CS255: Winter 2017

PRPs and PRFs

1. Abstract block ciphers: PRPs and PRFs,
2. Security models for encryption,
3. Analysis of CBC and counter mode

Dan Boneh, Stanford University

PRPs and PRFs

- Pseudo Random Function (**PRF**) defined over (K,X,Y):

$$F: K \times X \rightarrow Y$$

such that exists "efficient" algorithm to evaluate F(k,x)

- Pseudo Random Permutation (**PRP**) defined over (K,X):

$$E: K \times X \rightarrow X$$

such that:

1. Exists "efficient" algorithm to evaluate E(k,x)
2. The function $E(k, \cdot)$ is one-to-one
3. Exists "efficient" inversion algorithm D(k,x)

Running example

- <u>Example PRPs</u>: 3DES, AES, …

  AES-128: $K \times X \rightarrow X$ where $K = X = \{0,1\}^{128}$

  DES: $K \times X \rightarrow X$ where $X = \{0,1\}^{64}$, $K = \{0,1\}^{56}$

  3DES: $K \times X \rightarrow X$ where $X = \{0,1\}^{64}$, $K = \{0,1\}^{168}$

- Functionally, any PRP is also a PRF.
  - A PRP is a PRF where X=Y and is efficiently invertible
  - A PRP is sometimes called a ***block cipher***

Secure PRFs

- Let F: $K \times X \to Y$ be a PRF

  - Funs[X,Y]: the set of **all** functions from X to Y
  - $S_F = \{ F(k,\cdot) \text{ s.t. } k \in K \} \subseteq$ Funs[X,Y]

- Intuition: a PRF is **secure** if
  a random function in Funs[X,Y] is indistinguishable from
  a random function in S_F

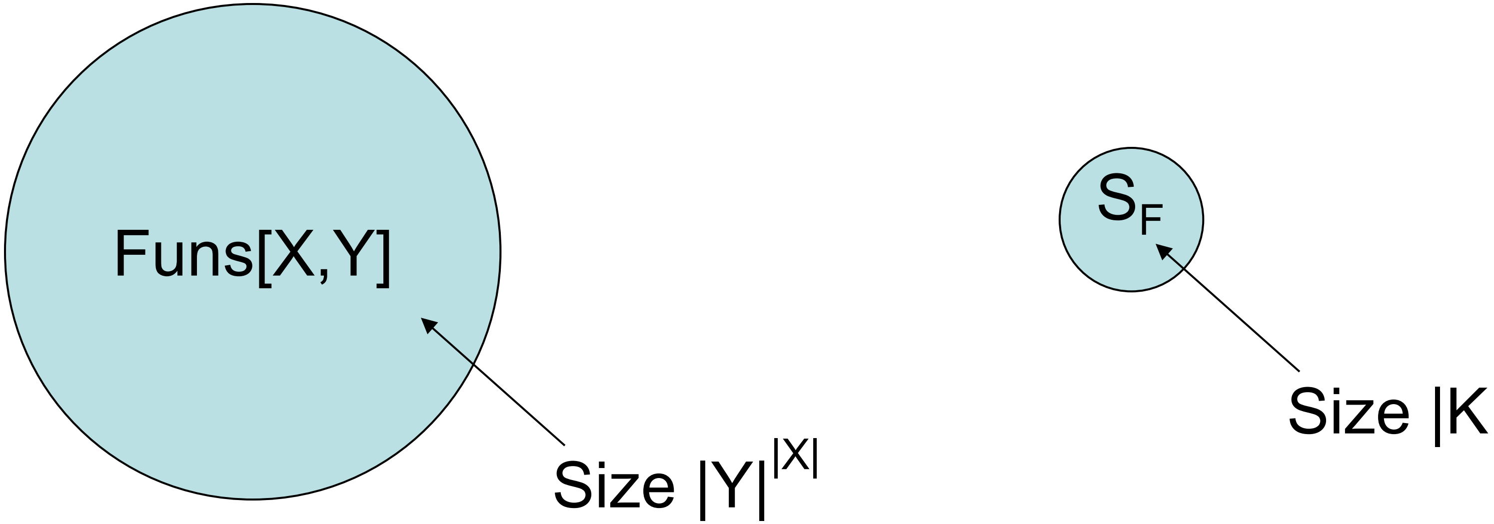

Secure PRFs

- Let $F: K \times X \to Y$ be a PRF

$$\begin{cases} \text{Funs}[X,Y]: & \text{the set of } \textbf{all} \text{ functions from } X \text{ to } Y \\ S_F = \{ F(k,\cdot) \text{ s.t. } k \in K \} & \subseteq \text{Funs}[X,Y] \end{cases}$$

- Intuition: a PRF is **secure** if
  a random function in Funs[X,Y] is indistinguishable from a random function in S_F

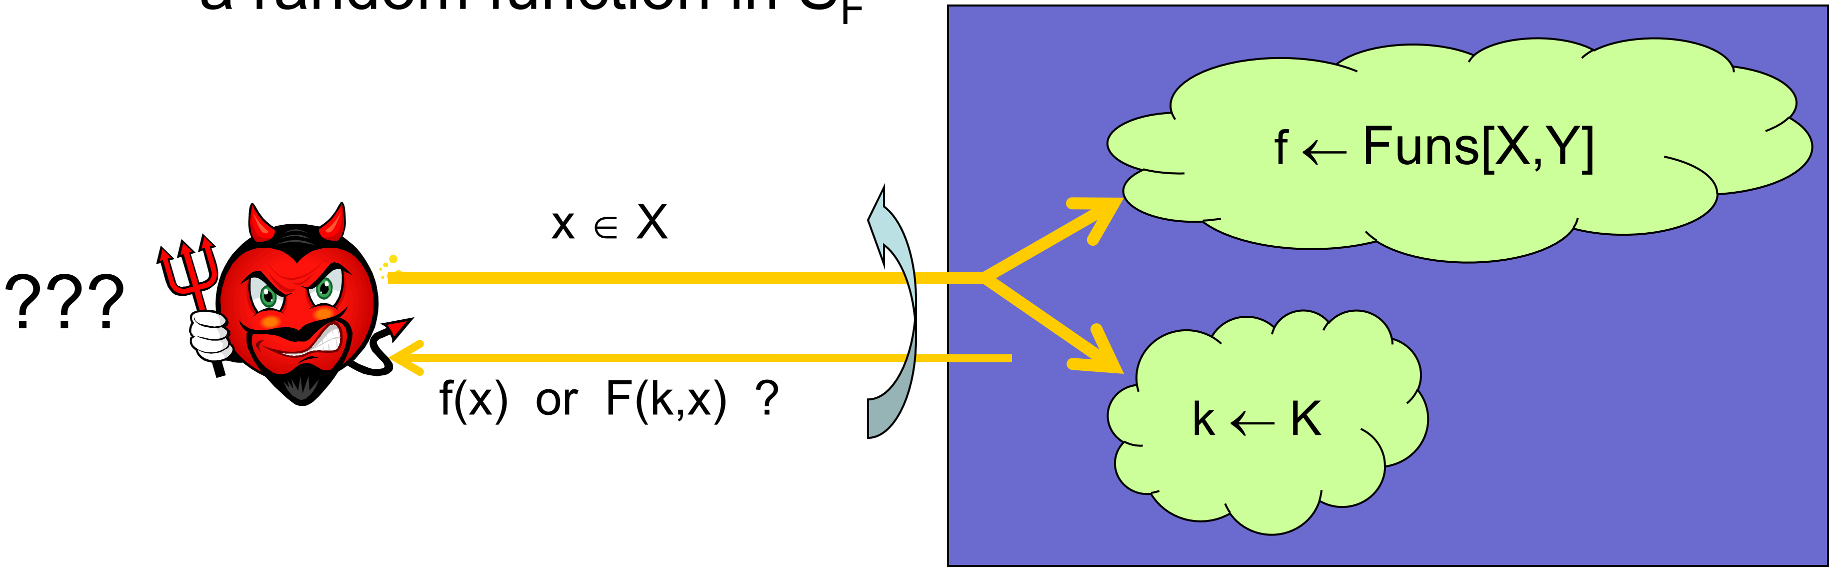

Secure PRF: defintion

- For b=0,1 define experiment EXP(b) as:

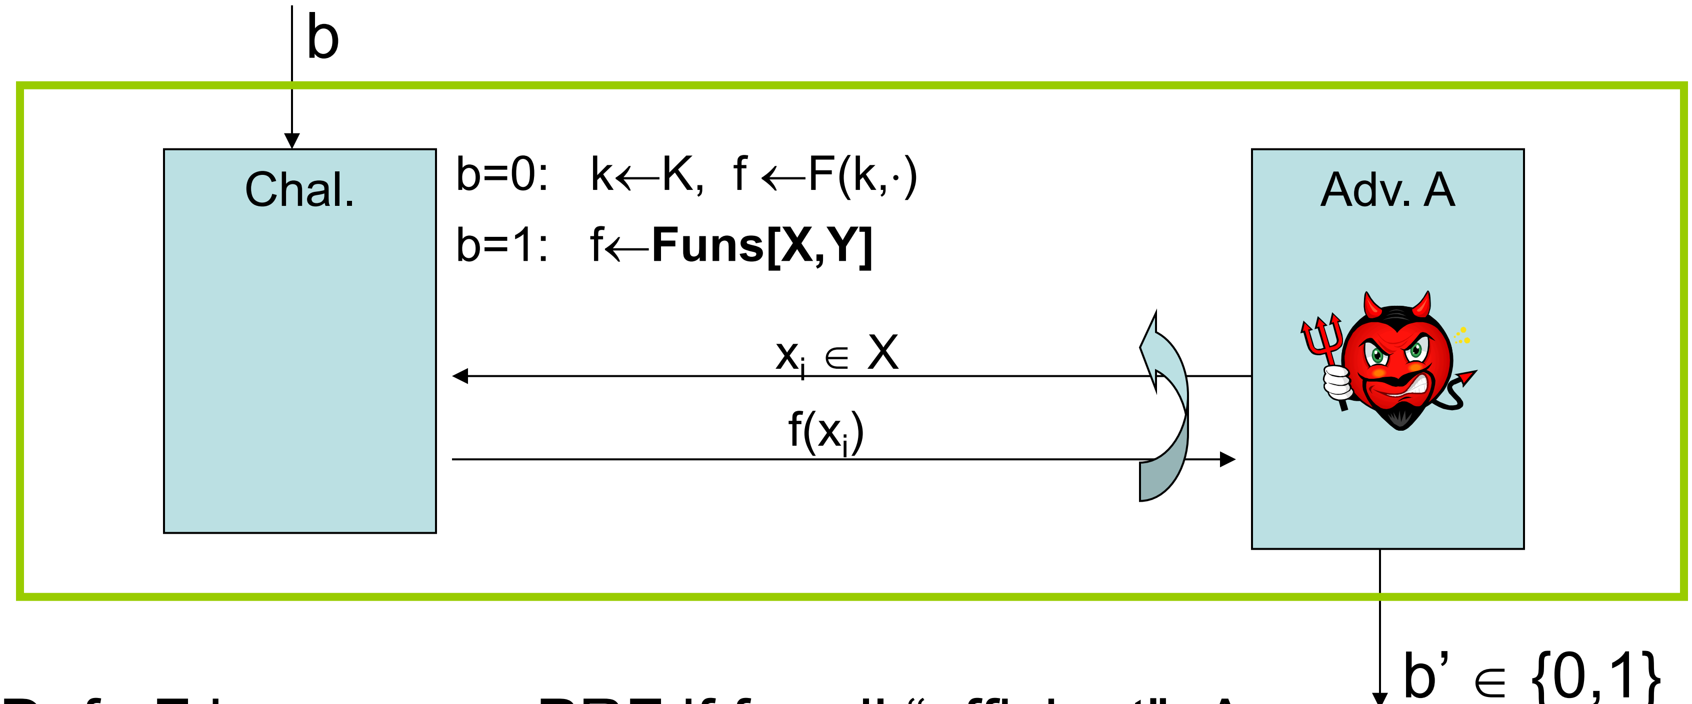

- Def: F is a secure PRF if for all "efficient" A:

$$Adv_{PRF}[A,F] = \left| Pr[EXP(0)=1] - Pr[EXP(1)=1] \right|$$

is "negligible."

Secure PRP

- For b=0,1 define experiment EXP(b) as:

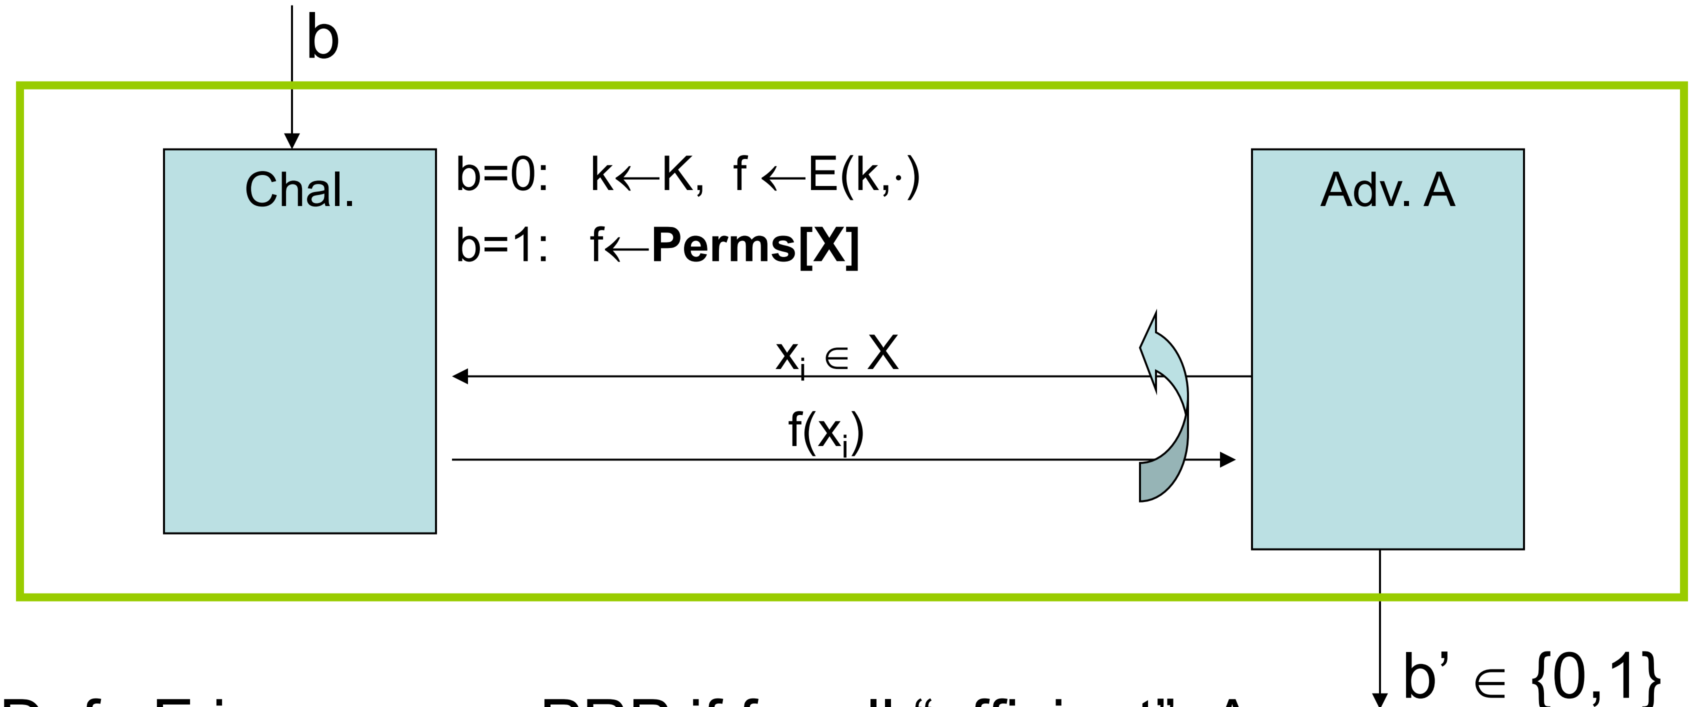

- Def: E is a secure PRP if for all "efficient" A:

$$\text{Adv}_{\text{PRP}}[A,E] = \left| \Pr[\text{EXP}(0)=1] - \Pr[\text{EXP}(1)=1] \right|$$

is "negligible."

Example secure PRPs

- Example secure PRPs: 3DES, AES, …

  AES_{256}: $K \times X \rightarrow X$ where $X = \{0,1\}^{128}$

  $K = \{0,1\}^{256}$

- AES_{256} PRP Assumption (example) :

  All explicit 2^{80}–time algs A have PRP Adv[A, $\mathbf{AES_{256}}$] $< 2^{-40}$

PRF Switching Lemma

Any secure PRP is also a secure PRF.

<u>Lemma</u>: Let E be a PRP over (K,X)

Then for any q-query adversary A:

$$\left| \mathrm{Adv}_{\mathrm{PRF}}[A,E] - \mathrm{Adv}_{\mathrm{PRP}}[A,E] \right| < q^2 / 2|X|$$

⇒ Suppose |X| is large so that $q^2 / 2|X|$ is "negligible"

Then $\mathrm{Adv}_{\mathrm{PRP}}[A,E]$ "negligible" ⇒ $\mathrm{Adv}_{\mathrm{PRF}}[A,E]$ "negligible"

Using PRPs and PRFs

- Goal: build "secure" encryption from a PRP.

- Security is always defined using two parameters:

  1. What "**power**" does adversary have?
     examples:
     - Adv sees only one ciphertext (one-time key)
     - Adv sees many PT/CT pairs (many-time key, CPA)

  2. What "**goal**" is adversary trying to achieve?
     examples:
     - Fully decrypt a challenge ciphertext.
     - Learn info about PT from CT (semantic security)

Incorrect use of a PRP

Electronic Code Book (ECB):

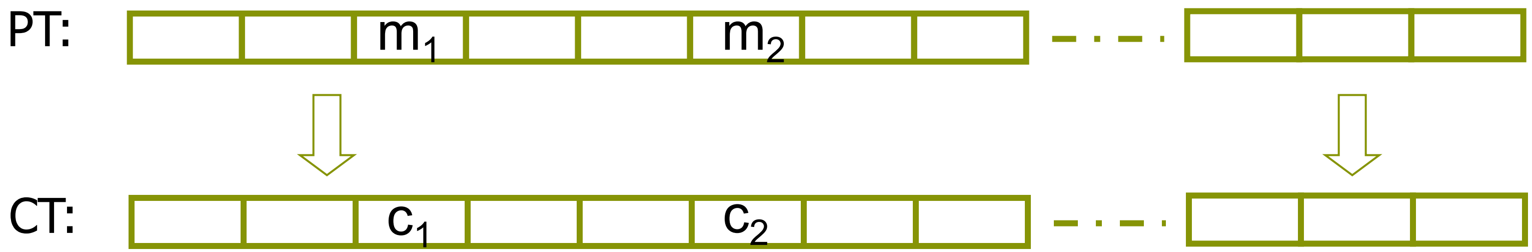

Problem:

– if $m_1 = m_2$ then $c_1 = c_2$

In pictures

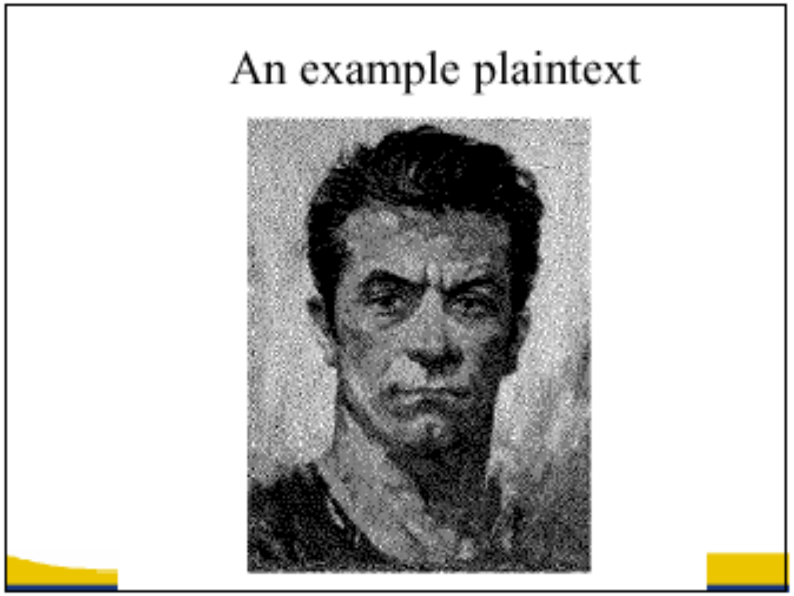

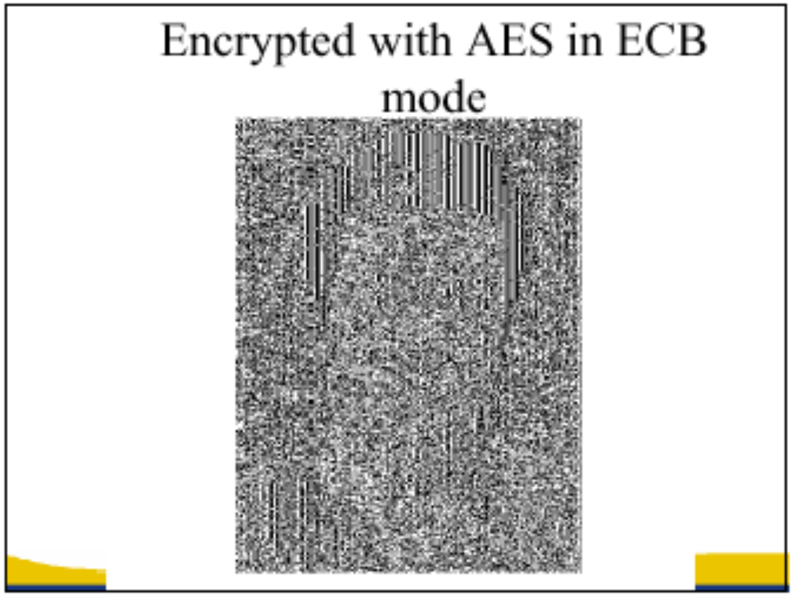

(courtesy B. Preneel)

Modes of Operation for One-time Use Key

Example application:

Encrypted email.   New key for every message.

Semantic Security for one-time key

- $\mathbb{E}$ = (E,D) a cipher defined over (K,M,C)
- For b=0,1 define EXP(b) as:

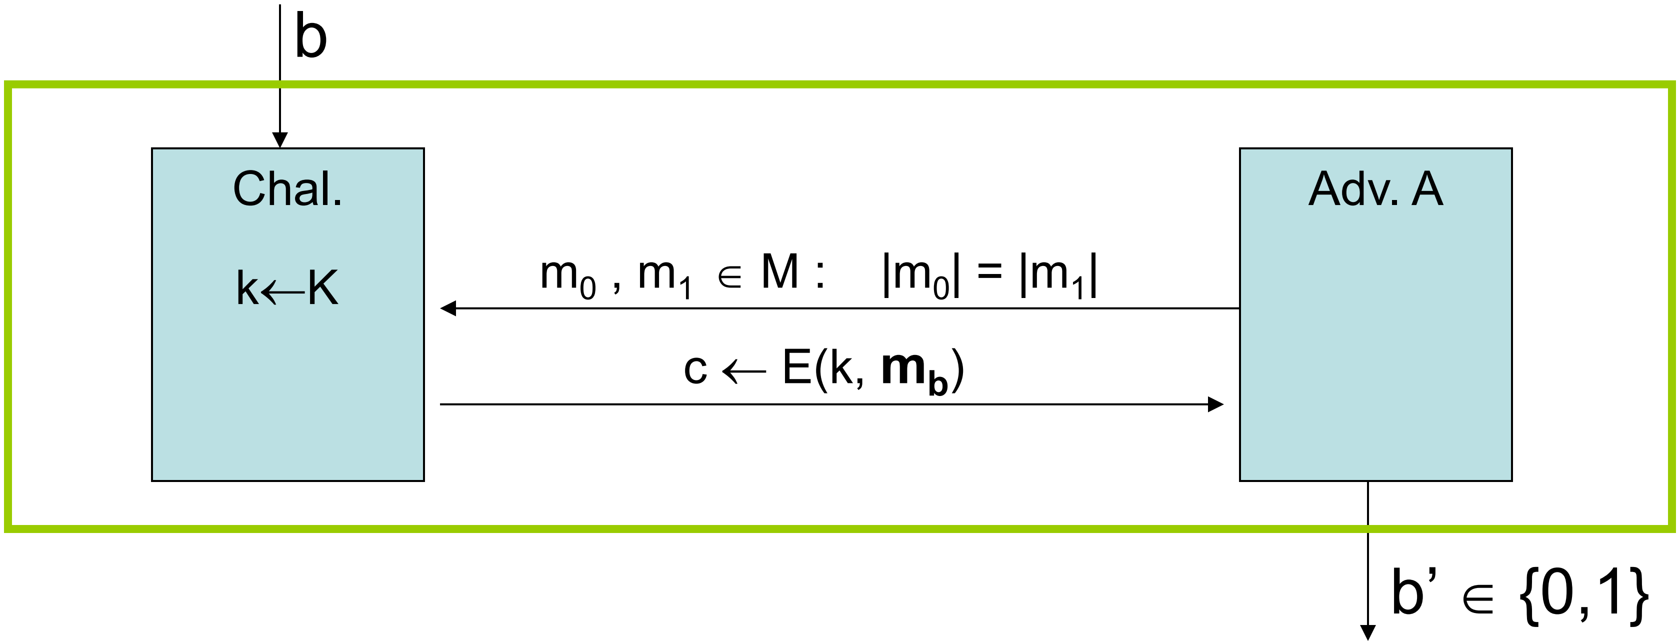

- Def: $\mathbb{E}$ is sem. sec. for one-time key if for all "efficient" A:

$$\text{Adv}_{SS}[A,\mathbb{E}] = \left| \Pr[\text{EXP}(0)=1] - \Pr[\text{EXP}(1)=1] \right|$$

is "negligible."

Semantic security (cont.)

Sem. Sec. ⇒ no "efficient" adversary learns info about PT from a **single** CT.

Example: suppose efficient A can deduce LSB of PT from CT.

Then $\mathbb{E}$ = (E,D) is not semantically secure.

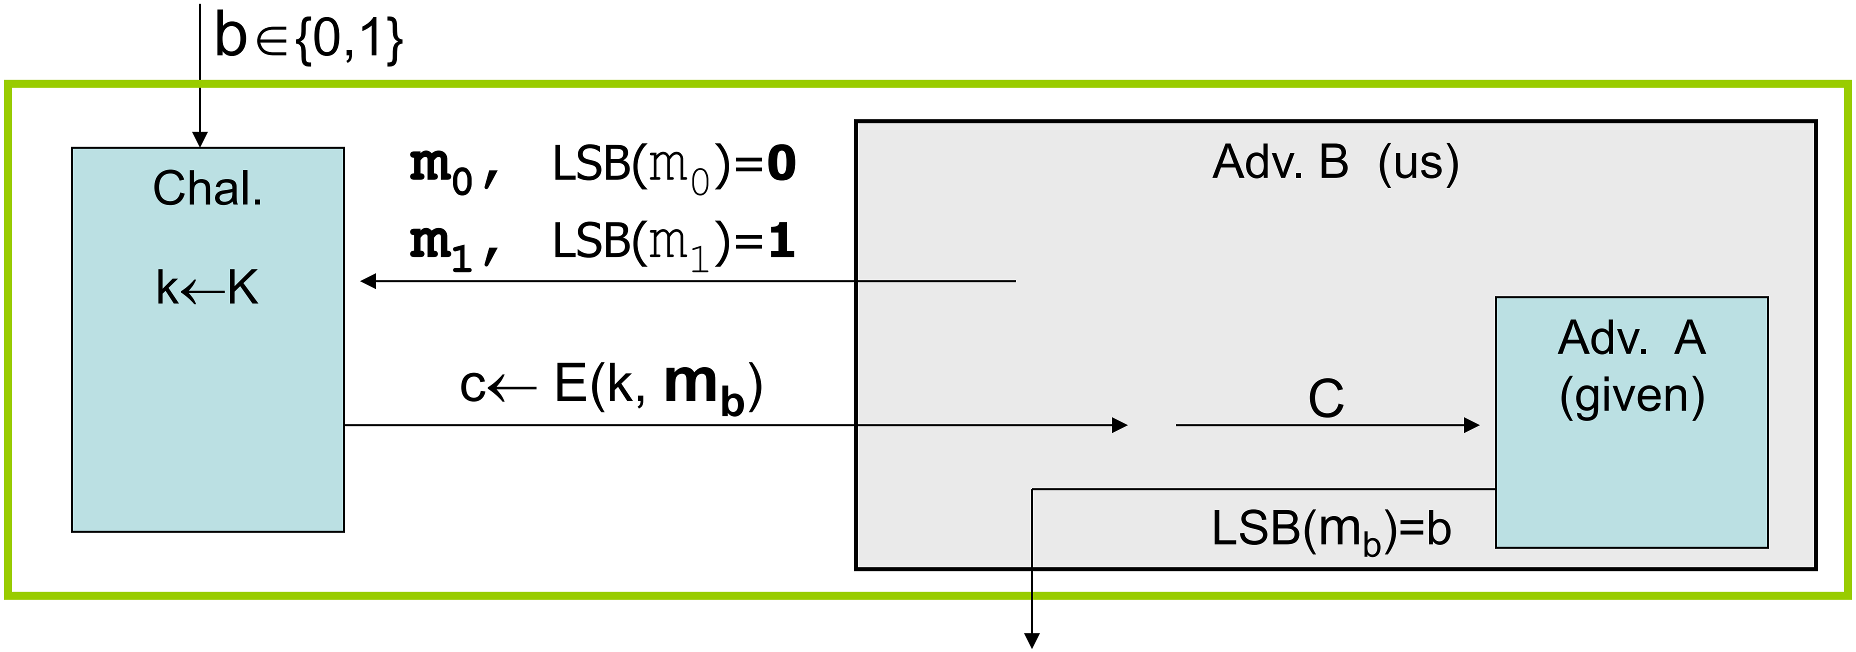

Then $\text{Adv}_{SS}[B, \mathbb{E}] = 1$ ⇒ $\mathbb{E}$ is not sem. sec.

Note: ECB is not Sem. Sec.

Electronic Code Book (ECB):

– Not semantically secure for messages that contain more than one block.

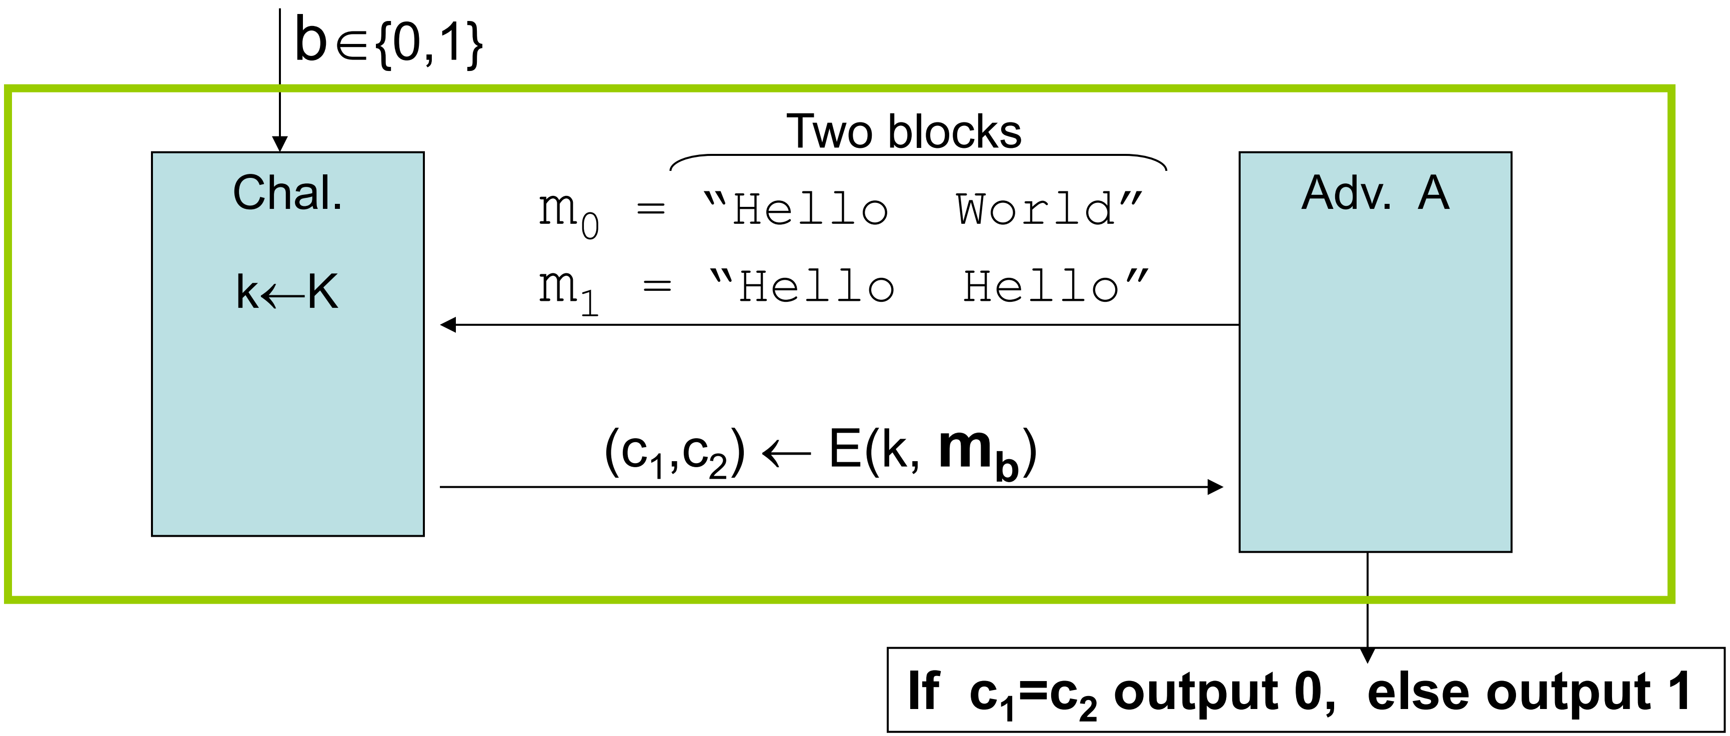

Then $\text{Adv}_{SS}[A, \text{ECB}] = 1$

Secure Constructions

Examples of sem. sec. systems:

1. $Adv_{SS}[A, OTP] = 0$ for **all** A

2. Deterministic counter mode from a PRF F :
   - $E_{DETCTR}(k,m)$ =

m[0]	m[1]	…	m[L]

$\oplus$

F(k,0)	**F(k,1)**	…	**F(k,L)**

---

c[0]	c[1]	…	c[L]

- Stream cipher built from PRF (e.g. AES, 3DES)

Det. counter-mode security

Theorem: For any $L>0$.

If F is a secure PRF over (K,X,X) then

E_{DETCTR} is sem. sec. cipher over (K,X^L,X^L).

In particular, for any adversary A attacking E_{DETCTR}

there exists a PRF adversary B s.t.:

$$Adv_{SS}[A, E_{DETCTR}] = 2 \cdot Adv_{PRF}[B, F]$$

$Adv_{PRF}[B, F]$ is negligible (since F is a secure PRF)

$\Rightarrow$ $Adv_{SS}[A, E_{DETCTR}]$ must be negligible.

Modes of Operation for Many-time Key

Example applications:

1. File systems: Same AES key used to encrypt many files.
2. IPsec: Same AES key used to encrypt many packets.

Semantic Security for many-time key (CPA security)

Cipher $\mathbb{E} = (E,D)$ defined over (K,M,C).
For $b=0,1$ define $EXP(b)$ as:

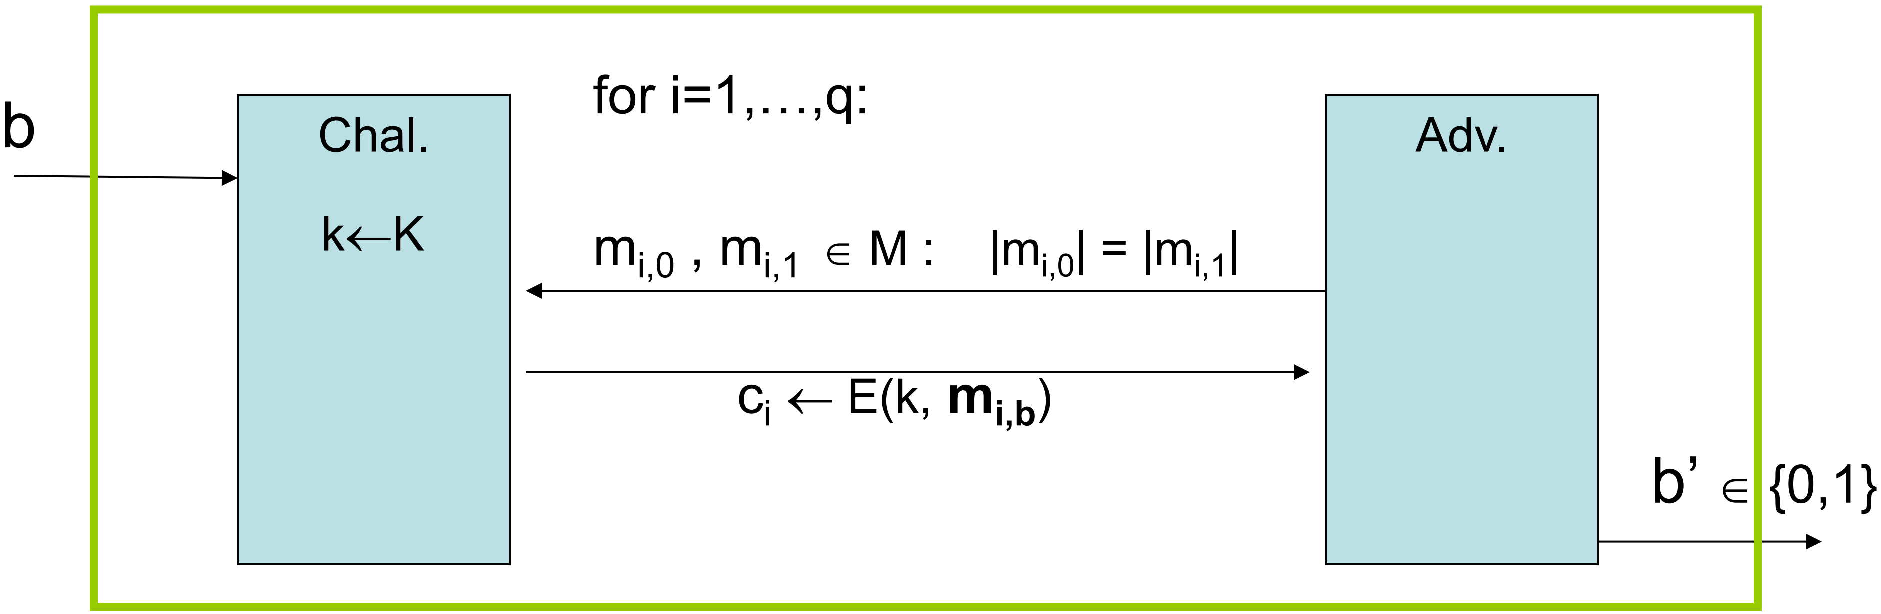

if adv. wants $c = E(k, m)$ it queries with $m_{j,0} = m_{j,1} = m$

Def: $\mathbb{E}$ is sem. sec. under CPA if for all "efficient" A:

$$\text{Adv}_{\text{CPA}}[A,\mathbb{E}] = \left| \Pr[EXP(0)=1] - \Pr[EXP(1)=1] \right|$$

is "negligible."

Security for many-time key

<u>Fact:</u> stream ciphers are insecure under CPA.

- More generally: if E(k,m) always produces same ciphertext, then cipher is insecure under CPA.

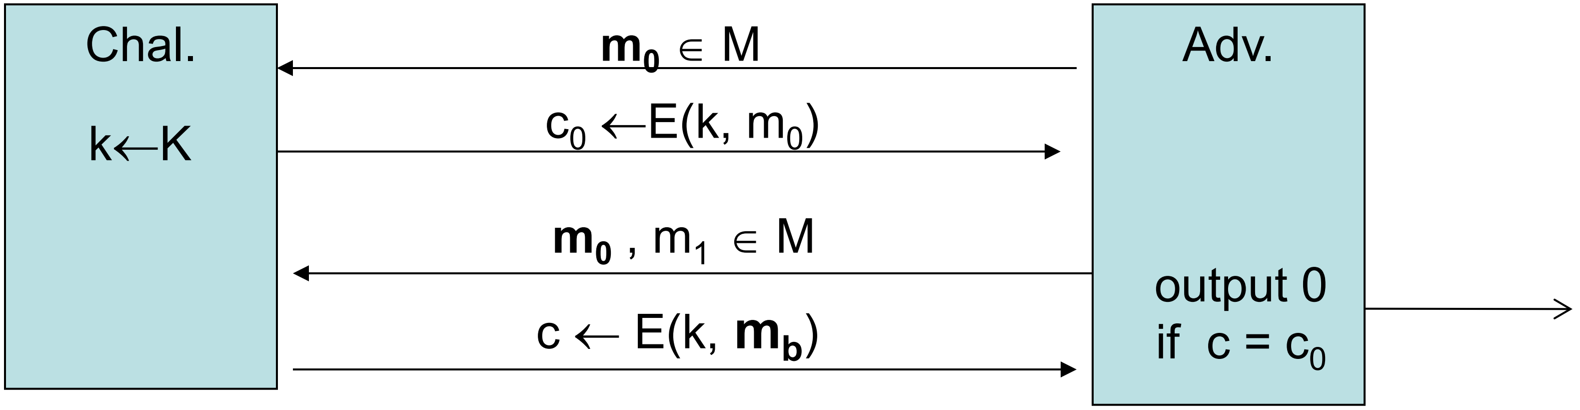

If secret key is to be used multiple times ⇒

given the same plaintext message twice,
the encryption alg. must produce different outputs.

Nonce-based Encryption

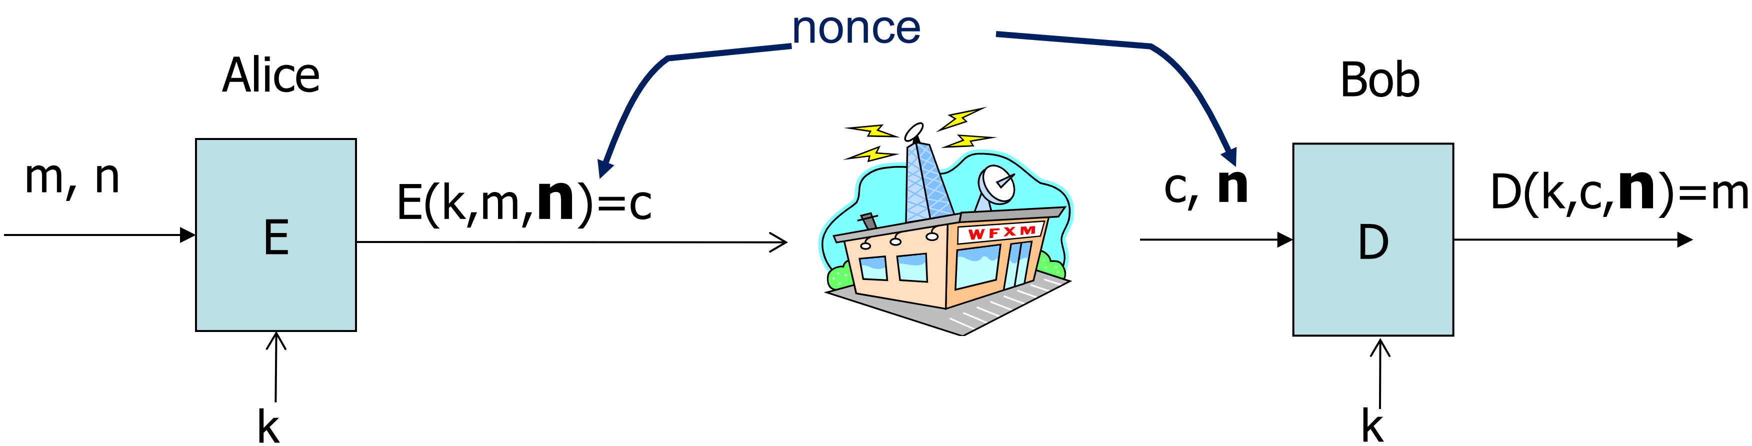

nonce n: a value that changes from msg to msg

(k,n) pair <u>never</u> used more than once

- <u>method 1</u>: encryptor chooses a random nonce, $n \leftarrow \mathcal{N}$
- <u>method 2</u>: nonce is a counter (e.g. packet counter)
  - used when encryptor keeps state from msg to msg
  - if decryptor has same state, need not send nonce with CT

Construction 1: CBC with random nonce

Cipher block chaining with a random IV (IV = nonce)

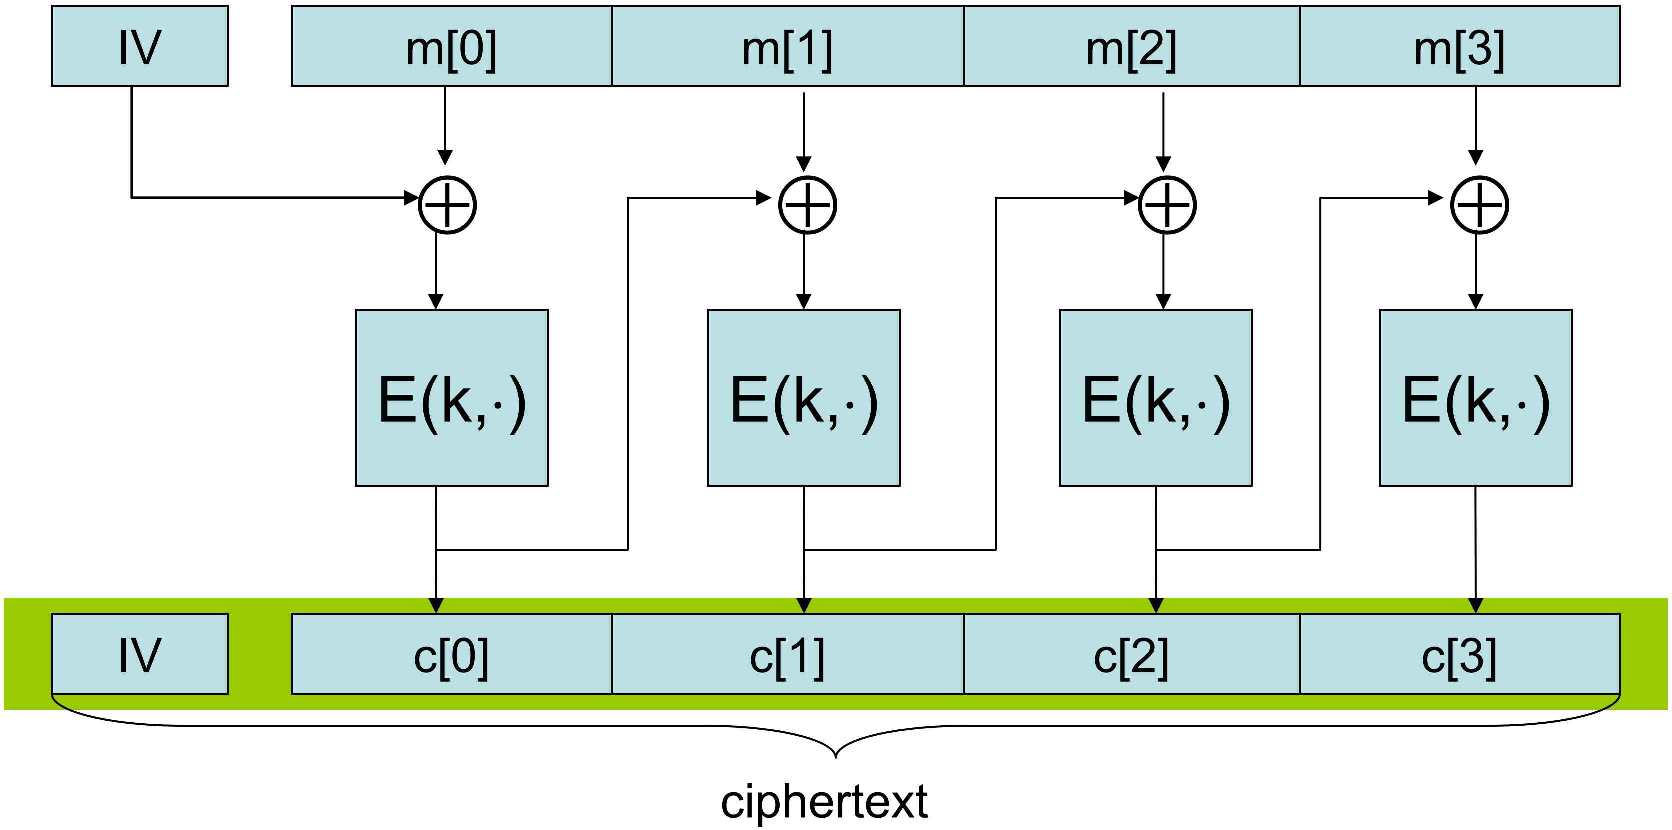

note: CBC where attacker can predict the IV is not CPA-secure. HW.

CBC: CPA Analysis

CBC Theorem: For any $L>0$,

If E is a secure PRP over (K,X) then

E_{CBC} is a sem. sec. under CPA over (K, X^L, X^{L+1}).

In particular, for a q-query adversary A attacking E_{CBC} there exists a PRP adversary B s.t.:

$$Adv_{CPA}[A, E_{CBC}] \leq 2 \cdot Adv_{PRP}[B, E] + 2\, q^2 L^2 / |X|$$

Note: CBC is only secure as long as $q^2L^2 \ll |X|$

Construction 1': CBC with **unique** nonce

Cipher block chaining with <u>unique</u> IV (IV = nonce)

unique IV means: (key,IV) pair is used for only one message

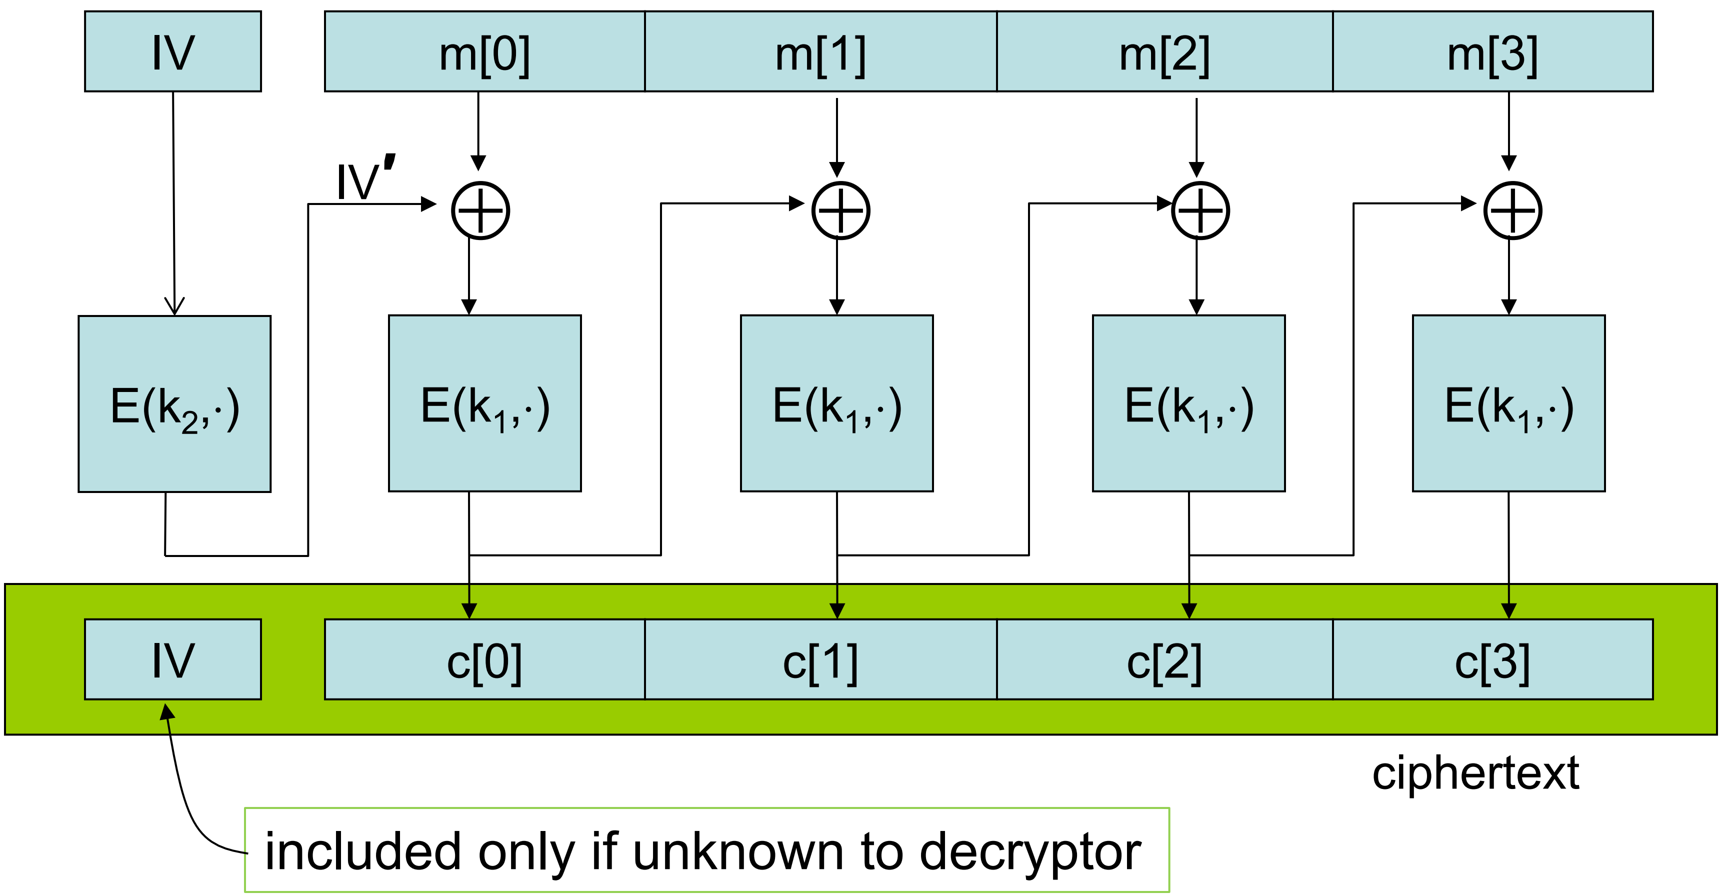

A CBC technicality: padding

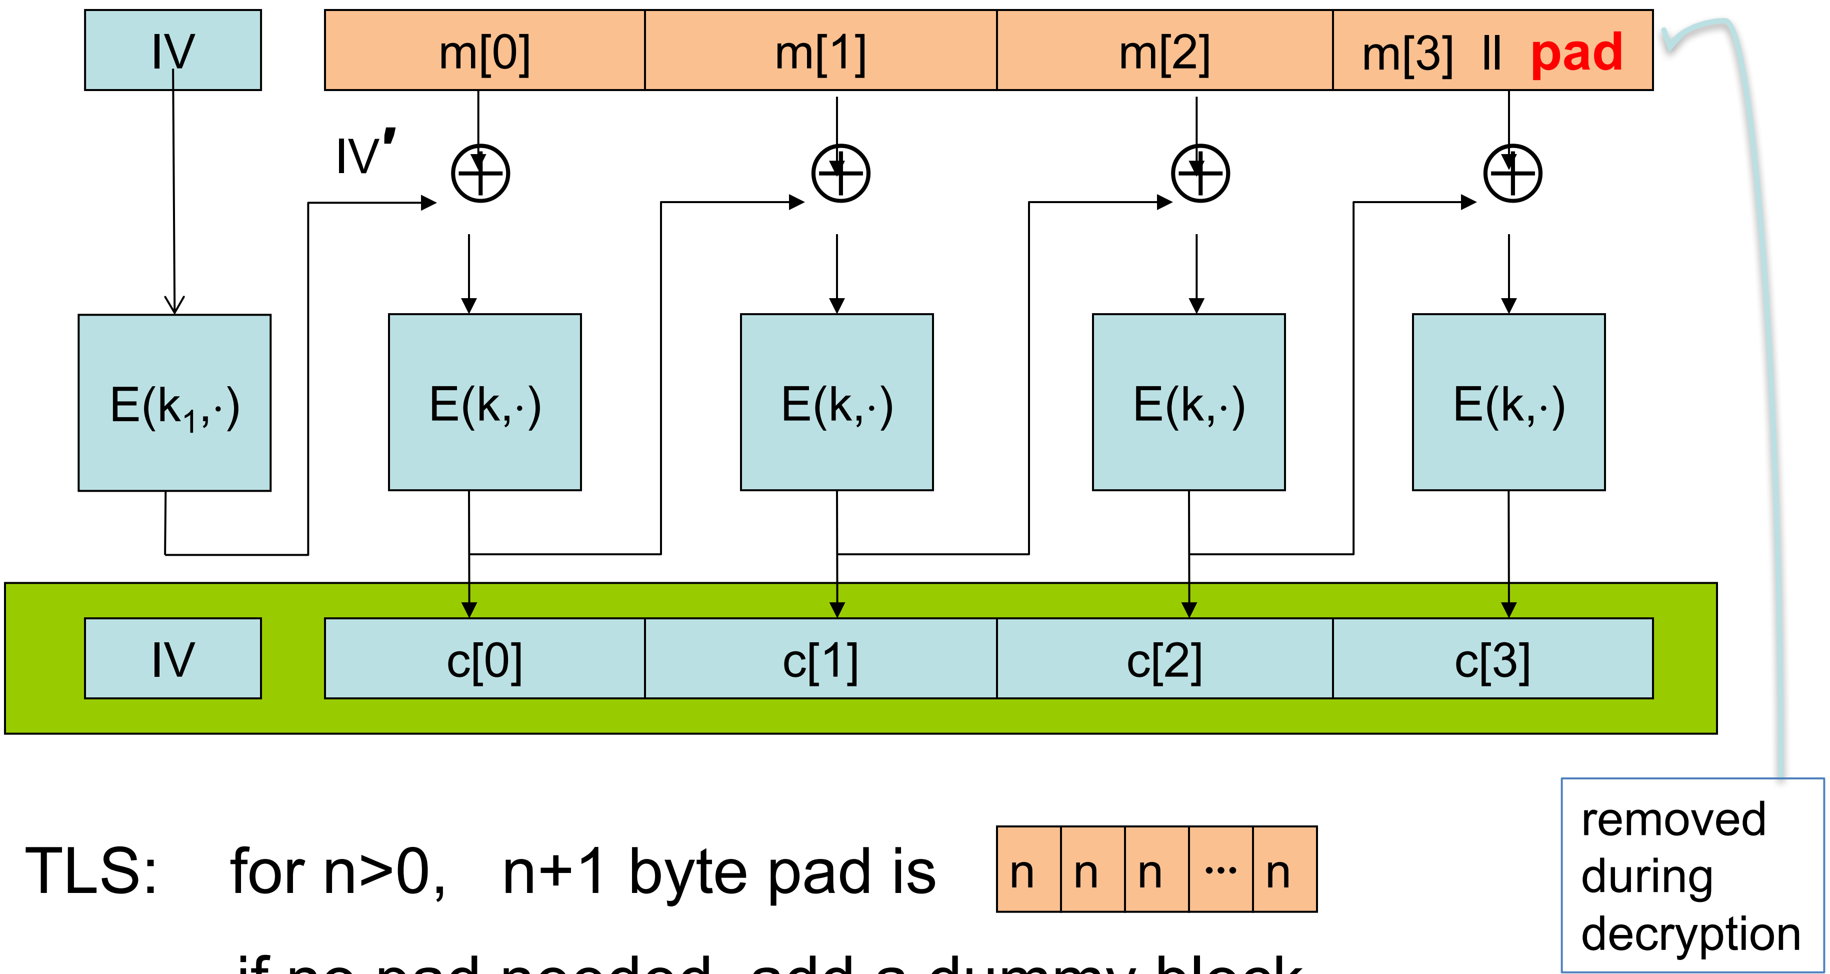

TLS: for n>0, n+1 byte pad is | n | n | n | ··· | n |

if no pad needed, add a dummy block

Construction 2: rand ctr-mode

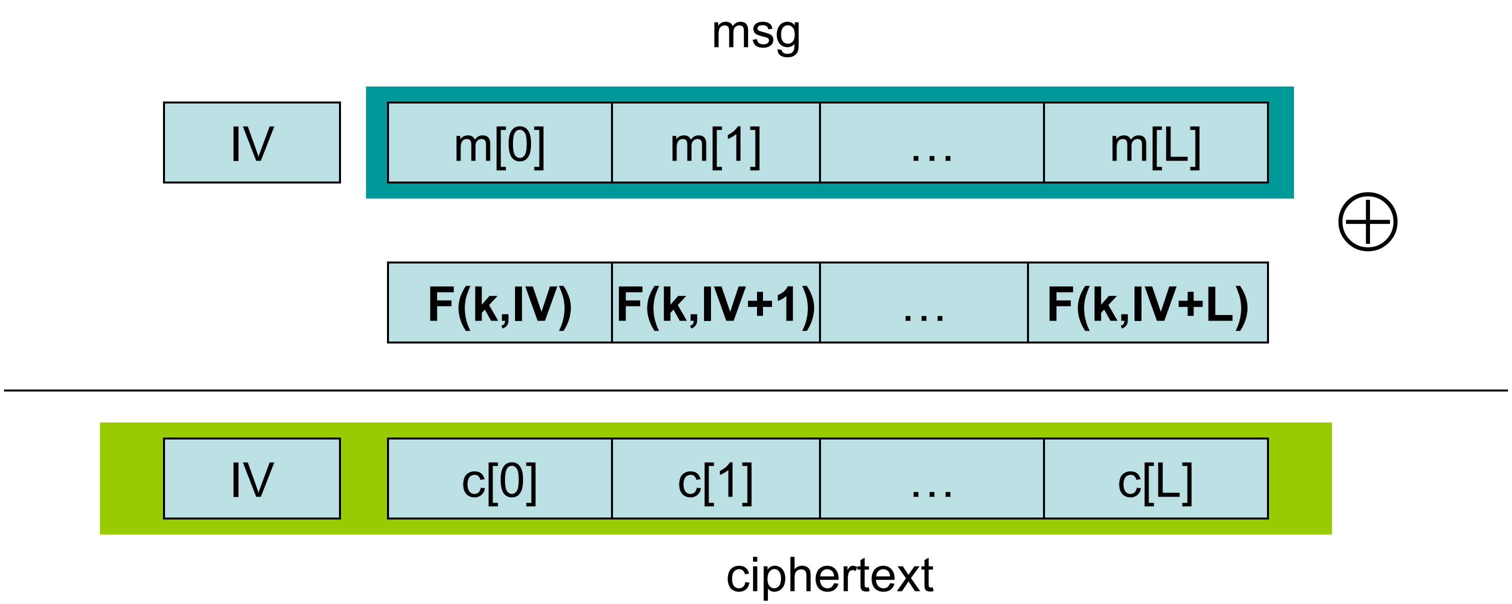

IV - chosen at random for every message

note: parallelizable (unlike CBC)

Construction 2': nonce ctr-mode

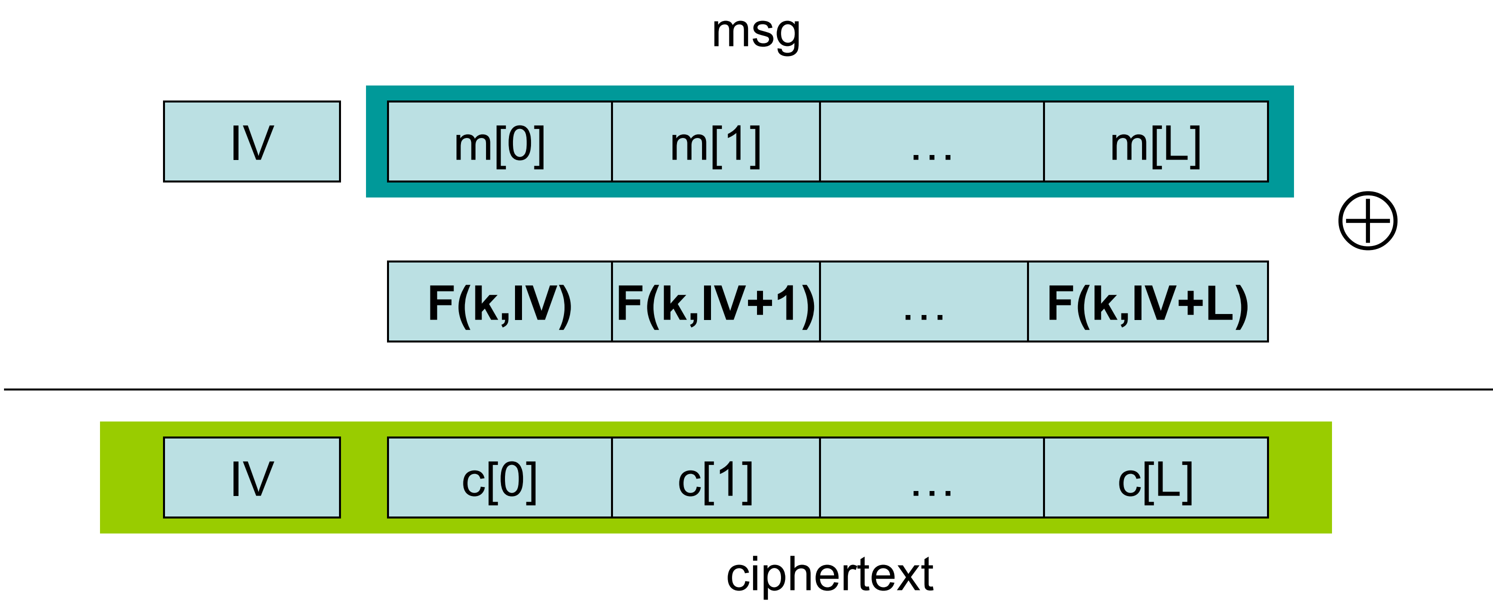

To ensure F(K,x) is never used more than once, choose IV as:

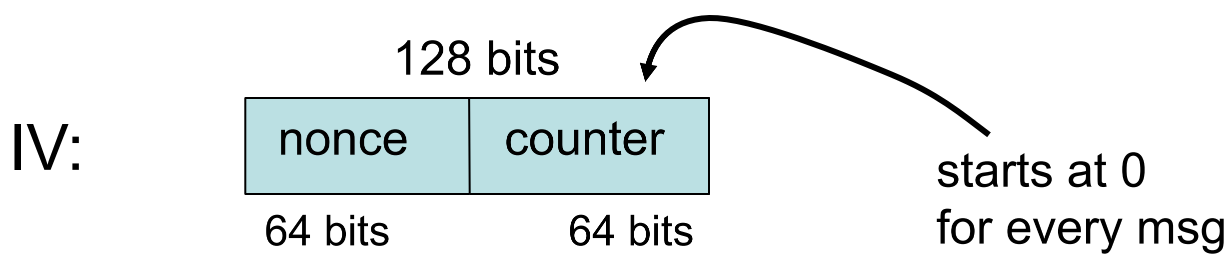

rand ctr-mode: CPA analysis

Randomized counter mode: random IV.

<u>Counter-mode Theorem</u>: For any L>0,

If F is a secure PRF over (K,X,X) then

E_{CTR} is a sem. sec. under CPA over (K, X^L, X^{L+1}).

In particular, for a q-query adversary A attacking E_{CTR} there exists a PRF adversary B s.t.:

$$\text{Adv}_{\text{CPA}}[A, E_{CTR}] \leq 2 \cdot \text{Adv}_{\text{PRF}}[B, F] + 2\, q^2 L / |X|$$

<u>Note</u>: ctr-mode only secure as long as $q^2 L \ll |X|$

Better then CBC !

An example

$$\text{Adv}_{\text{CPA}}[A, E_{\text{CTR}}] \leq 2 \cdot \text{Adv}_{\text{PRF}}[B, E] + \mathbf{2\, q^2\, L / |X|}$$

q = # messages encrypted with k , L = length of max msg

Suppose we want $\mathbf{Adv_{CPA}}[A, E_{CTR}] \leq q^2 L / |X| \leq \mathbf{1/ 2^{32}}$

- AES: $|X| = 2^{128} \Rightarrow q L^{1/2} < 2^{48}$

So, after 2^{32} CTs each of len 2^{32}, must change key

(total of 2^{64} AES blocks)

Comparison: ctr vs. CBC

	CBC	**ctr mode**
uses	PRP	PRF
parallel processing	No	Yes
Security of rand. enc.	q^2 L^2 << |X|	q^2 L << |X|
dummy padding block	Yes	No
1 byte msgs (nonce-based)	16x expansion	no expansion

(for CBC, dummy padding block can be avoided using ciphertext stealing)

Summary

PRPs and PRFs: a useful abstraction of block ciphers.

We examined two security notions:

1. Semantic security against one-time CPA.
2. Semantic security against many-time CPA.

Note: neither mode ensures data integrity.

Stated security results summarized in the following table:

Power / Goal	**one-time key**	**Many-time key (CPA)**	**CPA and CT integrity**
Sem. Sec.	steam-ciphers det. ctr-mode	rand CBC rand ctr-mode	later