

# ENGG5383

# Applied Cryptography



Sherman Chow  
Chinese University of Hong Kong  
Spring 2025  
Lecture 2: Symmetric-Key Primitives Part (I)



# Introduction & One-Time Pad

Chapters 0-1 of “The Joy of Cryptography”


# Security Parameter (& some notations)

- We want a “set” of cryptosystems parameterized by  $\lambda$
- Algo.’s run by all honest parties take the commonly agreed  $\lambda$
- They run in time *polynomial* in their input length  $\lambda$ ,  $\text{poly}(\lambda)$ :
  - *Sufficiently fast* (e.g., time) complexity, say, for honest parties, e.g.,  $\lambda^3$
  - Closure property:  $\text{poly}(\lambda) \cdot \text{poly}(\lambda)$  is still in  $\text{poly}(\lambda)$
- Security parameter of the system is  $1^\lambda$  (with length  $\lambda$  bits)
  - $1^\lambda$  denotes  $\lambda$  copies of 1's,  $|1^\lambda| = \lambda$ ,  $1^\lambda$  is in  $\{0, 1\}^\lambda$ , a set of  $\lambda$  bits (0/1)
  - If we put  $\lambda$  as an input, the length of (input)  $\lambda$ ,  $|\lambda|$ , is  $\log(\lambda)$  bits
- Brute-force attacks against the system should run in time  $2^\lambda$

# Probabilistic Polynomial Time (PPT) Algo.

- Deterministic algorithm:  $y = A(x)$
- Probabilistic: “flipping a coin” internally to randomize  $A()$
- $y \leftarrow A(x)$ ,  $y$  is the random variable corresponds to  $A$ 's output
- Or  $y = A(x; r)$ , where  $r$  denotes  $A$ 's “*coin tossing*”
  - when we had the need to specify the randomness explicitly
- All algorithms' details are public (Kerckhoffs' principle)

# Negligible Function



- A function  $v(\lambda)$  is called negligible, denoted **negl**( $\lambda$ )
- $(\forall c > 0) (\exists \lambda') (\forall \lambda \geq \lambda') [v(\lambda) \leq 1/\lambda^c]$
- Less than the inverse of any polynomial for large enough  $\lambda$

$$\leq 1/\lambda^c$$

$$\forall c > 0$$

$$(\forall \lambda \geq \lambda')$$

- Alt. def.:  $v$  is  $\text{negl}(\lambda)$  if for every poly  $p$ :  $\lim_{\lambda \rightarrow \infty} p(\lambda)v(\lambda) = 0$ .
- e.g.,  $2^{-\lambda/2}$  ( $= 1/2^{\lambda/2}$ ),  $2^{-\sqrt{\lambda}}$ ,  $2^{-\log^2 \lambda}$ ,  $n^{-\log \lambda}$
- Prob. of breaking a secure system should be negligible in  $\lambda$
- We have  $\text{poly}(\lambda) \cdot \text{negl}(\lambda) = \text{negl}(\lambda)$  (abusing notations)
- Definition:  $f \approx g$  if  $|f(\lambda) - g(\lambda)|$  is  $(\leq) \text{negl}(\lambda)$  //  $f, g: \mathbf{N} \rightarrow \mathbf{R}$

# What constitutes an encryption scheme?

- A crypto scheme/construction is a collection of algorithms
- *Syntax*: the inputs/outputs (the “function signature”) of all algorithms
- A symmetric-key encryption (SKE) scheme,  $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec})$ :
- $\text{KeyGen}(1^\lambda)$ : A randomized algorithm that outputs a key  $k \in \mathbf{K}$
- $\text{Enc}(k, m)$ : A (possibly randomized) algorithm that takes a key  $k \in \mathbf{K}$  and plaintext  $m \in \mathbf{M}$  as input, and outputs a ciphertext  $c \in \mathbf{C}$
- $\text{Dec}(k, c)$ : A deterministic algorithm that takes a key  $k \in \mathbf{K}$  and ciphertext  $c \in \mathbf{C}$  as input, and outputs a plaintext  $m \in \mathbf{M}$
- $\Sigma.\mathbf{K}$ ,  $\Sigma.\mathbf{M}$ ,  $\Sigma.\mathbf{C}$  denotes the key, message, ciphertext space, resp.

# (Perfect) Correctness

- For all ( $\forall$ )  $k \in \Sigma.\mathbf{K}$  and all  $m \in \Sigma.\mathbf{M}$ ,  
 $\Pr[\Sigma.\text{Dec}(k, \Sigma.\text{Enc}(k, m)) = m] = 1$
- The definition is written in terms of the probability since  $\text{Enc}()$  is allowed to be randomized.
- Counterexample:  $\text{Enc}(k, m) = 0^\lambda$
- One might relax the perfect correctness requirement
  - e.g., for efficiency
  - or for properties we do not know how to achieve otherwise
    - e.g., fully homomorphic encryption

# One-Time Pad (OTP) based on XOR

- eXclusive OR (XOR): For  $b_1 \oplus b_2$  ( $b_i$  is a bit), output as the table
  - a logical operator that returns 1 (true) if the number of 1 (true) inputs is odd
  - XOR is also addition modulo 2 ( $1 + 1 = 2, 2 \bmod 2 = 0$ )
- For bit-string operation  $S_1 \oplus S_2$ , just  $\oplus$  in a bit-wise manner
- $\text{OTP} = \{\text{KeyGen}, \text{Enc}, \text{Dec}\}$
- $\text{KeyGen}(1^\lambda)$ :
  - *uniformly* sample a  $\lambda$ -bit string  $k$
  - output  $k$
- $\text{Enc}(k, m) \rightarrow c = m \oplus k$ :
  - ( $m$  is  $\lambda$ -bit long)
- $\text{Dec}(k, c) \rightarrow m = c \oplus k$

KeyGen:

$k \leftarrow \{0, 1\}^\lambda$   
return  $k$

Enc( $k, m \in \{0, 1\}^\lambda$ ):  
return  $k \oplus m$

	$b_2$		
	XOR $\oplus$	0	1
0	0	0	1
1	1	1	0

$b_1$

Dec( $k, c \in \{0, 1\}^\lambda$ ):  
return  $k \oplus c$




# Fundamentals of “Provable Security”

- Security is a nebulous concept, but not if you took this course
- Provable:
  - We can formally define what it means to be secure
  - and then mathematically prove claims about security
  - e.g., logic of composing building blocks together in secure ways
- Security conference papers require “threat model.”
  - Is the threat too easy to defend? Is it too narrow/restrict?

# Modelling what the adversary sees

- We always treat the attacker as some (unspecified) process that receives output from an algorithm (eavesdrop here).
- Attacker can make the honest users
- run an algorithm with some input and
- see the outcome (not internal variables)
- Each time you call  $\text{eavesdrop}(m)$ ,
- you see a sample from the probability distribution

```
EAVESDROP( $m \in \{0, 1\}^\lambda$ ):  
-----  
 $k \leftarrow \{0, 1\}^\lambda$   
 $c := k \oplus m$   
return  $c$ 
```



# Attackers' Goal vs. Strength of Encryption

- Recover the plaintext  $m$
- Recover a part of the plaintext  $m$ 
  - (Weaker adversary)
  - To protect against a weaker adversary, a weaker scheme may suffice
  - The weaker scheme might be more efficient
- Recover the secret key
  - (Stronger adversary)
- “Deem” only a break when...
  - Whole  $m$  is recovered
    - (Weakest security level)
  - Some part of  $m$  is recovered
    - (Slightly stronger)
  - “1 bit information” of  $m$  is leaked
    - (Strongest)
    - May not be the actual bit of  $m$ 
      - Consider  $m$  is known to be “yes” or “no”



# Basic of Provable Security

Chapter 2 of “The Joy of Cryptography”

# Security Definition

*“Human ingenuity cannot concoct a cipher which human ingenuity cannot resolve.”*  
— Edgar Allan Poe,  
*“A Few Words on Secret Writing”*, 1841

- how to write a security definition
- how to understand & interpret security definitions
- how to prove security using the **hybrid** technique
- how to demonstrate insecurity using attacks
  - w.r.t. the definition
- Let's consider definitions for a more general encryption scheme.
  - (the key is still used for encrypting once)
  - (1 step at a time, we'll study “regular”/“multi-use” encryption soon)



# Provable Security: to define


- Defining an adversary via its goal and capabilities
- A “blueprint” defining security according to the syntax
- but not the internal details of the algorithms or the attacks
- Security always considers the attacker’s view of the system.
- What is the “interface” that honest users expose to the attacker by their use of the cryptography?
- And does that particular interface benefit the attacker?

# A General Encryption Algorithm Model

- There are two inputs to  $\text{Enc}()$ : the key and the plaintext
- The key is our source of randomness and hence security
- The key, generated according to  $\text{KeyGen}()$ , is kept secret
- For now, we still assume each key is used to encrypt once
- The view we said for OTP:
  - A general interface:

*nitty  
gritty*

```
EAVESDROP( $m \in \{0, 1\}^\lambda$ ):  
 $k \leftarrow \{0, 1\}^\lambda$   
 $c := k \oplus m$   
return  $c$ 
```



```
CTXT( $m \in \Sigma.\mathcal{M}$ ):  
 $k \leftarrow \Sigma.\text{KeyGen}$   
 $c := \Sigma.\text{Enc}(k, m)$   
return  $c$ 
```

*syntax*


# Library/Oracle: Interface for Attacking (SKE)

- Adversaries can submit plaintexts via an “encryption oracle”
- receiving their corresponding ciphertexts.
- An *oracle* is an abstract entity that responds to “queries”
  - a controlled environment simulating access to a specific cryptographic operation (like encryption or decryption@§9)
  - without revealing the full knowledge of the *underlying secret*
- A “pessimistic” choice  $\Rightarrow$  Giving more power to attackers
- If an SKE scheme is secure against a “powerful” attacker
- then it’s also secure in “more realistic scenarios”
  - where the attacker has some uncertainty about the plaintexts.



# Two Styles of Security Definition

- Real-or-Random: make sure ciphertexts look “nonsense”
- Left-or-Right: break the “linkage” to the 2 known possibilities
  - Both trying to achieve the “goal of secure encryption”
  - (There might be diff. styles or diff. ways to formalize for each style.)



# Real vs. Random (more specific)

- “An encryption scheme is a good one if its ciphertexts look like random junk to an attacker when ...”
- “each key is secret and used to encrypt only one plaintext, even when the attacker chooses the plaintexts.”
- “ $\Sigma$  is secure if, when you plug its KeyGen and Enc algorithms into the template of the `ctxt` subroutine, the below two implementations of `ctxt` **induce identical behavior in every calling program.**”

real

```
CTXT( $m \in \Sigma.\mathcal{M}$ ):  
-----  
 $k \leftarrow \Sigma.\text{KeyGen}$   
 $c := \Sigma.\text{Enc}(k, m)$   
return  $c$ 
```

vs.

```
CTXT( $m \in \Sigma.\mathcal{M}$ ):  
-----  
 $c \leftarrow \Sigma.\mathcal{C}$   
return  $c$ 
```

random

# A Simple Proof that OTP is RoR-secure

real

```
CTXT( $m \in \Sigma.\mathcal{M}$ ):  
-----  
 $k \leftarrow \Sigma.\text{KeyGen}$   
 $c := \Sigma.\text{Enc}(k, m)$   
return  $c$ 
```

vs.

```
CTXT( $m \in \Sigma.\mathcal{M}$ ):  
-----  
 $c \leftarrow \Sigma.\mathcal{C}$   
return  $c$ 
```

random

```
CTXT( $m$ ):  
-----  
 $k \leftarrow \{0, 1\}^\lambda$  // KeyGen of OTP  
 $c := k \oplus m$  // Enc of OTP  
return  $c$ 
```

vs.

```
CTXT( $m$ ):  
-----  
 $c \leftarrow \{0, 1\}^\lambda$  //  $\mathcal{C}$  of OTP  
return  $c$ 
```

By the properties of XOR ( $\oplus$ ),  
if  $k$  is chosen uniformly at random, so does  $c$ , no matter what  $m$  is

# Left vs. Right

RoR: “an encryption scheme is a good one if its ciphertexts look like random junk to an attacker when each key is secret and used to encrypt only one plaintext, even when the attacker chooses the plaintexts.”

“ $\Sigma$  is secure if, encryptions of  $m_L$  **look like** encryptions of  $m_R$  to an attacker, when each key is secret and used to encrypt only one plaintext, even when the attacker **chooses**  $m_L$  and  $m_R$ .”

EAVESDROP( $m_L, m_R \in \Sigma.\mathcal{M}$ ):

$k \leftarrow \Sigma.\text{KeyGen}$   
 $c := \Sigma.\text{Enc}(k, m_L)$   
return  $c$

(Exercise: Prove OTP is LoR-secure)

EAVESDROP( $m_L, m_R \in \Sigma.\mathcal{M}$ ):

$k \leftarrow \Sigma.\text{KeyGen}$   
 $c := \Sigma.\text{Enc}(k, m_R)$   
return  $c$

(See Exercise 2.15 for an alternative formalization)

# Overview of Hybrid (Proving) Technique

- proving security by demonstrating that a sequence of cryptographic systems are **interchangeable** with one another
- breaks the comparison into **smaller (manageable) steps**
  - instead of directly comparing two *significantly* different systems
- We start with a system (library) and make a **sequence of small modifications** to arrive at the desired system.
- Each modification must be **justified as having “no” effect** on the attacker’s ability to distinguish between the systems.
- Avoids *tedious probability calculations* by focusing on intermediate transitions between libraries.

# Construct a Sequence of Hybrids

- The **heart** of the hybrid technique:
- Designing a sequence of **intermediate libraries**, starting with the first target library and ending with the second.
- Each hybrid should introduce a **small, manageable change** compared to the previous one.
- Your “**bridges**” most likely come from one of two sources:
  - (i) A common technique: **interchangeable code** (p.23)
  - (ii) **Security of the underlying** building block, or **proven fact**

# Justify Each Transition

- The **crucial** part of the proof is to rigorously argue **why** each hybrid is interchangeable with the one before it.
  - demonstrating that no attacker, represented by any calling program, can distinguish between the two libraries based on their outputs.
- Common Justifications (corresponding to (i) and (ii) in the last slide):
- (i) They are the **same** “program”/”library”
- (ii) Applying **known** security properties (**two interchangeable libraries**)
  - leverage proven security properties of underlying cryptographic primitives
  - e.g., when you proving a bigger system using OTP,
  - you can justify the change from real OTP ciphertext to a random bitstring

# Programming-Like Terminologies

- A *library*  $\mathcal{L}$  is a collection of subroutines & *private/static* var.
- A library's *interface* consists of the names, argument types, and output type of all its subroutines
- $\mathcal{A} \diamond \mathcal{L}$ : a program  $\mathcal{A}$  includes calls to subroutines in  $\mathcal{L}$  (*linking*)
- $\mathcal{A} \diamond \mathcal{L} \Rightarrow z$ : denote the event that  $\mathcal{A} \diamond \mathcal{L}$  outputs the value  $z$

$\mathcal{A}$ :
$m \leftarrow \{0, 1\}^\lambda$
$c := \text{CTXT}(m)$
return $m \stackrel{?}{=} c$

An example of  $\mathcal{A}$ :  
choosing a random  $m$  and hoping that  $\text{ctxt}(m)$  is just  $m$ ?


$$\Pr[\mathcal{A} \diamond \mathcal{L} \Rightarrow \text{true}] = 1/2^\lambda$$

$\mathcal{L}$
<u>CTXT(<math>m</math>):</u>
$k \leftarrow \{0, 1\}^\lambda$
$c := k \oplus m$
return $c$



# Another Example (as in the textbook)

- A “challenge”/“game” asking the adversary to guess every single bit of a string picked uniformly at random:



# Interchangeability as Formal Security

- Let  $\mathcal{L}_0$  and  $\mathcal{L}_1$  be two libraries that have the same interface.
- $\mathcal{L}_0$  and  $\mathcal{L}_1$  are said to be *interchangeable*, denoted by  $\mathcal{L}_0 \equiv \mathcal{L}_1$
- if for all programs  $\mathcal{A}$  that output a boolean value (true/false)
  - $\Pr[\mathcal{A} \diamond \mathcal{L}_0 \Rightarrow \text{true}] = \Pr[\mathcal{A} \diamond \mathcal{L}_1 \Rightarrow \text{true}]$
  - We can also call  $\mathcal{A}$  as a *distinguisher*
- Lemmas about interchangeable “hybrid” libraries:
  - 1.  $(\mathcal{A} \diamond \mathcal{L}_1) \diamond \mathcal{L}_2 \equiv \mathcal{A} \diamond (\mathcal{L}_1 \diamond \mathcal{L}_2)$ ; i.e., “associativity” of  $\diamond$
  - 2. If  $\mathcal{L}_{\text{left}} \equiv \mathcal{L}_{\text{right}}$ , for any library  $\mathcal{L}^*$ , we have  $\mathcal{L}^* \diamond \mathcal{L}_{\text{left}} \equiv \mathcal{L}^* \diamond \mathcal{L}_{\text{right}}$

# RoR: (One-time) Uniform Ciphertexts

$$\begin{array}{l} \mathcal{L}_{\text{ots}\$-real}^{\Sigma} \\ \hline \text{CTXT}(m \in \Sigma.\mathcal{M}): \\ k \leftarrow \Sigma.\text{KeyGen} \\ c \leftarrow \Sigma.\text{Enc}(k, m) \\ \text{return } c \end{array}$$

$\equiv$

$$\begin{array}{l} \mathcal{L}_{\text{ots}\$-rand}^{\Sigma} \\ \hline \text{CTXT}(m \in \Sigma.\mathcal{M}): \\ c \leftarrow \Sigma.C \\ \text{return } c \end{array}$$

(uniform as a shorthand of uniformly random)

the superscript  $\Sigma$  means the library is “parameterized” by scheme  $\Sigma$

the “\$” symbol denotes “random” as in coin tossing

(One-time RoR-)Security of OTP:

$$\begin{array}{l} \mathcal{L}_{\text{otp-real}} \\ \hline \text{EAVESDROP}(m \in \{0, 1\}^{\lambda}): \\ k \leftarrow \{0, 1\}^{\lambda} \quad // \text{OTP.KeyGen} \\ \text{return } k \oplus m \quad // \text{OTP.Enc}(k, m) \end{array}$$

$\equiv$

$$\begin{array}{l} \mathcal{L}_{\text{otp-rand}} \\ \hline \text{EAVESDROP}(m \in \{0, 1\}^{\lambda}): \\ c \leftarrow \{0, 1\}^{\lambda} \quad // \text{OTP.C} \\ \text{return } c \end{array}$$

# Left-or-Right Style One-time Security




- Formally, an encryption scheme  $\Sigma$  has one-time secrecy if:

$$\begin{array}{|c} \mathcal{L}_{\text{ots-L}}^{\Sigma} \\ \hline \text{EAVESDROP}(m_L, m_R \in \Sigma.\mathcal{M}): \\ \hline k \leftarrow \Sigma.\text{KeyGen} \\ c \leftarrow \Sigma.\text{Enc}(k, m_L) \\ \text{return } c \end{array} \equiv \begin{array}{|c} \mathcal{L}_{\text{ots-R}}^{\Sigma} \\ \hline \text{EAVESDROP}(m_L, m_R \in \Sigma.\mathcal{M}): \\ \hline k \leftarrow \Sigma.\text{KeyGen} \\ c \leftarrow \Sigma.\text{Enc}(k, m_R) \\ \text{return } c \end{array}$$

Uniform ciphertext  $\Rightarrow$  Left-or-right secrecy  
the same argument in [2-otp-proof.pdf](#) remains valid for any one-time RoR-secure encryption

Left-or-right secrecy  $\not\Rightarrow$  Uniform ciphertext  
Make a “contrived” counterexample LoR-secure  $\text{OTP}'$   
 $\text{OTP}'$  where  $\text{OTP}'.\text{Enc}() := \text{OTP}.\text{Enc}() \parallel 01\}^{\lambda}$

# Demonstrating Insecurity with Attacks



( $\mathcal{L}_{\text{ots}\$-L}$  omitted)

KeyGen:

$k \leftarrow \{0, 1\}^\lambda$   
return  $k$

Enc( $k, m$ ):

return  $k \& m$  // bitwise-AND

- Left-hand side is algorithm  $\mathcal{A}$  we designed for attack.
- Right-hand side inserts the insecure algorithm  $\Sigma$  into 2 libraries/templates.
- $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{ots}\$-real}^\Sigma \Rightarrow \text{true}] = 1.$
- $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{ots}\$-rand}^\Sigma \Rightarrow \text{true}] = 2^{-\lambda}.$
- $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{ots}\$-L}^\Sigma \Rightarrow \text{true}] = 1.$
- $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{ots}\$-R}^\Sigma \Rightarrow \text{true}] = 2^{-\lambda}.$



# Computational Security

Chapter 4 of “The Joy of Cryptography”


# $\approx$ (essentially the same / diff. negligibly)

- $\Pr[E] \approx p$  means  $|\Pr[E] - p|$  is  $\text{negl}(\lambda)$
- $\Pr[X] \approx 0$  (resp.  $1$ )  $\Leftrightarrow$  event  $X$  almost never (resp. always) happens
- $\Pr[A] \approx \Pr[B]$ 
  - $\checkmark$  events  $A$  and  $B$  happen with essentially the same probability
  - $\times$  events  $A$  and  $B$  almost always happen together (cf. head/tail)
- $\approx$  is transitive:  $\Pr[A] \approx \Pr[B]$  and  $\Pr[B] \approx \Pr[C] \rightarrow \Pr[A] \approx \Pr[C]$ 
  - perhaps  $\Pr[A] - \Pr[C]$  is slightly larger, but still negligible
  - $\Pr[X_0] \approx \Pr[X_1], \Pr[X_1] \approx \Pr[X_2], \dots, \Pr[X_{j-1}] \approx \Pr[X_j] \rightarrow \Pr[X_0] \approx \Pr[X_j]?$
  - $\checkmark$  when  $j$  is polynomial (say, in  $\lambda$ , the parameter of interest)
  - $\times$  when  $j$  is exponential (say, in  $\lambda$ , e.g.,  $2^\lambda$ )



# Real vs. Random Indistinguishability

- We switch to **indistinguishability** for **computational security** instead of **interchangeability** for **perfect security** from now on
- The attacker interacts with “something”
- behind a wall via a limited interface
  - (Issuing query, getting response)
- In a real world: real crypto algorithm
- In an ideal world: “random” behavior,
- which is “perfectly secure”
- If no PPT algorithm can distinguish them
- then the real crypto algorithm is secure





# Indistinguishability (Computational Security)

- Let  $\mathcal{L}_0$  and  $\mathcal{L}_1$  be two libraries that have the same interface
- We consider  $\mathcal{A} \diamond \mathcal{L}_b \Rightarrow b'$ 
  - $\mathcal{A}$  is linked to  $(\diamond) \mathcal{L}_b$  for a unknown random bit  $b$
  - $\mathcal{A}$  outputs a bit  $b'$  to declares its guess of  $b$ .
- We say that  $\mathcal{L}_0$  and  $\mathcal{L}_1$  are **indistinguishable**, i.e.,  $\mathcal{L}_0 \approx \mathcal{L}_1$
- if for **all PPT programs**  $\mathcal{A}$  that output a boolean value
- $\Pr[\mathcal{A} \diamond \mathcal{L}_0 \Rightarrow 1] \approx \Pr[\mathcal{A} \diamond \mathcal{L}_1 \Rightarrow 1]$
- i.e., **no** PPT algorithms can differentiate between  $\mathcal{L}_0$  and  $\mathcal{L}_1$
- $|\Pr[\mathcal{A} \diamond \mathcal{L}_0 \Rightarrow 1] - \Pr[\mathcal{A} \diamond \mathcal{L}_1 \Rightarrow 1]|$  is also called advantage of  $\mathcal{A}$

# Indistinguishability is Transitive

- $\mathcal{L}_0$  and  $\mathcal{L}_1$  are **indistinguishable**, i.e.,  $\mathcal{L}_0 \approx \mathcal{L}_1$
- if for **all PPT programs**  $\mathcal{A}$  that output a boolean value
- $\Pr[\mathcal{A} \diamond \mathcal{L}_0 \Rightarrow 1] \approx \Pr[\mathcal{A} \diamond \mathcal{L}_1 \Rightarrow 1]$
  
- $\approx$  is transitive: we can do hybrid proofs
- If  $\mathcal{L}_0 \equiv \mathcal{L}_1$  then  $\mathcal{L}_0 \approx \mathcal{L}_1$
- If  $\mathcal{L}_0 \approx \mathcal{L}_1$  then  $\mathcal{L}^* \diamond \mathcal{L}_0 \approx \mathcal{L}^* \diamond \mathcal{L}_1$  for any poly.-time library  $\mathcal{L}^*$

# Simple Specific Examples of $\mathcal{A}()$ ; & Union Bound

**$\mathcal{A}_{obvious}$**

```
do q times:
  if PREDICT( $\theta^\lambda$ ) = true
    return 1
return 0
```

**$\mathcal{A}_{infinite\ power}$**

```
do  $\infty$  times:
  if PREDICT( $\theta^\lambda$ ) = true
    return 1
return 0
```

■  $\Pr[\mathcal{A}_{obvious} \diamond \mathcal{L}_{right} \Rightarrow 1] = 0$

■  $\Pr[\mathcal{A}_{obvious} \diamond \mathcal{L}_{left} \Rightarrow 1]$

$= 1 - \Pr[\mathcal{A}_{obvious} \diamond \mathcal{L}_{left} \Rightarrow 0]$  *ret. true if  $x = s$*

$= 1 - \Pr[q \text{ independent Predict() calls ret. false}]$

$(= 1 - (1 - 1 / 2^\lambda)^q$  (*// too many terms to list*))

$\leq \Pr[1^{st} \text{ call ret. true}] + \Pr[2^{nd} \text{ call ret. true}] + \dots$

$\leq q / 2^\lambda$  (*// a "loose" bound but suffice for us*)

**$\mathcal{L}_{left}$**

```
PREDICT(x):
  s  $\leftarrow$   $\{0, 1\}^\lambda$ 
  return x  $\stackrel{?}{=} s$ 
```

**$\mathcal{L}_{right}$**

```
PREDICT(x):
  return false
```

# Difference (“Bad-Event”) Lemma

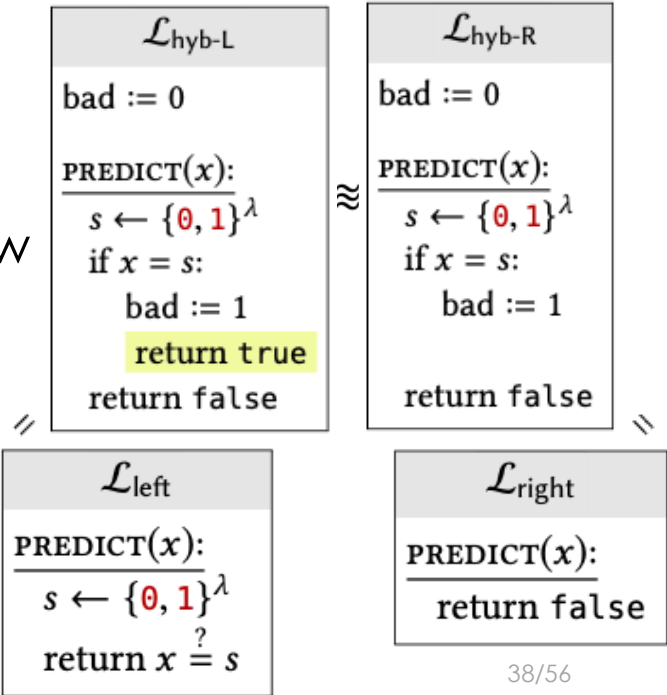
- Two libraries are *expected* to execute exactly the same
- until some *rare & exceptional* (“*bad*”) condition happens.
- Let  $\mathcal{L}_0$  and  $\mathcal{L}_1$  be libraries that each define a variable “*bad*” that is initialized to 0.
- If  $\mathcal{L}_0$  and  $\mathcal{L}_1$  have identical code, except for code blocks reachable only when *bad* = 1
  - (e.g., think of it as guarded by an “if (*bad* = 1)” statement)
- then  $|\Pr[\mathcal{A} \diamond \mathcal{L}_0 \Rightarrow 1] - \Pr[\mathcal{A} \diamond \mathcal{L}_1 \Rightarrow 1]| \leq \Pr[\mathcal{A} \diamond \mathcal{L}_0 \text{ sets } \textit{bad} = 1]$

# Motivation, Notations, Observation

- Why such lemma? Because it is a common situation
  - e.g., while using a crypto. primitive  $\Sigma$  for a certain bigger task
  - you want to argue the “whole task” is secure if  $\Sigma$  remains secure
- Let  $B_i$  be the event that  $\mathcal{A} \diamond \mathcal{L}_i$  sets *bad* to 1 at some point.
- $\neg B_i$  denotes the corresponding complement event.
  - (The notation in the textbook is  $\overline{B}_i$ )
- $\Pr[\mathcal{A} \diamond \mathcal{L}_i \Rightarrow 1] = \Pr[\mathcal{A} \diamond \mathcal{L}_i \Rightarrow 1 \mid B_i] \Pr[B_i] + \Pr[\mathcal{A} \diamond \mathcal{L}_i \Rightarrow 1 \mid \neg B_i] \Pr[\neg B_i]$ 
  - by definition (of *conditional probability*)
  - useful in proving the lemma (omitted, see the textbook)

# Using the Lemma: $\mathcal{L}_{\text{left}} \approx \mathcal{L}_{\text{right}} (\forall \text{ PPT } \mathcal{A})$

- We just show the indistinguishability of 1 calling program
- Define 2 simple hybrids:  $\mathcal{L}_{\text{hyb-L}}$  and  $\mathcal{L}_{\text{hyb-R}}$ :
- The only diff. of  $\mathcal{L}_{\text{hyb-L}}$  from  $\mathcal{L}_{\text{left}}$  is var. **bad**
  - It never reads from this variable.
  - This change has no effect.
- The only diff. of  $\mathcal{L}_{\text{hyb-R}}$  from  $\mathcal{L}_{\text{hyb-L}}$  is in yellow
- $|\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{hyb-L}} \Rightarrow 1] - \Pr[\mathcal{A} \diamond \mathcal{L}_{\text{hyb-R}} \Rightarrow 1]|$   
 $\leq \Pr[\mathcal{A} \diamond \mathcal{L}_{\text{hyb-L}} \text{ sets } \text{bad} = 1]$  (let it be  $p^*$ )
- 1 call to  $\text{Predict}()$ ,  $p^* = 1/2^\lambda$
- (poly)  $q$  calls to  $\text{Predict}()$   $\rightarrow q/2^\lambda$  is negl.
- $\mathcal{L}_{\text{hyb-R}}$  and  $\mathcal{L}_{\text{right}}$  both always return false.





# Pseudorandom Generators

Chapter 5 of “The Joy of Cryptography”

# Pseudorandom generator (PRG)



- Syntax: a deterministic function  $G$  whose outputs are longer than its inputs, *i.e.*,  $G: \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda + \ell}$ 
  - We call  $\ell$  the *stretch*;  $\ell = 0$  (or even negative) is a trivial PRG
- Security requirement:
  - When the input “seed” to  $G$  is chosen uniformly at random, it induces a “certain” distribution over the possible output
  - More formally, the distribution should be *pseudorandom*
  - *i.e.*, it is *indistinguishable* from the uniform distribution



“a randomness multiplier”



# Pseudorandomness, formally

- Let  $G: \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda + \ell}$  be a deterministic function with  $\ell > 0$ .
- We say that  $G$  is a secure PRG if  $\mathcal{L}_{\text{prg-real}}^G \approx \mathcal{L}_{\text{prg-rand}}^G$ :

Real world uses a real PRG: picking a seed  $s$ , returns  $G(s)$




Ideal world returns random: returning a random  $|G(s)|$ -bit bitstring

- This definition is kind of a “master” definition that encompasses all practical (polynomial) statistical tests

<https://dilbert.com/strip/2001-10-25><sub>41/56</sub>

# Interchangeability vs. Indistinguishability

- Only  $2^\lambda$  possible outputs cannot cover the range  $\{0, 1\}^{2\lambda}$
- Outputs of  $G()$  can't be perfectly random
  - $\because$  Unbounded (not PPT) adversaries can try all possible inputs
- But outputs of  $G()$  can be *indistinguishable* from uniform.
- Attackers in practice are PPT (in  $\lambda$ ) can only test a small fraction of the possible inputs/guesses.
- $\mathcal{L}_{\text{prg-real}}$  samples from distribution of **red dots**
- $\mathcal{L}_{\text{prg-rand}}$  directly samples the **uniform distribution** on  $\{0, 1\}^{2\lambda}$



# Teach me how to construct a PRG!

- If you can, you get a Ph.D. right away [\*\*]
  - If it were possible to prove that some function  $G$  is a secure PRG
  - it'd resolve the famous P vs. NP (nondeterministic poly. time) problem
- The next best thing that cryptographic research can offer are *candidate PRGs*, which are *conjectured* to be secure
  - In practice, those ones that have been subjected to *significant public scrutiny* and *resisted all attempts at attacks so far*
- The entire rest of this course (or the textbook) is based on cryptography that is only *conjectured* to be secure
- I thought this course is about the *science* of cryptography!?
- Provable security is rigorous but conditional on a few conjectures

# How NOT to build a PRG

- $G(s) := s || s$  ( $||$  denotes string concatenation)
- Every string exhibits a “discernible pattern”
  - its 1<sup>st</sup> half equal to its 2<sup>nd</sup> half
  - not likely for a uniform distribution

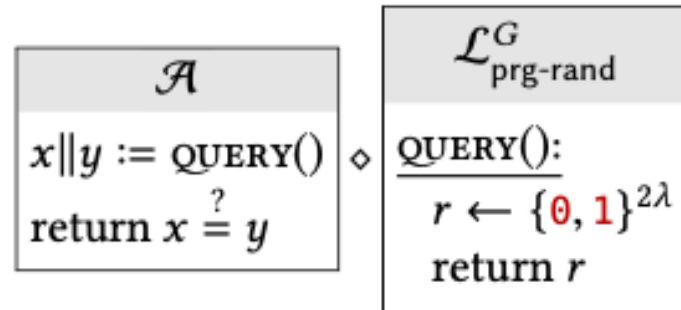
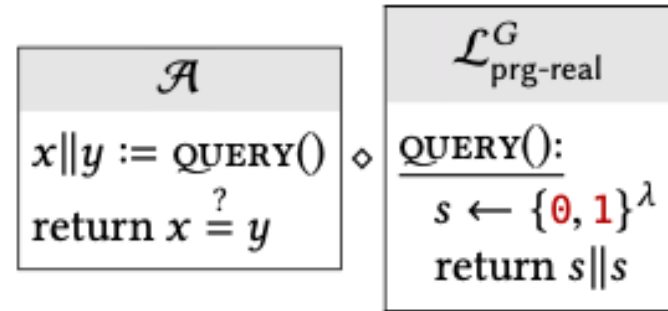
- $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{prg-real}}^G \Rightarrow 1] = 1$

- $\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{prg-rand}}^G \Rightarrow 1] = 1/2^\lambda$

- Advantage of  $\mathcal{A}$

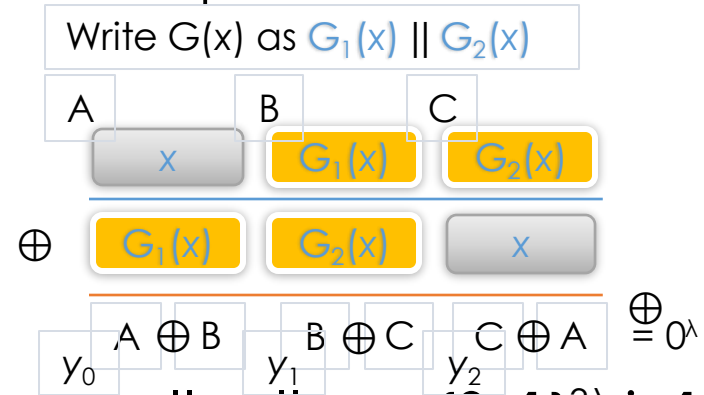
$$= |\Pr[\mathcal{A} \diamond \mathcal{L}_{\text{prg-real}}^G \Rightarrow 1] - \Pr[\mathcal{A} \diamond \mathcal{L}_{\text{prg-rand}}^G \Rightarrow 1]|$$

$$= 1 - 1/2^\lambda$$




# How NOT to build a PRG, encore

- Given  $G: \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$ , build  $G': \{0, 1\}^\lambda \rightarrow \{0, 1\}^{3\lambda}$  as follows.
- $G'(x) := (x \parallel G(x)) \oplus (G(x) \parallel x)$  // the output obscures the input?
- $\mathcal{A}$ :
  - $y_0 \parallel y_1 \parallel y_2 := \text{Query}() // y := G'(x)$ 
    - each  $y_i$  has length  $\lambda$
  - return  $(y_0 \oplus y_1 \oplus y_2) \stackrel{?}{=} 0^\lambda$
- $\mathcal{A}(G'(s))$  always returns true
- $\Pr[y_0 \oplus y_1 \oplus y_2 = 0^\lambda]$  for random  $y = y_0 \parallel y_1 \parallel y_2 \in \{0, 1\}^{3\lambda}$  is  $1/2^\lambda$
- (By the way,  $G''(x) := x \parallel G(x)$  is also insecure)




# Extending the Stretch of a PRG



- $G: \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\lambda + \ell}$ 
  - Stretch  $\ell$  measures how much longer its output is than its input
  - Can we build a PRG w/ larger stretch from one w/ smaller stretch?
- Suppose we have a length-doubling PRG  $G: \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$
- Can we make a length-tripling or quadrupling one?
- Are  $H_1/H_2$  secure?
- The longer
  - the merrier?
  - or the riskier?


$H_1(s):$

$x||y := G(s)$   
 $u||v := G(y)$   
return  $x||u||v$



$H_2(s):$

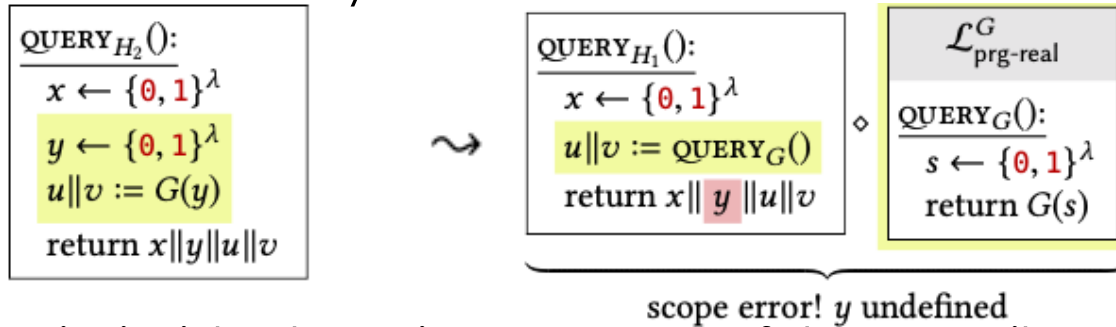
$x||y := G(s)$   
 $u||v := G(y)$   
return  $x||y||u||v$



[5-prg-feedback.pdf](#)

# Where the Proof Breaks Down for $H_2()$

- Let's try to "blindly" reproduce the security proof for  $H_1$  with  $H_2$
- We get stuck when we try to factor out the 2<sup>nd</sup> call to  $G$  via  $\mathcal{L}_{\text{prg-real}}$ :



- $s$  can only exist inside the private scope of the new library,
- while there still exists a "dangling reference"  $y$  in the original library.
- This particular proof strategy fails does not imply  $H_2()$  is insecure
  - although it is indeed insecure in this case (Exercise: concrete attack)

# More Discussions on the Failure

- A PRG's output is indistinguishable from random if
  1. its seed is uniform, and
  2. the seed is not used for anything else! (which breaks condition (1))
- This construction  $H_2$  violates condition (2)
  - Its output contains the "seed"  $y$ , so the seed is no longer random
- In the proof, we can only express a call to  $G$  in terms of  $\mathcal{L}_{\text{prg-real}}$  if the input to  $G$  is uniform and is used *nowhere else* (still uniform)
- Takeaway: These (subtle) issues are not limited to PRGs.
- Every hybrid security proof in this course includes steps where we factor out some statements in terms of some pre-existing library.
- Don't take these steps for granted!

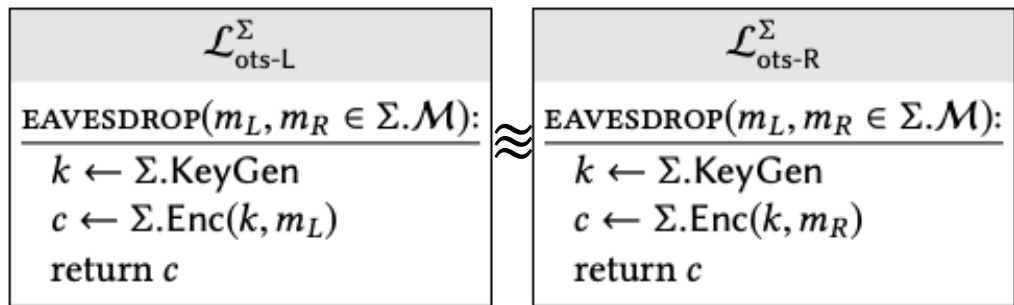


# (In)Security or Common Task in Crypto.

- We try to build “better” PRG from simpler PRG.
- “Security”: if the building blocks are secure then the construction is secure.
- To show insecurity, you shouldn’t directly attack the building blocks!
- We’ll be studying “fancy” higher-level constructions from “weaker” “primitives”
- “Insecurity”: secure building blocks don’t imply security for the whole thing.
- You should assume the building blocks are secure and attack the way that the building blocks are being used.

# Application: One-Time-Secret Encryption

- $\text{KeyGen}(1^\lambda)$ : output  $k$  sampled *uniformly* from  $\mathbf{K} = \{0, 1\}^\lambda$
- $\text{Enc}(k, m) \rightarrow c = m \oplus G(k) // \mathbf{M} = \{0, 1\}^{\lambda + \ell}$  if  $G: \mathbf{K} \rightarrow \{0, 1\}^{\lambda + \ell}$
- $\text{Dec}(k, c) \rightarrow m = c \oplus G(k) // \mathbf{C} = \{0, 1\}^{\lambda + \ell}$
  
- *Computational one-time secrecy* (of a general scheme  $\Sigma$ )



[5-potp-proof.pdf](#)

# Let's extend... indefinitely

- Indeed, even if the stretch is 1, we can further stretch it.
- The PRG-feedback construction can be generalized:
- We continue to feed part of  $G$ 's output into  $G$  again.
- Exercise for you: The proof still works similarly
  - the security of  $G$  is applied one at a time to each application of  $G$


$H_n(s)$ :

$s_0 := s$

for  $i = 1$  to  $n$ :


$s_i || t_i := G(s_{i-1})$

return  $t_1 || \dots || t_n || s_n$



# Stream Cipher


- A stream cipher  $G()$  **takes a seed  $s$  and length  $\ell$**  as input and outputs a string.
- It should satisfy the following requirements:
  1.  $G(s, \ell)$  is a string of length  $\ell$  (or multiple of  $\ell$  for simplicity)
  2. If  $i < j$ , then  $G(s, i)$  is a **prefix** of  $G(s, j)$ 
    - $G(s, n)$  is an infinitely-long string when  $n$  goes to infinity
  3. **For each  $n$** , the function  $G(\cdot, n)$  is a **secure PRG**
    - **$n$  is hardwired** to  $G$  and hence  **$G(\cdot; n)$**  only takes 1 input instead of 2 inputs
- **Simply** use our construction  $H_n(s)$  with  $n$  as  $\ell$
- **Keep outputting  $t_i$** 
  - not outputting  $s_n$  to keep the prefix property



$H_n(s)$   
 $s_0 := s$   
for  $i = 1$  to  $n$ :  
 $s_i || t_i := G(s_{i-1})$   
return  $t_1 || \dots || t_n || s_n$

# (Compromising) Secure Messaging

- Suppose Alice & Bob share a symmetric key  $k$  and are using a secure messaging app to exchange messages over a long period of time, so they worry  $k$  will be leaked
- Suppose an attacker eventually learns  $k$ .
- Then the attacker can decrypt all **past, present, and future** ciphertexts that it saw!
- Can we do **better**?



e.g.,  $k$  is leaked some time after  $t = 6$

Attacker can decrypt ctxt. created during  $t = 0$  to  $t = 5$  (& beyond)

# Forward-Secure Secure Messaging

- The attacker **can, of course, decrypt all future** ciphertexts
  - Why? Because the attacker “becomes” Bob since then
- There’s **hope** that the **past ciphertexts can’t be decrypted**
  - when the attacker gets the key in the present moment
- Forward secrecy: messages in the **present** are protected against a key-compromise that happens in the **future**



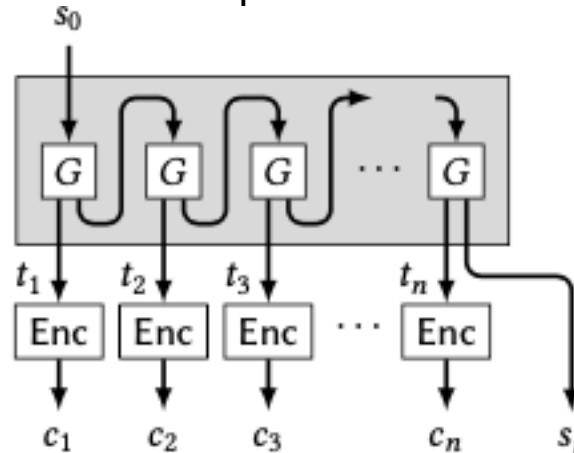
*Ctxt. created during time periods 0 to 6 remain secure if only  $s_6$  is leaked*

- We knew how to do that! (by “**evolving**” the “key” from  $s_n$  to  $s_{n+1}$ )
- This is our stream cipher, also known as *symmetric ratchet*
  - It is easy to **advance the key sequence** in the forward direction (from  $s_n$  to  $s_{n+1}$ ) but hard to reverse it (from  $s_{n+1}$  to  $s_n$ ).

# Theorem for Security of Ratchet

- If the symmetric ratchet is used with a **secure PRG G** and an encryption scheme  $\Sigma$  ( $\Sigma.K = \{0, 1\}^\lambda$ ) that **has uniform ciphertexts**, then the first  $n$  ciphertexts are **pseudorandom**, even to an eavesdropper who compromises the key  $s_n$ .

```
ATTACK( $m_1, \dots, m_n$ ):  
   $s_0 \leftarrow \{0, 1\}^\lambda$   
  for  $i = 1$  to  $n$ :  
     $s_i || t_i := G(s_{i-1})$   
     $c_i \leftarrow \Sigma.\text{Enc}(t_i, m_i)$   
  return ( $c_1, \dots, c_n, s_n$ )
```




# Proof

- Boxed part is indistinguishable to a PRG  $H_n$
- PRG  $H_n$  is indistinguishable to a truly random function
  - skipped (factoring out  $\mathcal{L}_{\text{prg-real}}$  and replacing it w/  $\mathcal{L}_{\text{prg-rand}}$ )
- Keys of  $\Sigma$  look truly random, used once, & nowhere else
  - so, we factor them out and
  - replace them w/ uniform ciphertexts

```

ATTACK( $m_1, \dots, m_n$ ):
 $s_0 \leftarrow \{0, 1\}^\lambda$ 
 $t_1 \parallel \dots \parallel t_n \parallel s_n := H_n(s_0)$ 
for  $i = 1$  to  $n$ :
   $c_i \leftarrow \Sigma.\text{Enc}(t_i, m_i)$ 
return ( $c_1, \dots, c_n, s_n$ )
    
```



```

ATTACK( $m_1, \dots, m_n$ ):
for  $i = 1$  to  $n$ :
   $t_i \leftarrow \{0, 1\}^\lambda$ 
 $s_n \leftarrow \{0, 1\}^\lambda$ 
for  $i = 1$  to  $n$ :
   $c_i \leftarrow \Sigma.\text{Enc}(t_i, m_i)$ 
return ( $c_1, \dots, c_n, s_n$ )
    
```

```

ATTACK( $m_1, \dots, m_n$ ):
for  $i = 1$  to  $n$ :
   $t_i \leftarrow \{0, 1\}^\lambda$ 
 $c_i \leftarrow \Sigma.\text{Enc}(t_i, m_i)$ 
 $s_n \leftarrow \{0, 1\}^\lambda$ 
return ( $c_1, \dots, c_n, s_n$ )
    
```

```

ATTACK( $m_1, \dots, m_n$ ):
for  $i = 1$  to  $n$ :
   $c_i \leftarrow \text{CTXT}(m_i)$ 
 $s_n \leftarrow \{0, 1\}^\lambda$ 
return ( $c_1, \dots, c_n, s_n$ )
    
```

```

 $\mathcal{L}_{\text{ots}\$-real}^\Sigma$ 
CTXT( $m \in \Sigma.\mathcal{M}$ ):
 $k \leftarrow \Sigma.\text{KeyGen}$ 
 $c \leftarrow \Sigma.\text{Enc}(k, m)$ 
return  $c$ 
    
```

```

ATTACK( $m_1, \dots, m_n$ ):
for  $i = 1$  to  $n$ :
   $c_i \leftarrow \Sigma.C$ 
 $s_n \leftarrow \{0, 1\}^\lambda$ 
return ( $c_1, \dots, c_n, s_n$ )
    
```