is being recorded

Substitution Only

Suppose we have S-boxes but no permutation.

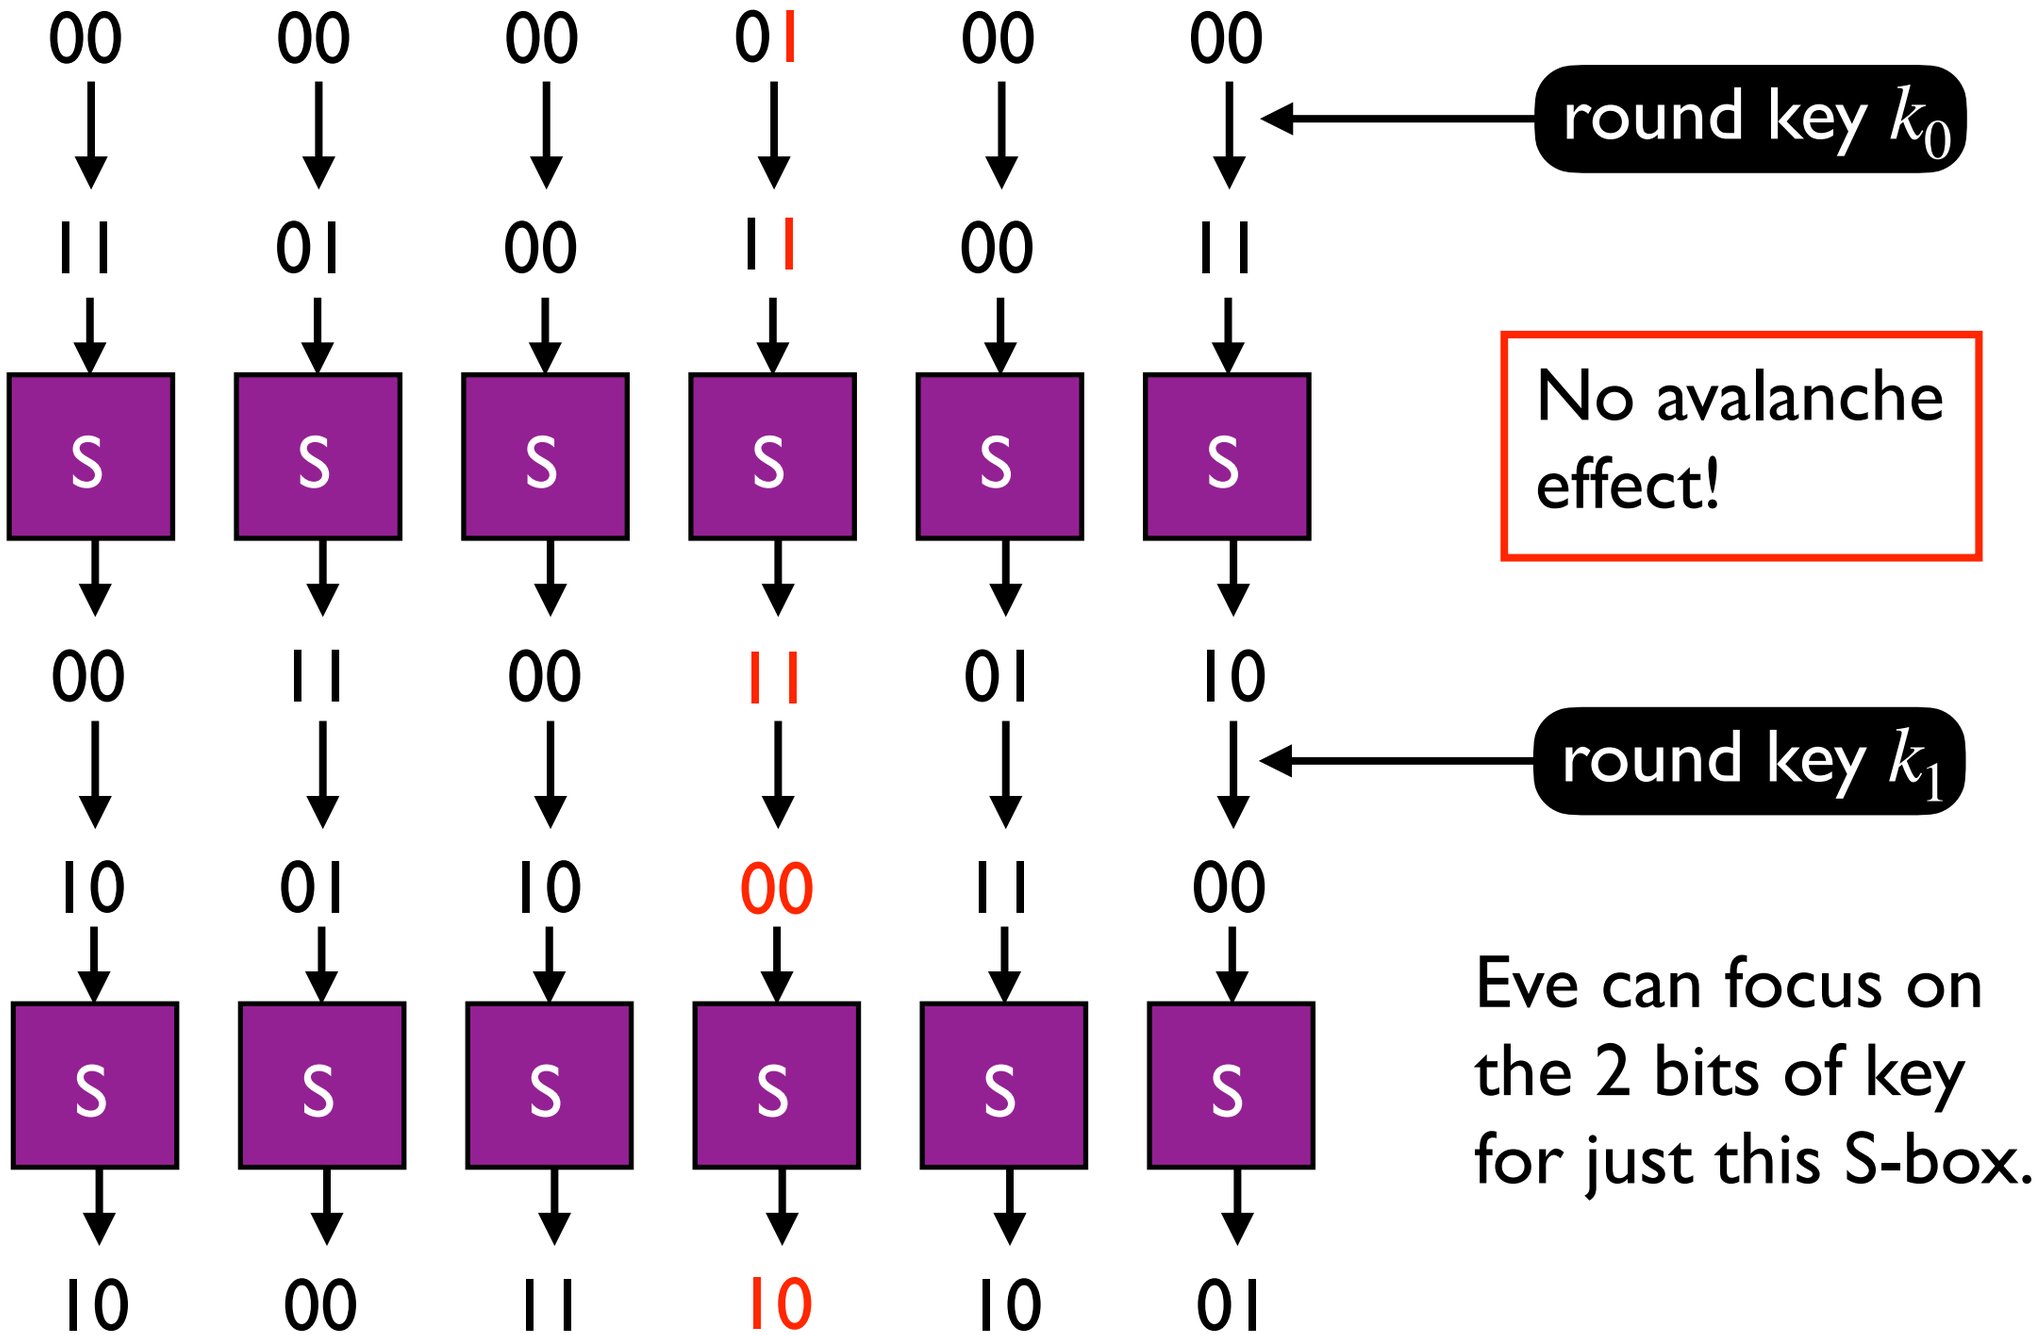

No avalanche effect!

Eve can focus on the 2 bits of key for just this S-box.

This class is being recorded

Permutation Only

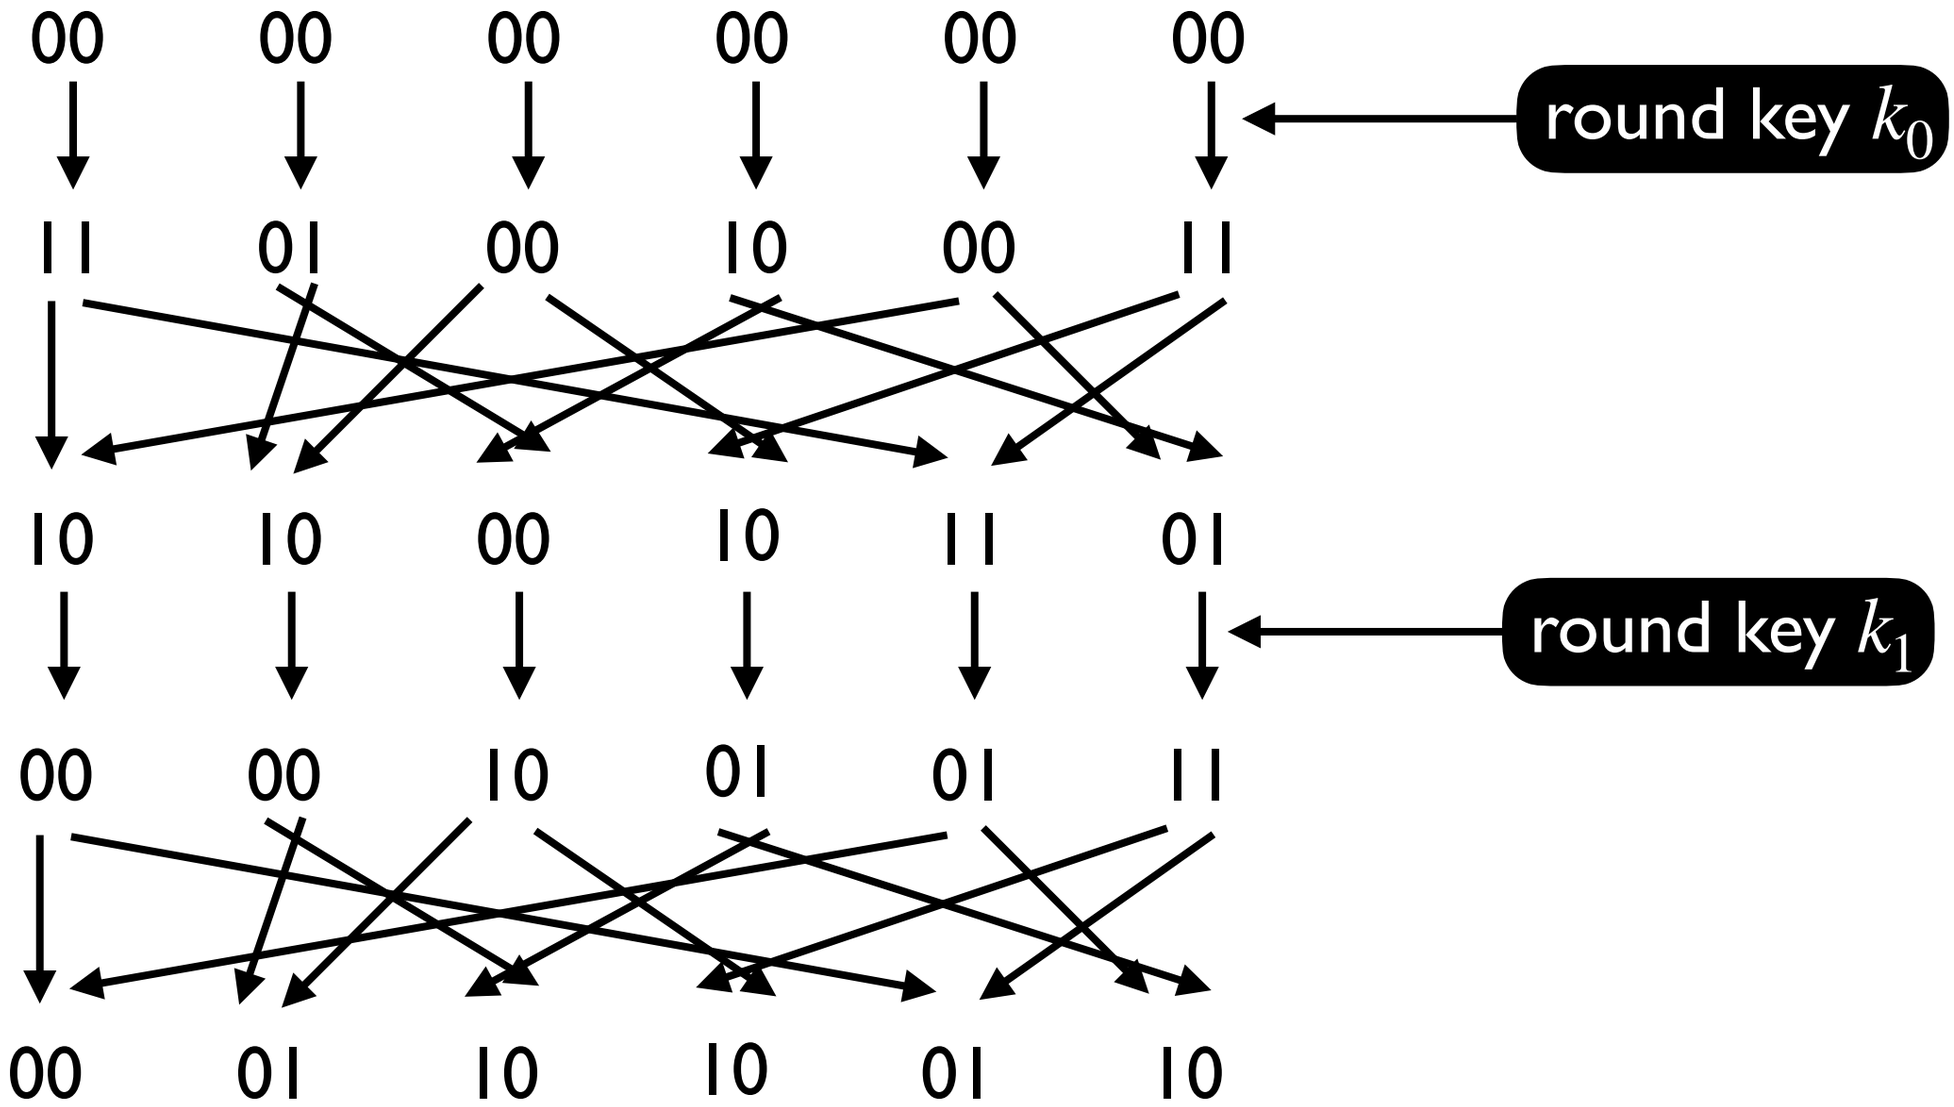

This class is being recorded

Permutation Only

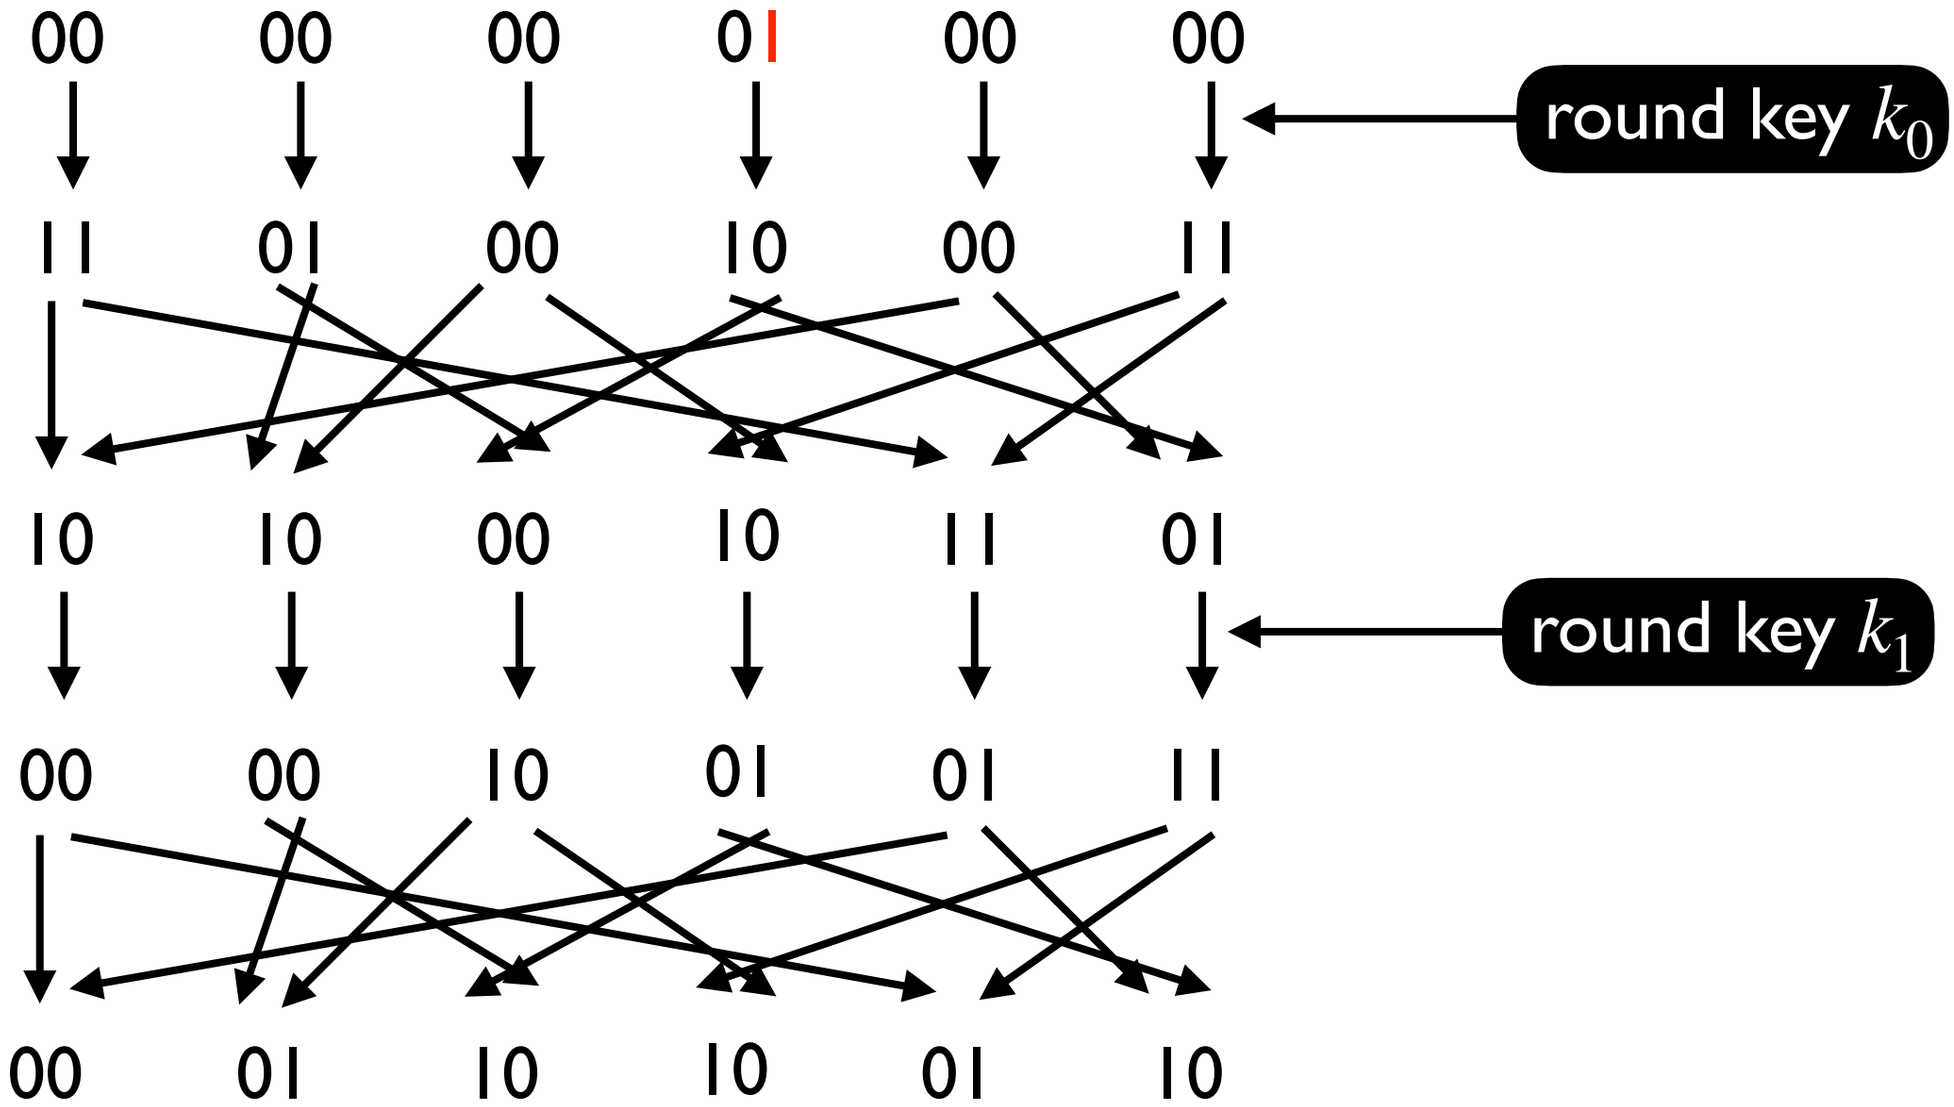

This class is being recorded

Permutation Only

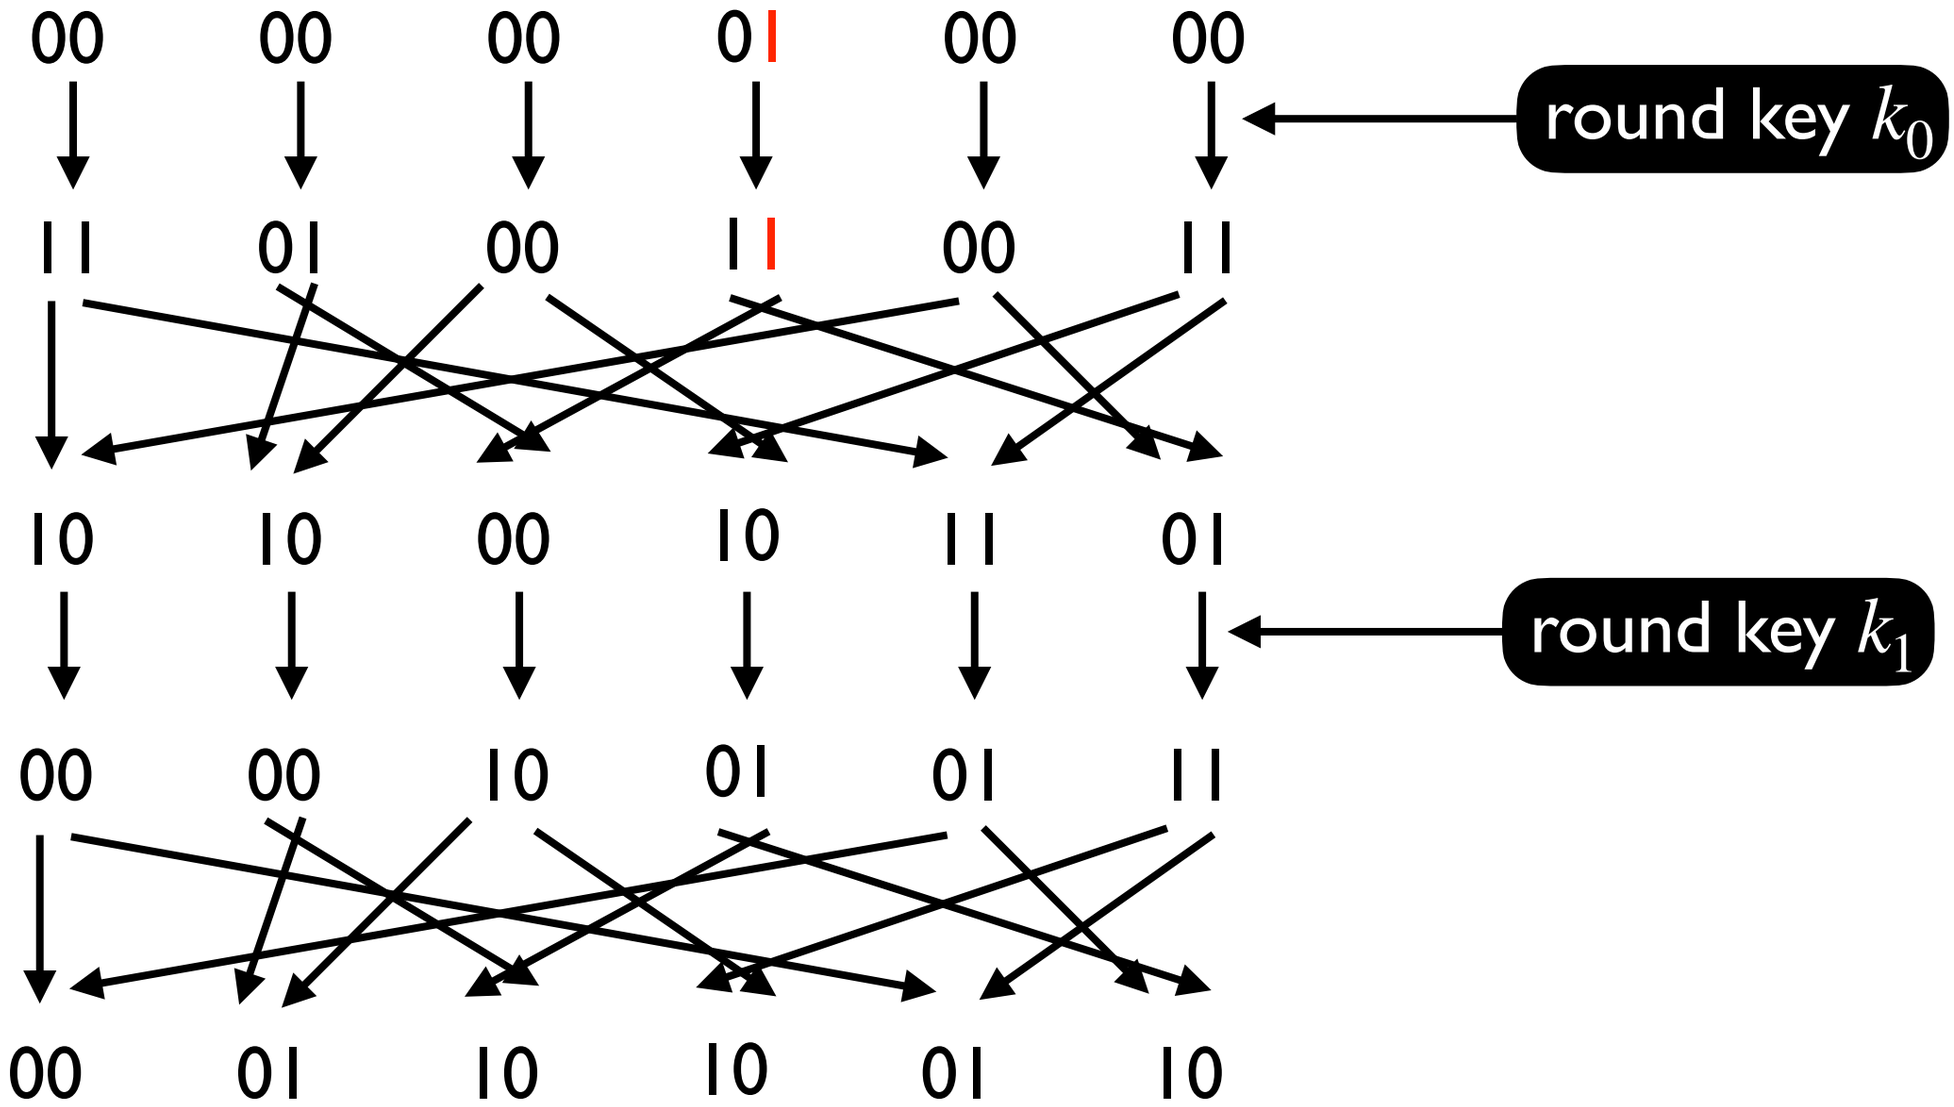

This class is being recorded

Permutation Only

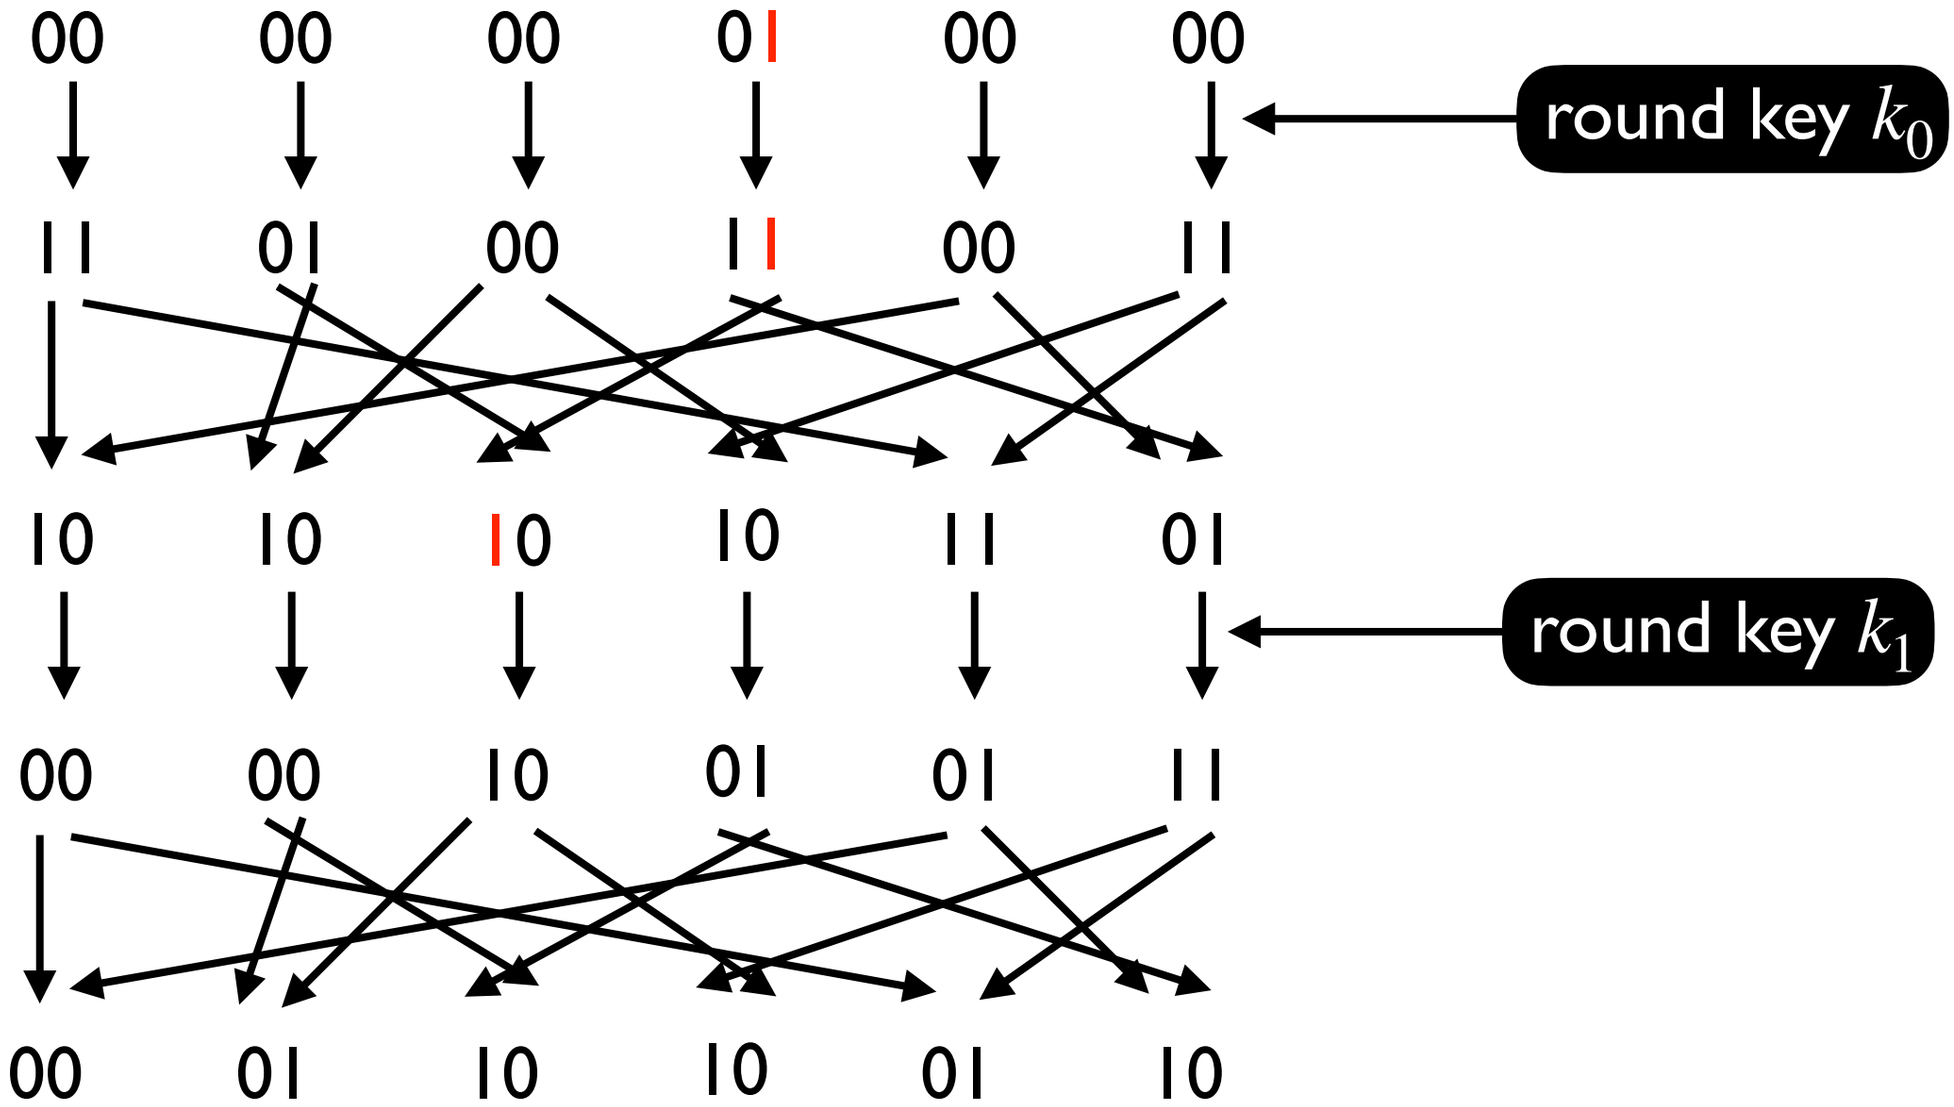

This class is being recorded

Permutation Only

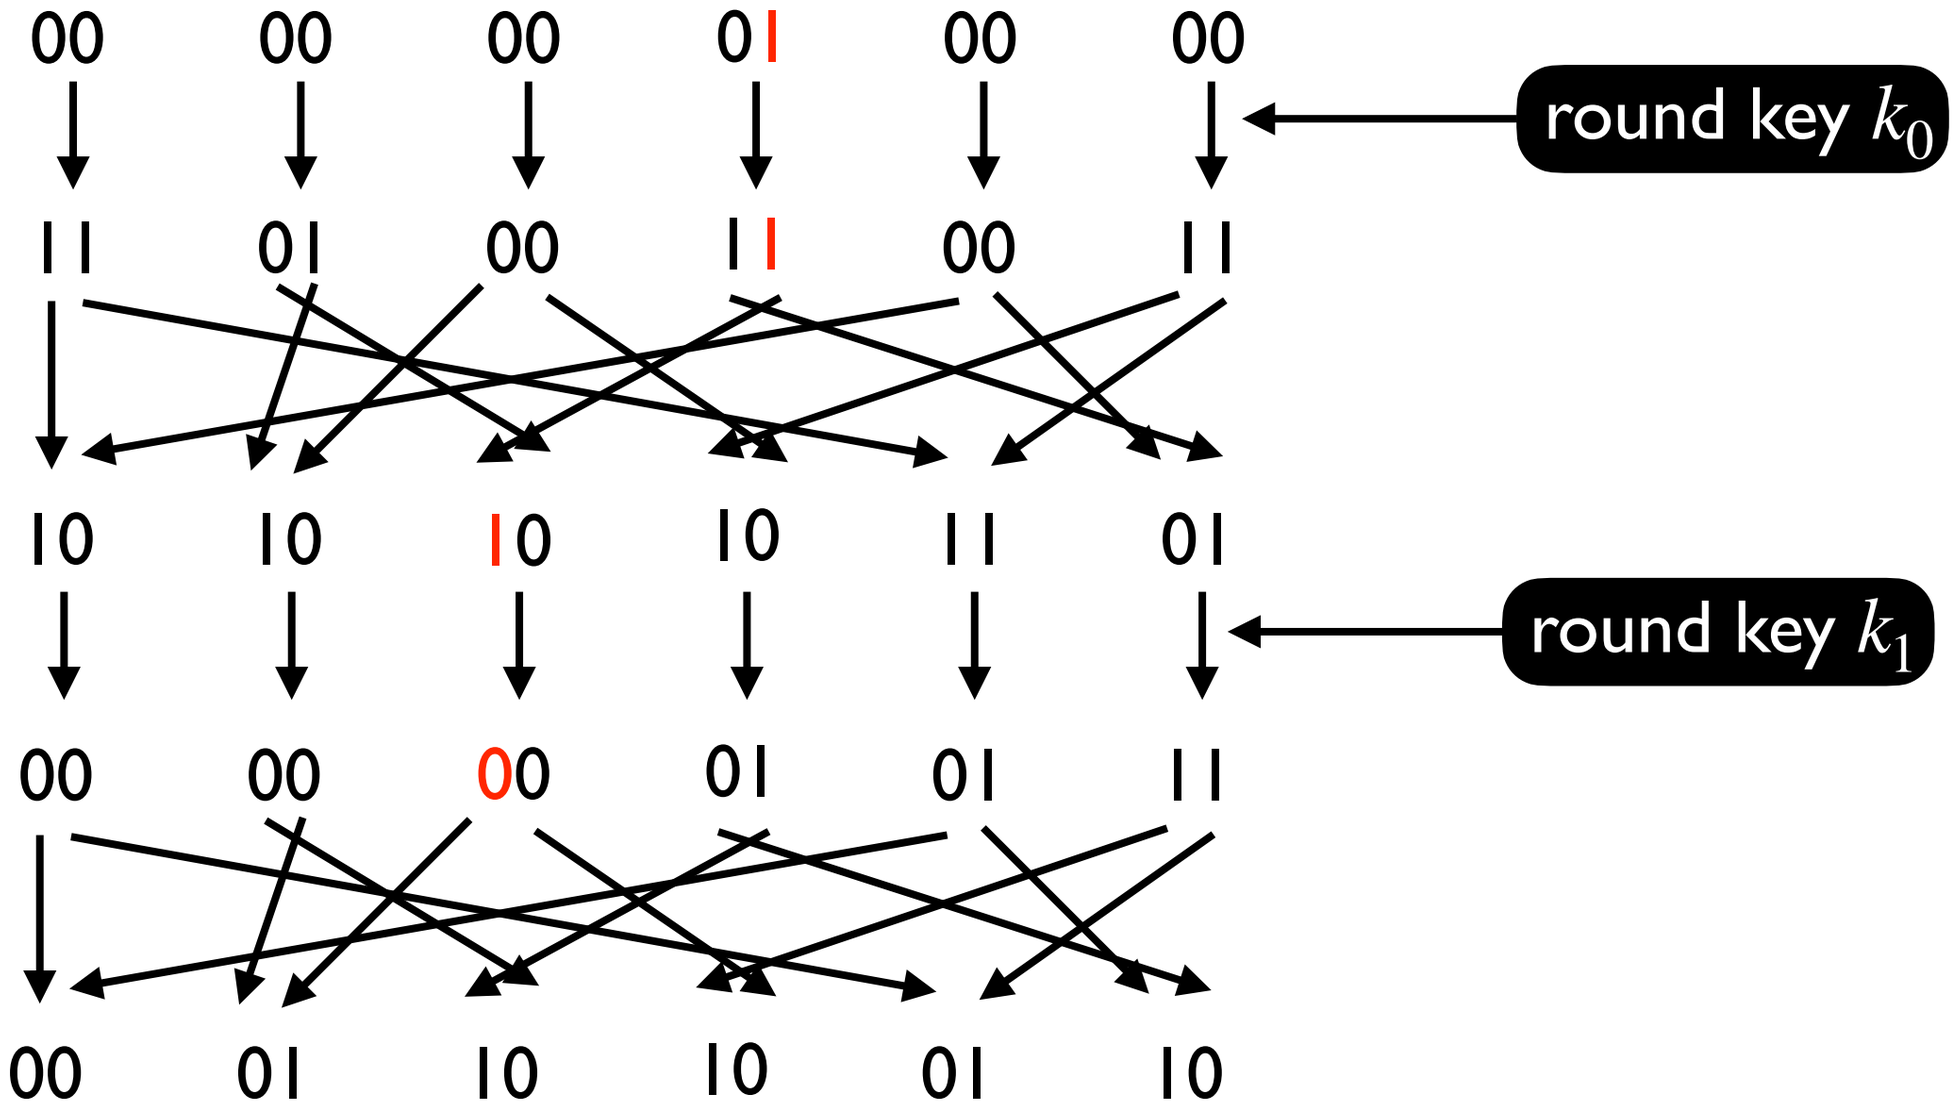

This class is being recorded

Permutation Only

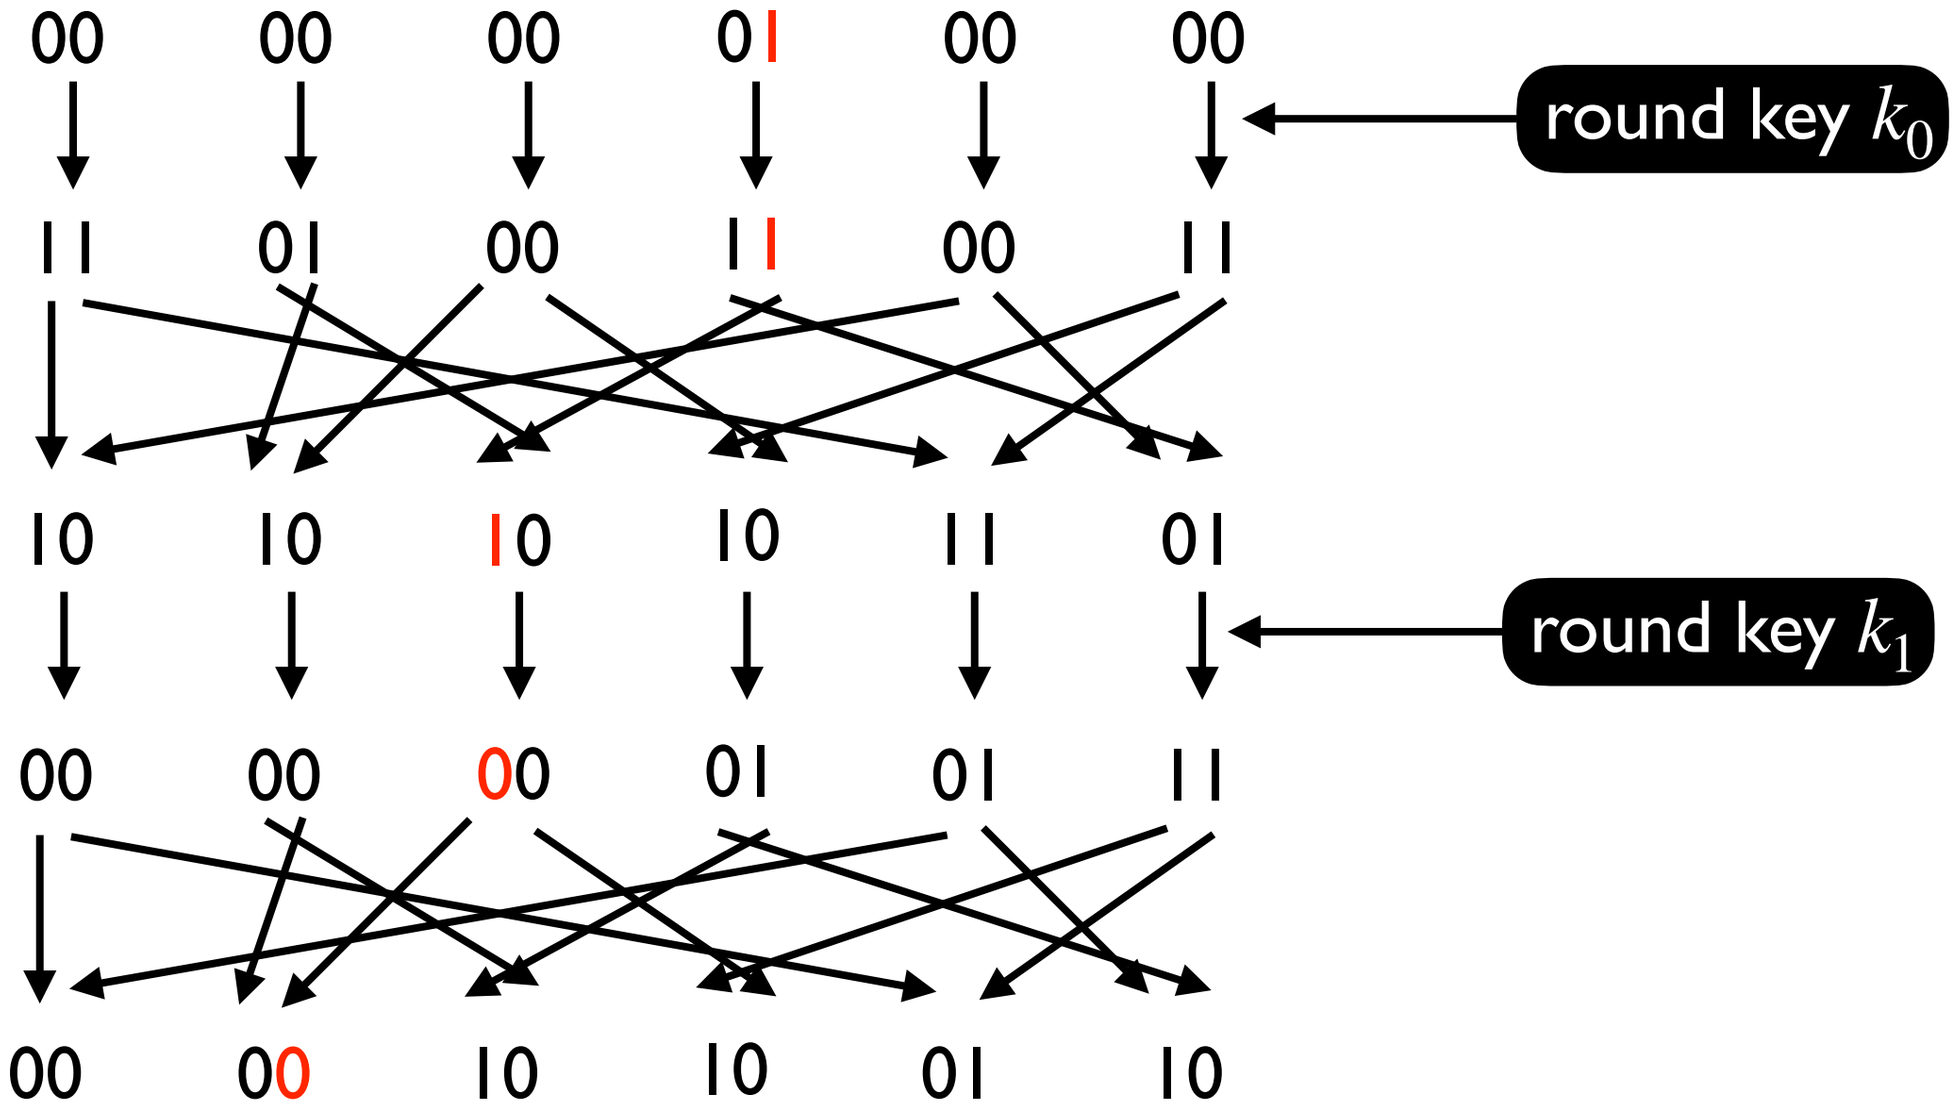

This class is being recorded

Permutation Only

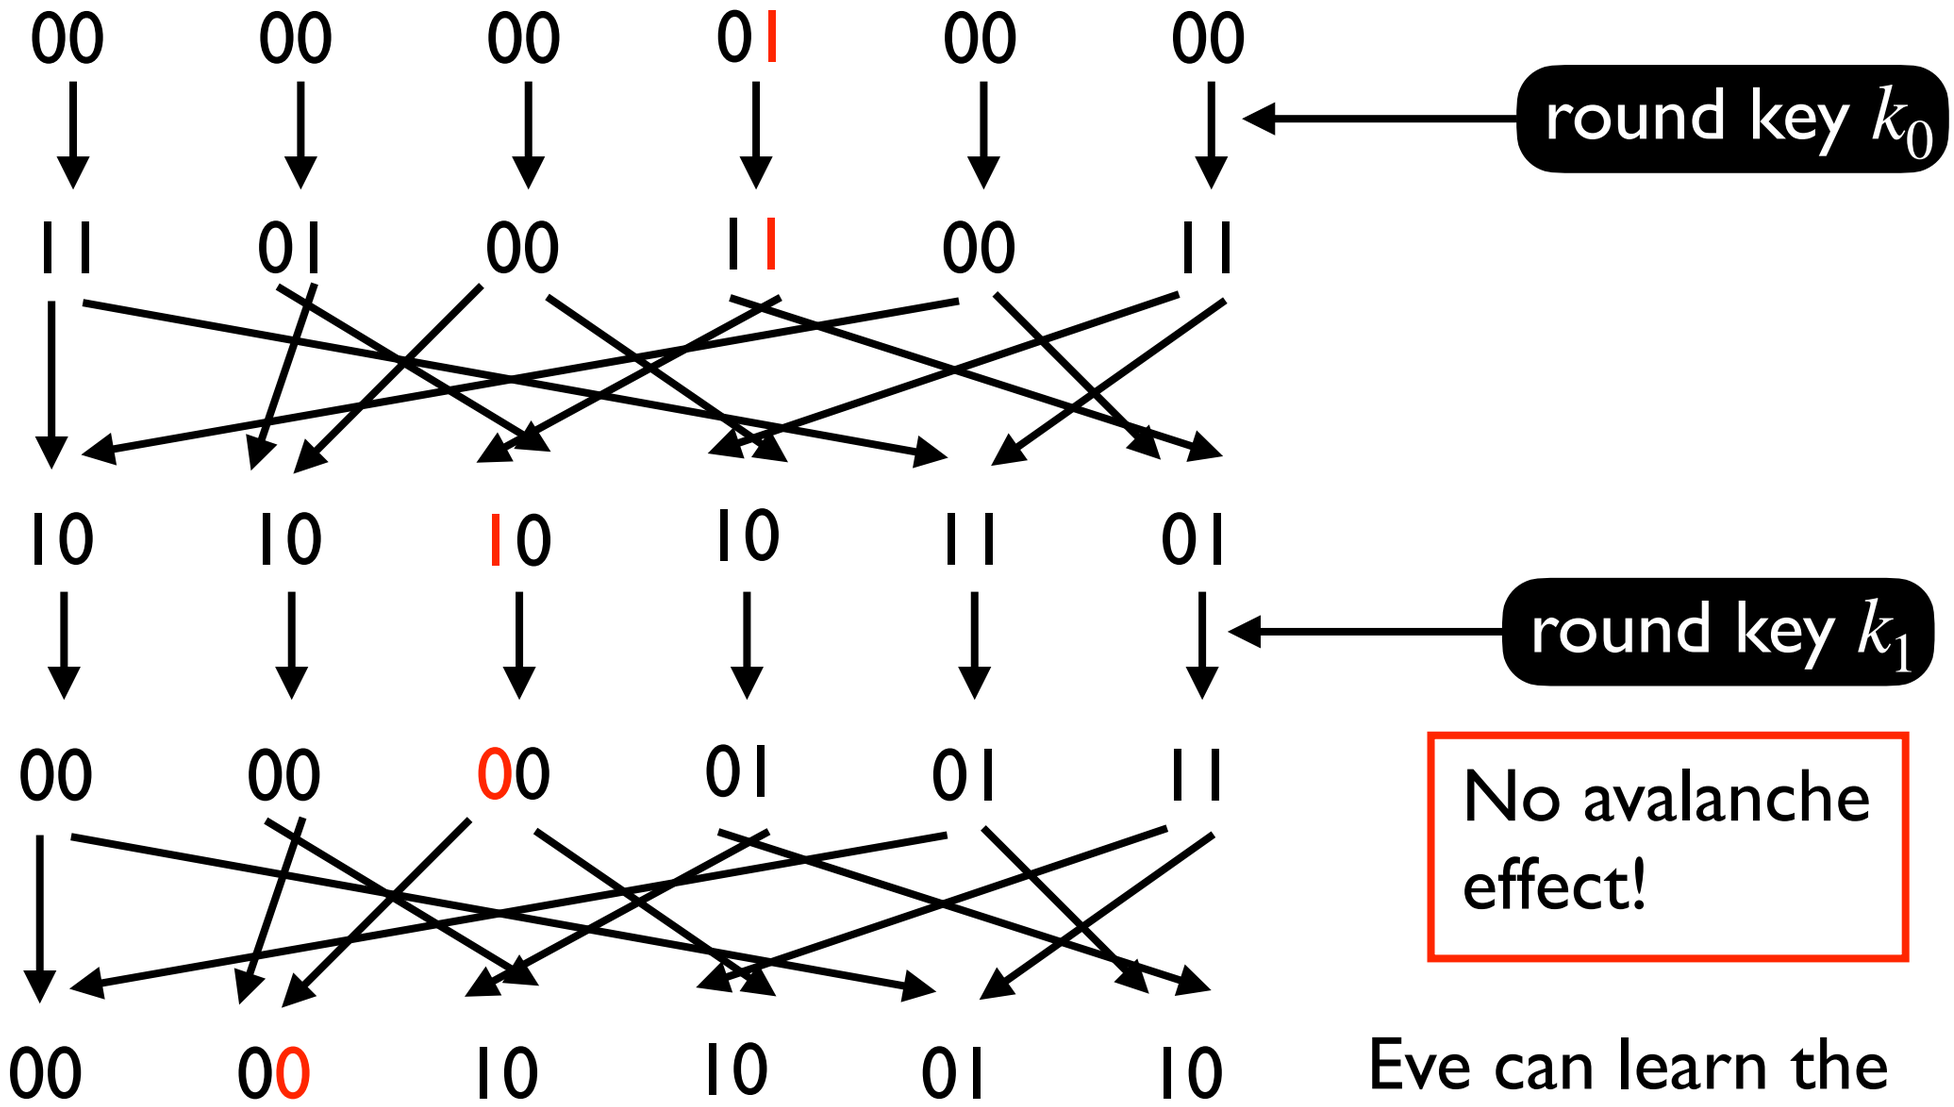

No avalanche effect!

Eve can learn the XOR of the bits of key corresponding to this path.

This class is being recorded

AES Overview

The AES permutation takes a 128-bit input represented as 4 x 4 matrix of bytes:

AES is basically a substitution-permutation network.

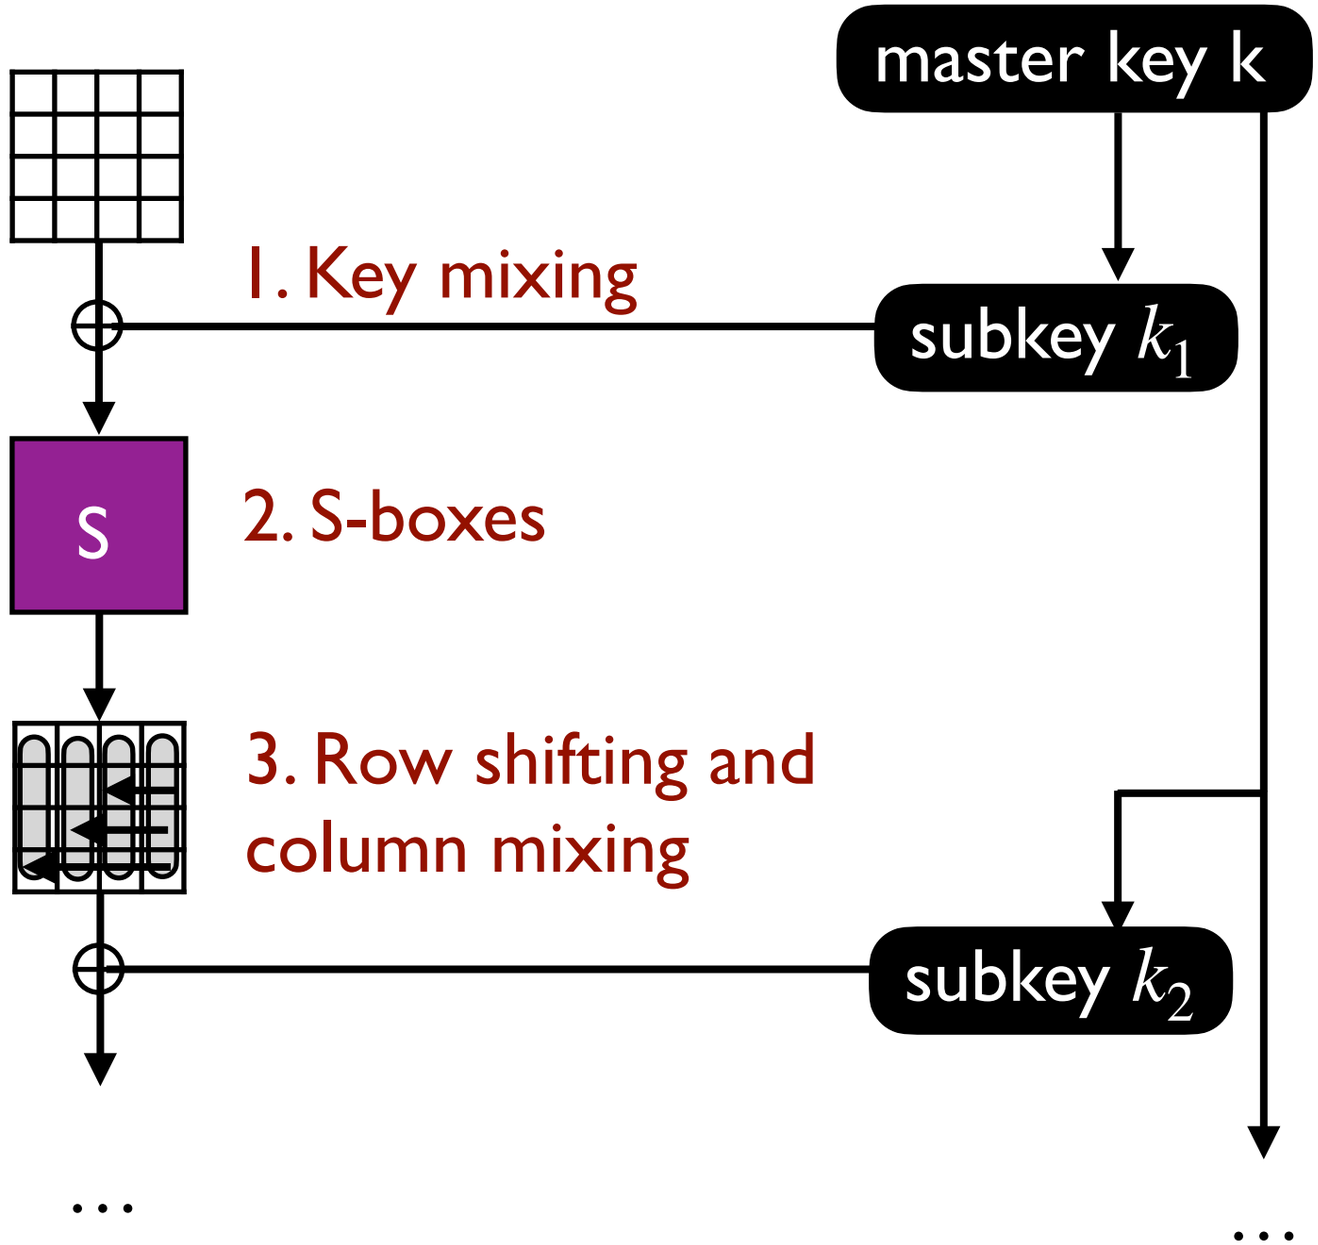

This class is being recorded

AES Key Schedule and Key Mixing

AES is standardized for 3 key lengths: 128 bits, 192 bits, and 256 bits.

Each subkey is 128 bits and is derived from the master key by more complex transformations than in DES. In particular, the later subkeys are derived using the AES S-box.

The number of rounds also depends on the key length (longer key = more rounds = more secure):

- 128-bit key: 10 rounds
- 192-bit key: 12 rounds
- 256-bit key: 14 rounds

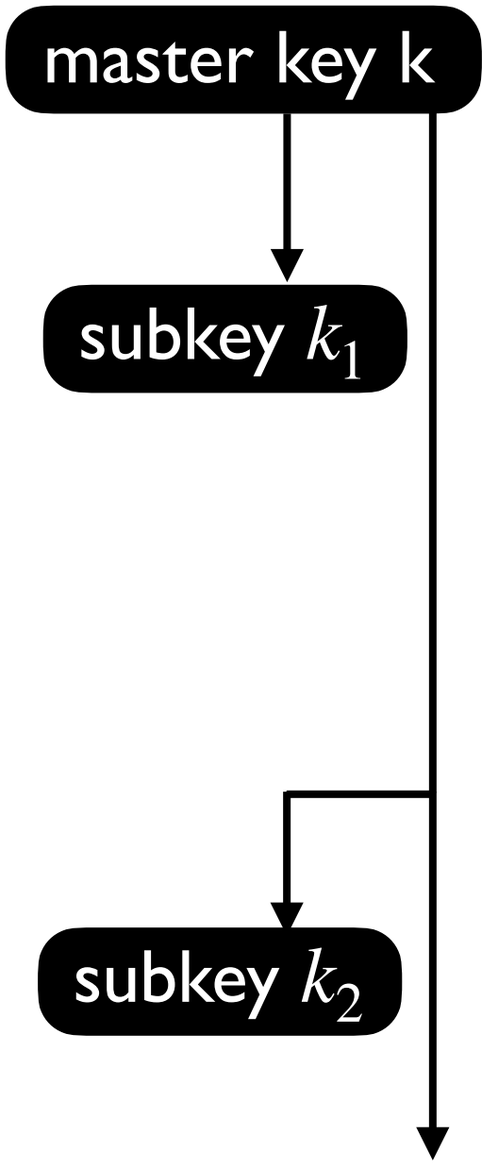

The 128-bit subkey is then XORed with the 128 bits of the state at the key mixing stages.

This class is being recorded

Applying AES S-Boxes

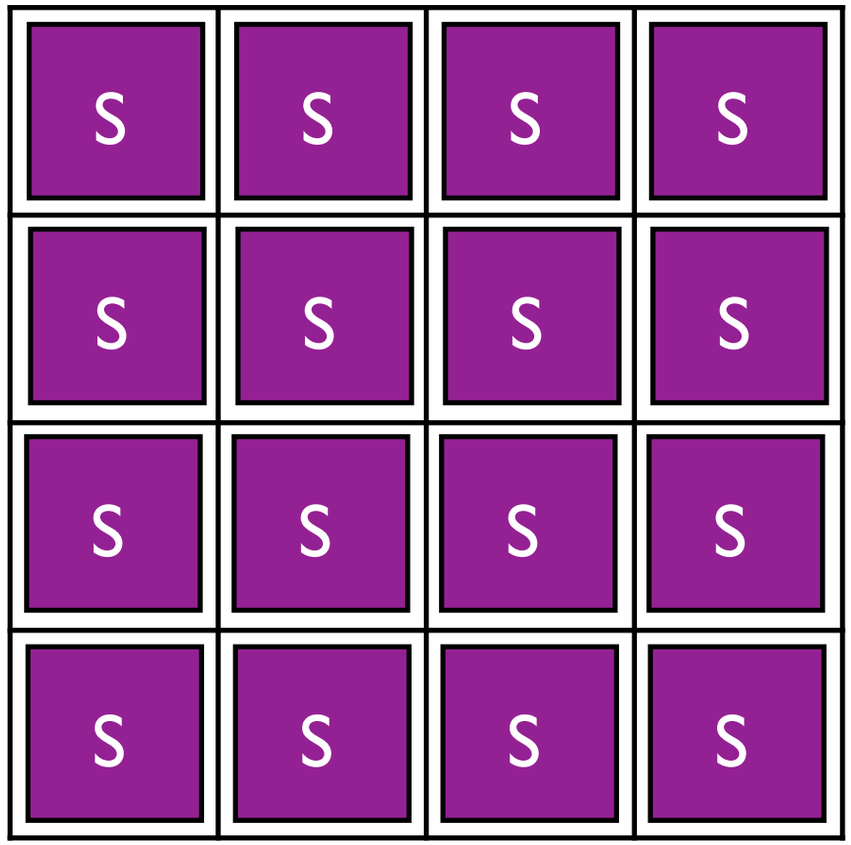

The AES S-box takes a 1-byte (8 bit) input and 1-byte output. It is invertible and again invoked via a table lookup.

The same S-box is applied to each entry of the matrix.

Again, the S-box is chosen to introduce confusion by magnifying small changes in the input.

This class is being recorded

Shifting Rows in AES

The diffusion step in AES consists of two pieces. In the first part, the rows are shifted independently.

In the ShiftRows step, the ith row is shifted cyclically by i spaces.

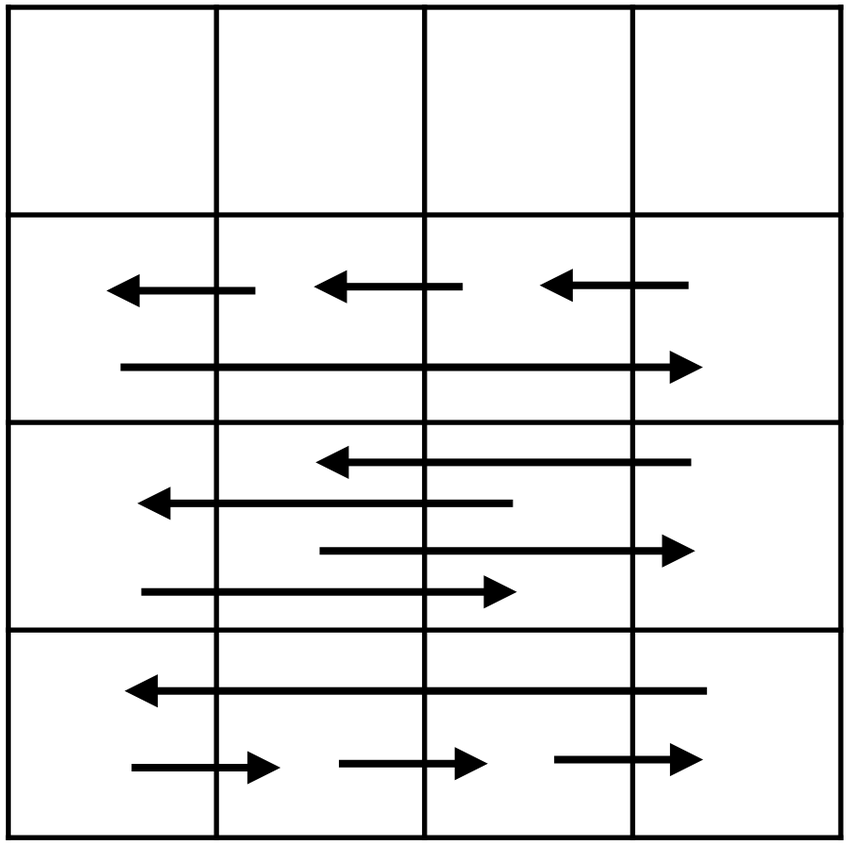

This class is being recorded

Column Mixing in AES

The MixColumns deviates from a substitution-permutation network by using a linear transformation on each column instead of permuting the bits. But it has a similar effect. (Linear transformations are also easy to invert, like permutations.)

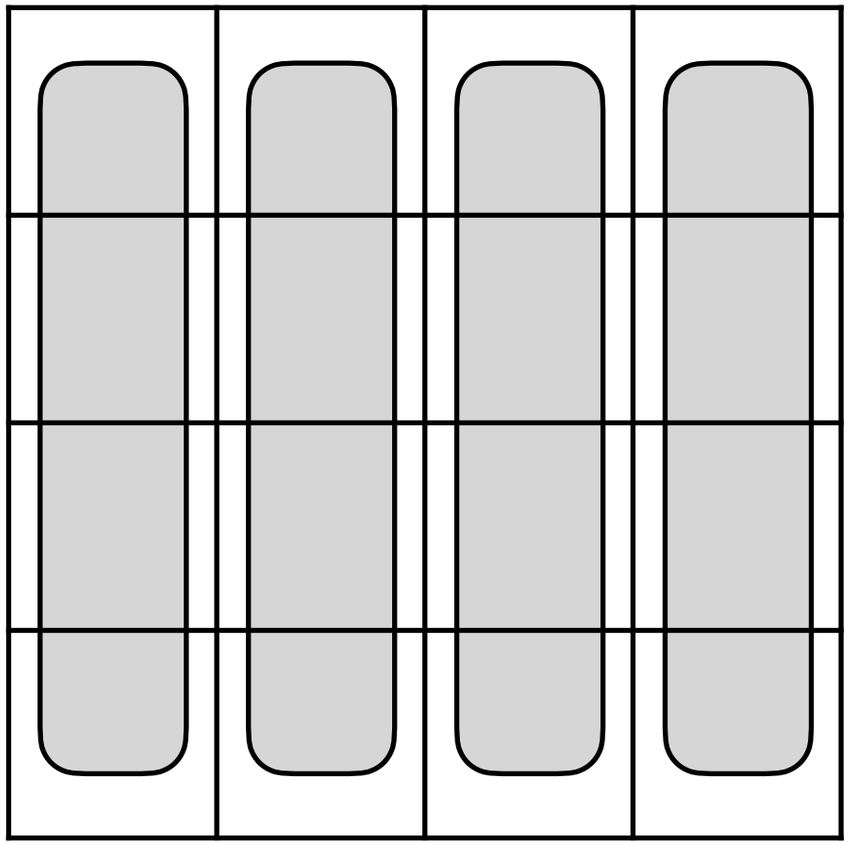

There would not be much point in ending with ShiftRows and MixColumns steps, since by themselves they can be easily inverted by Eve. Instead, the last round replaces MixColumns with another key mixing step (requiring an extra subkey).

This class is being recorded

Breaking AES

Vote: Is there a known way to break AES? (Yes/No/Other)

This class is being recorded

Breaking AES

Vote: Is there a known way to break AES? (Yes/No/Other)

All three answers are correct in a sense.

There is no (publicly) known practical attack against AES in its ideal implementation.

But … if it is not implemented properly, the encryption algorithm can sometimes leak additional information that can help narrow down the key.

This is known as a side-channel attack.

This seems like cheating — but yes, Eve cheats. If she can't win playing by your rules, she will try to change the rules.

This class is being recorded

Side Channel Attacks

There are a wide variety of known side-channel attacks. E.g.:

Timing Attacks: Some computations take longer than others. If we're not *very* careful, different keys or messages will lead to quicker or slower computations, and Eve can detect that.

Power Analysis Attacks: Similarly, some computations may draw more power (or produce more heat) than others. This can also be used by Eve to narrow down the key or message.

Cache Attacks: By monitoring cache use, an attacker may be able to determine information about what computation the encryption process is using.

Electromagnetic or Acoustic Attacks: EM radiation or sounds may leak from the computer, revealing some information about the encryption process.

This class is being recorded

How to Access a Side Channel

In some cases, it is possible to perform a side channel attack over a network. For instance, in a timing attack, Eve can interact with a server (or monitor Bob's interaction with the server) and measure the length of time it takes for the server to respond to each query. This reveals something about the encryption time.

In other cases, it may be necessary for Eve to monitor the physical vicinity of Alice. For instance, in an electromagnetic attack, Eve needs some way of seeing the leaked radiation.

In other cases, Eve may need a process on the same computer as Alice. For instance, Eve may manage to get some low-privilege malware on the machine which is unable to read Alice's message but can see how its own cache usage depends on the encryption. Or perhaps Eve's program is simply running on the same cloud server as Alice.

This class is being recorded

Side Channel Attacks on AES

AES's S-boxes and column mixing operations are defined via operations on finite fields. These can be computed, but to optimize speed, they are usually pre-computed into lookup tables, which are cached during an encryption.

This enables a cache side-channel attack.

These attacks seem doable in practice given a poor implementation and a process running on the same device (e.g., same cloud server) as Alice.

But there are also countermeasures, so in most cases, AES should be considered secure.

This class is being recorded

How Do We Use Symmetric Crypto?

The symmetric cryptosystems we have seen so far require Alice and Bob to share a key unknown to Eve. This is fine if they occasionally meet in person and communicate regularly, but what if Alice and Bob have not met before?

How can they establish their first key remotely?

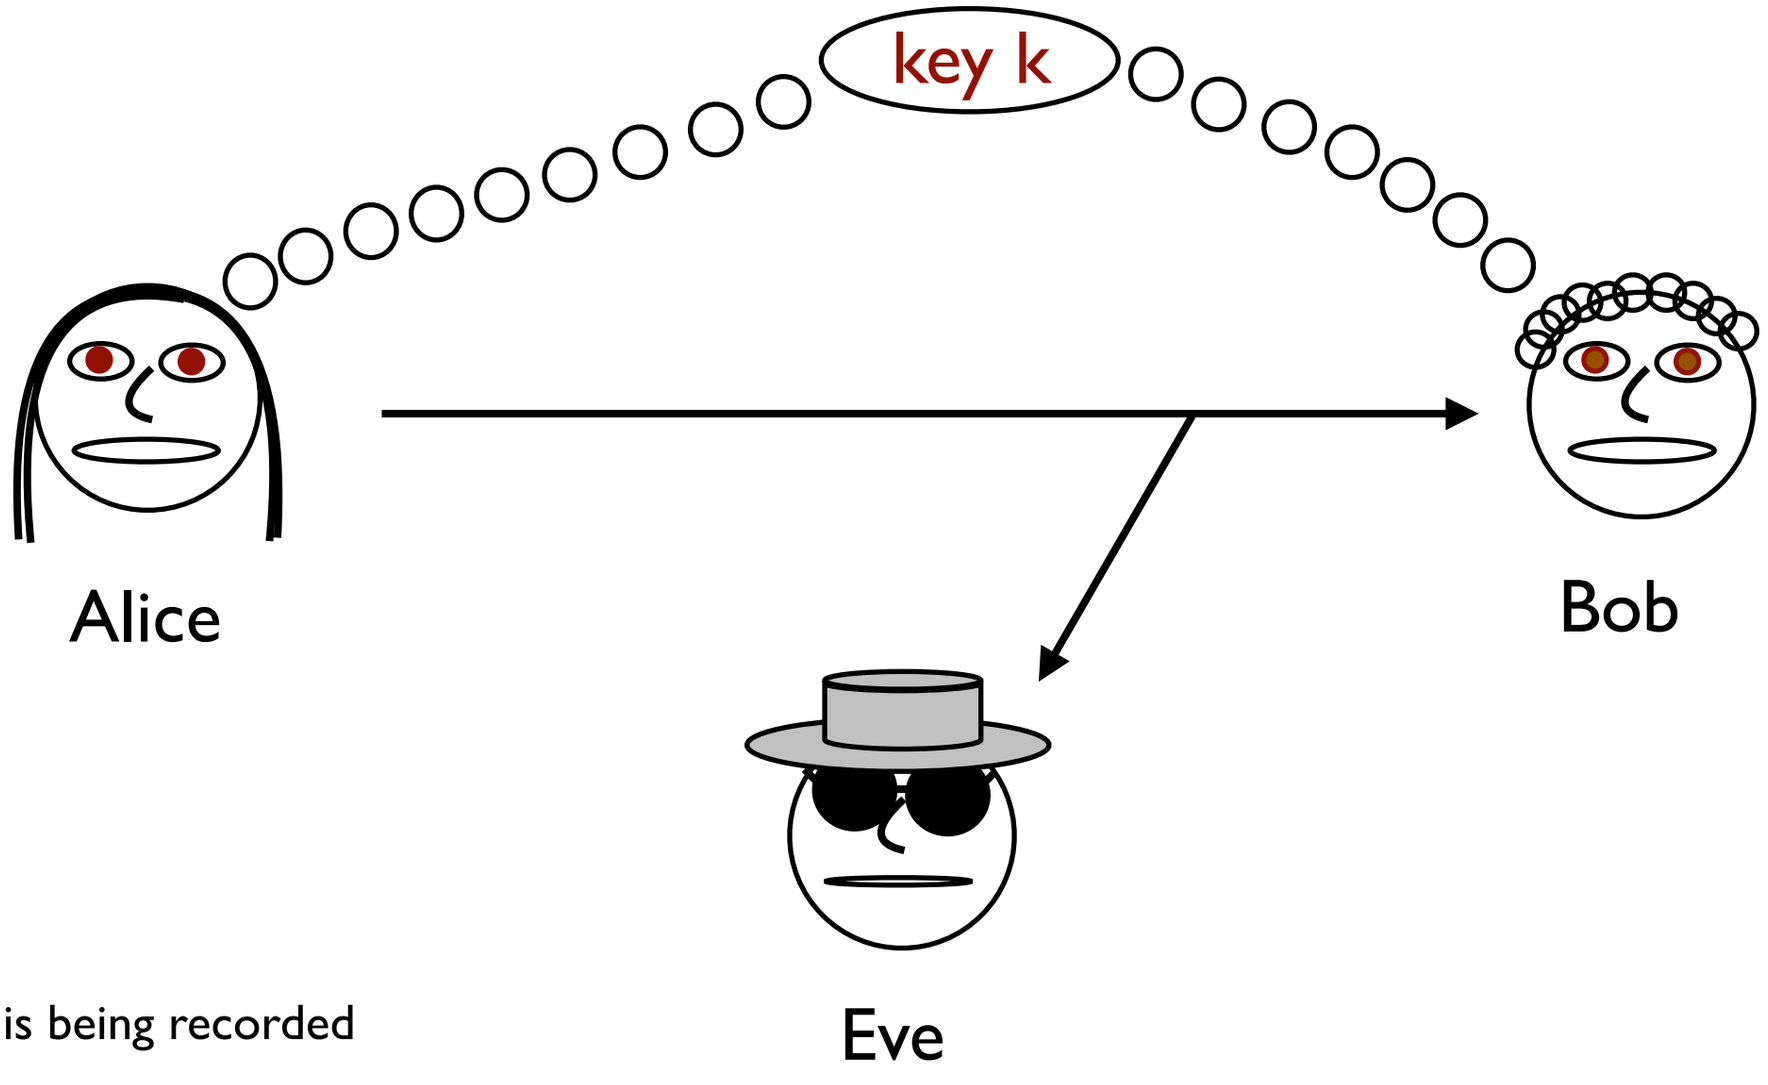

This class is being recorded

Public Key Encryption

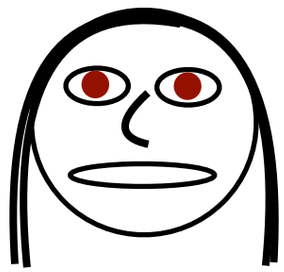

Alice

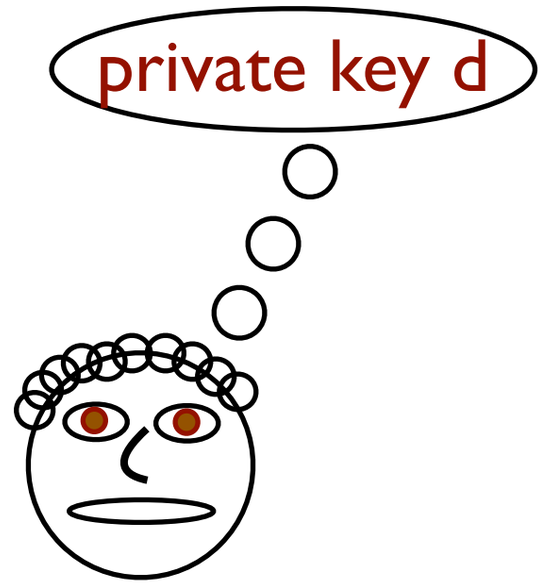

Bob

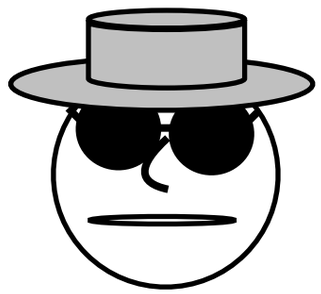

Eve

Public-key encryption is an asymmetric protocol.

This class is being recorded

Public Key Encryption

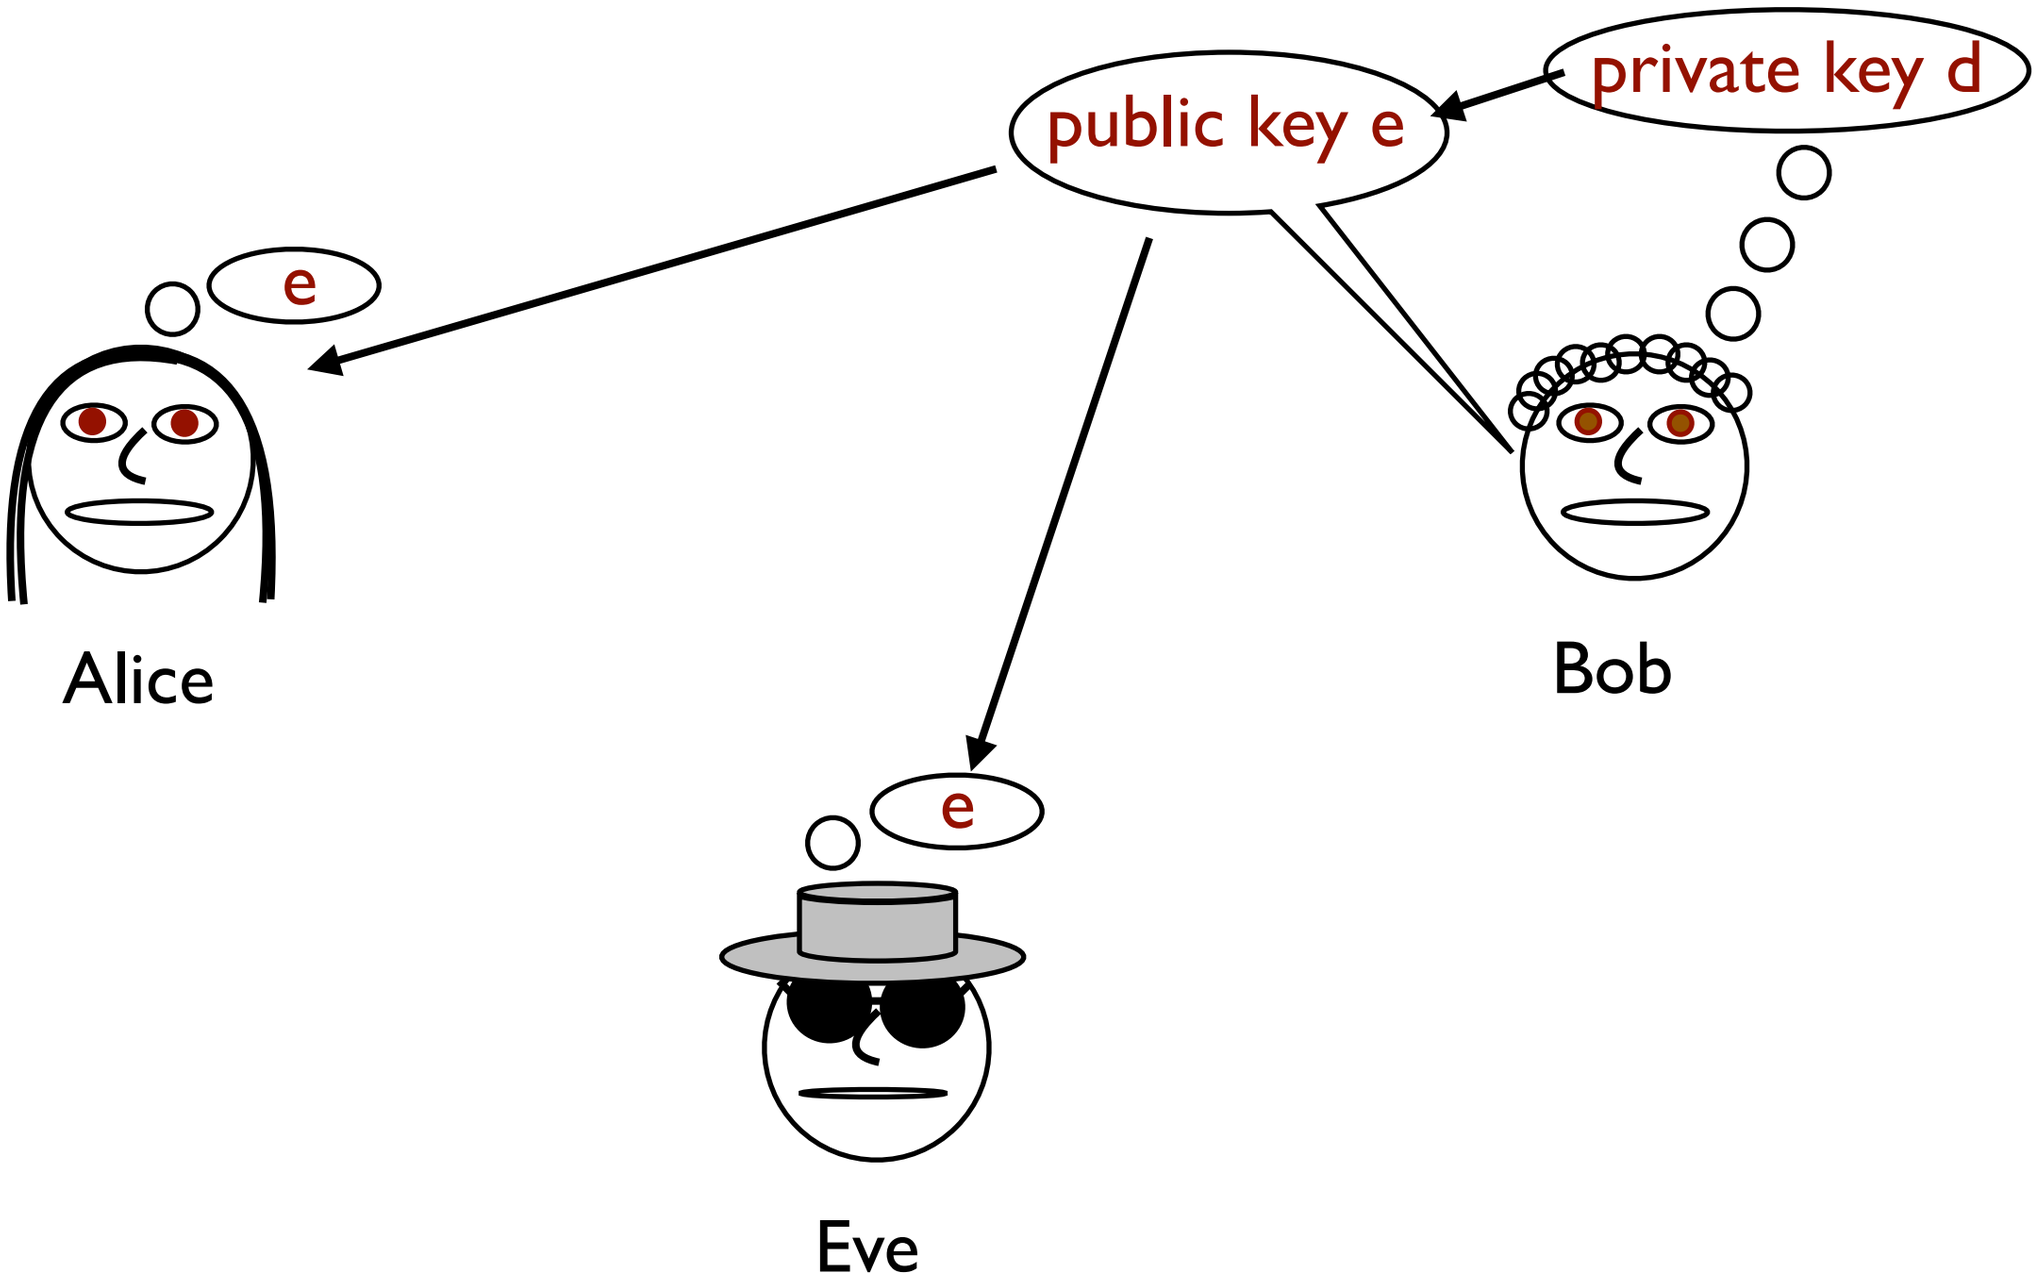

Public-key encryption is an asymmetric protocol.

This class is being recorded

Public Key Encryption

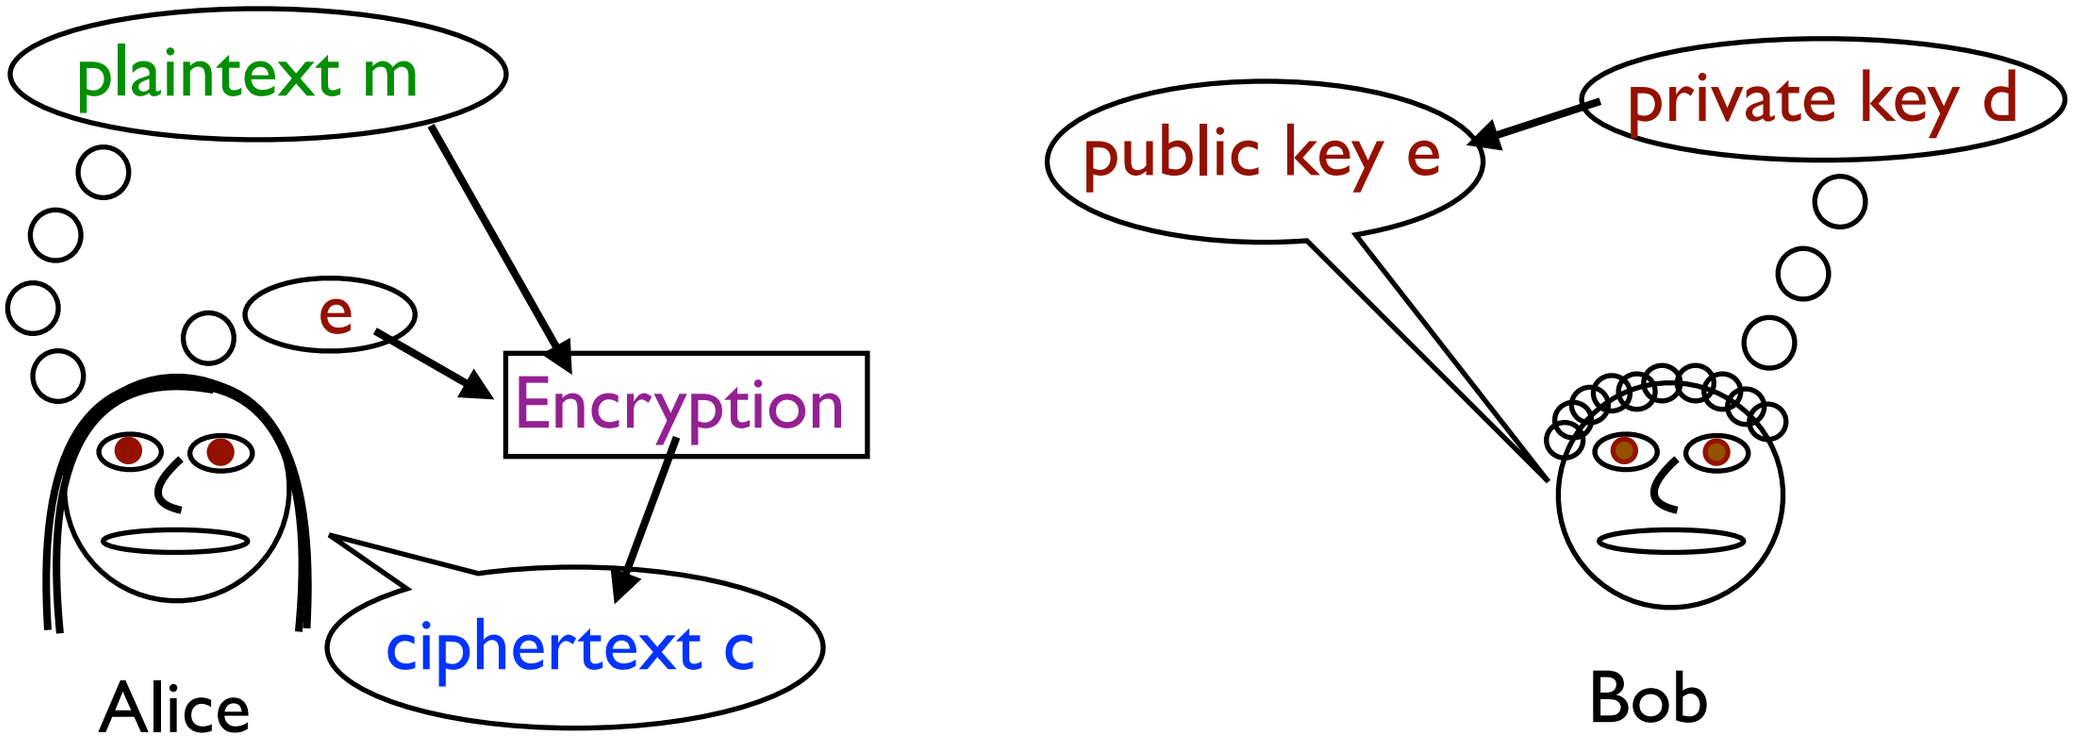

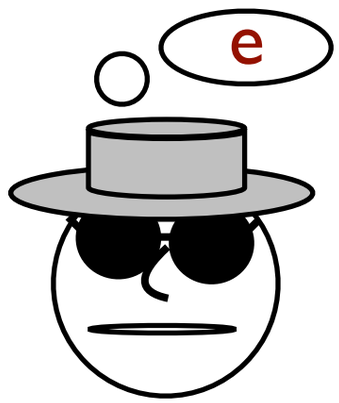

Eve

Public-key encryption is an asymmetric protocol.

This class is being recorded

Public Key Encryption

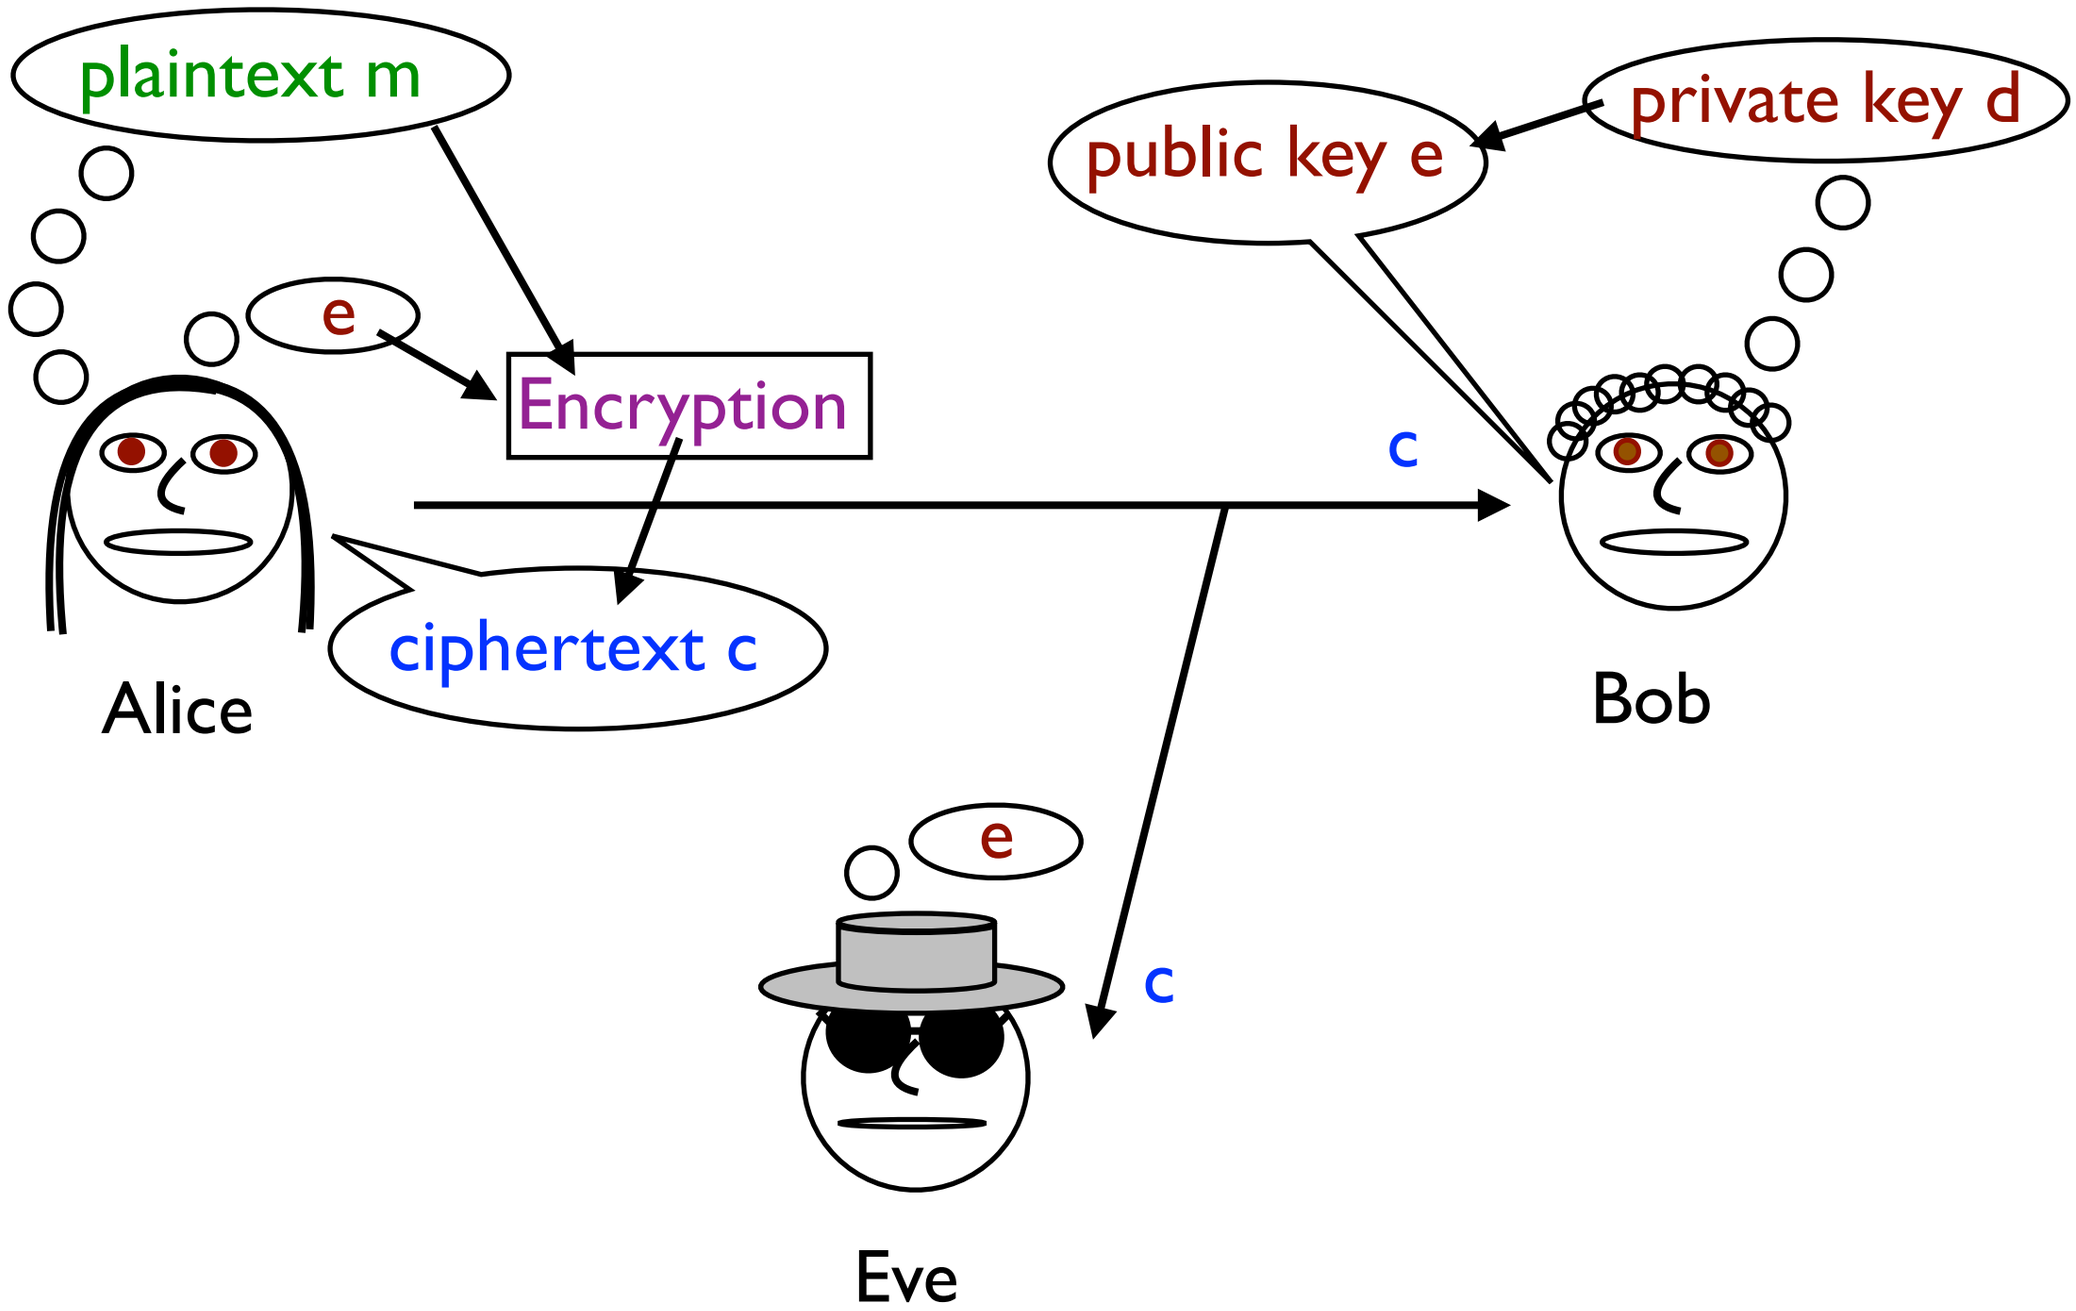

Public-key encryption is an asymmetric protocol.

This class is being recorded

Public Key Encryption

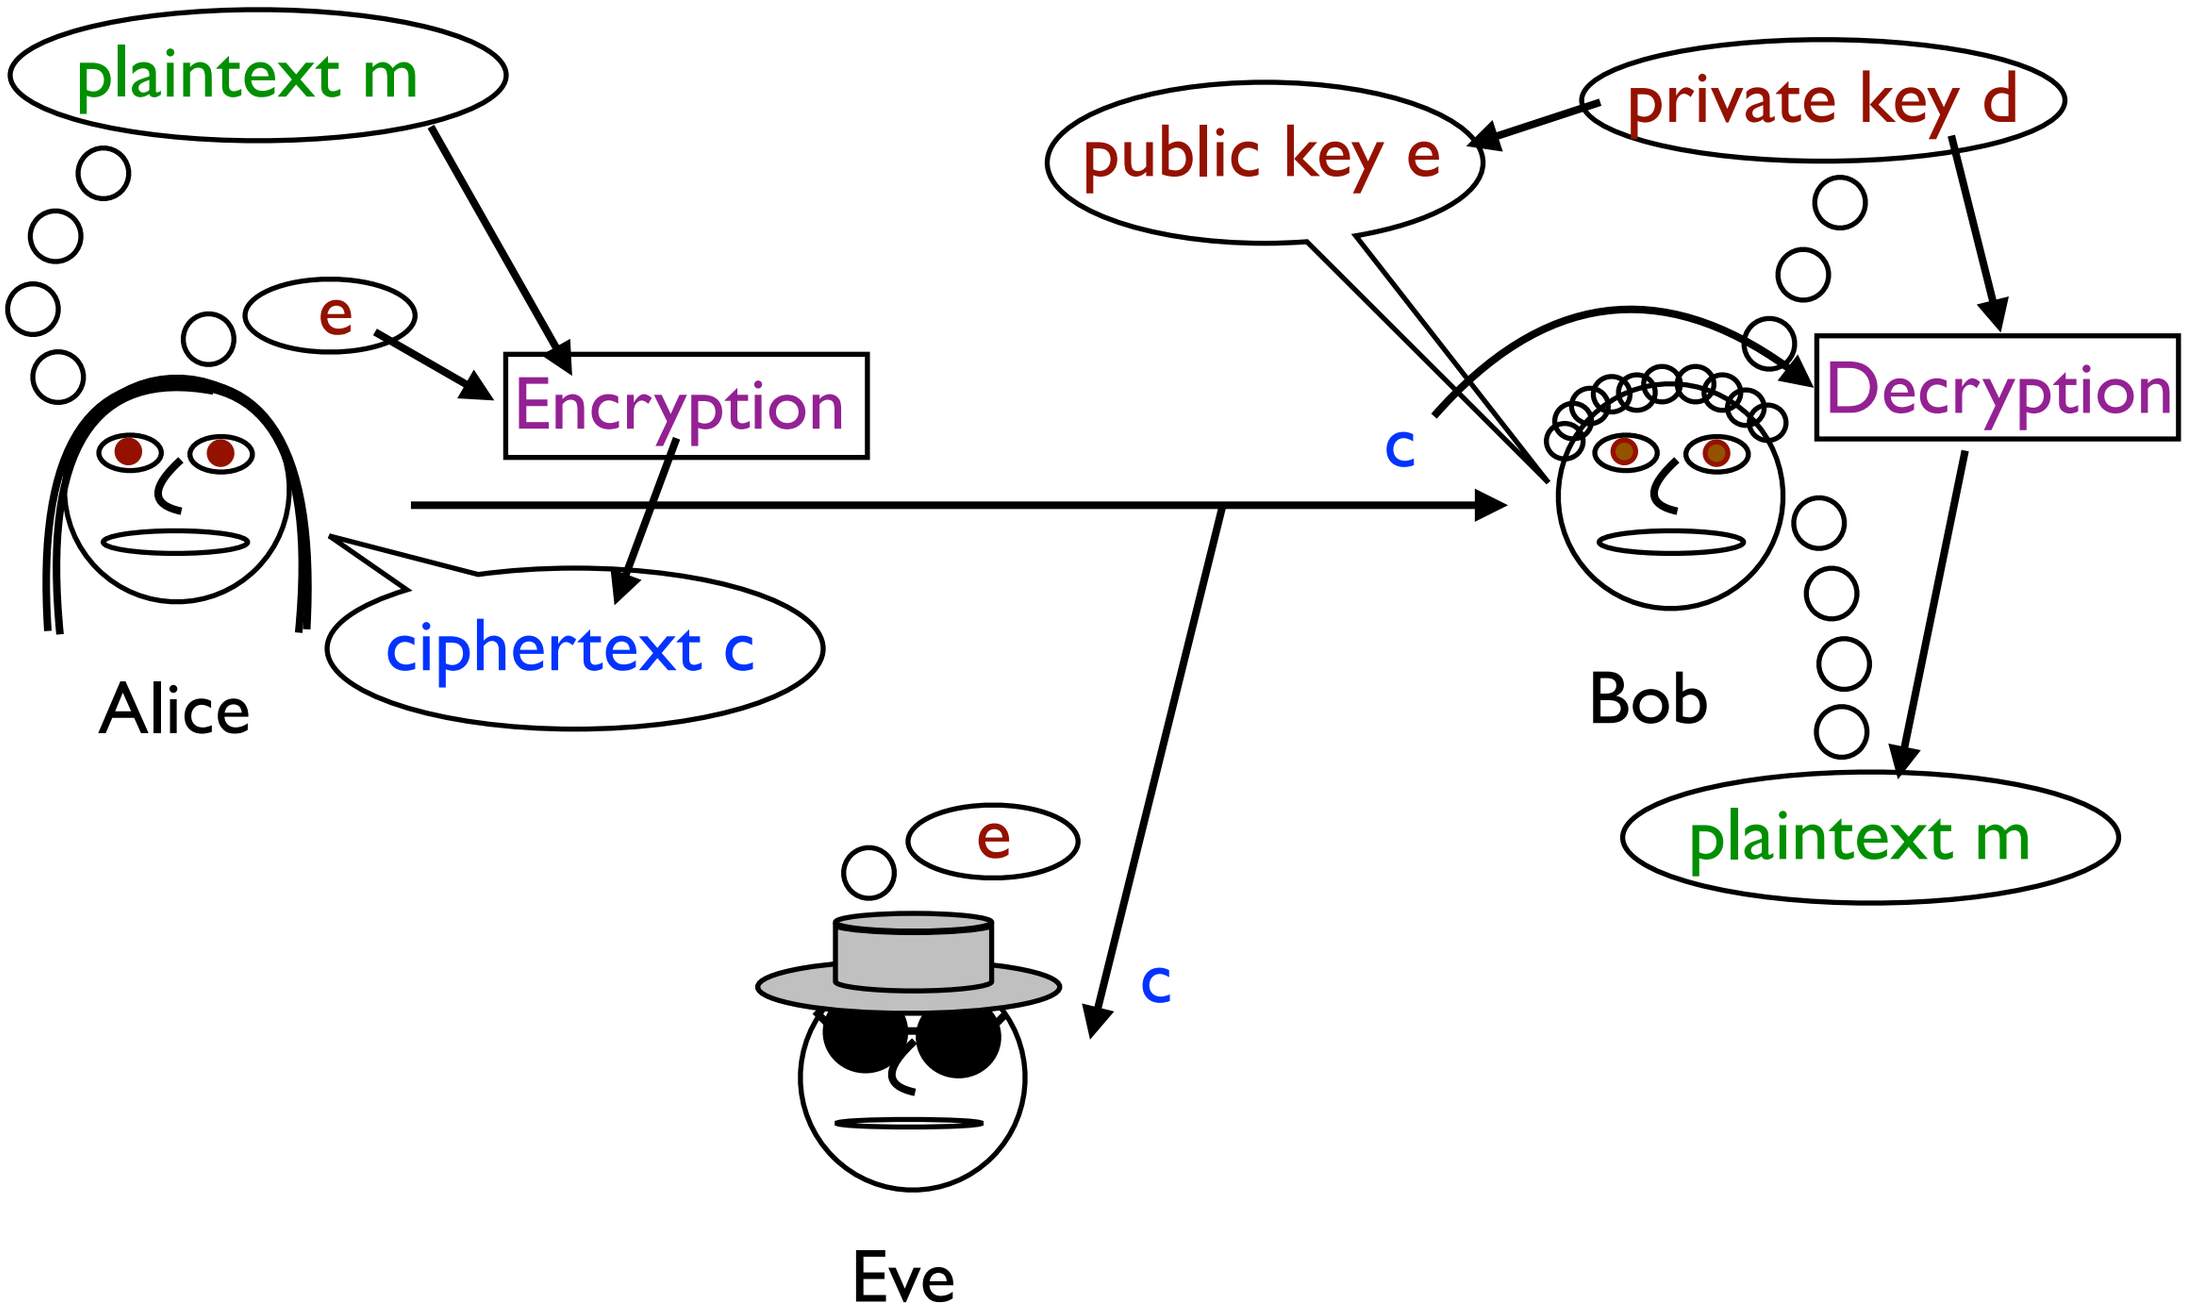

Public-key encryption is an asymmetric protocol.

This class is being recorded

Public Key Encryption

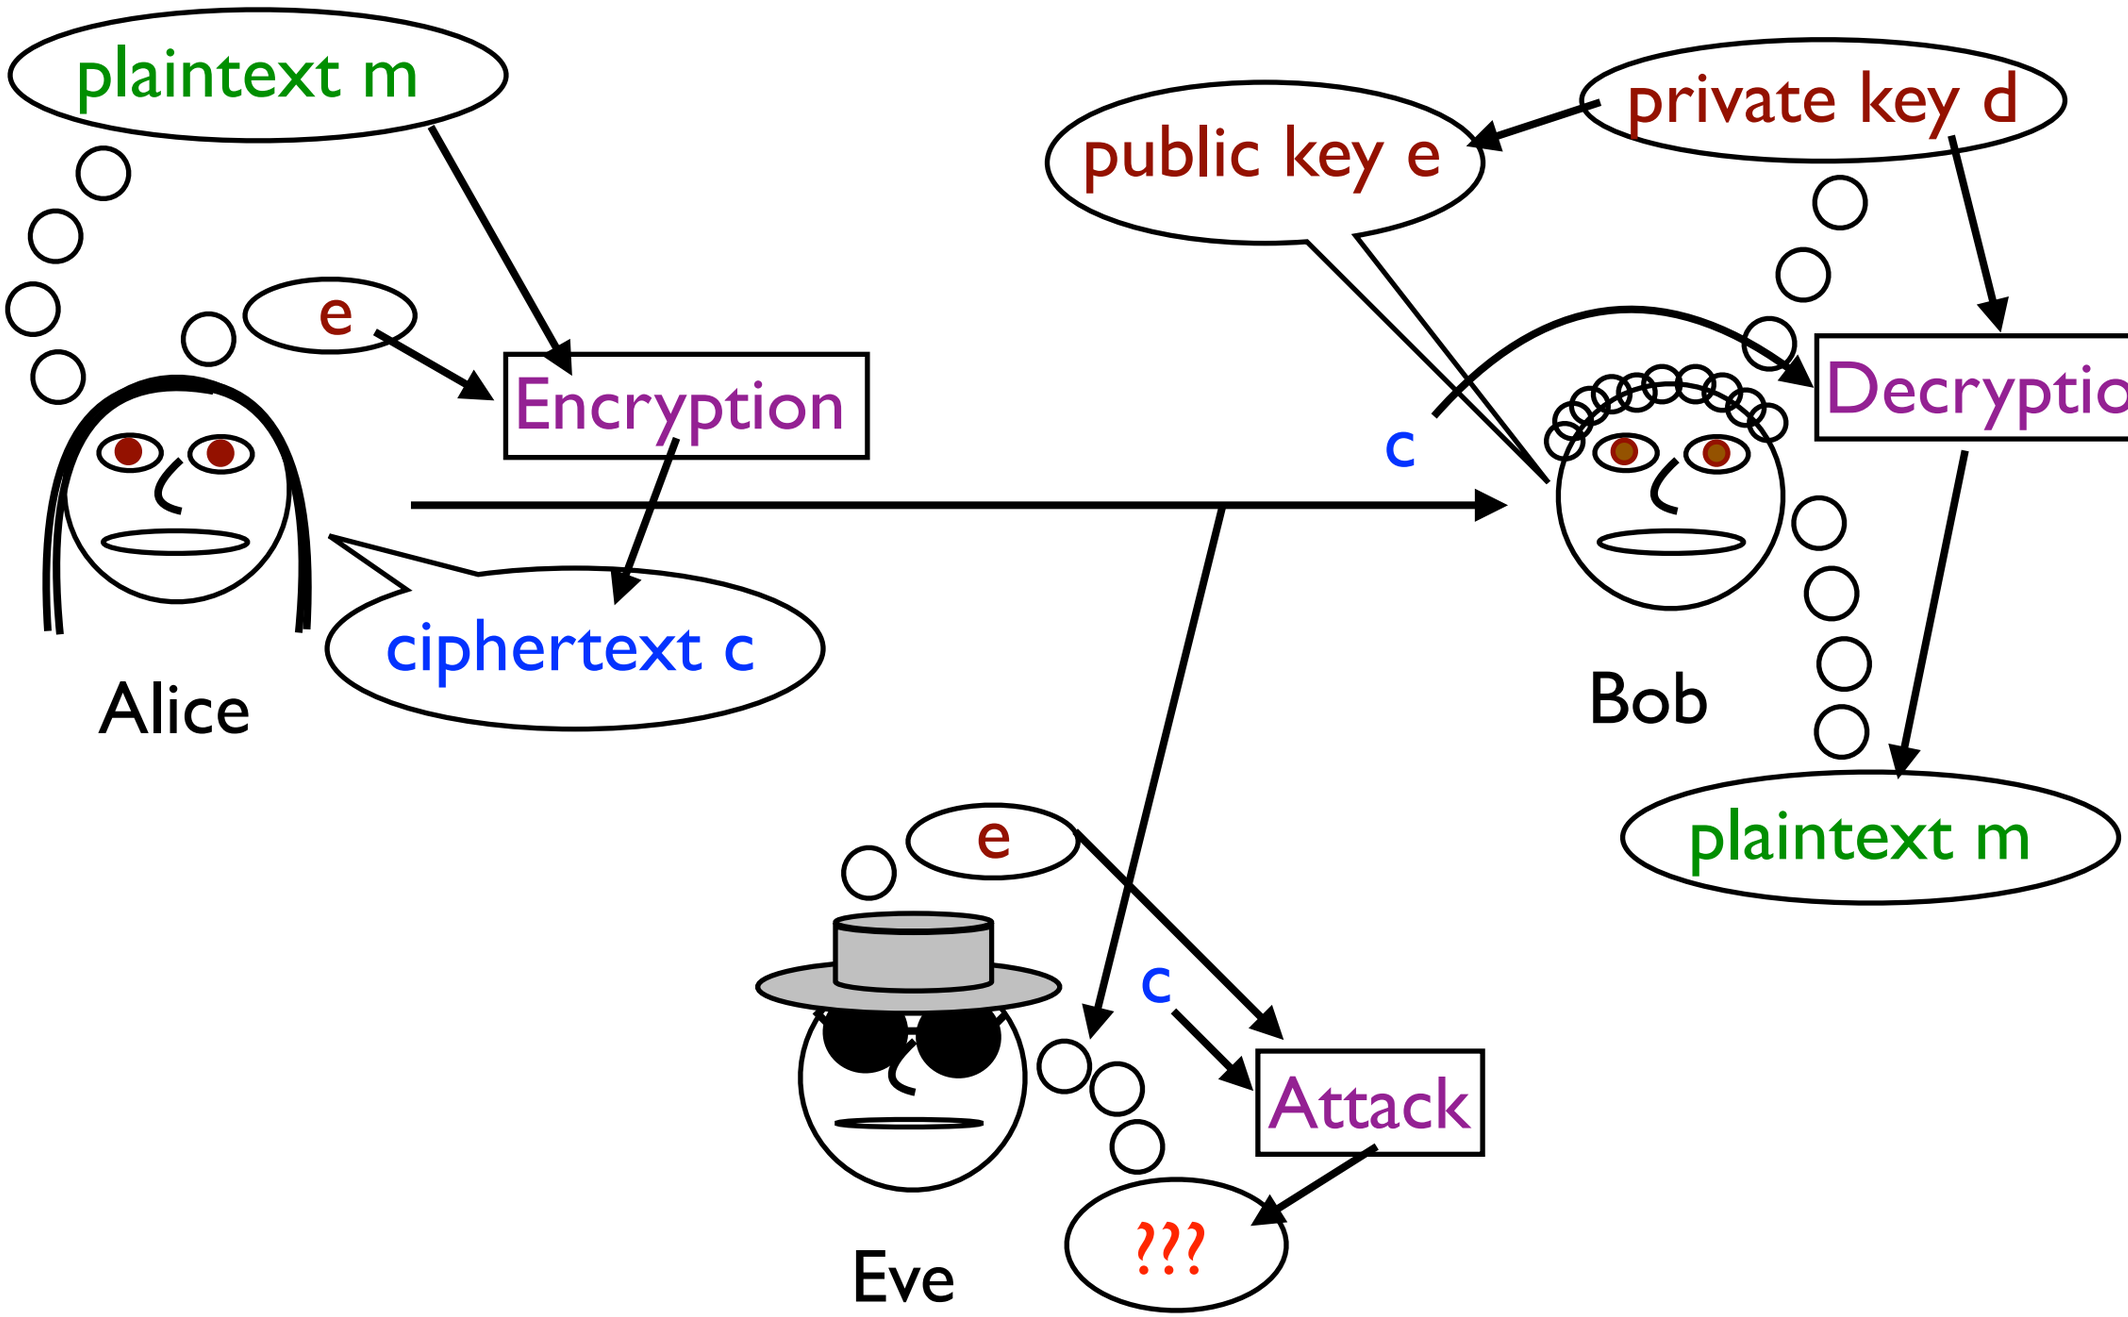

Public-key encryption is an asymmetric protocol.

This class is being recorded

Public Key Encryption

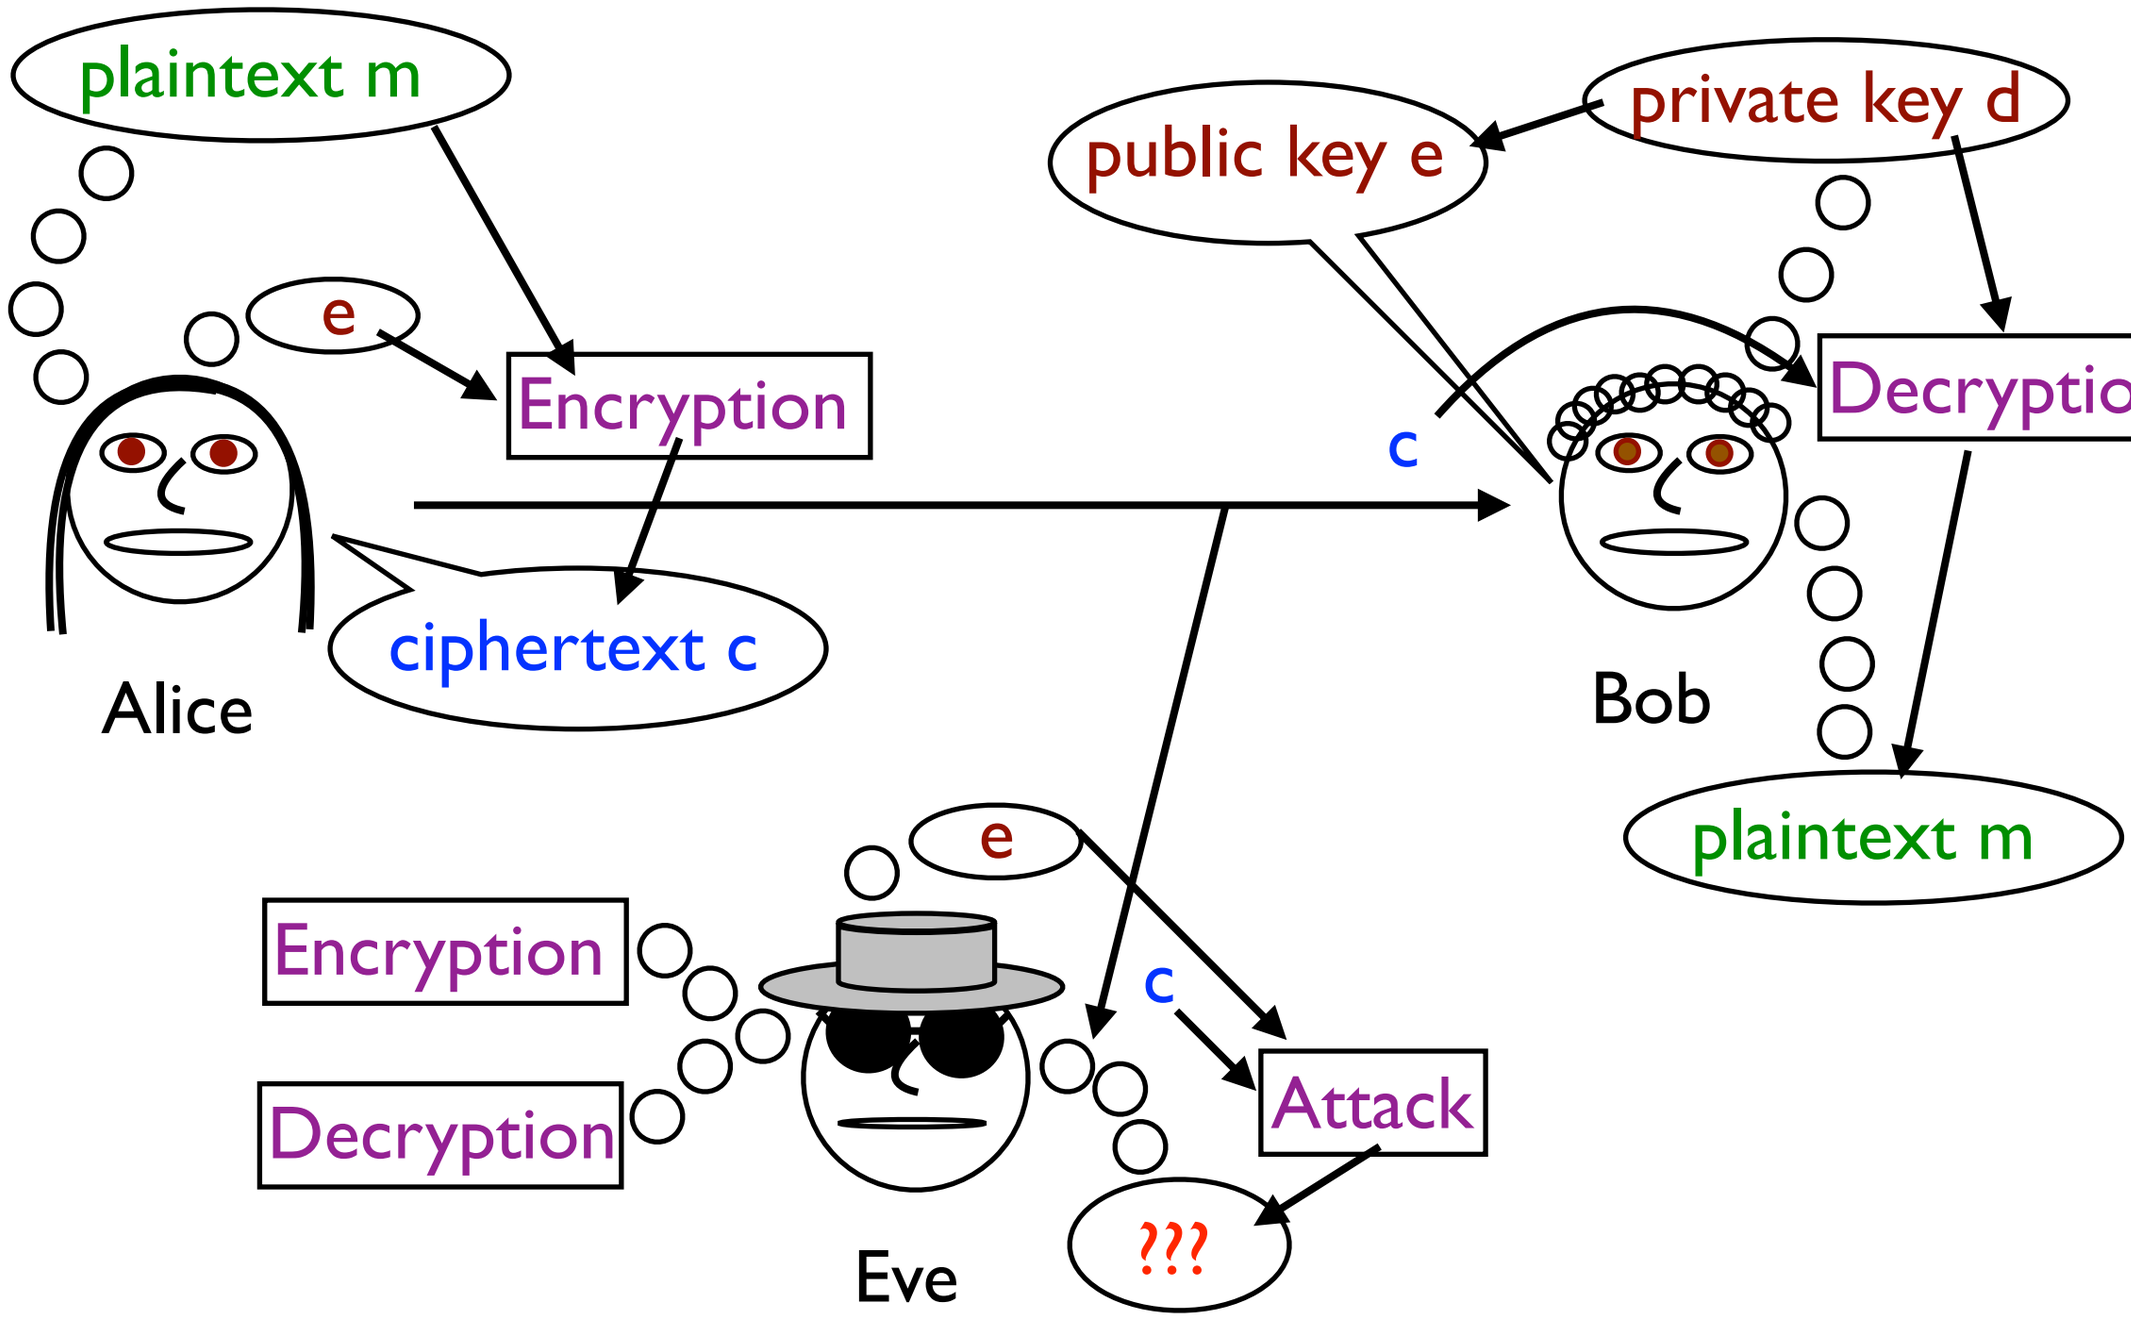

Public-key encryption is an asymmetric protocol.

This class is being recorded