AMERICAN UNIVERSITY OF BEIRUT

Applied Cryptography

CMPS 297AD/396AI

Fall 2025

Part 1: Provable Security

1.2: One-Time Pad & The Provable Security Mindset

Nadim Kobeissi

https://appliedcryptography.page

How it's made

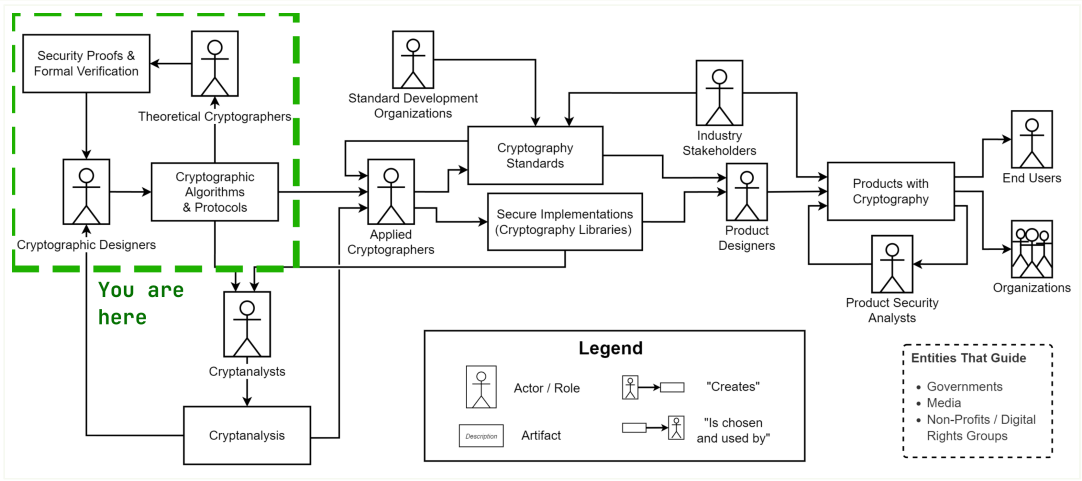

Fischer et al., The Challenges of Bringing Cryptography from Research Papers to Products: Results from an Interview Study with Experts, USENIX Security 2024

Thinking about secrecy

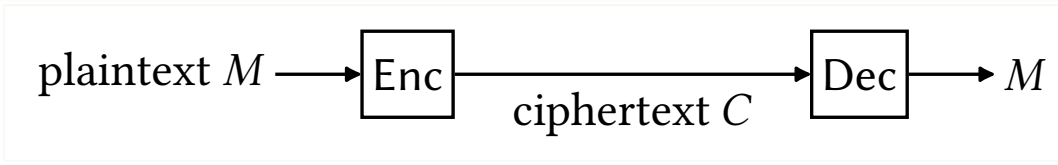

Source: The Joy of Cryptography

Thinking about secrecy

- Keep the whole design secret?
- **“Advantages”**:
  - Attacker doesn’t know how our cipher (or system, more generally,) works.
- **Disadvantages**:
  - Figuring out how the thing works might mean a break.
  - Can’t expose cipher to scrutiny.
  - Everyone needs to invent a cipher.

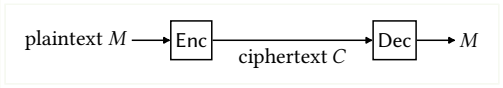

Source: The Joy of Cryptography

Kerckhoff's principle

- *"A cryptosystem should be secure even if everything about the system, except the key, is public knowledge."* — Auguste Kerckhoffs, 1883
- **Why it matters**:
  - No "security through obscurity"
  - The key is the only secret: the rest can be audited, tested, trusted
  - Encourages open standards and peer review
  - If your system's security depends on nobody knowing how it works, it's not secure.

Thinking about secrecy

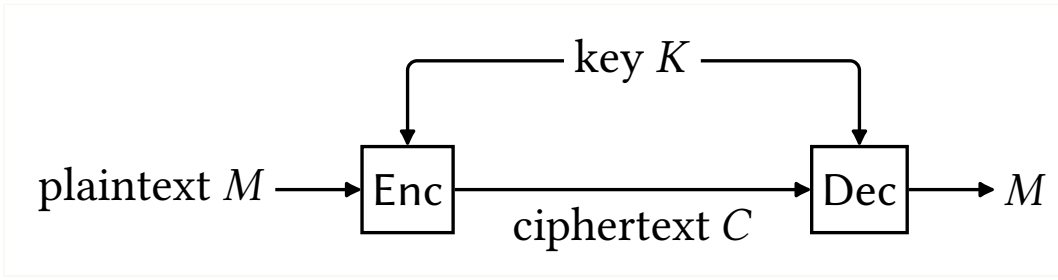

Concentrate all the need for secrecy in the key!

Thinking about secrecy

- Cipher can be scrutinized, used by anyone.
- Design can be shown to hold so long as the key is secret.
- This is how virtually all cryptography is designed today.

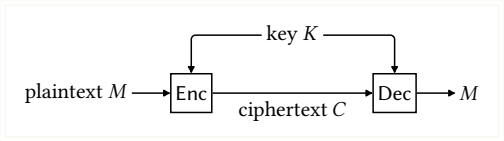

Source: The Joy of Cryptography

One-time pad

First look at a symmetric cipher

$\underline{\textsc{Enc}(K, M):}$

$C := K \oplus M$

return C

$\underline{\textsc{Dec}(K, C):}$

$M := K \oplus C$

return M

XOR (Exclusive OR) operation

A	B	A ⊕ B
0	0	0
0	1	1
1	0	1
1	1	0

Table: Truth table for XOR operation

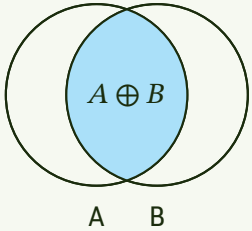

- XOR returns 1 when inputs differ
- XOR returns 0 when inputs are the same
- Key property: $x \oplus x = 0$ and $x \oplus 0 = x$
- Self-inverse: $(M \oplus K) \oplus K = M$

One-time pad

First look at a symmetric cipher

$$
\begin{array}{ll}
 & 11101111101111100011 \quad M \\
\oplus & 00011001110000111101 \quad K \\
\hline
= & 11110110011111011110 \quad C = \mathsf{Enc}(K, M)
\end{array}
$$

(We're encoding the message and key as bits)

One-time pad

First look at a symmetric cipher

$$
\begin{array}{ll}
 & 11110110011111011110 \quad C \\
\oplus & 00011001110000111101 \quad K \\
\hline
= & 11101111101111100011 \quad M = \mathsf{Dec}(K, C)
\end{array}
$$

(We're encoding the message and key as bits)

Key derivation

Uniform distribution

- How to derive K?
- K is ideally random.
- True randomness isn't practical, so K is in practice pseudo-random.
- We need a pseudo-random uniform distribution:
- If $\mathcal{S}$ is a set of m items, then the uniform distribution over $\mathcal{S}$ assigns probability $\frac{1}{m}$ to each item $x \in \mathcal{S}$
- In practice, this just means we need the bits to be random, unpredictable, uniformly distributed in terms of probability
- Sampling a K from a pseudo-random uniform distribution is written as $K \twoheadleftarrow \{0,1\}^n$

Adversary's access to oracle

- "Victim" chooses their key.
- Adversary chooses the message and receives the ciphertext.
- We say that **the adversary has access to an encryption oracle**.

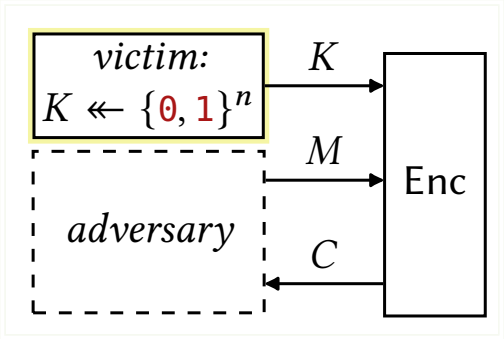

Source: The Joy of Cryptography

Adversary's access to oracle

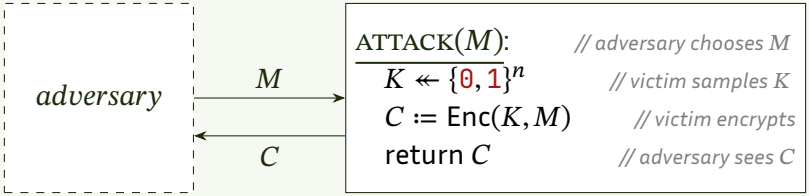

Adversary's access to oracle

- Adversary can query oracle an unbounded number of times.
- Two queries with same M may return different (C, C'), since victim may use different K.
- K is always chosen correctly (pseudo-random uniform sampling)
- "Randomized oracle"
- Attacker cannot see K.

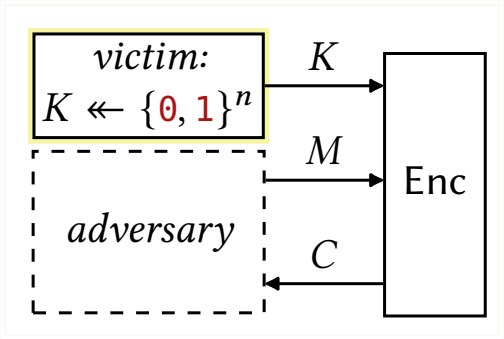

Source: The Joy of Cryptography

Adversary's access to oracle

- When we prove security, we prove what is or isn't possible by the attacker calling Attack(M).

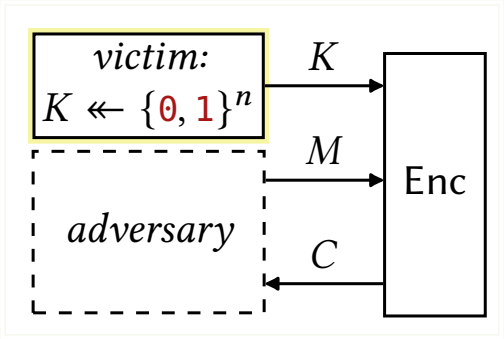

Source: The Joy of Cryptography

“If I use OTP according to the attack scenario (I sample keys uniformly and use each key to encrypt just one ciphertext), then no matter how the plaintexts are chosen, and no matter how the ciphertext is subsequently used, I can enjoy a certain security guarantee.”

One-time pad

Correctness proof

- $\forall(n > 0,\ K \in \{0,1\}^n,\ M \in \{0,1\}^n),\ \mathsf{Dec}(K, \mathsf{Enc}(K, M)) = M$
- For all positive n, any key of n bits and message of n bits will decrypt back to the same plaintext if encrypted into a ciphertext.
- **Proof**:

$$\begin{aligned}
\mathsf{Dec}(K, \mathsf{Enc}(K, M)) &= \mathsf{Dec}(K, K \oplus M) \\
&= K \oplus (K \oplus M) \\
&= (K \oplus K) \oplus M \\
&= 0^n \oplus M \\
&= M \qquad \square
\end{aligned}$$

One-time pad

How do we prove security?

- **Generally**: a cipher is secure if the adversary can't distinguish the output of calls to $ATTACK$ from random junk.
- **Formally**: For all positive integers n and all choices of plaintext $M \in \{0,1\}^n$, the output of the following subroutine is uniformly distributed:

$\textsc{Attack}(M)$:

$K \twoheadleftarrow \{0,1\}^n$

$C := K \oplus M$

return C

One-time pad

How do we prove security?

- If the key is random, the output will be uniformly distributed!
- Suppose $M = \texttt{01}$:
  - $K = \texttt{00}$ is chosen with probability $1/4$: $C = K \oplus M = \texttt{00} \oplus \texttt{01} = \texttt{01}$.
  - $K = \texttt{01}$ is chosen with probability $1/4$: $C = K \oplus M = \texttt{01} \oplus \texttt{01} = \texttt{00}$.
  - $K = \texttt{10}$ is chosen with probability $1/4$: $C = K \oplus M = \texttt{10} \oplus \texttt{01} = \texttt{11}$.
  - $K = \texttt{11}$ is chosen with probability $1/4$: $C = K \oplus M = \texttt{11} \oplus \texttt{01} = \texttt{10}$.

$\textsc{Attack}(M)$:

$K \twoheadleftarrow \{\texttt{0},\texttt{1}\}^n$

$C := K \oplus M$

return C

One-time pad

How do we prove security?

- What if this is true only for $M = \texttt{01}$?
- Fine, let's pick any $M, C \in \{\texttt{0},\texttt{1}\}^n$.
- What is $\Pr[\text{Attack(M)} = \texttt{C}]$?
- Answer: Exactly when $C = \text{Enc}(K, M) = K \oplus M$.
- ...which occurs for exactly one K.
- Since K is chosen uniformly from $\{\texttt{0},\texttt{1}\}^n$, the probability of choosing that K is $\frac{1}{2^n}$. □

$\textsc{Attack}(M)$:

$K \twoheadleftarrow \{\texttt{0},\texttt{1}\}^n$

$C := K \oplus M$

return C

One-time pad

From the adversary's perspective...

$\textsc{Attack}(M)$:

$K \twoheadleftarrow \{0,1\}^n$

$C := K \oplus M$

return C

$\approx$

(indistinguishable from)

$\textsc{Junk}(M)$:

$C \twoheadleftarrow \{0,1\}^n$

return C

“Real or random?”

Limitations of security proofs

Part 1

- Rigor and the real world famously don't mix.
- Security proofs are good for rigor but address very little regarding real-world concerns:
  - How can Alice & Bob obtain a secret key, which only they know?
  - How can they keep K secret?
  - How can a computer sample from the uniform distribution?
  - How can Alice ensure that C is sent reliably to Bob?

Limitations of security proofs

Part 2

- More questions proofs don't address:
  - How can Alice hide the fact that she is talking to Bob (rather than hide only the content)?
  - How can Alice be sure that she is communicating with Bob, not an impostor?
  - How can we incentivize Alice and Bob to use encryption?
  - Should the government be allowed to obtain a warrant to read encrypted communications?
- Security proofs are about specific properties within specific models.
- Real-world security depends on many factors beyond what our models capture.
- Having a security proof is necessary but not sufficient for real-world security.

The value of security proofs

- Despite limitations, security proofs provide important benefits:
  - **Precise guarantees**: Clearly define what security properties are achieved.
  - **Confidence**: When properly structured, proofs ensure no obvious attacks exist.
  - **Foundation for composition**: Proven components can be securely combined.
  - **Precise terminology**: Forces us to clearly define our terms and assumptions.
- Security proofs help identify the *boundaries* of security:
  - What assumptions are necessary?
  - What threats are addressed vs. unaddressed?
  - What conditions must hold for security to be maintained?

The provable security mindset

- Building systems with provable security in mind:
  - Start with clear security goals and adversary model.
  - Design systems whose security can be formally analyzed.
  - Identify and document necessary assumptions.
  - Distinguish between proven properties and conjectures.
- Good practical security requires both:
  - Rigorous proofs for core mechanisms.
  - Practical engineering to address real-world constraints.

OTP: security assumptions & constraints

Part 1

- Our security proofs rely on specific assumptions about the adversary:
  - **Key reuse**: Keys are never intentionally reused (though may repeat by chance).
  - **Observation only**: Adversary passively observes ciphertext but doesn't tam.per with it
  - **Message independence**: Choice of message M is independent of key K.
  - **Key secrecy**: Adversary learns nothing about the key.
  - **No sampling influence**: Adversary cannot influence how the key is sampled.
- These constraints are *necessary* for the security proofs to hold!

OTP: security assumptions & constraints

Part 2

- Side-channel attacks violate our model:
  - We assume adversary cannot coerce victim to run a different algorithm.
  - Cannot observe execution details:
    - CPU timing information (clock cycles).
    - Memory access patterns.
    - Cache hits/misses.
    - Power consumption during encryption.
- Real-world security requires considering these additional attack vectors.
- Our security proofs address a *specific threat model* that may not capture all real-world threats.

One-time pad

What's so special about XOR?

- Let's replace $\oplus$ with $\wedge$. What would happen?
- Output no longer uniform!

A	B	A $\wedge$ B
0	0	0
0	1	0
1	0	0
1	1	1

Table: Truth table for AND operation

$\underline{\text{ATTACK}(M):}$

$K \twoheadleftarrow \{0,1\}^n$

$C := K \wedge M$

return C

One-time pad

What about $\% n$?

- Let's replace $\oplus$ with $\% n$. What would happen?
- Still good!
- Can you prove correctness and security?

$\underline{\textsc{Attack}(M)}$:

$K \twoheadleftarrow \mathbb{Z}_n$

$C := (K + M) \% n$

return C

AMERICAN UNIVERSITY OF BEIRUT

Applied Cryptography

CMPS 297AD/396AI

Fall 2025

Part 1: Provable Security

1.2: One-Time Pad & The Provable Security Mindset

Nadim Kobeissi

https://appliedcryptography.page