



Computer Security and Privacy (COM-301)


Introduction to applied cryptography

Plan of the Day

- Module 1 – Introduction
- Module 2 – Adversaries in Cryptography
- Module 3 – Perfect Secrecy and One-Time-Pad
- Module 4 – From OTP to Stream Ciphers
- Module 5 – Public Key Cryptography

Module 1: Introduction

Why cryptography matters?



Data in transit



Data at rest

- **Data in Transit:** Securing communication (e.g., HTTPS, VPNs).
- **Data at Rest:** Securing stored information (e.g., encrypted backups).

What can we do with cryptography

Cryptography is also the building block for:

- **Integrity:** Ensuring data hasn't been tampered with (Hashing/Signatures).
- **Authentication:** Verifying identity (Protocols).
- But also **Anonymity**, etc.

In this first lecture we will focus on the basics of using cryptography for confidentiality.

The Core Problem For Today

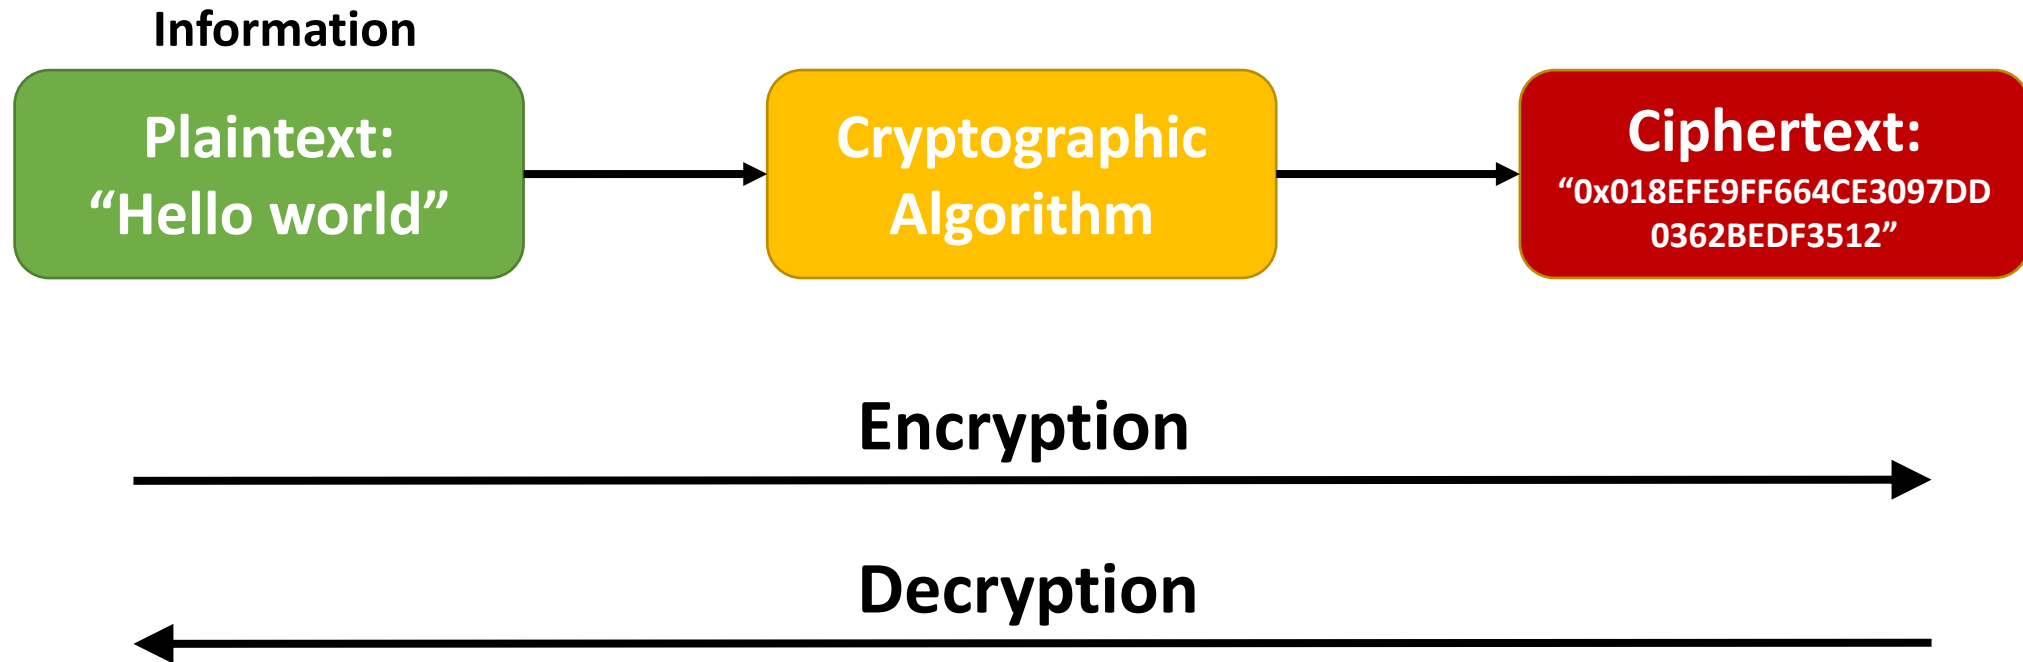
Secure Communication over an Insecure Channel

We need a way for Alice to send a message to Bob, so that Eve (the adversary) cannot read it.

Note: Confidentiality for data at rest can be seen as a particular case, where Alice = Bob.

The origins of cryptography: the quest for confidentiality

Confidentiality: information cannot be accessed by unauthorized parties



Encryption cannot be reversed without a **KEY**

Cryptography as Functions

Encryption/Decryption

- **Plaintext (M):** Original message.
- **Key (K):** Secret value.
- **Ciphertext (C):** Scrambled message.

$$C = E_K(M) \quad \text{and} \quad M = D_K(C)$$

Cryptographic algorithms for confidentiality

1. Get a key shared between Alice and Bob.

2. Encrypt message $m \rightarrow \text{Enc}(k,m)$



3. Send encrypted message $\text{Enc}(k,m)$

4. Decrypt message $\text{Dec}(k,m) \rightarrow m$



The Core Requirement - Invertibility

The decryption function D_K *must* be the **invert** of E_K :

$$\forall K, M, D_K(E_K(M)) = M$$

If not, Bob cannot recover the original message M .

SECURITY REQUIREMENT: To provide security, these functions must also be **hard to invert** *without* knowing the key K .


Hard to Invert: The Key Space

- **Concrete Meaning:** To find the plaintext (M), in an ideal cryptosystem, the only viable attack for the adversary (Eve) must be to **try every single possible secret key K** (a **brute-force attack**).
- **Goal:** A secure scheme forces the attacker into this situation.


Example 1: Caesar's Cipher

- **Encryption:** Shift each letter by a fixed number (K).
- **Decryption:** Shift each letter by a fixed number ($-K$).
- **Key Space:** Only 25 possible keys.

Plaintext: ATTACK AT DAWN $\xrightarrow{K=3}$ Ciphertext: DWWDFN DW GDZQ



Encrypt



Decrypt

Security Insight: $\log_2(25) \approx 4.6$ bits of security. This is **far too small** for real-world use.

Cryptanalysis: The Flaw

- Eve knows that Alice wrote a long-enough message for Bob in English.
- She knows that English most frequent letter is 'e'.
- She can then simply identify the most frequent letter in the ciphertext: it likely corresponds to the letter 'e'.
- She can then directly recover the key (a number between 0 and 25). This is a **frequency analysis attack**.

More than 25 Keys: The Substitution Cipher

- Alternative encryption/decryption strategy: **Permuting the Alphabet**
- **Operation:** Each letter is mapped to a unique, different letter, defined by a **permutation** of the 26 letters.

Key: ABCDEFGHIJKLMNOPQRSTUVWXYZ
HOWBUGIACRYEVZXPJQMSNTFDKL

Encrypt/Decrypt: substitute by opposite letter

hello world → au**e**ex**f**xq**e**b

- **Key Space:** 26! (26 factorial) possible keys.

From 25 Keys to $\approx 4 \times 10^{26}$ Keys

- **26! (Number of Keys):** ≈ 403 septillion keys (4.03×10^{26}).
- **Bit Security:** $\log_2(26!) \approx 88.4$ bits of security?
- **Insight:** This key space might look like it is big enough, but this cipher is **still easily broken** by the same frequency analysis.

Cryptanalysis: The Flaw, Again

- The same **frequency analysis attack** is effective against the Substitution Cipher.
- **Vulnerability:** The cipher preserves the statistical properties of the plaintext language (e.g., single-letter frequencies).
- If Eve (the attacker) has a long enough ciphertext C , statistical analysis will break the cipher, rendering the full $26!$ key space **irrelevant**.

- **Goal of the adversary:** Notice that the goal of the adversary might not be necessarily to recover the full key, depending on the scenario, it might only attempt to recover part of the key (the mapping for a subset of letters).
- Sometimes the goal of the attacker is even smaller: Eve might want only to know whether Alice sent the message “Yes” or “No” to Bob.

Bits of Security Versus Key Space Size

- **Goal:** An N -bit key should offer security as close to N bits as possible (i.e., require 2^N attempts).
- **The Break:** If a 1024-bit key is broken with only 2^{40} attempts (40 bits of security), the algorithm is considered broken.
- **Caesar Lesson:** The 25-key space was cleverly reduced to **1 possibility**, similarly the ≈ 88 bits of the substitution cipher was reduced to just a few tries.

Module 2: Adversaries in Cryptography

Eve's Power: Passive Eavesdropping

In Caesar's (and the substitution) Cipher, Eve's capabilities were very limited:

- She observed a long-enough ciphertext C that corresponds to the encryption of a piece of English.
- (And some compute power to test a few keys).

This is not the only possible adversary model.

Other Security Models: Types of Adversaries

The security of a scheme is always measured relative to the power of the **adversary (Eve)**.

- **Passive Eavesdropper:** Eve can only *read* the ciphertext (the model we had until now).
- **More Active Attacker:** Eve can somehow influence the system by for example convincing Alice to send to Bob very specific messages, observing the precise time it takes Bob to run its algorithm, etc... anything that can go beyond just observing communication.

Active Model 1: Known Plaintext Attack (KPA)

- **Scenario:** Eve is given access to multiple pairs (M, C) corresponding to messages and their corresponding ciphertext, all encrypted with the secret key K .
- **Goal:** From these pairs, she is trying to guess the key K to decrypt other messages.

Why is KPA a Realistic Model?

- **Standard Headers:** Many encrypted messages contain predictable headers (e.g., “From:”, timestamps).
- **Malicious Guessing:** If Eve guesses a message (or a part of a message) and observes its encryption, she has an M/C pair.

KPA and the Substitution Cipher

- **Vulnerability:** Under a KPA, the substitution cipher becomes **even more broken**.
- **Why?** If Eve knows that $M = \text{'THE QUICK BROWN FOX'}$, she immediately recovers the mappings for multiple letters, drastically reducing the effective key space, without even needing to do a statistical analysis study!

Active Model 2: Chosen Plaintext Attack (CPA)

- **Scenario:** Eve can somehow convince Alice (or the system) to encrypt arbitrary messages (chosen by Eve) with the secret key.
- **More Abstractly:** For any chosen plaintext m , Eve can get access to $E_K(m)$.
- We often say that Eve is assumed to have access to an **Encryption Oracle**.

Why is CPA a realistic model?

We can imagine two simple scenarios.

- **Alternative Scenario:** Think of a secure messaging service (like Signal or WhatsApp) that uses the sender's key to encrypt the message. If Eve gets limited access to Alice's phone for a few minutes, she can maybe send **any** piece of data she chooses to Bob, to generate an (M, C) pair.
- **Encryption API:** Eve might interact with an encrypted API endpoint where she provides data and receives the encrypted output.

CPA and the Substitution Cipher

- **Complete Break:** The substitution cipher is **COMPLETELY broken** by a CPA.
- **Attack:** Eve chooses a plaintext that reveals the entire key: $M = \text{'abcdefghijklmnopqrstuvwxyz'}$.
- **Result:** The resulting ciphertext C is the entire key (the full permutation mapping), in just *one* query!

Practical Threat Model: Side-Channel Attacks


- **Scenario:** Eve is allowed to not only access an ideal oracle, but also measure physical characteristics of the actual implementation while it is encrypting/decrypting.
- **Timing Attack:** Eve measures **how long** it takes for a given message to get encrypted or a ciphertext to be decrypted.
- **Power Analysis:** Eve observes the **energy consumed** by the device (e.g., a credit card) doing the crypto.

Why are Side-Channel a Realistic Model?

- **Malicious Ownership:** You often own a device (like a phone, smart card, or embedded chip) that holds keys belonging to another principal (e.g., Google, your bank) that are meant to be unreadable to you.
- **Concrete Example (DRM):** The decryption key for Digital Rights Management (DRM) content (like a movie or song) is embedded in your media player or console. By measuring power or timing signals, an attacker could potentially extract that secret key.

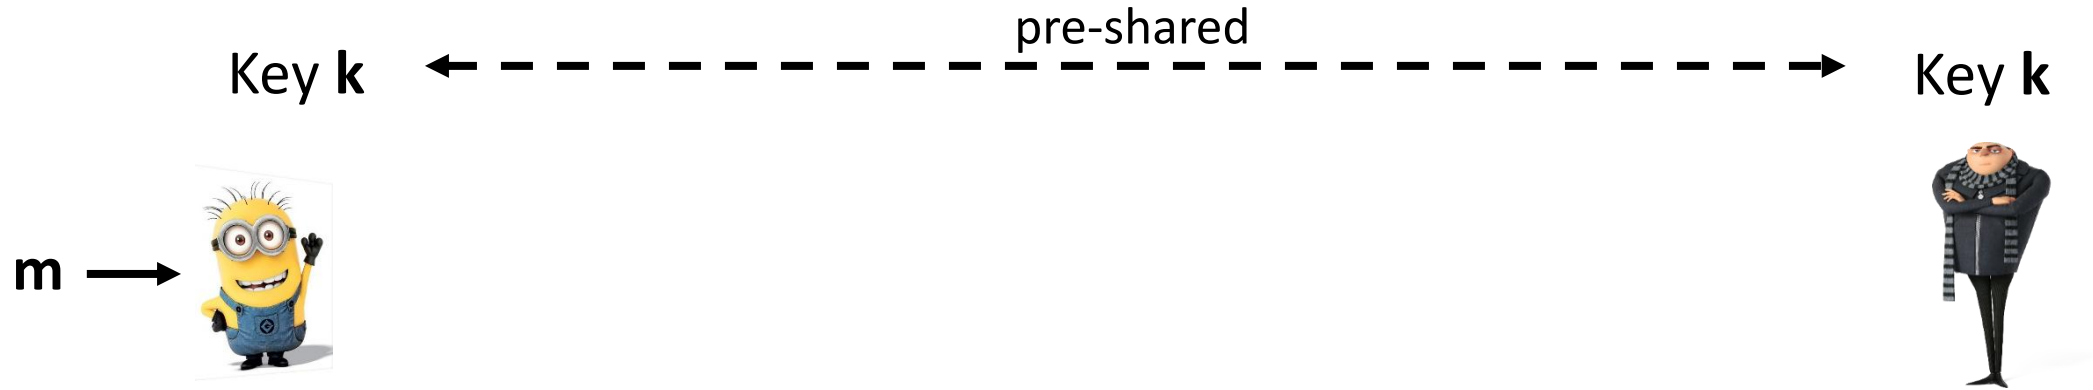
Module 3: Perfect Secrecy and One-Time-Pad

Obtaining perfect secrecy: One Time Pad (OTP)



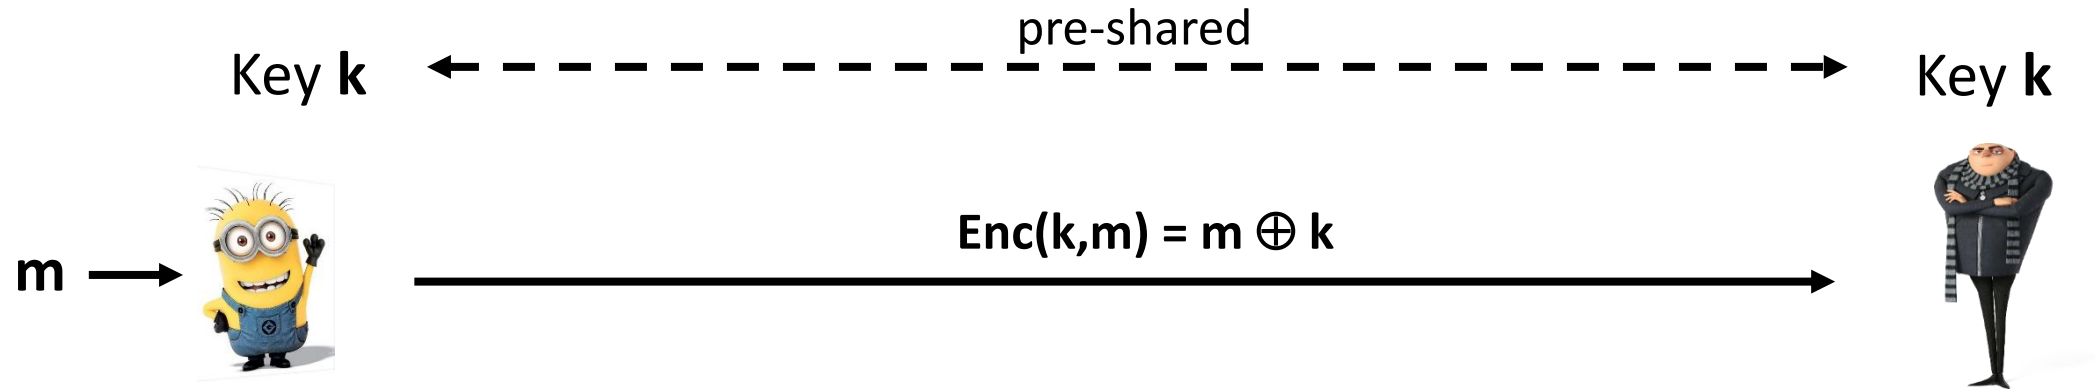
Obtaining perfect secrecy: One Time Pad (OTP)

Key = string k of random bits **as long as the message**



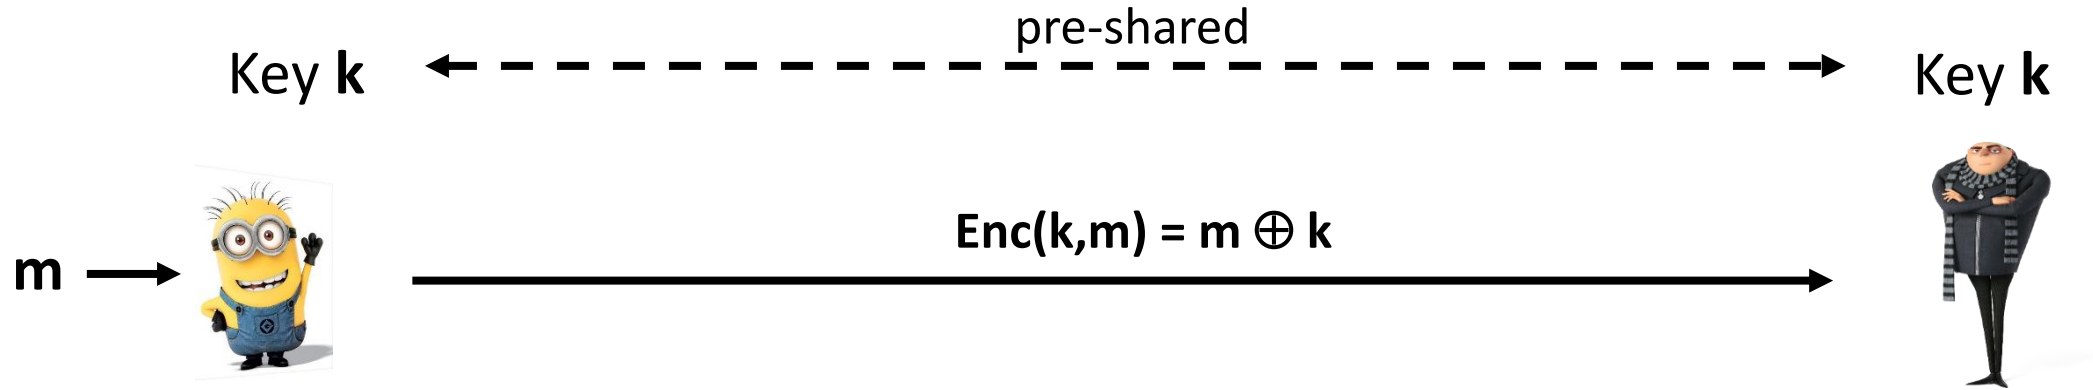
Obtaining perfect secrecy: One Time Pad (OTP)

Key = string k of random bits **as long as the message**



Obtaining perfect secrecy: One Time Pad (OTP)

Key = string k of random bits **as long as the message**




Message
Binary
OTP-Key
Encryption

```
YEAH (ASCII Hex: 59454148)
01011001010001010100000101001000
⊕ 01110101000111010100101001001010
---
00101100010110000000101100000010
```

Obtaining perfect secrecy: One Time Pad (OTP)

Key = string k of random bits **as long as the message**




Message
Binary
OTP-Key
Encryption

```
YEAH (ASCII Hex: 59454148)
01011001010001010100000101001000
⊕ 01110101000111010100101001001010
-----
00101100010110000000101100000010
⊕ 01110101000111010100101001001010
-----
01011001010001010100000101001000
```

same

Obtaining perfect secrecy: One Time Pad (OTP)

Key = string k of random bits **as long as the message**

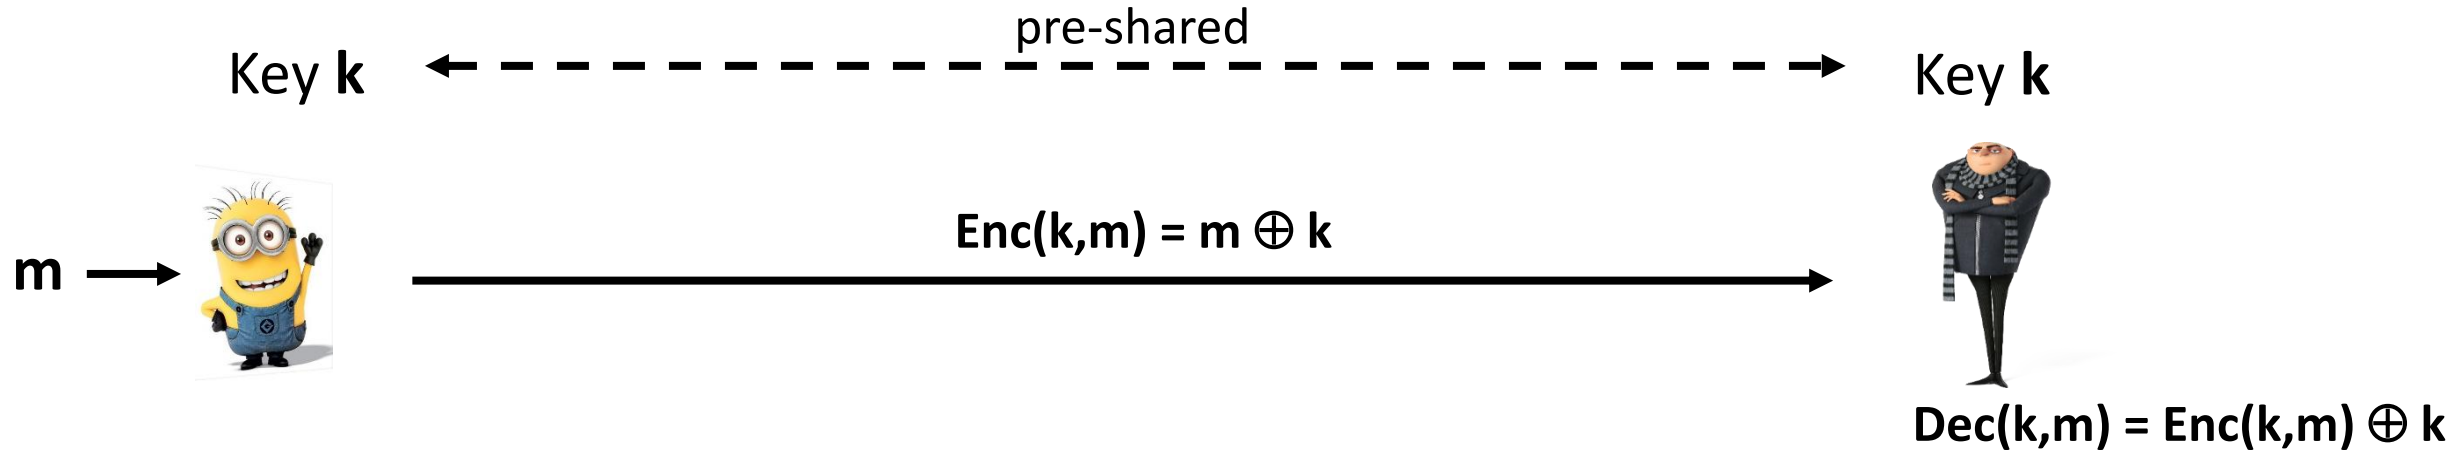


Message
Binary
OTP-Key
Encryption

	YEAH (ASCII Hex: 59454148)		NOPE (ASCII Hex: 4e4f5045)
	01011001010001010100000101001000		01001110010011110101000001000101
\oplus	01110101000111010100101001001010	← same →	\oplus 01110101000111010100101001001010
	<hr/> 0010110001011000000101100000010		<hr/> 00111011010100100001101000001111

Obtaining perfect secrecy: One Time Pad (OTP)

Key = string k of **random** bits as long as the message



Message
Binary (ASCII)
OTP-Key
Encryption

YEAH

01011001010001010100000101001000

01110101000111010100101001001010

00101100010110000000101100000010

same

different

YEAH

01001110010011110101000001000101

11100001010000010111101011010001

1010111000011100010101010010100

Proof of Security (Back to AICC1)

Why is the OTP *Perfectly* Secure?

- For any ciphertext C , **every possible plaintext M is equally likely.**
- The key K is uniformly random, so C gives the adversary **zero new information** about M .

Formally (Perfect Secrecy): $\forall m c, P(M = m | E_K(m) = c) = P(M = m)$

$$\begin{aligned} P(M = m | E_K(m) = c) &= \frac{P(M=m \cap E_K(m)=c)}{P(E_K(m)=c)} = \frac{P(M=m \cap (m \text{ xor } K)=c)}{P(E_K(m)=c)} \\ &= \frac{P(M = m \cap K = m \text{ xor } c)}{P(K = m \text{ xor } c)} \\ &= \frac{P(M = m) * P(K = m \text{ xor } c)}{P(K = m \text{ xor } c)} \\ &= P_K(M = m) \end{aligned}$$

The Key Problem of OTP

The Practical Limitation

OTP requires keys that are **Truly Random, Message-Length, and Used Only Once**.

The Showstopper:

Securely pre-distributing message-length keys for real-world applications is **impractical**.

Cryptanalysis Question: Key Reuse

Suppose Alice reuses the key K to encrypt **two single-word english messages of length 5**, M_1 and M_2 , producing C_1 and C_2 .

$$C_1 = M_1 \oplus K \quad \text{and} \quad C_2 = M_2 \oplus K$$

Your task: Show that Eve can "likely" recover the key and then the two messages M_1 and M_2 .

OTP: Integrity is NOT Included

The Integrity Flaw

- OTP guarantees **Confidentiality** (nobody can read it).
- It **DOES NOT** guarantee **Integrity** (nobody can tamper with it).

Attack Example (Eve flips the first bit of M):

1. Eve flips the first bit of C to get C' .
2. Bob decrypts C' : $M' = C' \oplus K = (M \oplus K \oplus \Delta) \oplus K = M \oplus \Delta$.


Result: Eve successfully changed the message (M' is M with one bit flipped) **without knowing K .**

Module 4: From OTP to Stream Ciphers

Two Main Types of Symmetric Schemes

Encryption of plaintext and decryption of ciphertext are done using

THE SAME KEY



- **Block Ciphers:** Operate on fixed-size blocks (e.g., 128 bits).
 - *Example: AES (Advanced Encryption Standard).*
- **Stream Ciphers:** Operate one bit/byte at a time, like a pseudo-OTP.
 - *Example: ChaCha20.*

We will dive deeper into **Stream Ciphers**.

What is a symmetric cryptographic key?

Fixed-size input to symmetric cryptographic primitives.

The size of the key influences the level of security provided

Key properties

Known to both parties

Partners must agree on the key **before** starting using the primitive

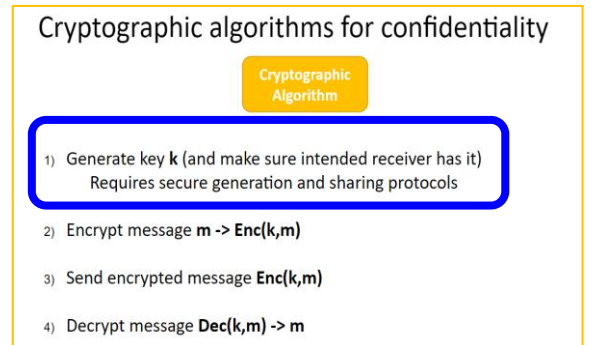
It is reused

The key is pre-shared once* and then reused

* keys do have a “duration”

It must be secret

Revealing the key eliminates any protection provided by the primitive



Stream Ciphers: The Pseudo-OTP

A stream cipher attempt to emulate the OTP's security while solving the key-length problem.

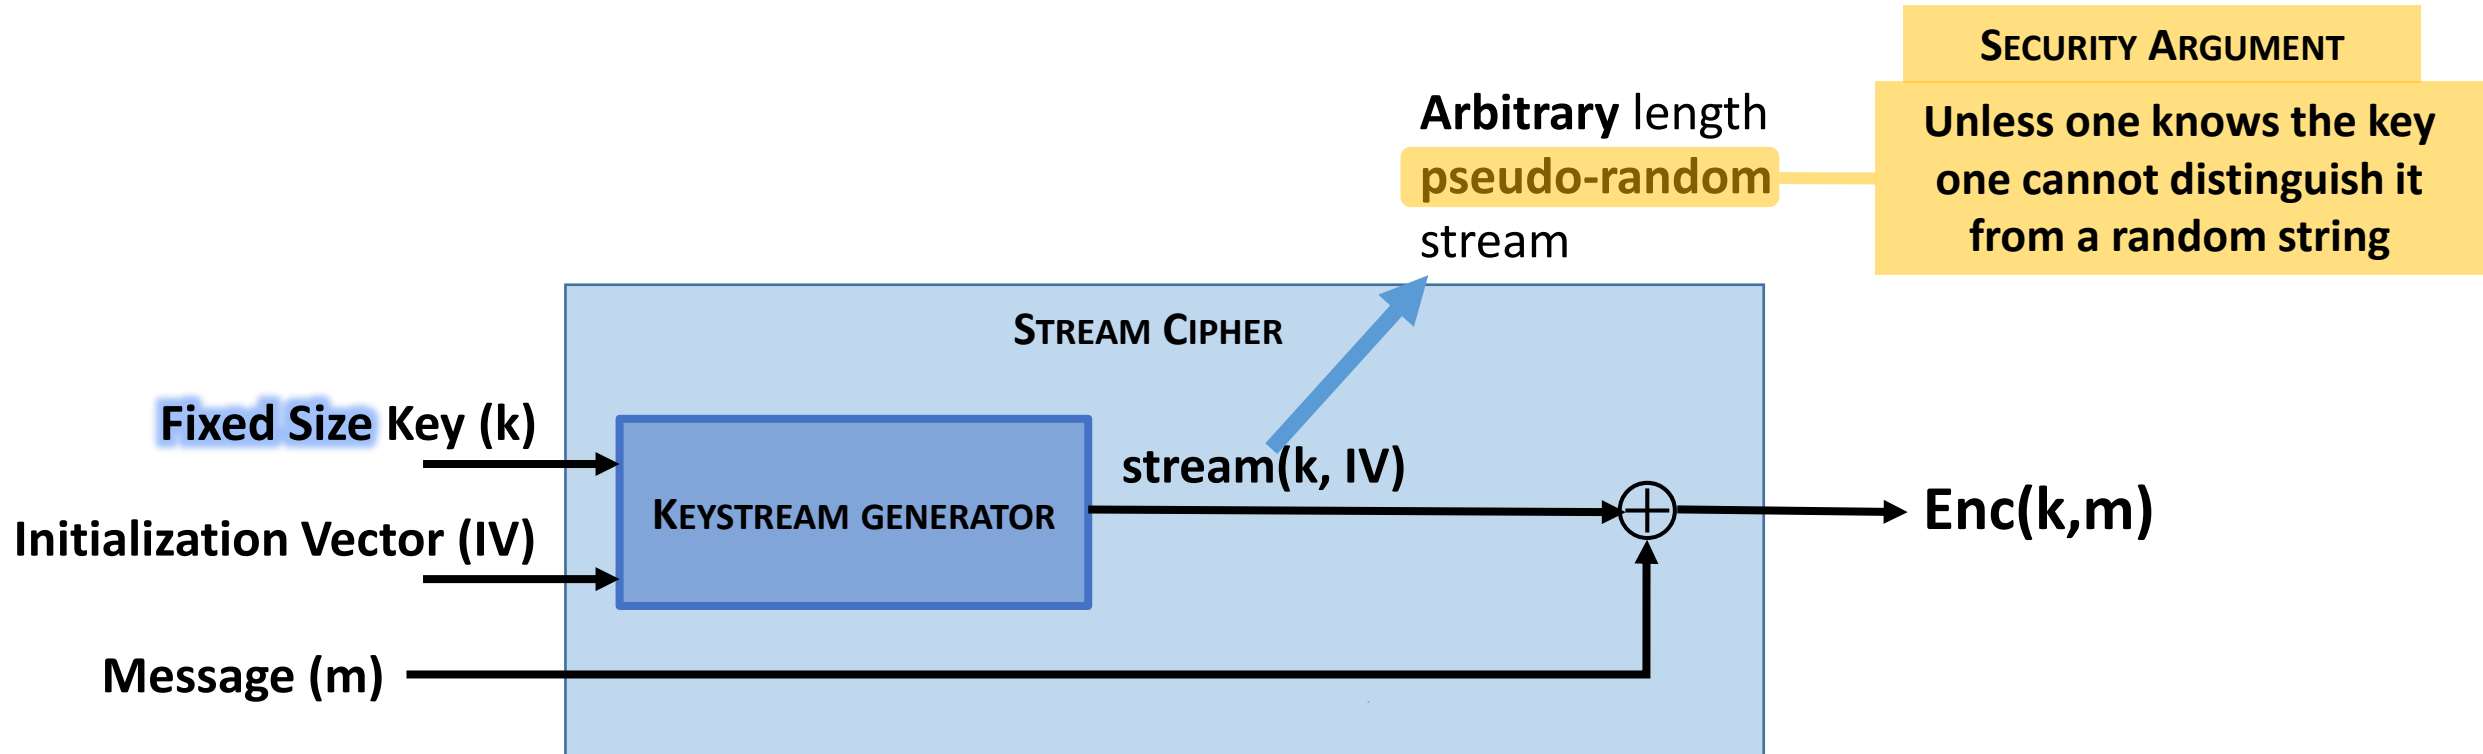
1. **Short Key (K):** Alice and Bob share a small, manageable key.
2. **Key Stream Generator:** Uses K and an **Initialization Vector (IV)** to produce an arbitrary long, *pseudo-random* bit stream (S).
3. **Encryption:** $C = M \oplus S$.

Question: Why does the Key Stream Generator takes in also an IV, and not just the key?

KSG: Finite Secret, Infinite “Randomness”?

- We are generating **potentially infinitely long randomness** (the stream S) from a **finite secret** (K).
- **Security:** Depends entirely on S being computationally indistinguishable from a truly random OTP key.

Stream Ciphers



Strengths and Weaknesses

Speed	Algorithms are linear in time and constant in space.
Low Error Propagation	Errors in one bit do not affect subsequent symbols.
Low Diffusion	A change in one plaintext bit only affects one ciphertext bit.
Susceptibility to Modification	Low diffusion makes it easier to tamper with the message (like the OTP).

Cryptanalysis: The Problem with Periodicity

- Key Stream Generators which have finite state are eventually **periodic**.
- If the message is longer than the period, the stream *repeats*.
- **The Attack:** If the period is identified, and too small, the key stream is known for the entire message.

Interlude On Building KSG

You are not expected to fully understand the following next 3 slides, they are here to illustrate some design patterns for KSG, and challenges involved in building secure KSG, and hint at the math needed to study these objects.

The Linear Feedback Shift Register Building Block for KSG

- **Linear Recurrence Relation**, The next output bit is a linear function (XOR) of previous bits, apply a recurrence relation to extend the initial random sequence of bits, into an arbitrary length sequence of (maybe?) randomly looking bits:
- For example, start with an initial state of 4 bits, a_0, \dots, a_3 , and
$$a_n = a_{n-3} \oplus a_{n-4}$$
- **Output Sequence (starting from 1, 0, 0, 0):** 1,0,0,0,1,0,0,1,1,...
- **Advantage:** Uses only **shift registers** (for state storage) and **XOR gates** (for feedback), so it is extremely fast and efficient to implement in hardware.

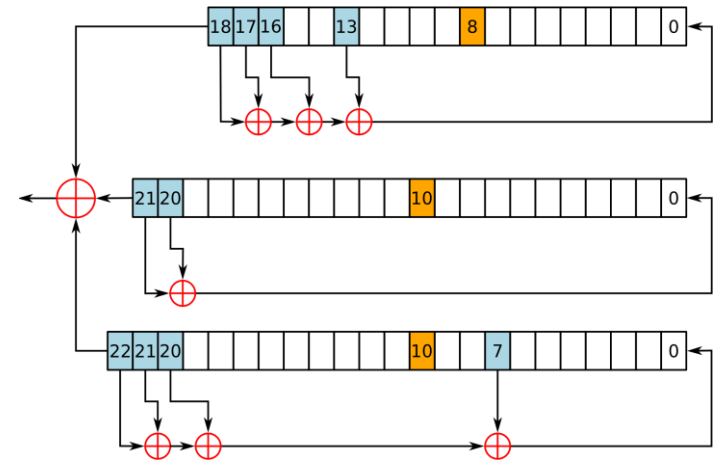
Randomness of LFSR

- **Maximal Period:** If the characteristic polynomial (AICCC1) of the recurrence relation of the LFSR is primitive (MATH-215, Algebra 3), the maximum possible number of states before repeating: For an L-bit register, the maximum period is $2^L - 1$ states.
- **Distribution Property:** As a consequence, the sequence generated exhibit good distribution properties. Specifically, every possible non-zero state appears exactly once in the cycle, so over one full cycle of values, the output sequence are highly balanced. For example, the sequence will contain exactly 2^{L-1} ones and $2^{L-1} - 1$ zeros.

LFSR - Math Danger

- Despite passing some statistical tests, the underlying operation of an LFSR is purely **linear**.
- **Consequence:** Techniques like the Berlekamp-Massey algorithm (beyond the scope of the class) can recover the entire LFSR state and the key stream by observing only a short segment of the output, breaking the cipher.
- **LFSRs should not be used directly as a Key Stream Generators**, though it has been traditionally used as a building block of Key Stream Generators.


The A5/1 Cipher



- **A5/1:** Used historically to secure GSM mobile phone communications.
- It was built by **combining multiple LFSRs** in a non-linear way to try and overcome the linearity flaw while still being cheap and having the nice statistical properties.
- **The Break:** Despite the increased complexity, mathematical weaknesses were found, allowing the cipher to be broken much faster than brute force.

Non-broken Stream Ciphers

- There are alternative, modern stream ciphers for which there are no known practical cryptanalysis.
- *Examples: **Trivium*** (based on NLFSRs, variation of LFSRs) and **Salsa20** (using different techniques).



Bit Security and Consequences

- **Cost of Breaking:** Exploiting a mathematical property or bug can break the scheme much faster than $2^{|K|}$ attempts, **reducing the *bit security***.
- **The Difficulty:** Designing mathematically secure and subtly resistant primitives is incredibly hard.
- **Conclusion:** Never design your own stream cipher, (block cipher, or hash function).
- **Best Practice:** Use **well-vetted, peer-reviewed, standard algorithms and their implementation** (like AES, ChaCha20, SHA-3).

Module 5: Public Key Cryptography

Recap: The Key Distribution Problem

Symmetric ciphers are fast, but require a **pre-shared secret key**.

Question: How can Alice and Bob agree on a secret key if all their communication is observed by Eve?

The Magic of Asymmetric Crypto

- **The Paradox:** Before the 1970s, it was considered impossible to communicate securely without a pre-shared secret.
- **The Revelation:** Public-Key Cryptography allows two parties who have never met to establish a secure channel.

Diffie-Hellman Setup

Alice, Bob (and Eve) know large public parameters: a prime number p and a generator g .

Alice's Actions

1. Chooses **private** secret a .
2. Computes **Public Value** $A = g^a \pmod{p}$.
3. Sends A to Bob.

Bob's Actions

1. Chooses **private** secret b .
2. Computes **Public Value** $B = g^b \pmod{p}$.
3. Sends B to Alice.

Note: Eve sees A , B , g , and p , but cannot find a or b due to the Discrete Log Problem.

The Shared Secret

Both Alice and Bob can compute a **shared secret** K :

$$\begin{aligned} K &= B^a \pmod{p} = (g^b)^a \pmod{p} \\ &= g^{ba} \pmod{p} \\ &= g^{ab} \pmod{p} \\ (g^b)^a \pmod{p} &= A^b \pmod{p} \end{aligned}$$

Where:

$K = B^a \pmod{p}$ is computable by Alice (she has the secret a)

$K = A^b \pmod{p}$ is computable by Bob (he has the secret b)

But Eve cannot compute K , because she has neither a nor b .

Trapdoor Functions

- This example of DH is one of the many examples of how to do public key cryptography to establish a shared secret (a shared key) to then to symmetric cryptography.
- They boil down to mathematical functions that are **easy to compute** in one direction.
- But **hard (high complexity) to do in the other direction**
- **Underlying Difficulty:** For DH, it is the **Discrete Logarithm Problem** .

Beyond DH

There are other ways to set a shared secret (or even to directly communicate), all involve interesting mathematics.

- **RSA**: Based on the Factoring problem.
- **Elliptic Curve Cryptography (ECC)**: Can do Discrete Log on ECC, instead of on public-key crypto.
- **Post-Quantum Cryptography**: New primitives (like **Lattice Crypto**) to resist quantum computers.

Man-In-The-Middle (MITM) (1/2)

Vanilla DH is Vulnerable to Active Attack

- DH only guarantees key *agreement*, not *authenticity* (verifying *who* you are agreeing with).
- **Eve Intercepts:** Eve intercepts A from Alice and B from Bob.

MITM (2/2)

Eve Impersonates

1. Eve sends E_B to Alice. Alice computes K_{AE} .
 2. Eve sends E_A to Bob. Bob computes K_{BE} .
- Eve shares a key K_{AE} with Alice and a separate key K_{BE} with Bob.
 - She can read, modify, and relay all communication.

The Need for Trust (Authentication)

- DH solves **Key Distribution**, but creates the **Identity Verification Problem**.
- **Solution:** DH must be paired with a mechanism (like **Digital Signatures**) to verify identity. (We will discuss that next week)
- **Setup:** Parties still need to trust a third party (a Certificate Authority) for public key ownership.