

- Modern stream ciphers (eSTREAM project: 2004-2008)

- Salsa20 (2005) \rightsquigarrow ChaCha (2008)

\hookrightarrow core design maps 256-bit key, 64-bit nonce, 64-bit counter onto a 312-bit output

enables using same key (and different nonces) to encrypt multiple messages (will discuss later)

allows random access into the stream

Design is more complex:
- relies on a sequence of rounds
- each round consists of 32-bit additions, XORs, and bit-shifts

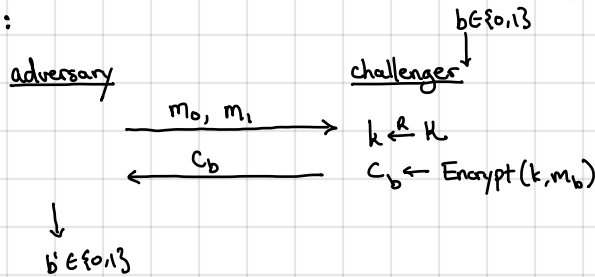
\hookrightarrow very fast even in software (4-14 CPU cycles/output byte) - used to encrypt TLS traffic between Android and Google services

Recall: the one-time pad is not reusable (i.e., the two-time pad is totally broken)

NEVER REUSE THE KEY TO A STREAM CIPHER!

But wait... we "proved" that a stream cipher was secure, and yet, there is an attack?

Recall security game:



Observe: adversary only sees one ciphertext key is only used once

\Rightarrow Security in this model says nothing about multiple messages / ciphertexts

Problem: If we want security with multiple ciphertexts, we need a different or stronger definition (CPA security)

Reusable security: security against chosen-plaintext attacks (CPA-security)

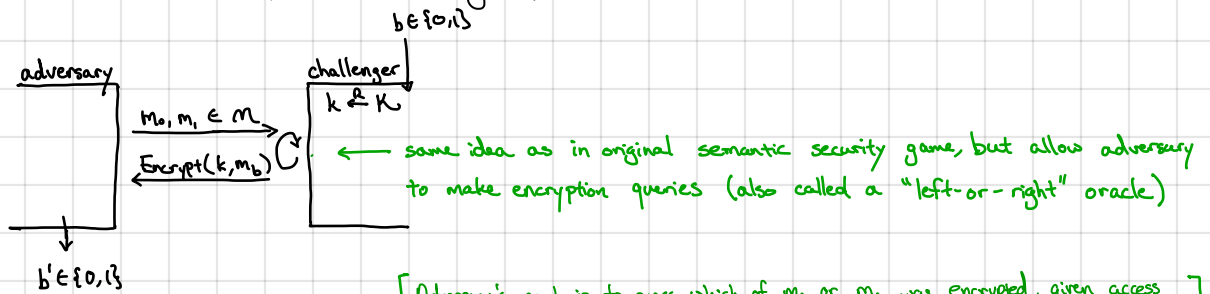
adversary does not just passively observe, it can choose the messages to be encrypted!

- \hookrightarrow semantic security should hold even if adversary sees multiple encrypted messages of its choosing
- \hookrightarrow captures many settings where adversary might know the message that is encrypted (e.g., predictable headers or site content in web traffic) or be able to influence it (e.g., client replies to an email sent by adversary)
- \hookrightarrow goal is to capture as broad of a range of attacks as possible

Definition: An encryption scheme $\Pi_{SE} = (\text{Encrypt}, \text{Decrypt})$ is secure against chosen-plaintext attacks (CPA-secure) if for all efficient adversaries A :

$$\text{CPAAdv}[A, \Pi_{SE}] = |\Pr[W_0 = 1] - \Pr[W_1 = 1]| = \text{negl.}$$

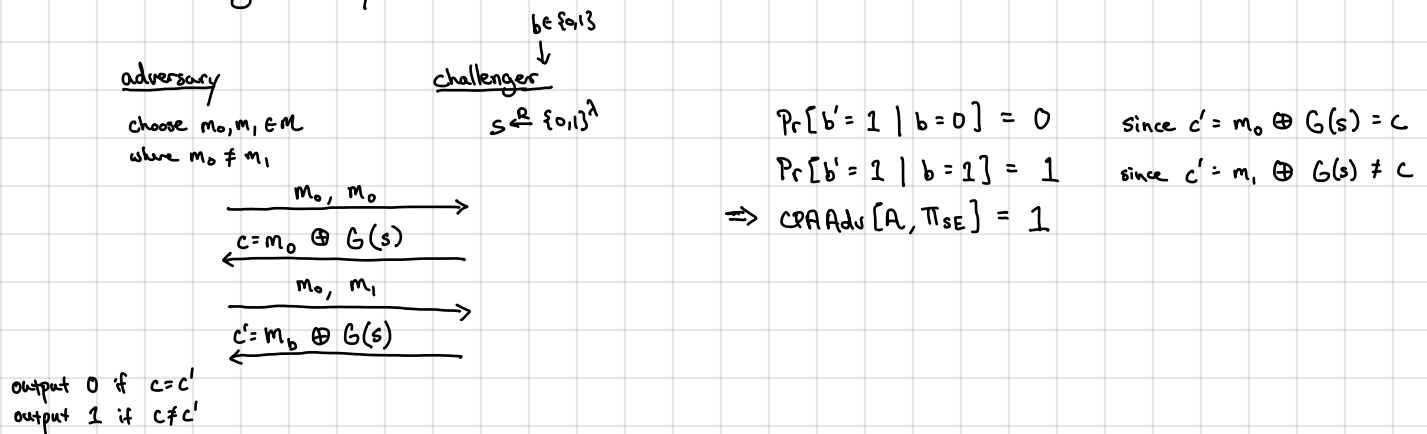
where W_b ($b \in \{0, 1\}$) is the output of the following experiment:



Adversary's goal is to guess which of m_0 or m_1 was encrypted, given access to an encryption oracle (i.e., adversary gets to see encryptions of messages of its choice).

Claim. A stream cipher is not CPA-secure.

Proof. Consider the following adversary:



Observe: Above attack works for any deterministic encryption scheme.

\Rightarrow CPA-secure encryption must be randomized!

\Rightarrow To be reusable, cannot be deterministic. Encrypting the same message twice should not reveal that identical messages were encrypted.

To build a CPA-secure encryption scheme, we will use a "block cipher"

- Block cipher is an invertible keyed function that takes a block of n input bits and produces a block of n output bits

- Examples include 3DES (key size 168 bits, block size 64 bits)


AES (key size 128 bits, block size 128 bits)

Will define block ciphers abstractly first: pseudorandom functions (PRFs) and pseudorandom permutations (PRPs)

\rightarrow General idea: PRFs behave like random functions

PRPs behave like random permutations

Definition. A function $F: K \times X \rightarrow Y$ with key-space K , domain X , and range Y is a pseudorandom function (PRF) if for all efficient adversaries A , $|W_0 - W_1| = \text{negl}$, where W_b is the probability the adversary outputs 1 in the following experiment:



$$\text{PRFAdv}[A, F] = |W_0 - W_1| = |\Pr[A \text{ outputs } 1 \mid b=0] - \Pr[A \text{ outputs } 1 \mid b=1]|$$

Intuitively: input-output behavior of a PRF is indistinguishable from that of a random function (to any computationally-bounded adversary)

3DES: $\{0,1\}^{168} \times \{0,1\}^{64} \rightarrow \{0,1\}^{64}$	$ K = 2^{168}$	$ \text{Funs}[X, Y] = (2^{64})^{(2^{64})}$	} space of random functions is exponentially-larger than key-space
AES: $\{0,1\}^{128} \times \{0,1\}^{128} \rightarrow \{0,1\}^{128}$	$ K = 2^{128}$	$ \text{Funs}[X, Y] = (2^{128})^{(2^{128})}$	

Definition: A function $F: K \times X \rightarrow X$ is a pseudorandom permutation (PRP) if

- for all keys k , $F(k, \cdot)$ is a permutation and moreover, there exists an efficient algorithm to compute $F^{-1}(k, \cdot)$:

$$\forall k \in K : \forall x \in X : F^{-1}(k, F(k, x)) = x$$

- for $k \xleftarrow{R} K$, the input-output behavior of $F(k, \cdot)$ is computationally indistinguishable from $f(\cdot)$ where $f \xleftarrow{R} \text{Perm}[X]$ and $\text{Perm}[X]$ is the set of all permutations on X (analogous to PRF security)

Note: a block cipher is another term for PRP (just like stream ciphers are PRGs)

Observe that a block cipher can be used to construct a PRG:

$$F: \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}^n \text{ be a block cipher}$$

Define $G: \{0,1\}^n \rightarrow \{0,1\}^{ln}$ as

$$G(k) = F(k, 1) \parallel F(k, 2) \parallel \dots \parallel F(k, l)$$

← this stream cipher allows random access!


↑ string concatenation ↑ write input as an n-bit string
(just require that $n > \log l$)

we said PRP above
(will revisit this) ↘

Theorem. If F is a secure PRF, then G is a secure PRG.

Proof. As usual, we show the contrapositive: if G is not a secure PRG, then F is not a secure PRF.

Suppose we have efficient adversary A for G . We use A to build adversary for F :



1. If $l = \text{poly}$, then B is efficient

2. If $b=0$: B sends $G(k)$ to A
where k is a uniformly random key

If $b=1$: B sends uniformly random string (f is random function) to A

$$3. \text{PRFAdv}[B, F] = |\Pr[b'=1 | b=0] - \Pr[b'=1 | b=1]|$$

$$= |\Pr[A \text{ outputs } 1 | b=0] - \Pr[A \text{ outputs } 1 | b=1]|$$

$$= \text{PRGAdv}[A, G]$$

which is non-negligible by assumption.

But... we used a block cipher (PRP) in our construction above. Does the proof still go through?

Not quite...

for a random function, $f(1) = f(2)$ with probability $\frac{1}{2^n}$

for a random permutation, $f(1) = f(2)$ with probability 0

} but 2^{-n} might be very very small...
adversary won't notice unless it sees a "collision" [i.e., two values x, y where $f(x) = f(y)$]

PRF Switching Lemma. Let $F: K \times X \rightarrow X$ be a secure PRP. Then, for any Q -query adversary A :

$$|\text{PRPAdv}[A, F] - \text{PRFAdv}[A, F]| \leq \frac{Q^2}{2|X|}$$

Proof Idea. Adversary essentially cannot tell the difference unless it sees a collision. If there is no collision, then it is just seeing random values. How many queries before there is a collision? Birthday paradox: $Q \sim \sqrt{|X|}$

Take-away: If $|X|$ is large (e.g., exponential), then we can use a PRP as a PRF.

- 3DES: $n = 64$ so $|X| = 2^{64}$ [if adversary makes $\ll 2^{32}$ queries, then can use it as a PRF]

- AES: $n = 128$ so $|X| = 2^{128}$ [if adversary makes $\ll 2^{64}$ queries, then can use it as a PRF]