Resistance Against Adaptive Iterated Distinguishers by Decorrelation

Aslı Bay [1] Atefeh Mashatan[2] Serge Vaudenay[1]

Speaker: Petr Sušil[1]

Ecole Polytechnique Fédéral de Lausanne (EPFL), Switzerland

Canadian Imperial Bank of Commerce (CIBC), Canada

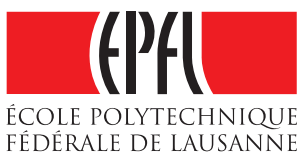

Outline

1. Decorrelation Theory
   - The Luby-Rackoff Model
   - Advantage of an adversary $\mathcal{A}$
   - Distribution matrix of a block cipher and its link with the advantage of $\mathcal{A}$
2. Adaptive distinguishers
   - Quantifying the security of block ciphers against adaptive iterated distinguishers

Decorrelation Theory

- Proposed by Vaudenay as a tool for proving resistance of block ciphers against a wide range of statistical attacks:
  - Differential attacks, linear attacks, truncated differential attacks, etc.
- A tool to design secure block ciphers against even undiscovered cryptanalytic attacks meeting certain broad criteria
- Proves the security of several block ciphers such as:
  - AES candidate DFC, NUT (n-Universal Transformation) families of block ciphers, the block cipher C, and KFC

The Luby-Rackoff Model

We consider a d-limited adversary $\mathcal{A}$ in the Luby-Rackoff Model

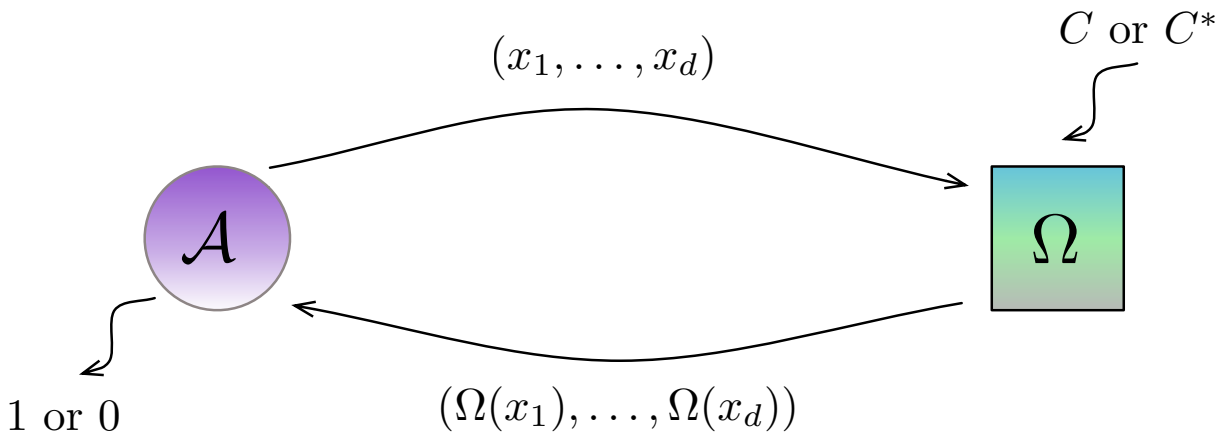

The Luby-Rackoff Model

We consider a d-limited adversary $\mathcal{A}$ in the Luby-Rackoff Model

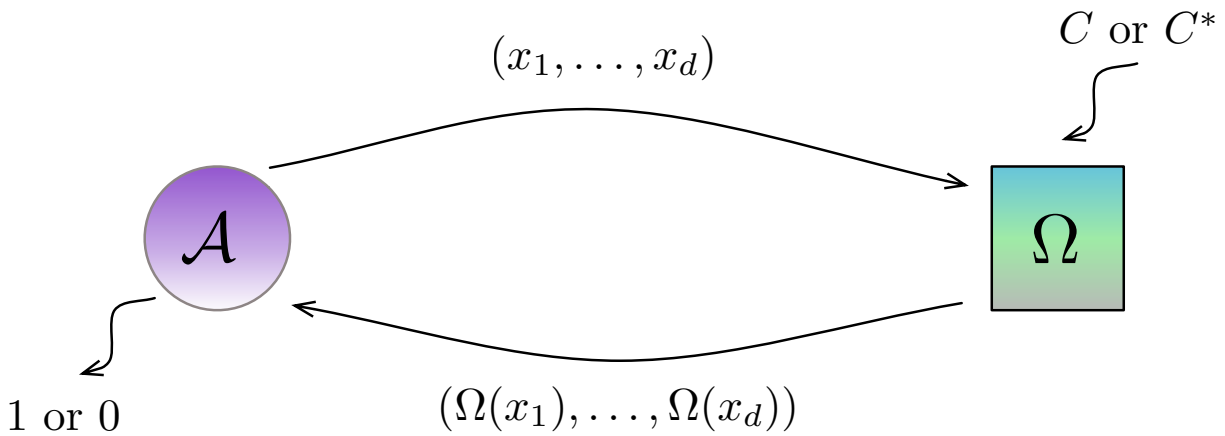

$$\mathrm{Adv}_{\mathcal{A}}(C, C^*) = \left| \Pr[\mathcal{A}(C) = 1] - \Pr[\mathcal{A}(C^*) = 1] \right|$$

- When d inputs are chosen **at once**, $\mathcal{A}$ is **non-adaptive**
- When d inputs are chosen depending on the answers to the previous queries, $\mathcal{A}$ is **adaptive**
- If advantage is **negligible** for all adversaries $\mathcal{A}$, then the cipher C is considered to be **secure**

Computing Advantage of $\mathcal{A}$ Using Decorrelation Theory

- Computing advantage is **not** an easy task in general
- Decorrelation Theory provides tools for computing the advantage of $\mathcal{A}$:

$$\mathrm{BestAdv}_{\zeta}(C, C^*) = \max_{\mathcal{A} \in \zeta} \mathrm{Adv}_{\mathcal{A}}(C, C^*)$$

Computing Advantage of $\mathcal{A}$ Using Decorrelation Theory

The best advantage of a **(non)-adaptive** distinguisher $\mathcal{A}$ is computed by the means of d**-wise distribution matrices:**

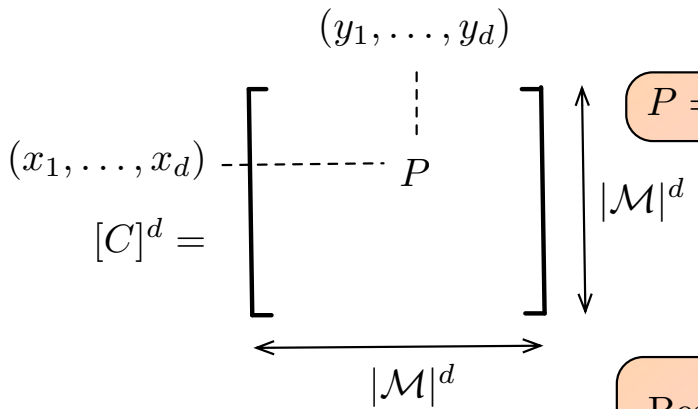

$$P = \Pr[C(x_1) = y_1, \ldots, C(x_d) = y_d]$$

$$\mathrm{BestAdv}_{\zeta}(C, C^*) = \frac{1}{2}\|[C]^d - [C^*]^d\|_{\diamond}$$

$$\|M\|_{\infty} = \max_{x_1,\ldots,x_d} \sum_{y_1} \cdots \sum_{y_d} |M_{(x_1,\ldots,x_d),(y_1,\ldots,y_d)}|$$

$$\|M\|_{A} = \max_{x_1} \sum_{y_1} \cdots \max_{x_d} \sum_{y_d} |M_{(x_1,\ldots,x_d),(y_1,\ldots,y_d)}|$$

Computing Advantage of $\mathcal{A}$ Using Decorrelation Theory

The best advantage of a **(non)-adaptive** distinguisher $\mathcal{A}$ is computed by the means of d**-wise distribution matrices:**

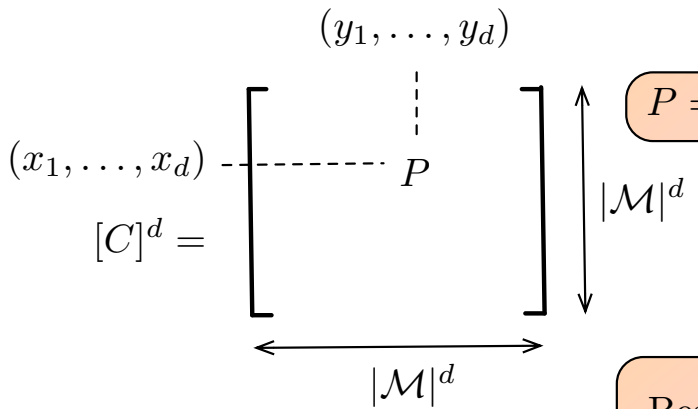

$$P = \Pr[C(x_1) = y_1, \ldots, C(x_d) = y_d]$$

$$\mathrm{BestAdv}_{\zeta}(C, C^*) = \frac{1}{2}\|[C]^d - [C^*]^d\|_{\diamond}$$

$$\|M\|_{\infty} = \max_{x_1,\ldots,x_d} \sum_{y_1} \cdots \sum_{y_d} |M_{(x_1,\ldots,x_d),(y_1,\ldots,y_d)}|$$ ⟸ Non-adaptive

$$\|M\|_{A} = \max_{x_1} \sum_{y_1} \cdots \max_{x_d} \sum_{y_d} |M_{(x_1,\ldots,x_d),(y_1,\ldots,y_d)}|$$ ⟸ Adaptive

Distinguishing Problems

The oracle Ω either implements:

- a random block cipher C or the perfect cipher C^*
- a MAC or the perfect function F^*
- the perfect function F^* or the perfect cipher C^*

Example: Distinguishing F^* from C^*

$$F^* : \{0,1\} \to \{0,1\} \text{ and } C^* : \{0,1\} \to \{0,1\}$$

$$[F^*]^2 = \begin{bmatrix} 1/2 & 0 & 0 & 1/2 \\ 1/4 & 1/4 & 1/4 & 1/4 \\ 1/4 & 1/4 & 1/4 & 1/4 \\ 1/2 & 0 & 0 & 1/2 \end{bmatrix}, [C^*]^2 = \begin{bmatrix} 1/2 & 0 & 0 & 1/2 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 & 0 \\ 1/2 & 0 & 0 & 1/2 \end{bmatrix}$$

$$[F^*]^2 - [C^*]^2 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1/4 & -1/4 & -1/4 & 1/4 \\ 1/4 & -1/4 & -1/4 & 1/4 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Best advantage of distinguishing F^* and C^* both non-adaptively and adaptively:

Example: Distinguishing F^* from C^*

$$F^* : \{0,1\} \to \{0,1\} \text{ and } C^* : \{0,1\} \to \{0,1\}$$

$$[F^*]^2 = \begin{bmatrix} 1/2 & 0 & 0 & 1/2 \\ 1/4 & 1/4 & 1/4 & 1/4 \\ 1/4 & 1/4 & 1/4 & 1/4 \\ 1/2 & 0 & 0 & 1/2 \end{bmatrix}, [C^*]^2 = \begin{bmatrix} 1/2 & 0 & 0 & 1/2 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 & 0 \\ 1/2 & 0 & 0 & 1/2 \end{bmatrix}$$

$$[F^*]^2 - [C^*]^2 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1/4 & -1/4 & -1/4 & 1/4 \\ 1/4 & -1/4 & -1/4 & 1/4 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Best advantage of distinguishing F^* and C^* both non-adaptively and adaptively:

$$\mathrm{BestAdv}(F^*, C^*) = \frac{1}{2}\|[F^*]^2 - [C^*]^2\|_\infty = \frac{1}{2}\|[F^*]^2 - [C^*]^2\|_A = \frac{1}{2}$$

A Non-adaptive Iterated Distinguisher of Order d

Iteration of a d-limited **non-adaptive** distinguisher $\mathcal{A}$ **"n times"**

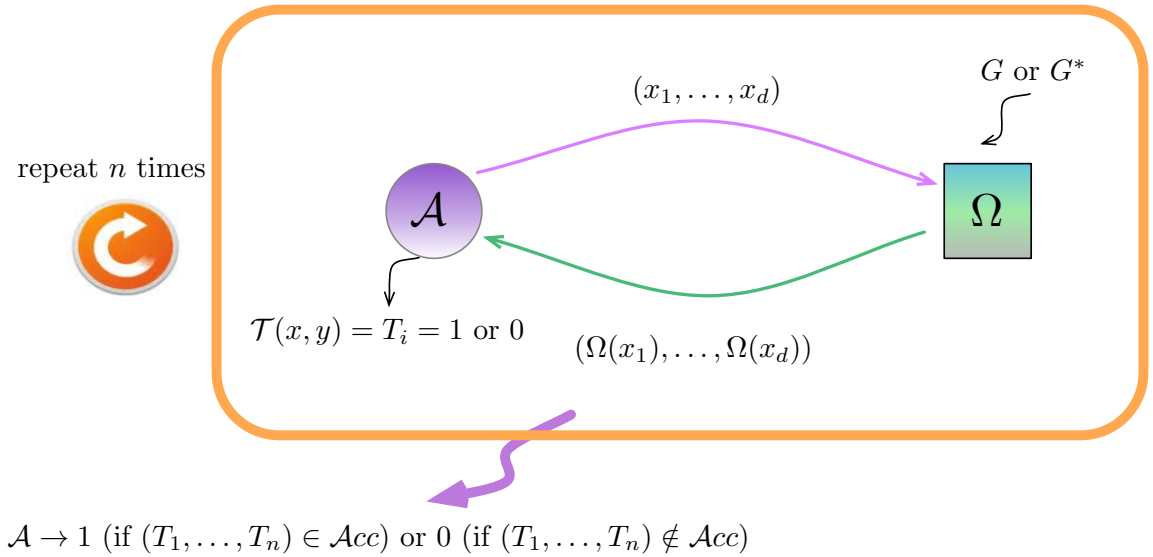

Parameters: an integer n, a set X, a distribution on X, a test $\mathcal{T}$, a set $\mathcal{A}cc$
for $i = 1$ to n **do**
    pick $x = (x_1, \ldots, x_d)$ at random
    get $y = (\Omega(x_1), \ldots, \Omega(x_d))$ $T_i = \mathcal{T}(x, y) \in \{0, 1\}$
end for
if $(T_1, \ldots, T_n) \in \mathcal{A}cc$ **then**
    output 1
else
    output 0
end if

Example 1: Linear Distinguisher

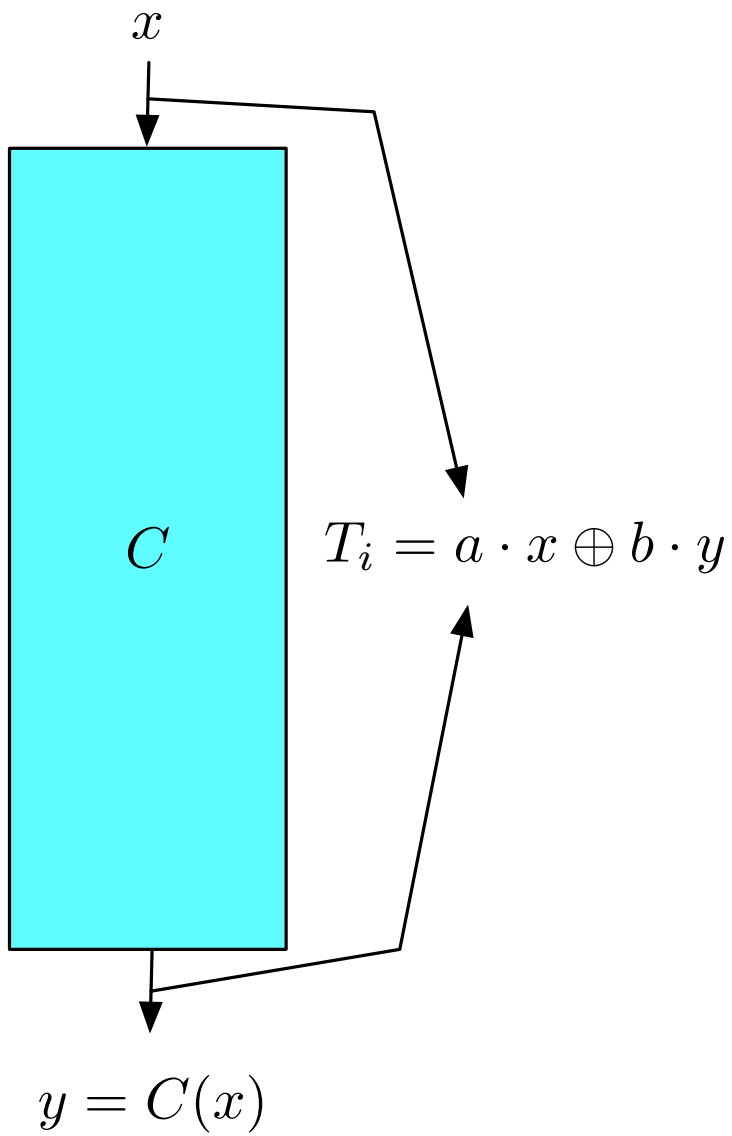

A non-adaptive iterated attack of order 1

Parameters: an integer n, a set X, a distribution $\mathcal{X}$ on X, a set I, masks a and b

for $i = 1$ **to** n
  Pick x at random from $\mathcal{X}$
  Set $y = c(x)$
  Set $T_i = a \cdot x \oplus b \cdot y$
end for

if $T_1 + \cdots + T_n \in I$ **then**
  Output 1
else
  Output 0

Example 2: Differential Distinguisher

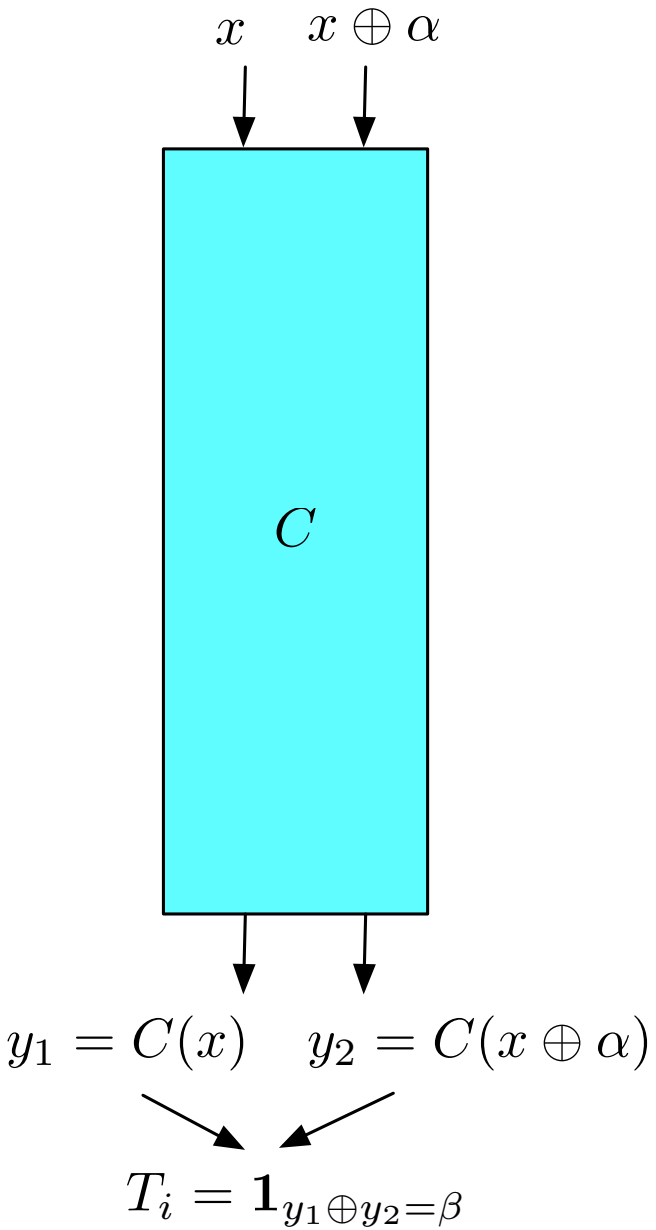

A non-adaptive iterated attack of order 2

Parameters: an integer n, a set X, a distribution $\mathcal{X}$ on X, differences α and β

for $i = 1$ **to** n
  Pick x_1 at random from $\mathcal{X}$
  Set $x_2 = x_1 \oplus \alpha$
  Set $y_1 = c(x_1)$, $y_2 = c(x_2)$
  Set $T_i = \mathbf{1}_{y_1 \oplus y_2 = \beta}$
end for

if $T_1 + \cdots + T_n \neq 0$ **then**
  Output 1
else
  Output 0

Adaptive Adversaries: Defining a Useful Function G

Let $\mathcal{G}$ be the set of functions G such that

$$G(x,b) = \begin{cases} G_0(x) \text{ (\textbf{encryption algorithm})}, & \text{if } b = 0, \\ G_1(x) \text{ (\textbf{decryption algorithm})}, & \text{if } b = 1, \end{cases}$$

where $G_0(x) = G(x,0)$ and $G_1(x) = G(x,1)$

G: a random element of $\mathcal{G}$
G^*: a uniformly distributed element of $\mathcal{G}$

G determines which operation will be performed

A d-limited Adaptive Distinguisher

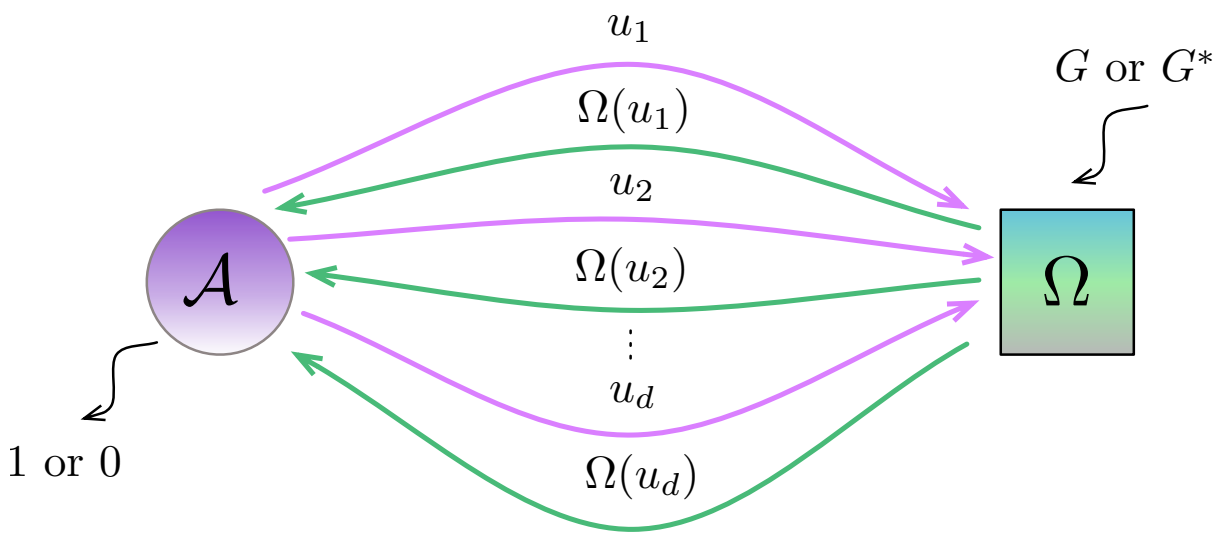

Parameters: a function $\mathcal{F}$, a test $\mathcal{T}$, a distribution $\mathcal{R}$ on $\{0,1\}^*$

Pick $r \in \{0,1\}^*$ at random from $\mathcal{R}$
Set $u_1 = (a_1, b_1) \leftarrow \mathcal{F}(\cdot; r)$
Set $v_1 = \Omega(u_1)$
Set $u_2 = (a_2, b_2) \leftarrow \mathcal{F}(v_1; r)$
Set $v_2 = \Omega(u_2)$
...
Set $u_d = (a_d, b_d) \leftarrow \mathcal{F}(v_1, \ldots, v_{d-1}; r)$
Set $v_d = \Omega(u_d)$

Output $\mathcal{T}(v_1, \ldots, v_d; r)$

An Adaptive Iterated Distinguisher of Order d

Iteration of a d-limited **adaptive** distinguisher $\mathcal{A}$ **"n times"**

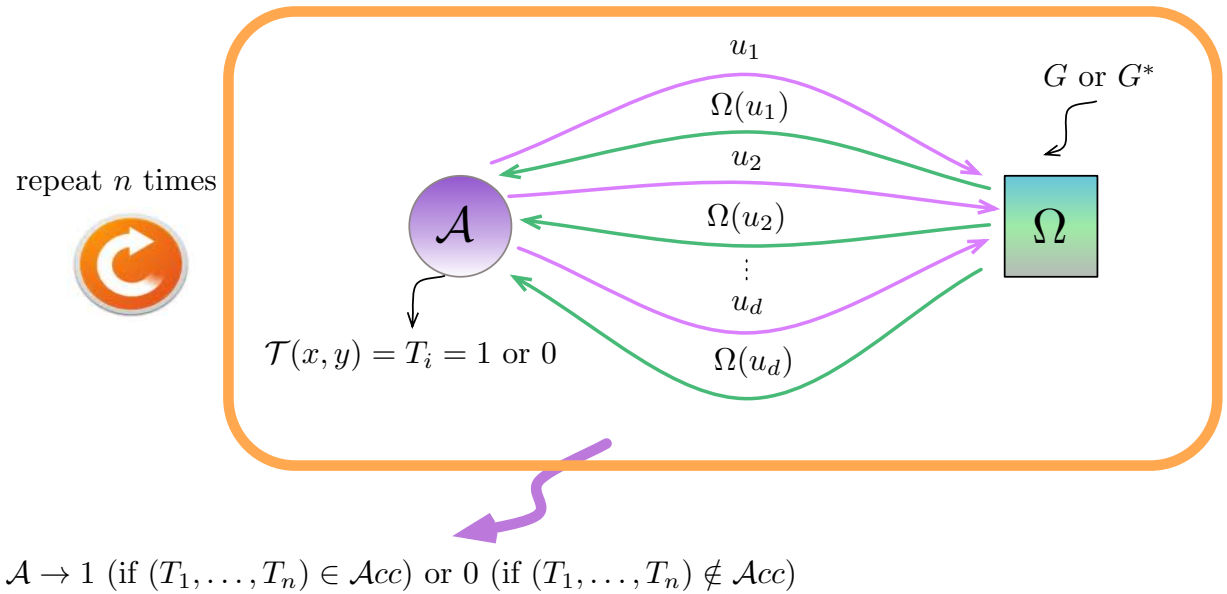

Parameters: an integer n, a function $\mathcal{F}$, a test $\mathcal{T}$, a set $\mathcal{A}cc$, a distribution $\mathcal{R}$ on $\{0,1\}^*$

for $k = 1$ **to** n

Set T_k (with independent coins) $\leftarrow$ output of adaptive distinguisher

end for

Output $1_{\mathcal{A}cc}(T_1, \ldots, T_n)$

Example: Boomerang Distinguishers

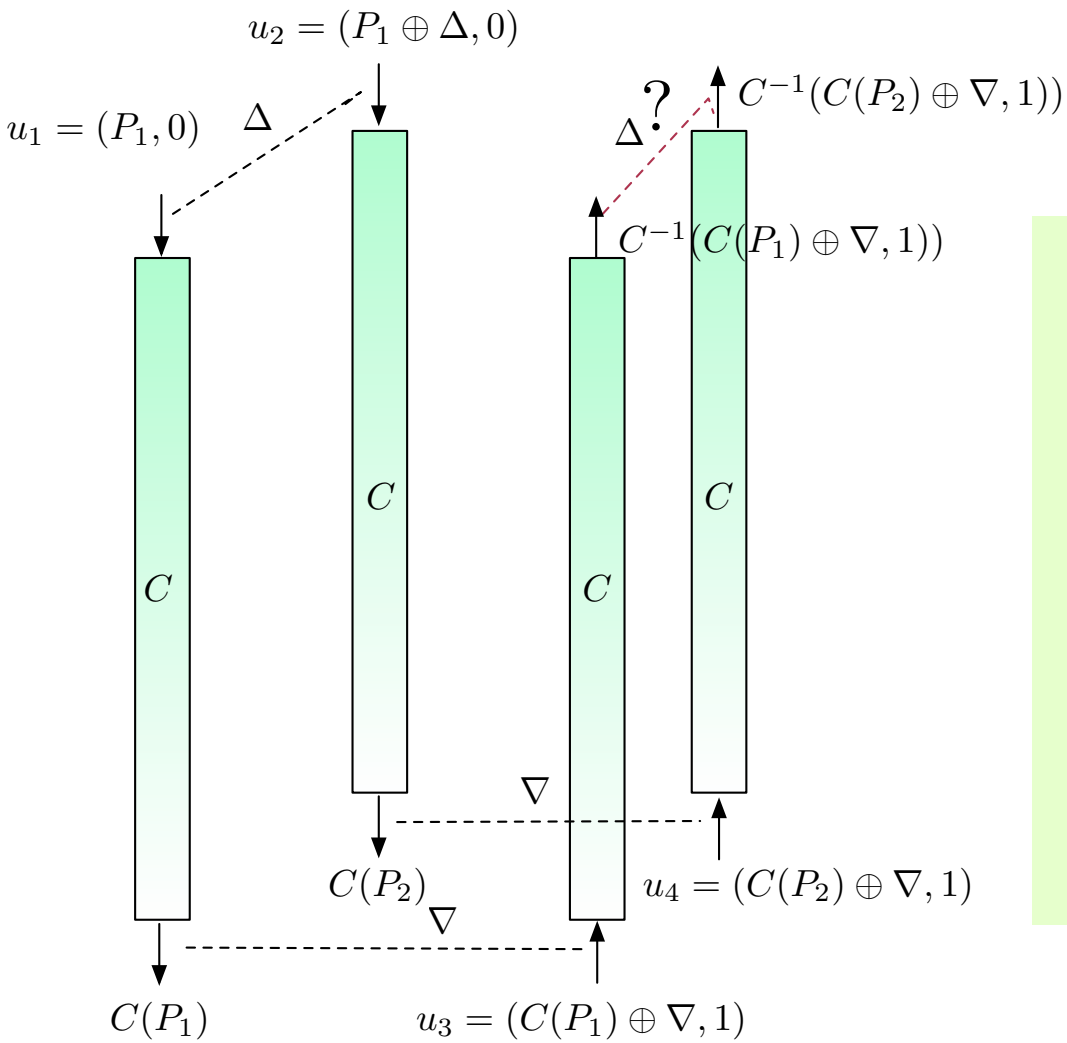

Parameters: an integer n, a set X, differences Δ and ∇

for $k = 1$ **to** n
  Pick P_1 uniformly at random from X
  Set $P_2 = P_1 \oplus \Delta$
  Set $C_1 = C(P_1)$, $C_2 = C(P_2)$
  Set $C_3 = C_1 \oplus \nabla$, $C_4 = C_2 \oplus \nabla$
  Set $P_3 = C^{-1}(C_3)$, $P_4 = C^{-1}(C_4)$
  Set $T_k = \mathbf{1}_{P_3 \oplus P_4 = \Delta}$

end for
if $T_1 + \cdots + T_n \neq 0$ **then**
  Output 1
else
  Output 0

Definitions

- A cipher C is d-**decorrelated** (ε negligible) if

$$\|[C]^d - [C^*]^d\|_\diamond \leq \varepsilon$$

- A cipher C is **perfect** d-**decorrelated** ($\varepsilon = 0$) if

$$\|[C]^d - [C^*]^d\|_\diamond = 0$$

Previous Work: Security against Non-adaptive Iterated Distinguishers of Order d

Theorem (Vaudenay)

An upper bound on the advantage of a **non-adaptive** *iterated distinguisher $\mathcal{A}$ of order d against a $2d$-*__decorrelated__ *cipher C with $\|[C]^{2d} - [C^*]^{2d}\|_\infty \leq \varepsilon$ is*

$$\mathsf{Adv}_{\mathcal{A}} \leq 5\sqrt[3]{\left(2\delta + \frac{5d^2}{2M} + \frac{3\varepsilon}{2}\right)n^2} + n\varepsilon$$

- M is the cardinality of the message space
- δ is the probability that any two iterations have at least one query in common
- n is the number of iterations

Our Contribution

- Security against Adaptive Iterated Distinguishers has not been dealt with yet
- We consider these distinguishers: **Adaptive Plaintext-Ciphertext Iterated Distinguishers**
  - Examples: Boomerang Distinguishers, Rectangle Distinguishers,
- We provide a bound against these adversaries
- We get a looser bound than Vaudenay's Theorem (for non-adaptive distinguishers) has (not surprising)

Advantage of Adaptive Plaintext-Ciphertext Iterated Distinguishers of Order d

Theorem

An upper bound on the advantage of an iterated adaptive distinguisher $\mathcal{A}_{\mathsf{AI}(d)}$ of order d against a $2d$-decorrelated function $G \in \mathcal{G}$ with $||[G]^{2d} - [G^*]^{2d}||_A \leq \varepsilon$ is

$$\mathsf{Adv}_{\mathcal{A}_{\mathsf{AI}(d)}} \leq 5\sqrt[3]{\left(2\theta + e^{8d^2/M} + \frac{2d^2}{M} + \frac{3\varepsilon}{2} - 1\right)n^2} + n\varepsilon,$$

- M is the cardinality of the message space
- θ is the probability that any two different iterations have at least one query in common for a given G
- n is the number of iterations

Question 1 : Is Extension for Decorrelation of Order $2d-1$ Possible?

- No!
- Similar to the first open problem posed by Vaudenay's Theorem
- Bay et al. solved this in [BMV12] which can also be applied to our case

- Decorrelation of order $2d$ is necessary
- Non-adaptive iterated distinguisher of order d against a cipher:
- 3-round Feistel construction decorrelated to the order $2d-1$, that is $\|[C]^{2d-1} - [C^*]^{2d-1}\|_A \leq 2(2d-1)^2/q$, where q is the cardinality of the finite field GF(q)

Question 2: Is Extension for High θ Possible?

- No!
- Similar to the second open problem posed by Vaudenay's Theorem
- Bay et al. solved this in [BMV12] which can also be applied to our case

- For high θ, we give an iterated distinguisher of order 1 on a cipher:
- 3-round Feistel construction decorrelated to the order $2d$ such that $\|[C]^{2d} - [C^*]^{2d}\|_A \leq 8d^2/2^k$, where 2^k is the number of elements in GF(2^k)

Summary

- We consider more general distinguishers: iterated adaptive distinguishers
- They have not been analyzed using Decorrelation Theory
- We quantify the security of these distinguishers and provide a bound for them
- We answer two questions related to our theorem

Thank You!