

COS433/Math 473: Cryptography

Mark Zhandry

Princeton University

Spring 2020

Announcements/Reminders

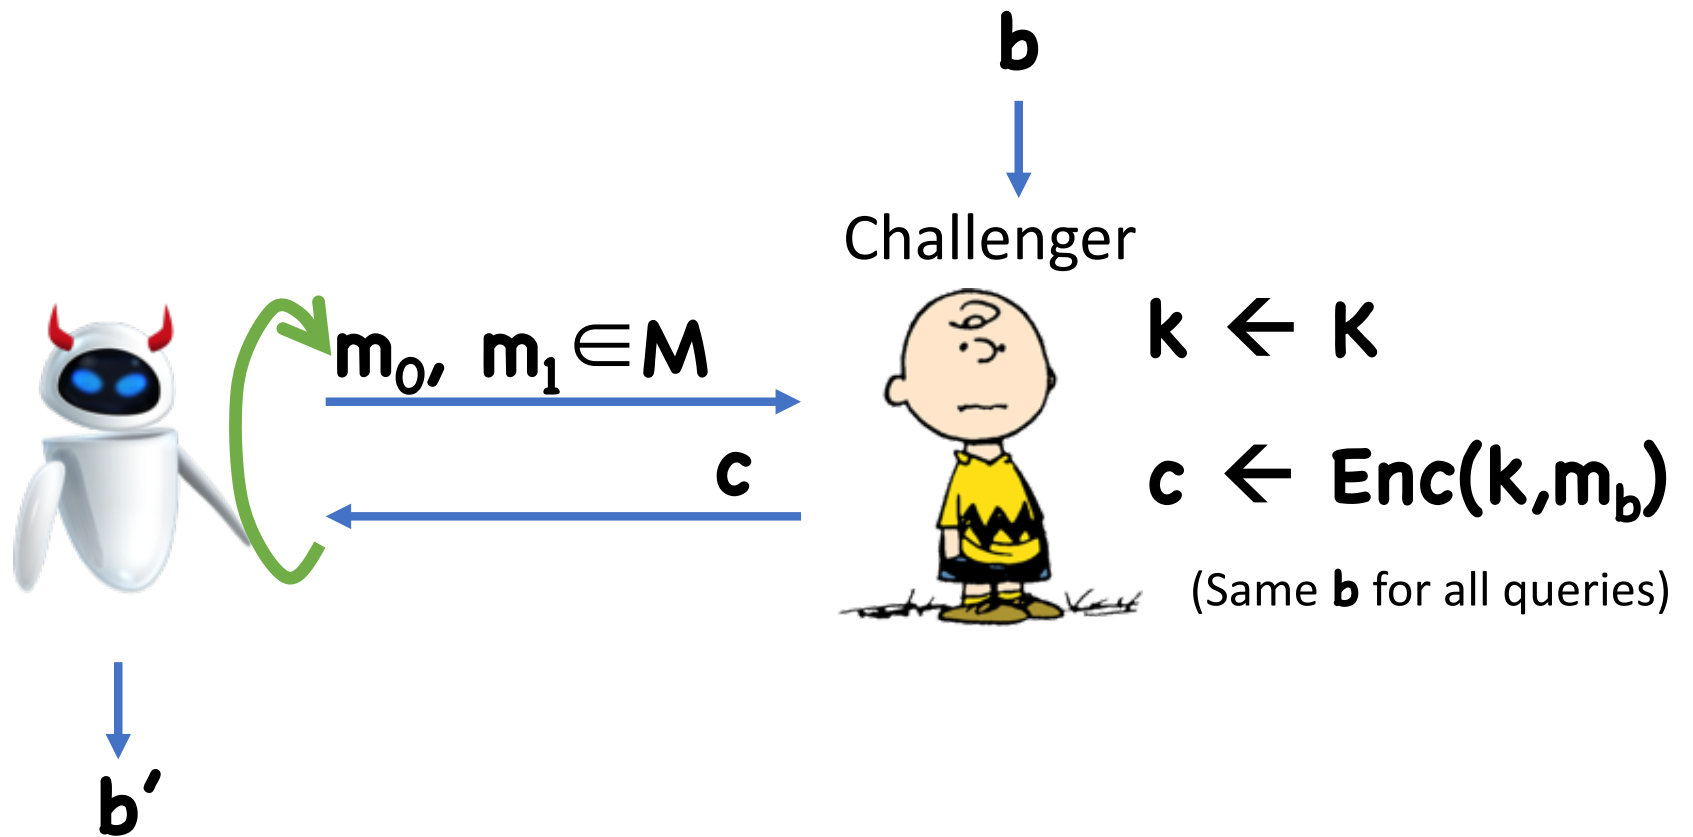
HW2 Due TODAY

HW3 Due March 5th

PR1 Due March 10th


Previously on COS 433...

Left-or-Right Experiment



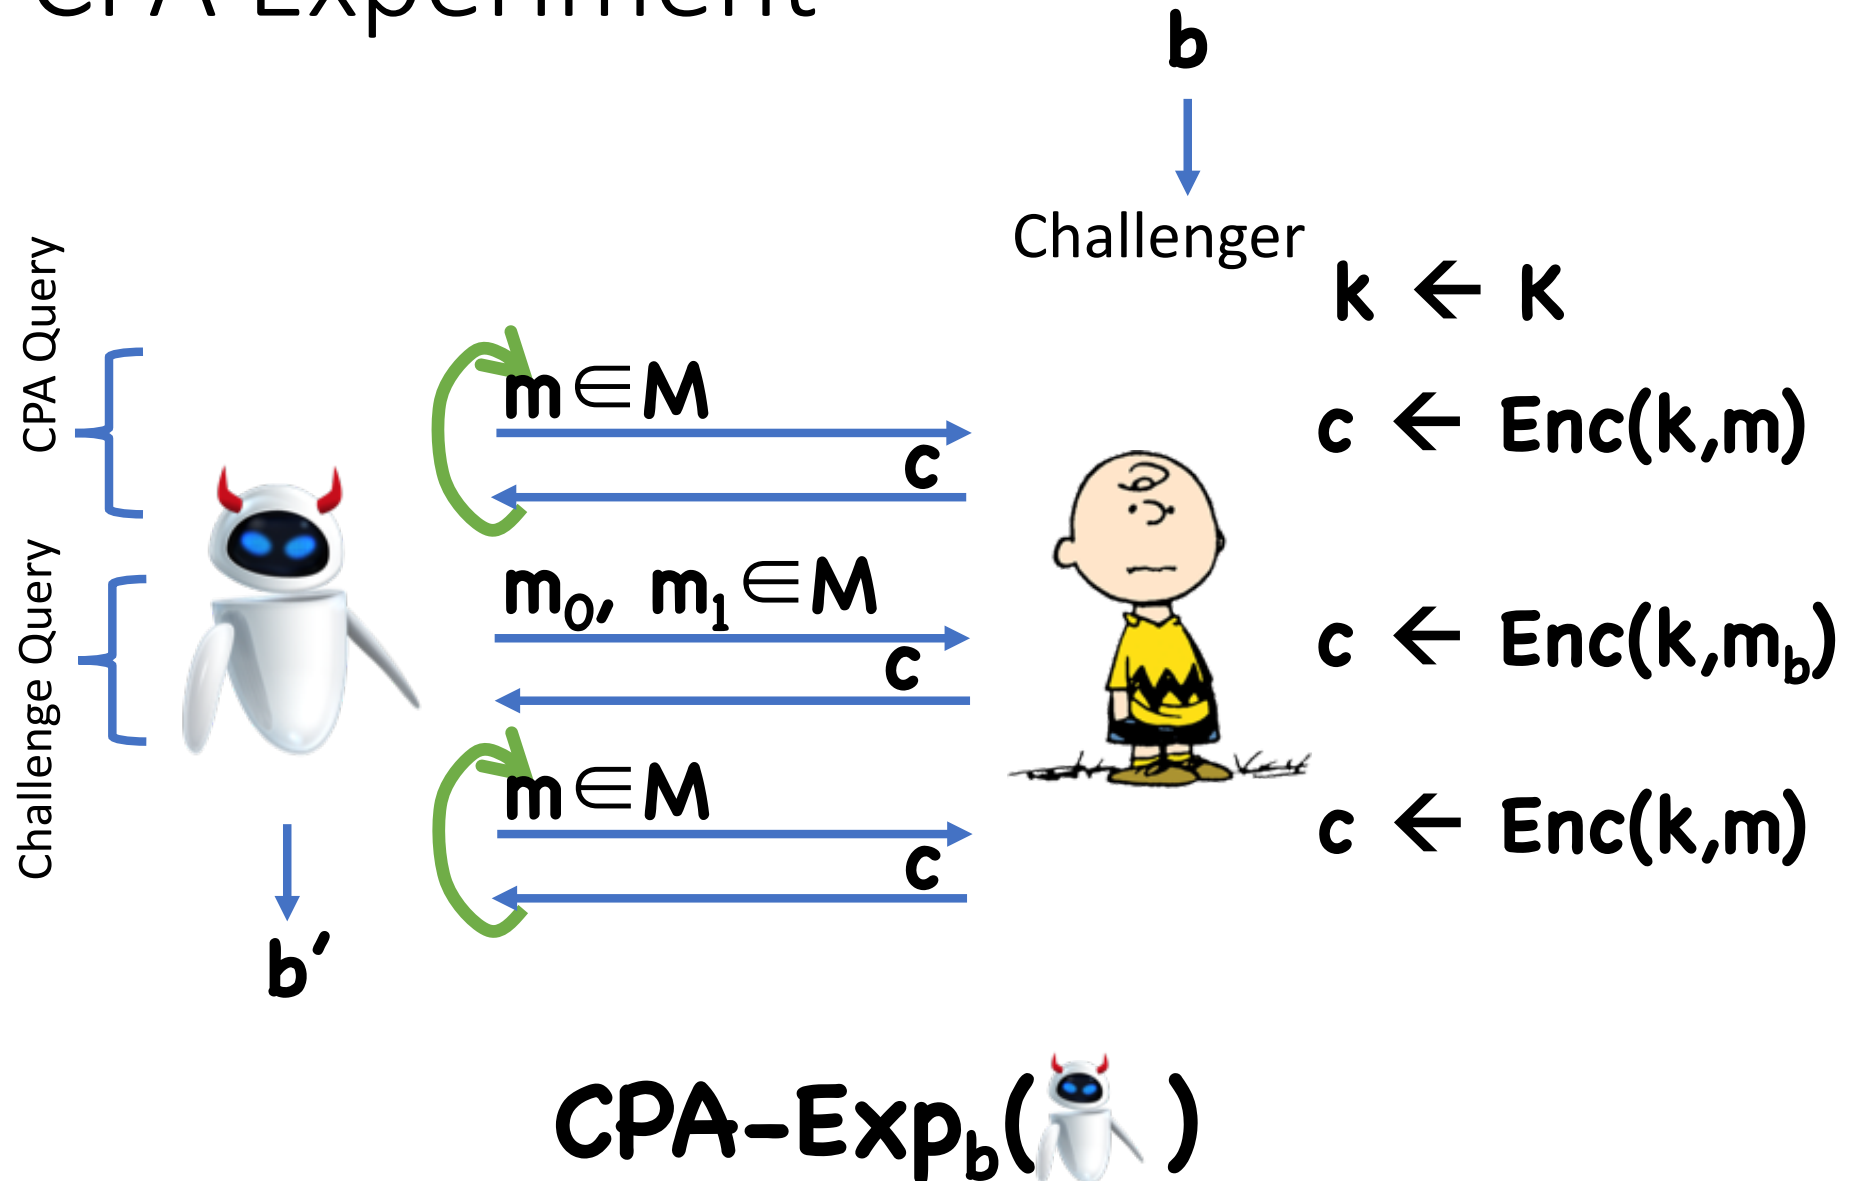
$\text{LoR-Exp}_b(\text{robot}, \lambda)$

LoR Security Definition

Definition: (Enc, Dec) has **Left-or-Right indistinguishability** if, for all  running in polynomial time, \exists negligible ϵ such that:

$$\left| \Pr[1 \leftarrow \text{LoR-Exp}_0(\text{robot}, \lambda)] - \Pr[1 \leftarrow \text{LoR-Exp}_1(\text{robot}, \lambda)] \right| \leq \epsilon(\lambda)$$

CPA Experiment



Generalized CPA Experiment

b



Challenger

$$k \leftarrow K$$

$$c \leftarrow \text{Enc}(k, m)$$


$$c \leftarrow \text{Enc}(k, m_b)$$

$$c \leftarrow \text{Enc}(k, m)$$



b'

Queries in any order



$$\text{GCPA-Exp}_b(\text{robot}, \lambda)$$

Equivalences

Theorem:

Left-or-Right indistinguishability



CPA-security



Generalized CPA-security

Therefore, you can use whichever notion you like best

Pseudorandom Functions

Functions that “look like” random functions


Syntax:

- Key space \mathbf{K}_λ
- Domain \mathbf{X}_λ
- Co-domain/range \mathbf{Y}_λ
- Function $\mathbf{F}:\mathbf{K}_\lambda \times \mathbf{X}_\lambda \rightarrow \mathbf{Y}_\lambda$

Correctness: \mathbf{F} is a function (deterministic)

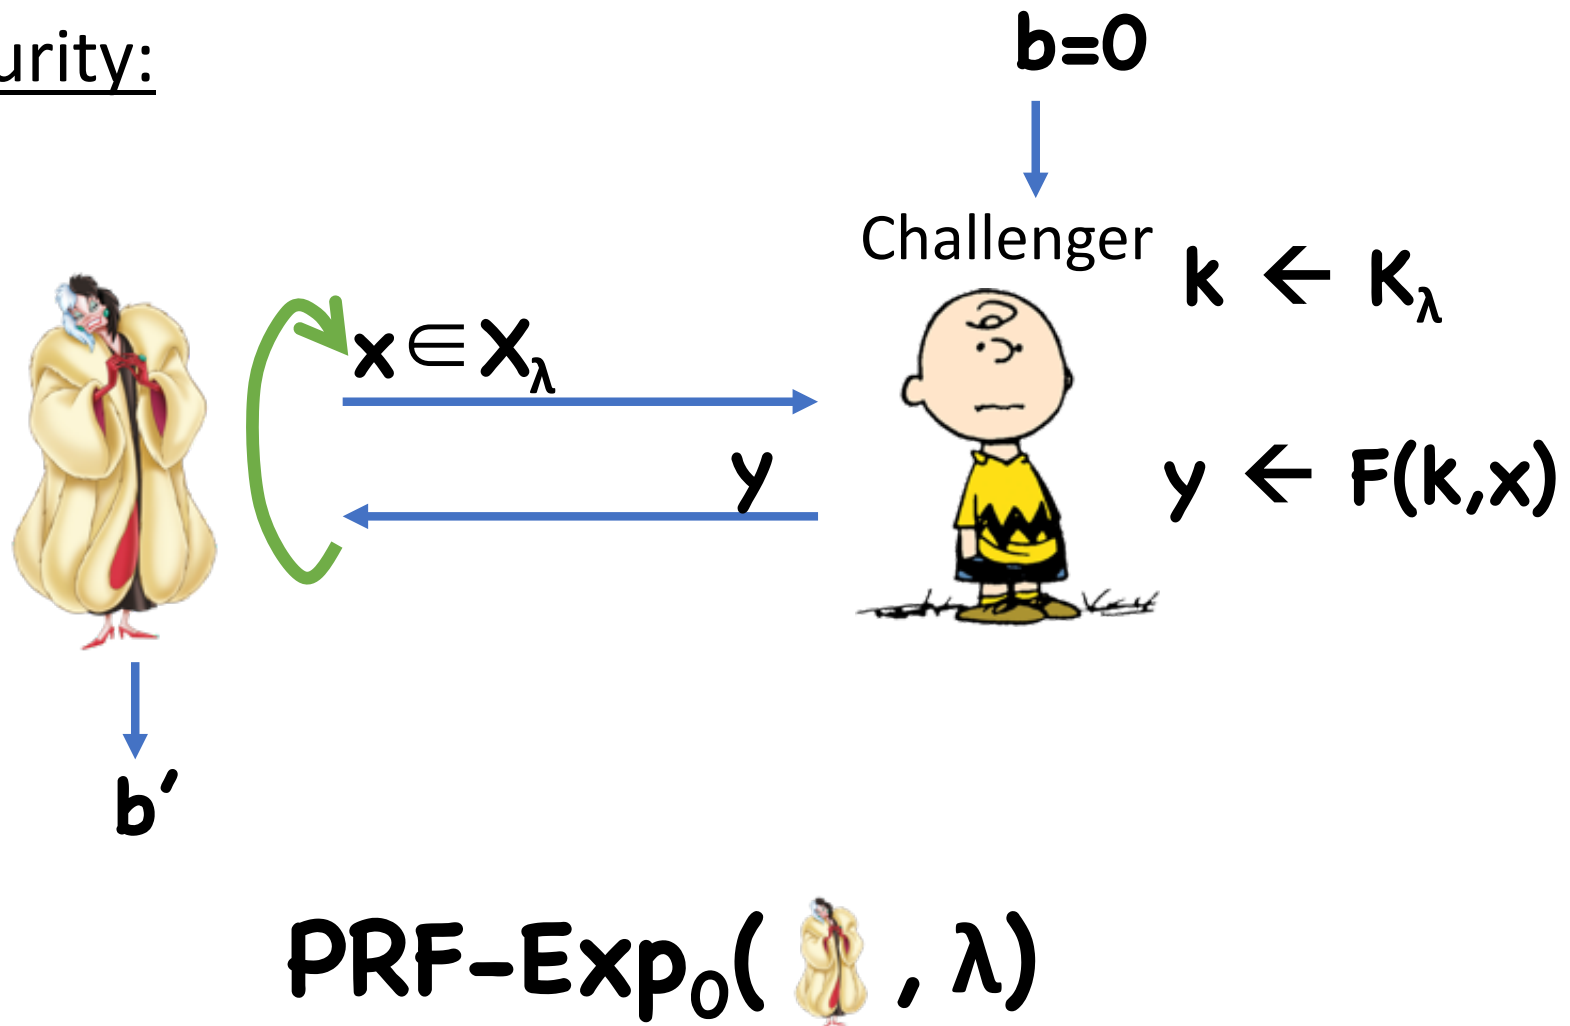
Pseudorandom Functions

Security:



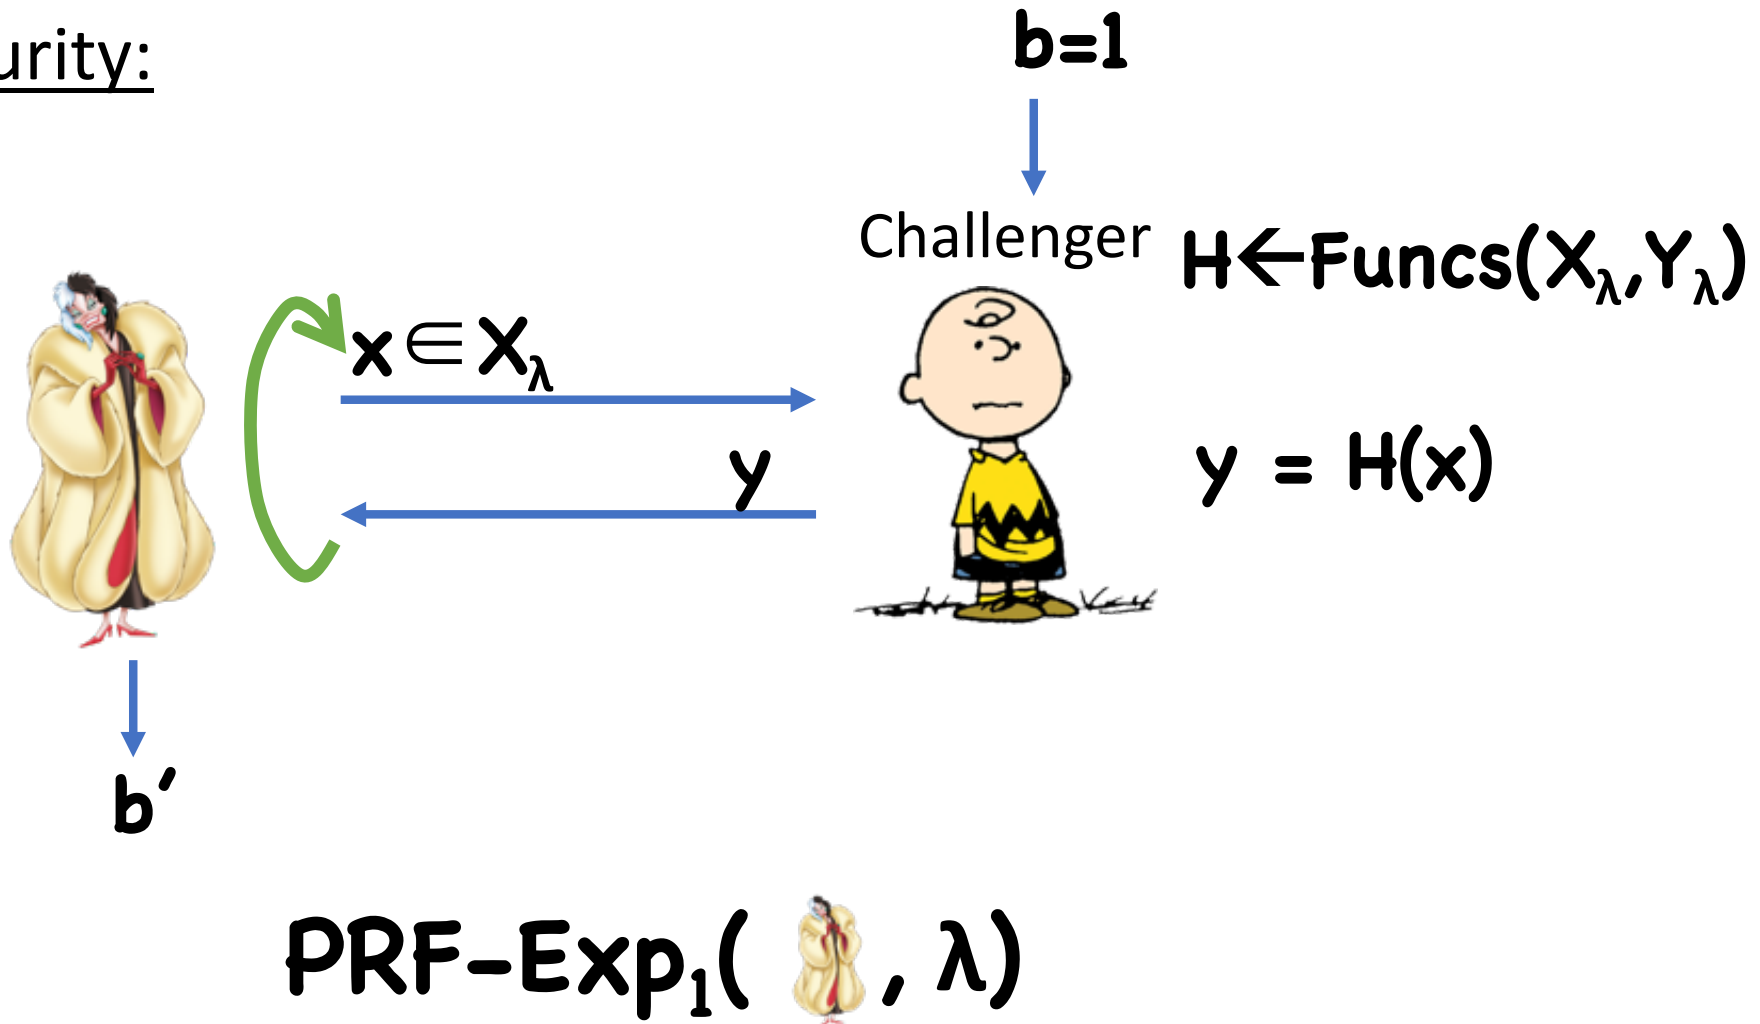
Pseudorandom Functions

Security:



Pseudorandom Functions

Security:



Using PRFs to Build Encryption

Enc(k, m):

- Choose random $r \leftarrow X_\lambda$
- Compute $y \leftarrow F(k, r)$
- Compute $c \leftarrow y \oplus m$
- Output (r, c)


Correctness:

- $y' = y$ since F is deterministic
- $m' = c \oplus y = y \oplus m \oplus y = m$

Dec(k, (r, c)):

- Compute $y' \leftarrow F(k, r)$
- Compute and output $m' \leftarrow c \oplus y'$

Using PRFs to Build Encryption



Today: More on PRFs

Security

Theorem: If \mathbf{F} is a secure PRF with domain \mathbf{X}_λ and $|\mathbf{X}_\lambda|$ is superpoly, then **(Enc,Dec)** is **LoR** secure.


Proof

Assume toward contradiction that there exists a  breaking **(Enc, Dec)**

Hybrids...

Proof

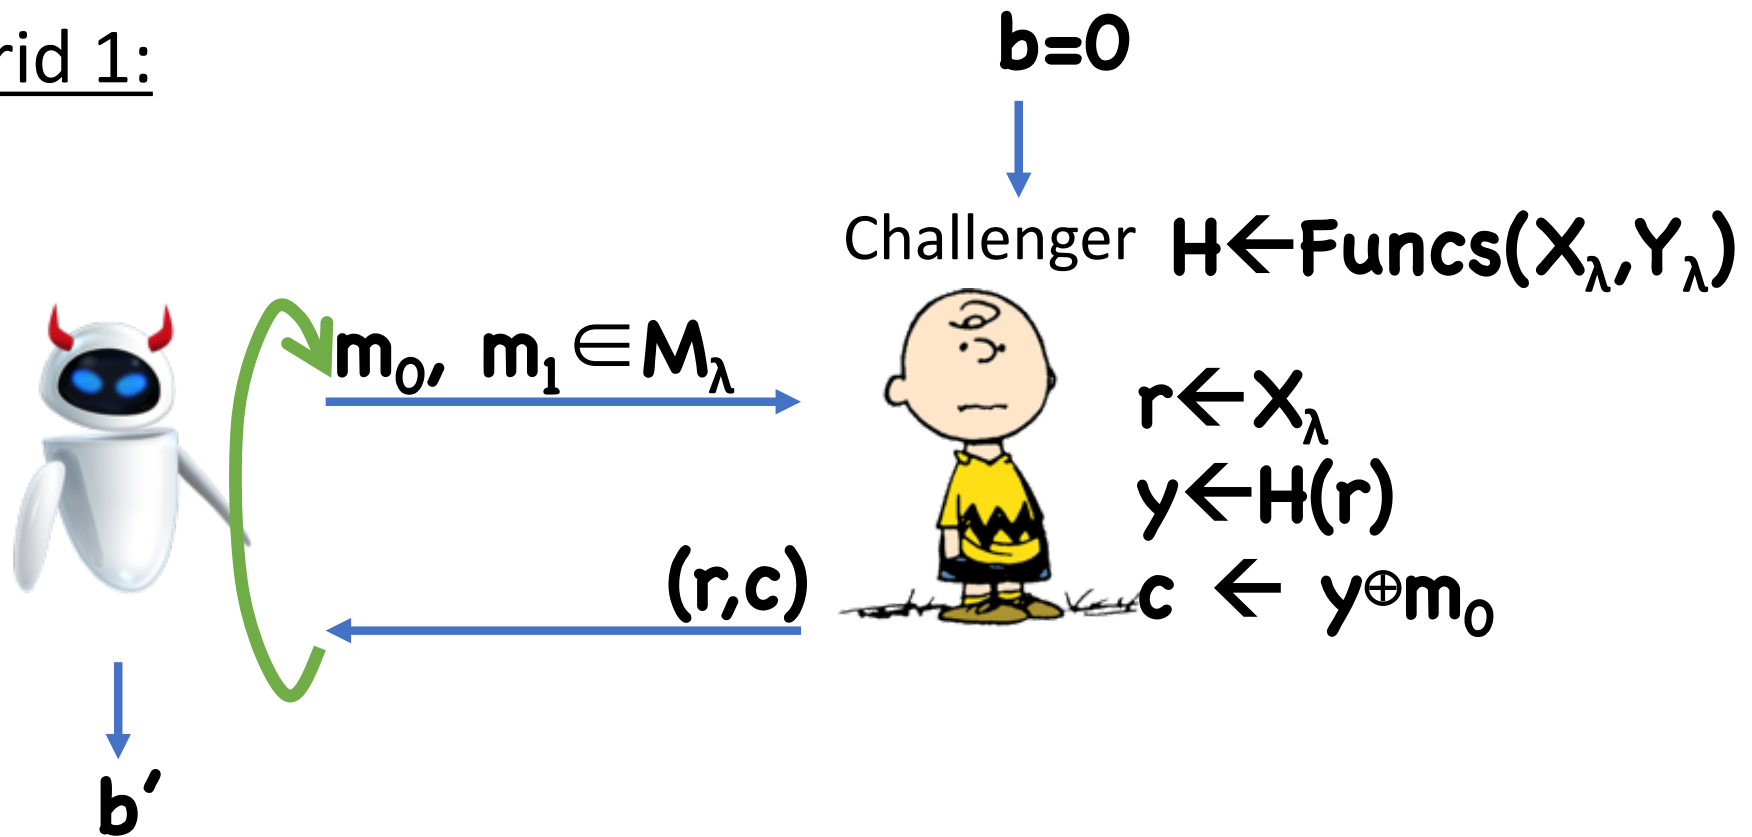
Hybrid 0:



$\text{LoR-Exp}_0(\text{robot}, \lambda)$

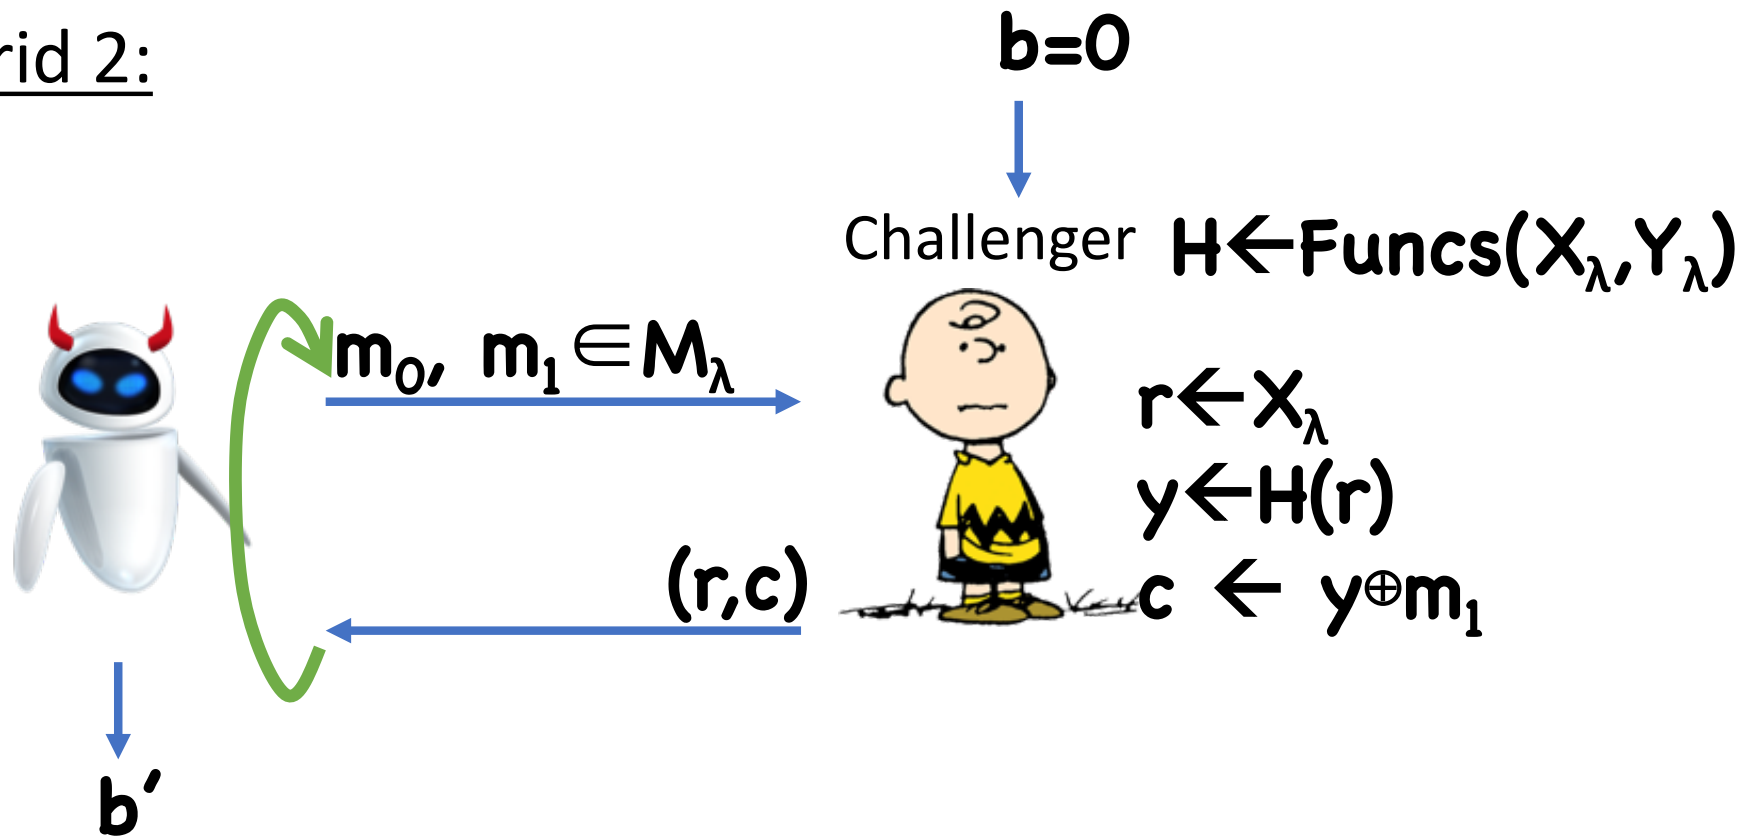
Proof

Hybrid 1:



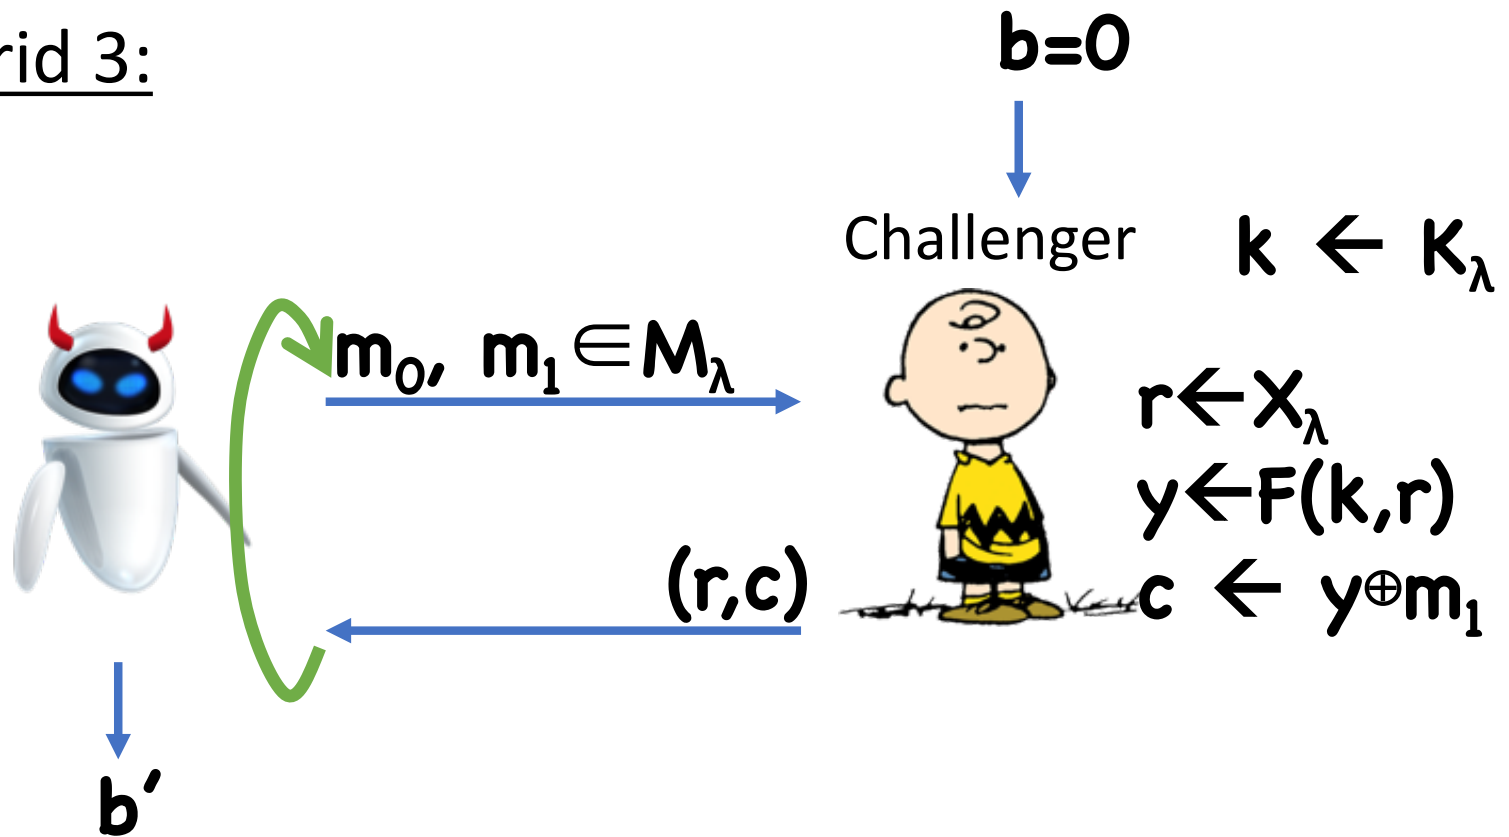
Proof

Hybrid 2:




Proof

Hybrid 3:



$\text{LoR-Exp}_1(\text{robot}, \lambda)$

Proof


Assume toward contradiction that there exists a  with advantage ϵ in breaking **(Enc, Dec)**


 distinguishes Hybrid 0 from Hybrid 3 with advantage ϵ , so either 

- Dist. Hybrid 0 from Hybrid 1 with adv. $\epsilon/2 - q^2/4|X|$
- Dist. Hybrid 1 from Hybrid 2 with adv. $q^2/2|X|$
- Dist. Hybrid 2 from Hybrid 3 with adv. $\epsilon/2 - q^2/4|X|$

Proof

Suppose  distinguishes Hybrid 0 from Hybrid 1



Construct 



Proof

Suppose  distinguishes Hybrid 0 from Hybrid 1

Construct 

- **PRF-Exp₀**(, λ) corresponds to Hybrid 0
- **PRF-Exp₁**(, λ) corresponds to Hybrid 1

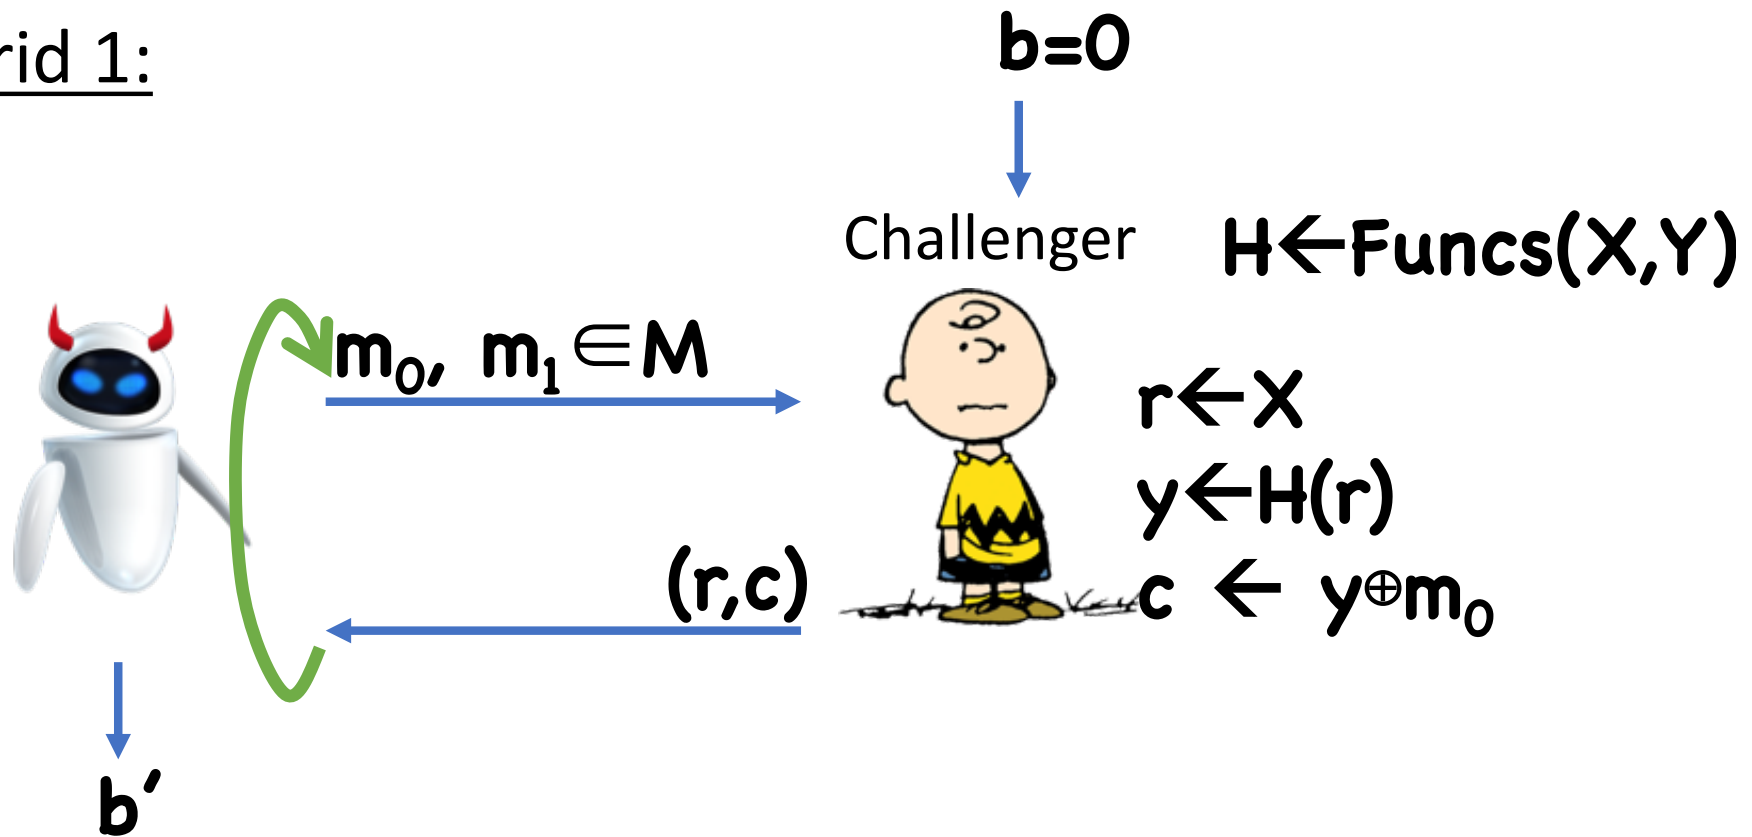
Therefore,  has advantage $\epsilon/2 - q^2/4|X|$
 \Rightarrow contradiction

Proof

Suppose  distinguishes Hybrid 1 from Hybrid 2

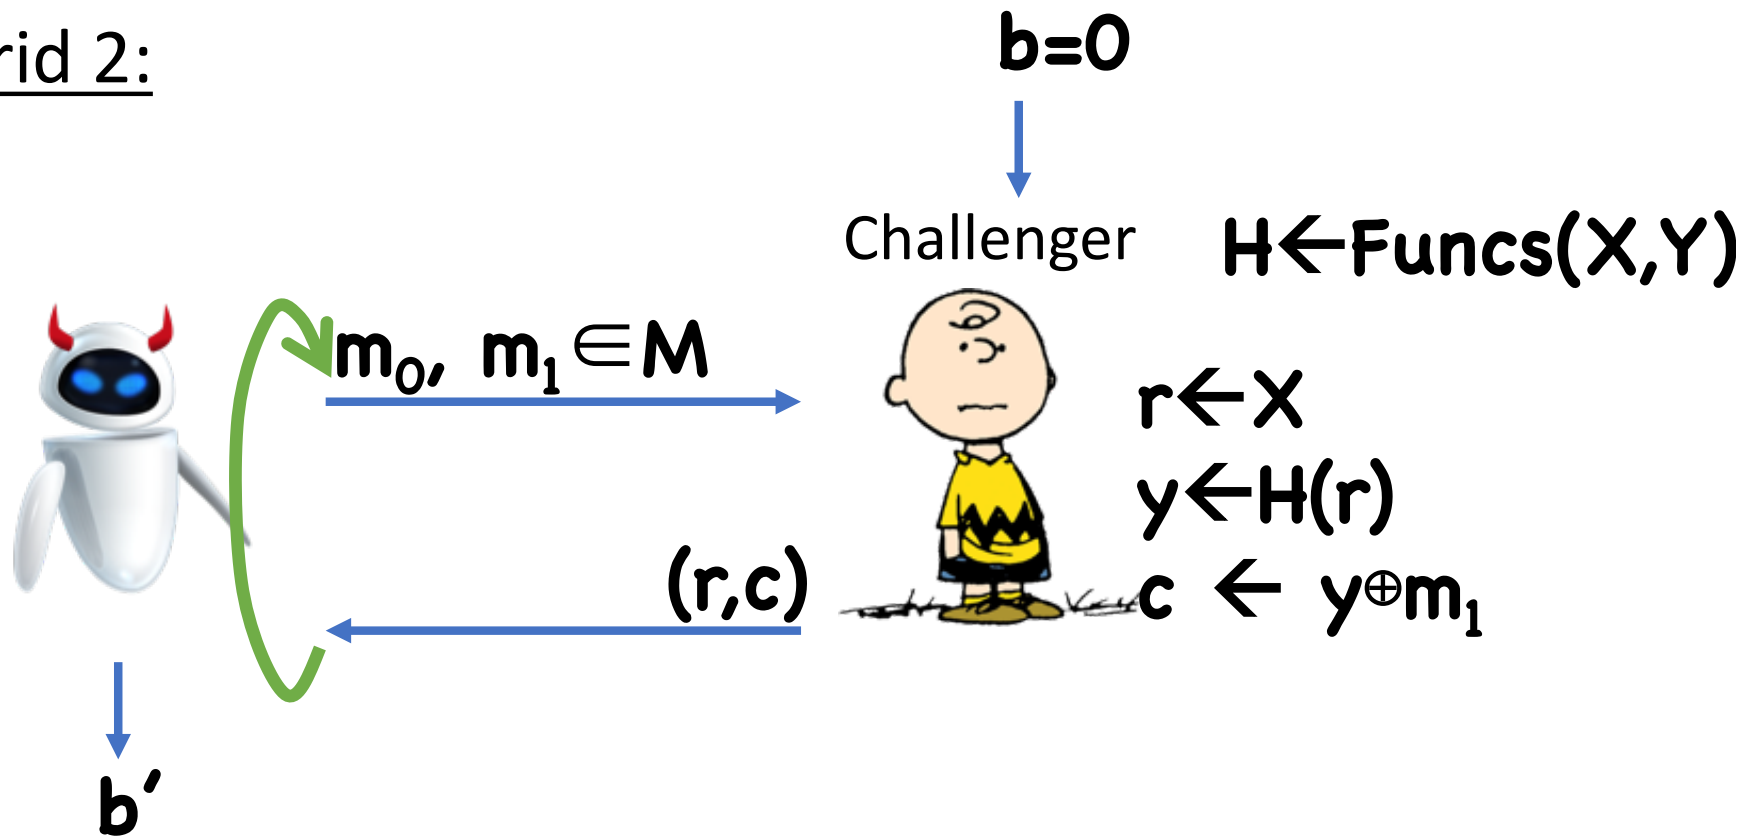
Proof

Hybrid 1:



Proof

Hybrid 2:



Proof

Suppose  distinguishes Hybrid 1 from Hybrid 2

As long as the \mathbf{r} 's for every query are distinct, the \mathbf{y} 's for each query will look like truly random strings

In this case, encrypting \mathbf{m}_0 vs \mathbf{m}_1 will be perfectly indistinguishable

- By OTP security

Proof

Suppose  distinguishes Hybrid 1 from Hybrid 2

Therefore, advantage is $\leq \Pr[\text{collision in the } \mathbf{r}'\text{'s}]$
 $< q^2/2|X|$

Proof

Suppose  distinguishes Hybrid 2 from Hybrid 3

Almost identical to the 0/1 case...


Using PRFs to Build Encryption

So far, scheme had fixed-length messages

- Namely, $\mathbf{M}_\lambda = \mathbf{Y}_\lambda$

Now suppose we want to handle arbitrary-length messages

Security for Arbitrary-Length Messages



Theorem: Given any CPA-secure **(Enc, Dec)** for fixed-length messages (even single bit), it is possible to construct a CPA-secure **(Enc, Dec)** for arbitrary-length messages

Construction

Let **(Enc,Dec)** be CPA-secure for single-bit messages

Enc'(k,m):

For $i=1, \dots, |m|$, run $c_i \leftarrow \text{Enc}(k, m_i)$

Output $(c_1, \dots, c_{|m|})$

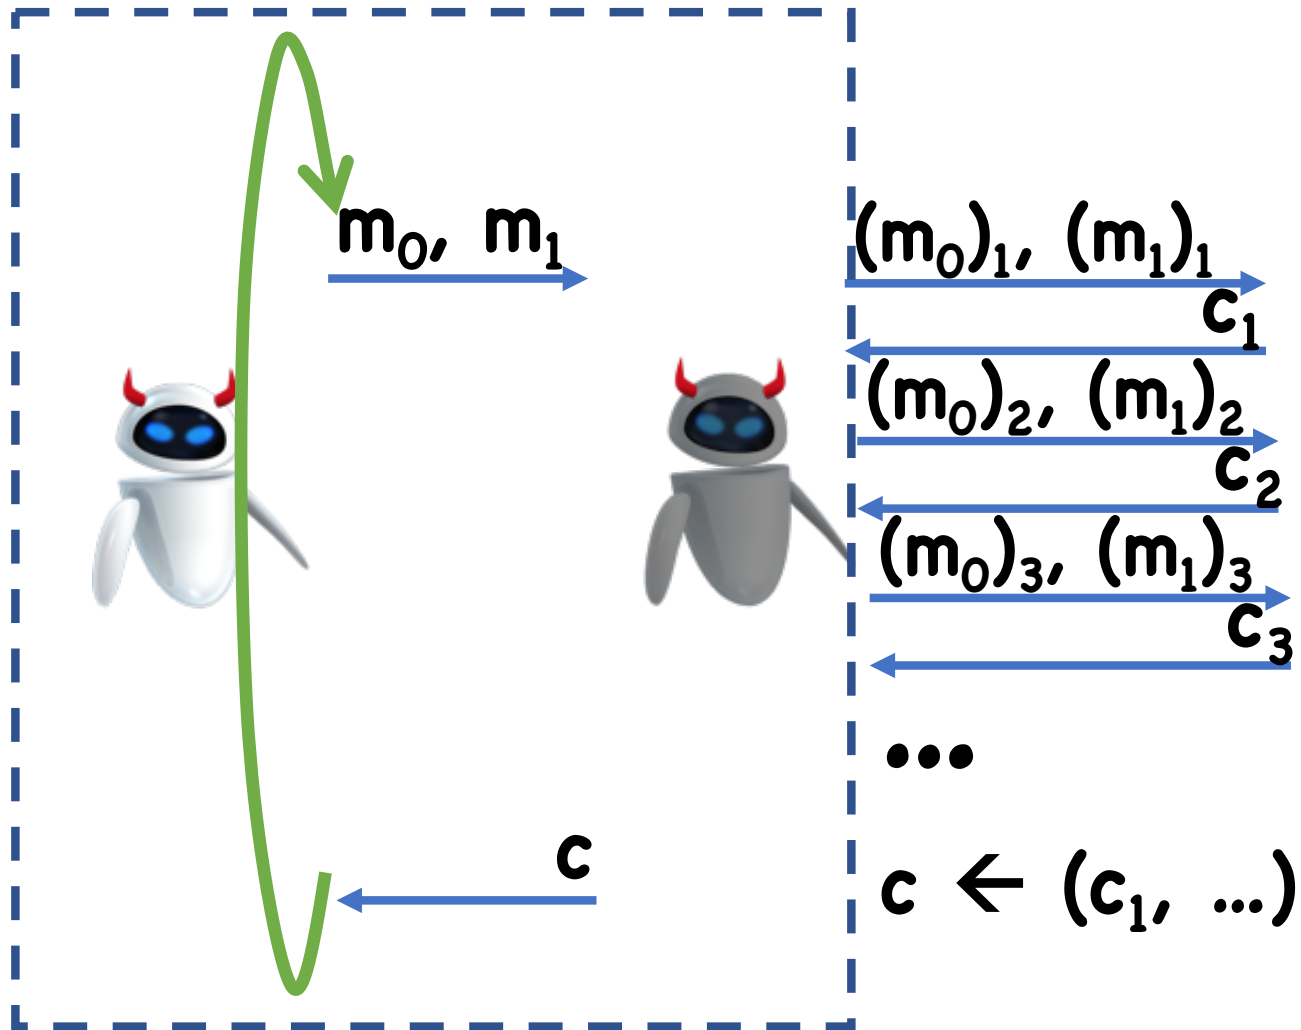
Dec'(k, (c₁, ..., c_l)):

For $i=1, \dots, l$, run $m_i \leftarrow \text{Dec}(k, c_i)$

Output $m = m_1 m_2 \dots m_l$

Theorem: If (Enc, Dec) is LoR secure, then $(\text{Enc}', \text{Dec}')$ is LoR secure

Proof (sketch)




Better Constructions Using PRFs

In PRF-based construction, encrypting single bit requires $\lambda+1$ bits

\Rightarrow encrypting l -bit message requires $\approx \lambda l$ bits

Ideally, ciphertexts would have size $\approx \lambda+1$

Solution 1: Add PRG/Stream Cipher



Solution 2: Counter Mode

Enc(k, m):


- Choose random $\mathbf{r} \leftarrow \{0,1\}^{\lambda/2}$
 - For $i=1,\dots,|m|$,
 - Compute $\mathbf{y}_i \leftarrow \mathbf{F}(\mathbf{k}, \mathbf{r} \parallel i)$
 - Compute $\mathbf{c}_i \leftarrow \mathbf{y}_i \oplus \mathbf{m}_i$
 - Output (\mathbf{r}, \mathbf{c}) where $\mathbf{c} = (\mathbf{c}_1, \dots, \mathbf{c}_{|m|})$
- Write i as $\lambda/2$ -bit string

Dec(k, (r,c)):

- For $i=1,\dots,l$,
 - Compute $\mathbf{y}_i \leftarrow \mathbf{F}(\mathbf{k}, \mathbf{r} \parallel i)$
 - Compute $\mathbf{m}_i \leftarrow \mathbf{y}_i \oplus \mathbf{c}_i$
- Output $\mathbf{m} = \mathbf{m}_1, \dots, \mathbf{m}_l$

Handles any message of length at most $2^{\lambda/2}$

Solution 2: Counter Mode



Block ciphers/Pseudorandom
Permutations

Pseudorandom Permutations

(also known as block ciphers)

Functions that “look like” random **permutations**


Syntax:

- Key space \mathbf{K}_λ
- Domain=Range= \mathbf{X}_λ
- Function $\mathbf{F}:\mathbf{K}_\lambda \times \mathbf{X}_\lambda \rightarrow \mathbf{X}_\lambda$
- Function $\mathbf{F}^{-1}:\mathbf{K}_\lambda \times \mathbf{X}_\lambda \rightarrow \mathbf{X}_\lambda$

Correctness: $\forall \mathbf{k}, \mathbf{x}, \mathbf{F}^{-1}(\mathbf{k}, \mathbf{F}(\mathbf{k}, \mathbf{x})) = \mathbf{x}$

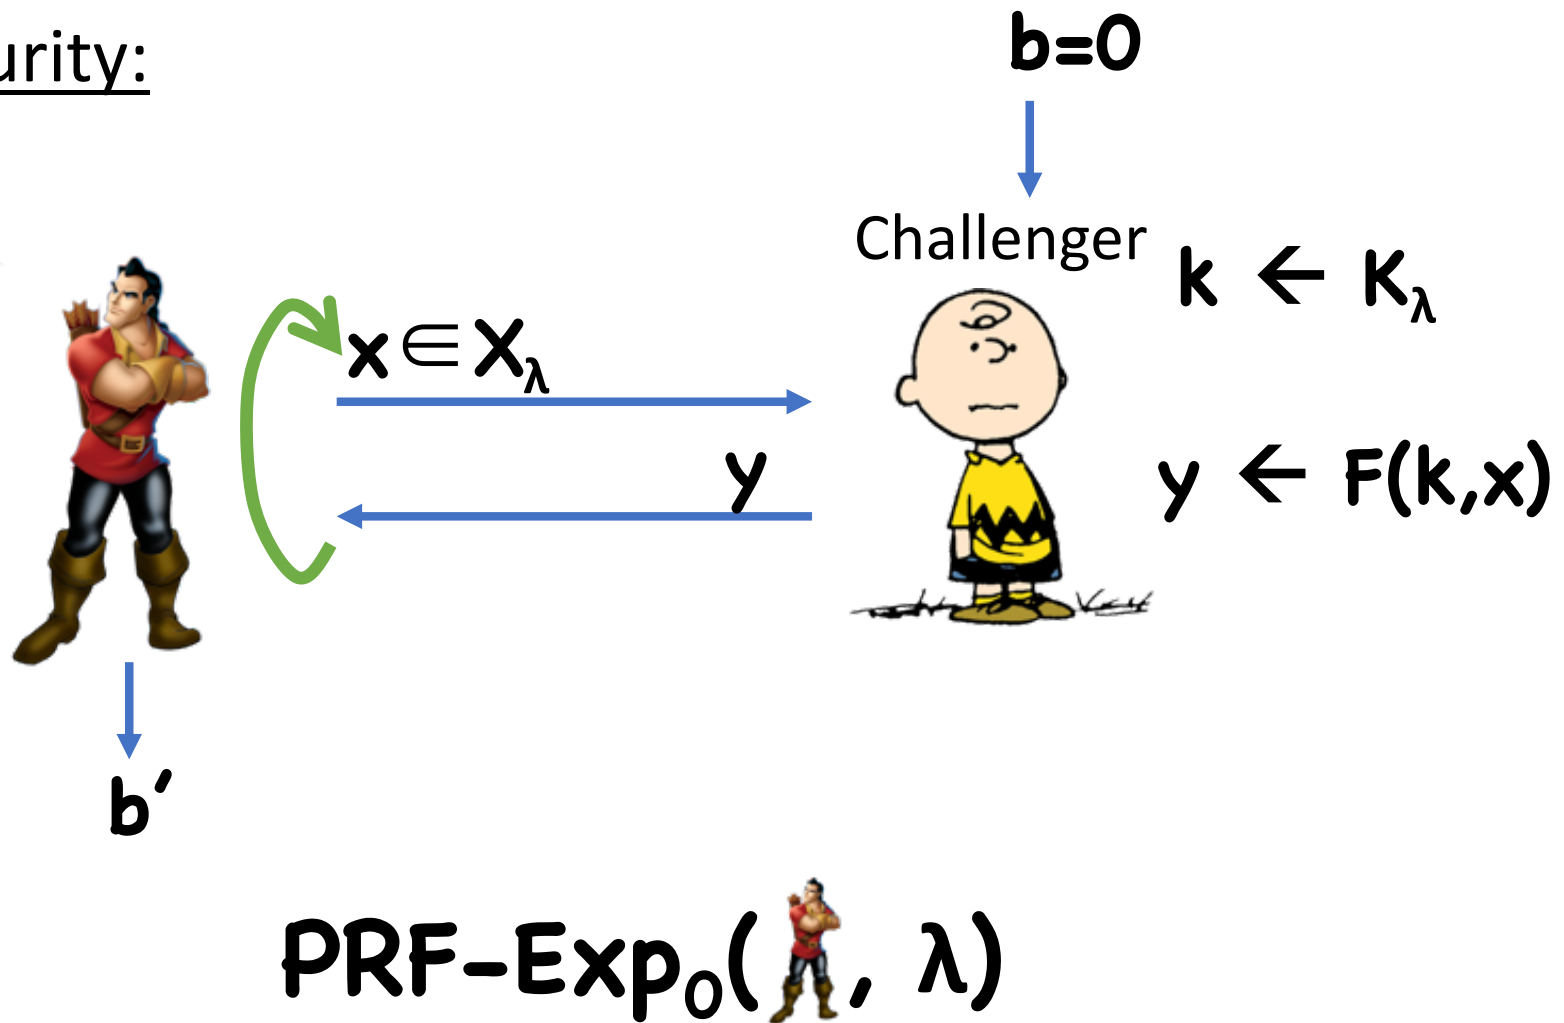
Pseudorandom Permutations

Security:



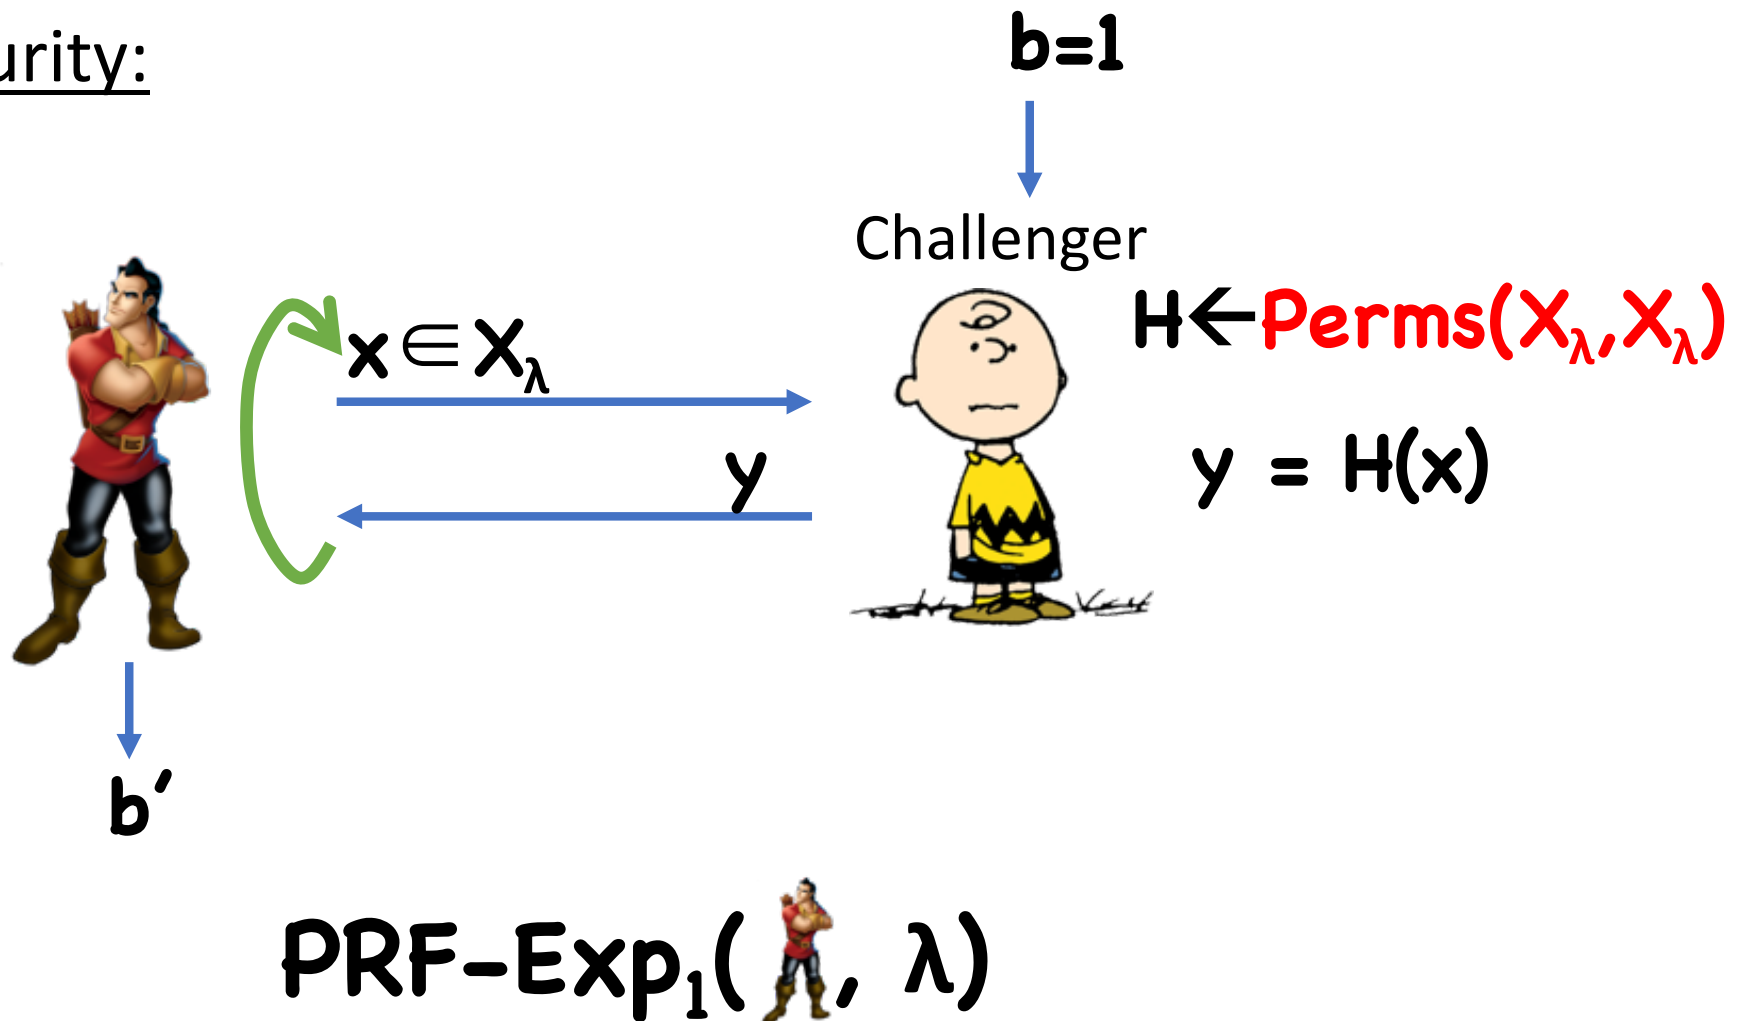
Pseudorandom Permutations

Security:




Pseudorandom Permutations

Security:



PRF Security Definition

Definition: F is a secure PRP if, for all  running in polynomial time, \exists negligible ϵ such that:

$$\left| \Pr[1 \leftarrow \text{PRF-Exp}_0(\text{img alt="superhero" data-bbox="470 530 495 600"}, \lambda)] - \Pr[1 \leftarrow \text{PRF-Exp}_1(\text{img alt="superhero" data-bbox="545 630 570 700"}, \lambda)] \right| \leq \epsilon(\lambda)$$


Theorem: Assuming $|X_\lambda|$ is super-polynomial, a PRP (F, F^{-1}) is secure iff F is secure as a PRF

Proof

Secure as PRP \Rightarrow Secure as PRF

- Assume , hybrids

Hybrid 0:



Proof

Secure as PRP \Rightarrow Secure as PRF

- Assume , hybrids

Hybrid 1:



b'



Challenger $H \leftarrow \text{Perms}(X, X)$



$y \leftarrow H(x)$

Proof

Secure as PRP \Rightarrow Secure as PRF

- Assume , hybrids

Hybrid 2:



b'



Challenger $H \leftarrow \text{Funcs}(X, X)$



$y \leftarrow H(x)$



Proof

Secure as PRP \Rightarrow Secure as PRF

- Assume , hybrids

Hybrids 0 and 1 are indistinguishable by PRP security

Hybrids 1 and 2?

- In Hybrid 1,  sees random **distinct** answers
- In Hybrid 2,  sees random answers
- Except with probability $\approx q^2/2|X_\lambda|$, random answers will be distinct anyway

Proof

Secure as PRF \Rightarrow Secure as PRP

- Assume , hybrids

Proof essentially identical to other direction

Suppose (F, F^{-1}) is a secure PRP

Is (F^{-1}, F) also a secure PRP?

Counter Example


Suppose (F, F^{-1}) is a secure PRP. Assume $X = \{0, 1\}^n$

Define (H, H^{-1}) as follows:


- Given k , let i be smallest input such that $F^{-1}(i)$ begins with a 0
- Let $x_0 = F^{-1}(0^n)$, $x_1 = F^{-1}(i)$
- $H(k, x) = \begin{cases} 0^n & \text{if } x = x_1 \\ i & \text{if } x = x_0 \\ F(k, x) & \text{otherwise} \end{cases}$

How to use block ciphers for encryption


Counter Mode (CTR)



Electronic Code Book (ECB)



ECB Decryption



Security of ECB?


Is ECB mode CPA secure?

Is ECB mode *one-time* secure?


Security of ECB



Plaintext




Ciphertext




Ideal

Cipher Block Chaining (CBC) Mode



(For now, assume all messages are multiples of the block length)

CBC Mode Decryption



Theorem: If (F, F^{-1}) is a secure pseudorandom permutation and $|X_\lambda|$ is super-polynomial, then CBC mode encryption is CPA secure.

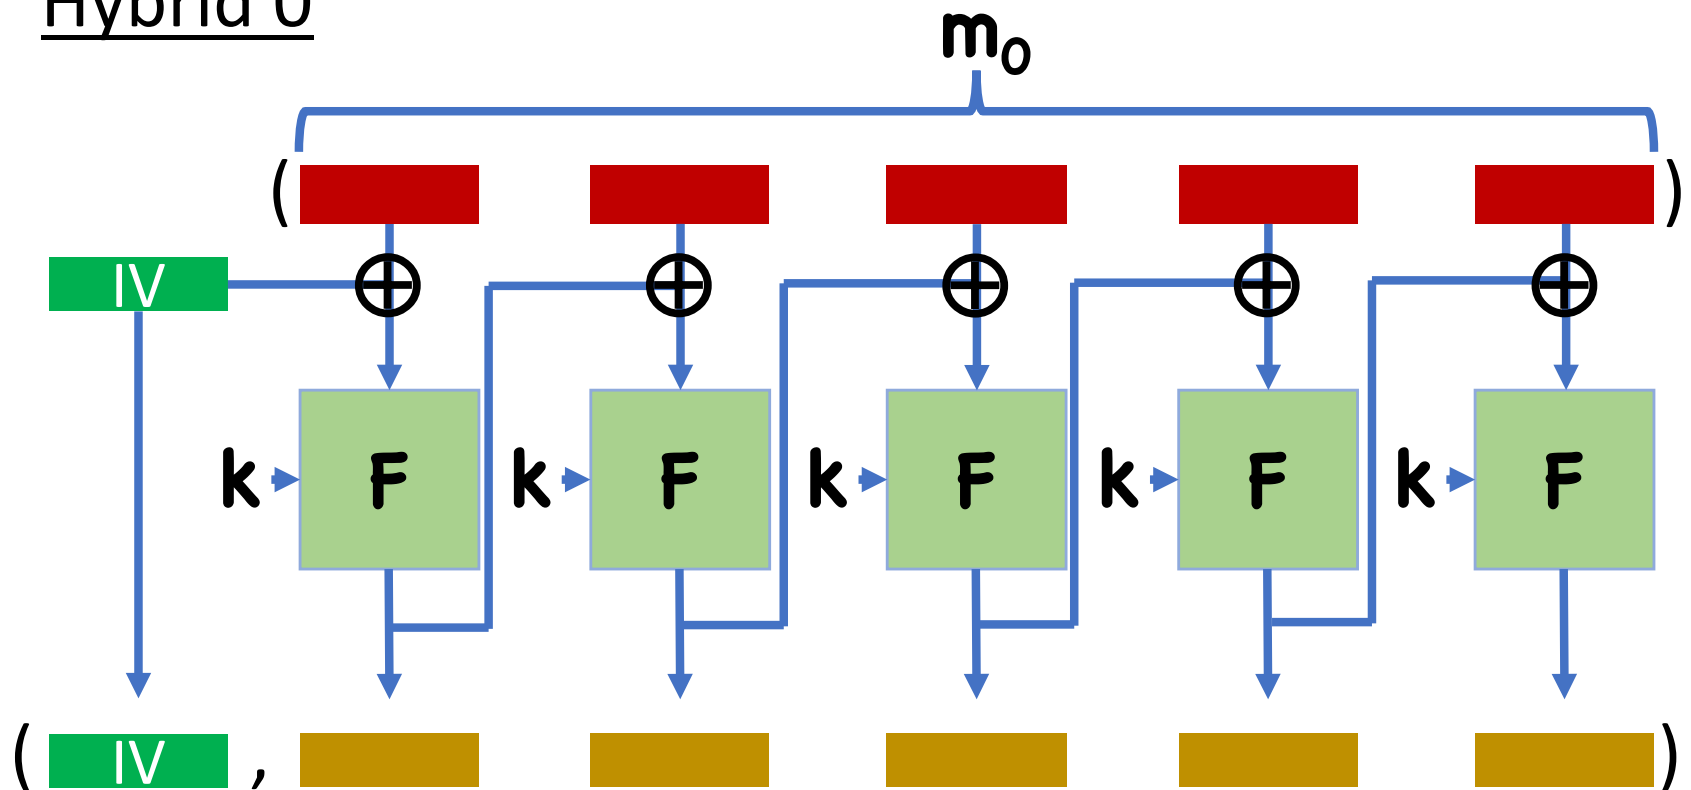
Proof Sketch

Assume toward contradiction an adversary  for CBC mode

Hybrids...


Proof Sketch

Hybrid 0




Proof Sketch

Hybrid 1



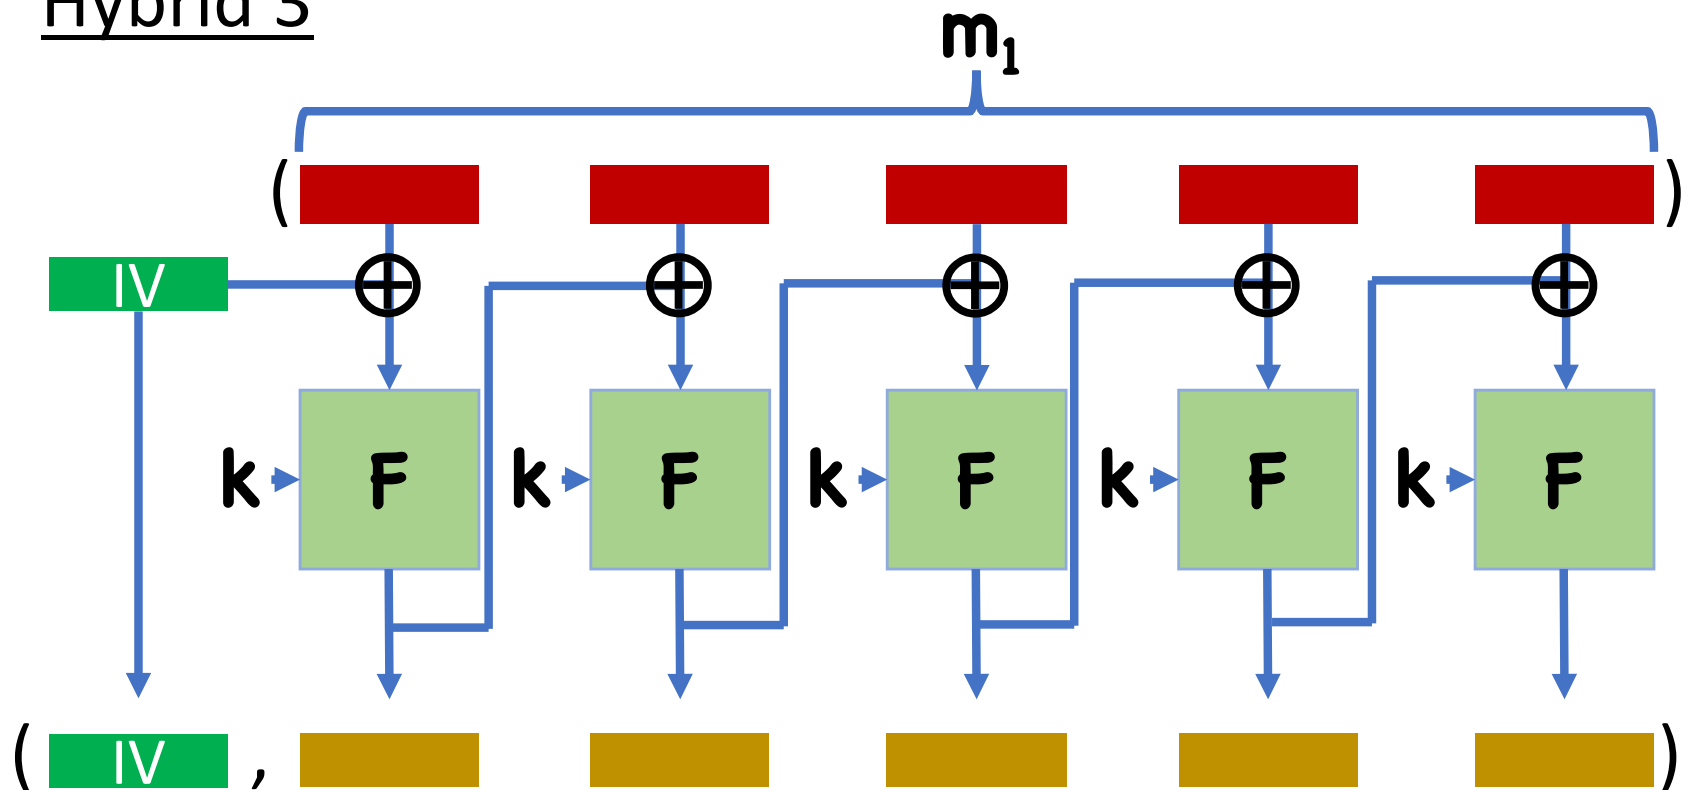
Proof Sketch

Hybrid 2



Proof Sketch

Hybrid 3



Proof Sketch

Hybrid 0,1 differ by replacing calls to \mathbf{F} with calls to random permutation \mathbf{H}

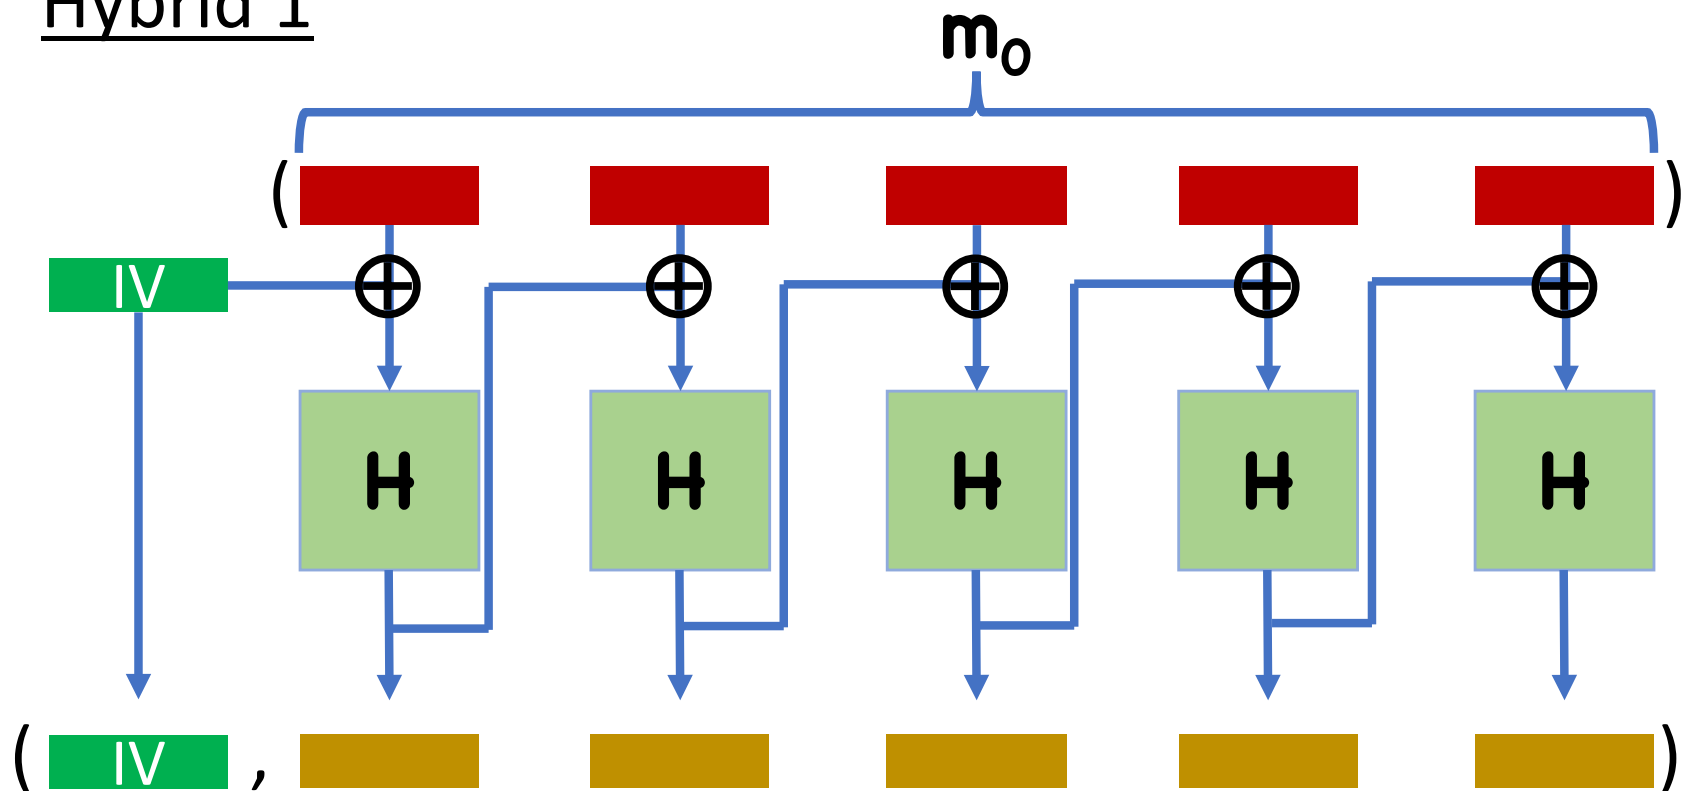
- Indistinguishable by PRP security

Same for Hybrids 2,3

All that is left is to show indistinguishability of 1,2


Proof Sketch

Hybrid 1




Proof Sketch

Hybrid 2



Proof Sketch

Idea:

- As long as, say, the sequence of left messages queried by  does not result in two calls to **F** on the same input, all outputs will be random (distinct) outputs
- For each message, first query to **F** will be uniformly random
- Second query gets XORed with output of first query to **F** $\Rightarrow \approx$ uniformly random

Proof Sketch

Idea:

- Since queries to \mathbf{F} are (essentially) uniformly random, probability of querying same input twice is exponentially small
- Ciphertexts will be essentially random
- True regardless of encrypting \mathbf{m}_0 or \mathbf{m}_1

Reminders

HW2 Due TODAY

HW3 Due March 5th

PR1 Due March 10th