Resistance Against Iterated Attacks by Decorrelation Revisited

Aslı Bay   Atefeh Mashatan   Serge Vaudenay

Ecole Polytechnique Fédéral de Lausanne (EPFL)

Outline

1. Decorrelation Theory
   - The Luby-Rackoff Model
   - Advantage of a non-adaptive adversary $\mathcal{A}$
   - Distribution matrix of a block cipher and its link with the advantage of the adversary $\mathcal{A}$
2. Solving two open problems
   - Necessary conditions for the security of block ciphers
   - Effects of the input distribution on the advantage of the adversary $\mathcal{A}$

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.
- Even provides the proof of security against not-yet discovered attacks
- Proves the security of several block ciphers such as:
  - 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

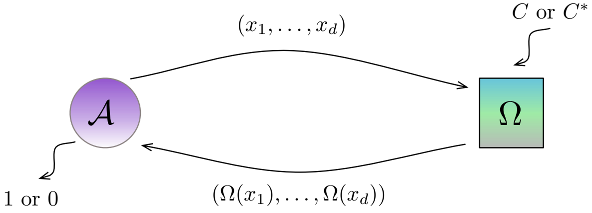

The Luby-Rackoff Model

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

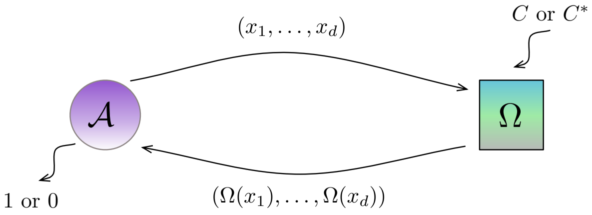

$$\mathrm{Adv}_{\mathcal{A}}(C, C^*) = \left| \Pr[\mathcal{A}(C) = 1] - \Pr[\mathcal{A}(C^*) = 1] \right|$$

When the d inputs are chosen **at once**, $\mathcal{A}$ is **non-adaptive**

— If advantage is negligible for all adversaries $\mathcal{A}$, then the cipher C is considered as **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 best 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 d**-wise distribution matrices**

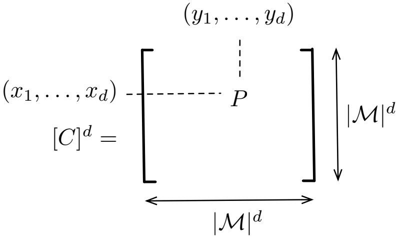

$$P = \Pr[C(x_1) = y_1, \ldots, C(x_d) = y_d]$$

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

$$\|A\|_\infty = \max_{x_1,\ldots,x_d} \sum_{y_1,\ldots,y_d} |A_{(x_1,\ldots,x_d)(y_1,\ldots,y_d)}|$$

A Non-adaptive Iterated Distinguisher of Order d

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

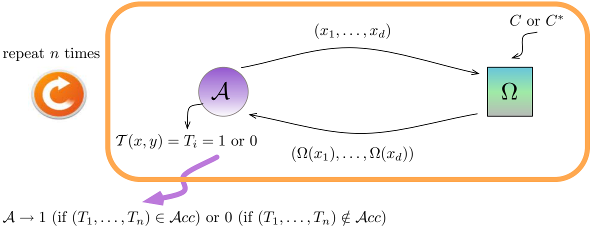

Examples:

Linear attacks have order $d = 1$

Differential attacks have order $d = 2$

Parameters: n, a distribution for 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

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 $2d$-decorrelared cipher C with $\|[C]^{2d} - [C^*]^{2d}\|_\infty \leq \varepsilon$ is

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

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

Two Open Problems

Two long-lasting open problems were posed by the previous Theorem

Problem 1: Could we extend to decorrelation of order $2d - 1$?

Two Open Problems

Two long-lasting open problems were posed by the previous Theorem

Problem 1: Could we extend 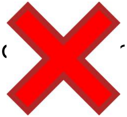 rrelation of order $2d - 1$?

Two Open Problems

Two long-lasting open problems were posed by the previous Theorem

Problem 1: Could we extend [illegible] rrelation of order $2d-1$?

Problem 2: Could we extend with a high δ ?

Two Open Problems

Two long-lasting open problems were posed by the previous Theorem

Problem 1: Could we extenc … rrelation of order $2d - 1$?

Problem 2: Could we extenc … igh δ ?

A 3-round Feistel Scheme C

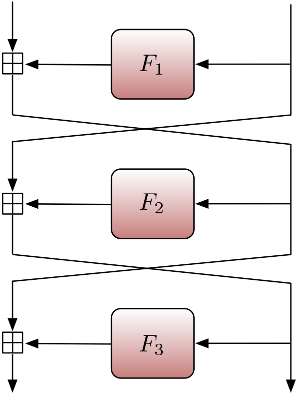

- $F_i(x) = a^i_{\kappa-1}x^{\kappa-1} + a^i_{\kappa-2}x^{\kappa-2} + \cdots + a^i_0$ over a finite field GF(p^k), $(a^i_{\kappa-1}, a^i_{\kappa-2}, \ldots, a^i_0) \in_U$ GF$(p^k)^\kappa$
- F_1, F_2 and F_3 are perfect κ-decorrelated functions
- C is a κ-decorrelated cipher with $\varepsilon = 2\kappa^2/p^k$ (Luby-Rackoff)

Solution of Problem 1: A cipher decorrelated to the order $2d-1$ may be broken by a non-adaptive iterated attack of order d

In this presentation: $d = 2$

How does the Distinguisher work?

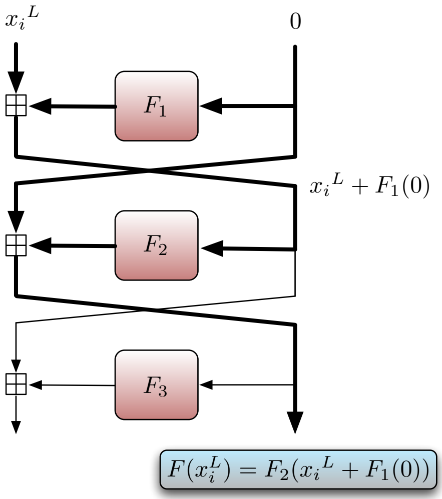

- Previous construction with $\kappa = 3$ over GF(p^k), $p > 2$
- We focus on F to distinguish the cipher C
- F is a random function: $F(x) = F_2(x + F_1(0))$, a polynomial degree ≤ 2

How does the Distinguisher work?

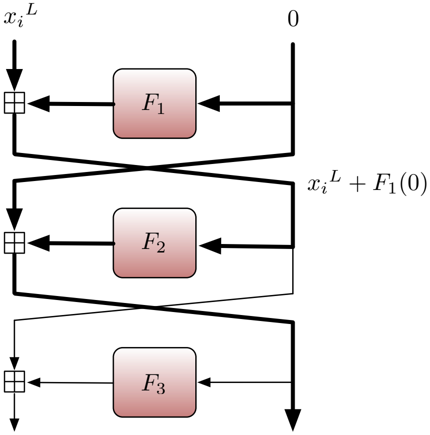

$$F(x_i^L) = F_2(x_i{}^L + F_1(0))$$

In each iteration, we have chosen plaintexts (x_1, x_2):

- $x_1 = x_1^L \| 0$ and $x_2 = x_2^L \| 0$
- $x_1^L + x_2^L = 0$
- $x_1^L \neq x_2^L$

How does the Distinguisher work?

Idea: Recovering a_1 of $F(x) = a_2x^2 + a_1x + a_0$

How does the Distinguisher work?

Idea: Recovering a_1 of $F(x) = a_2x^2 + a_1x + a_0$

- Send (x_1, x_2) s.t. $x_1^L + x_2^L = 0$ and $x_1^L \neq x_2^L$,
- Get $(y_1, y_2) = (\Omega(x_1), \Omega(x_2))$
- Solve

$$\left.\begin{array}{l} a_2(x_1^L)^2 + a_1x_1^L + a_0 = y_1^R \\ a_2(x_2^L)^2 + a_1x_2^L + a_0 = y_2^R \end{array}\right\} \Rightarrow a_1 = (y_1^R - y_2^R)(x_1^L - x_2^L)^{-1}$$

By only two iterations, F is distinguishable from F^* with high advantage

Solution of Problem 2: A cipher decorrelated to the order $2d$ may be broken by a non-adaptive iterated attack of order 1 (with high δ)

In this presentation: $d = 1$

How does the Distinguisher work?

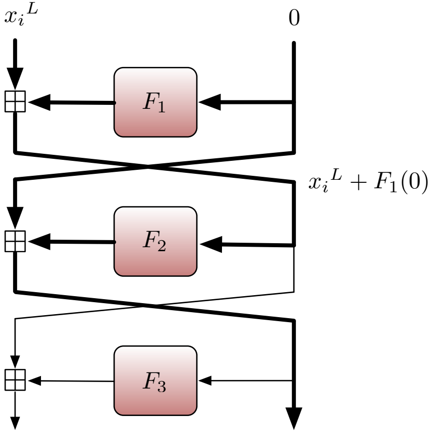

$F(x_i^L) = F_2(x_i{}^L + F_1(0))$

- Previous construction with $\kappa = 2$ over $\mathrm{GF}(2^k)$
- **Adversary's choice of the set of plaintexts is SMALL**:

  $S = \{x_1, x_2, x_3, x_4\}$

  - $x_i = x_i^L \| 0$, $1 \leq i \leq 4$
  - x_i's are pairwise distinct
  - $x_1^L + x_2^L + x_3^L + x_4^L = 0$
- In each iteration, a chosen plaintext x is taken from S

$$\delta = \frac{1}{4}$$

How does the Distinguisher work?

- **Reminder**: The **trace** of an element $\beta \in \mathsf{GF}(2^k)$ is defined as

$$\mathsf{Trace}(\beta) = \beta + \beta^2 + \cdots + \beta^{2^{k-1}}$$

- **A distinguishing property of F**:

$$\sum_{i=1}^{4} \mathsf{Trace}(F(x_i^L)) = 0, \quad \text{when } x_i = x_i^L \| 0 \in S,\ 1 \leq i \leq 4$$

How does the Distinguisher work?

- **Reminder**: The **trace** of an element $\beta \in \mathsf{GF}(2^k)$ is defined as

$$\mathsf{Trace}(\beta) = \beta + \beta^2 + \cdots + \beta^{2^{k-1}}$$

- **A distinguishing property of F**:

$$\sum_{i=1}^{4} \mathsf{Trace}(F(x_i^L)) = 0, \text{ when } x_i = x_i^L \| 0 \in S,\ 1 \leq i \leq 4$$

There is an **even** number of $F(x_i^L)$'s s.t. $\mathsf{Trace}(F(x_i^L)) = 1$

How does the Distinguisher work?

- Compute $T_i = \mathsf{Trace}(y_i^R)$ in each iteration
- Calculate the **average** $\bar{T} = \frac{1}{n}(T_1 + \cdots + T_n)$
- $E(\bar{T})$ and $E(\bar{T}^*)$:

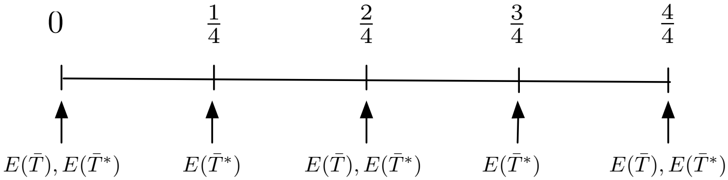

How does the Distinguisher work?

- Compute $T_i = \text{Trace}(y_i^R)$ in each iteration
- Calculate the **average** $\bar{T} = \frac{1}{n}(T_1 + \cdots + T_n)$
- $E(\bar{T})$ and $E(\bar{T}^*)$:

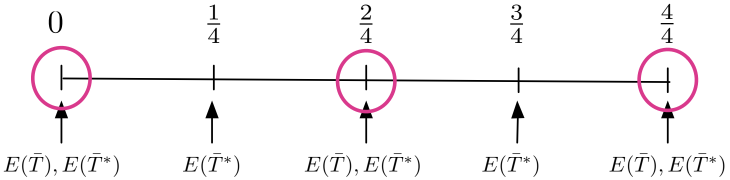

Distinguishing Set: $K = \bigcup_{m=0}^{2} \left(\frac{2m}{4} - \varepsilon, \frac{2m}{4} + \varepsilon \right)$, $\varepsilon > 0$

With 1000 iterations, F is distinguishable from F^* with high advantage

Conclusion

Two long-lasting open problems in Decorrelation Theory were settled:

- The $2d - 1$ decorrelation degree is not sufficient for a cipher to resist against a non-adaptive iterated distinguisher of order d
- When the probability of having a common query between different iterations is high, the advantage of the distinguisher **can** be high, too

Thanks...