

On the Optimality of Linear, Differential, and Sequential Distinguishers

Pascal Junod

Security and Cryptography Laboratory
Swiss Federal Institute of Technology
CH-1015 Lausanne, Switzerland
`pascal.junod@epfl.ch`

Abstract. In this paper, we consider the statistical decision processes behind a linear and a differential cryptanalysis. By applying techniques and concepts of statistical hypothesis testing, we describe precisely the shape of optimal linear and differential distinguishers and we improve known results of Vaudenay concerning their asymptotic behaviour. Furthermore, we formalize the concept of “sequential distinguisher” and we illustrate potential applications of such tools in various statistical attacks.

Keywords: Distinguishers, Statistical Hypothesis Testing, Linear Cryptanalysis, Differential cryptanalysis

1 Introduction

Historically, statistical procedures are indissociable of cryptanalytic attacks against block ciphers. One of the first attack exploiting statistical correlations in the core of DES [24] is Davies and Murphy’s attack [9]. Biham and Shamir’s differential cryptanalysis [1,2,3], Matsui’s attack against DES [17,18], Vaudenay’s statistical and χ^2 cryptanalysis [29], Harpes and Massey’s partitioning cryptanalysis [13], and Gilbert-Minier stochastic cryptanalysis [21] are attacks using statistical procedures in their core.

To the best of our knowledge, Murphy *et al.*, in an unpublished report [22], proposed for the first time a general statistical framework for the analysis of block ciphers using the technique of *likelihood estimation*. Other examples can be found in the field of cryptology: recently, Coppersmith, Halevi and Jutla [7] have devised a general statistical framework for analysing stream ciphers; they use the concept of statistical hypothesis testing for systematically distinguishing a stream cipher from a random function. Other examples (this list being non-exhaustive) include Maurer’s analysis of Simmon’s authentication theory [19,20] and Cachin’s theoretical treatment of steganography [4,5].

In a parallel way, some attempts of formalizing resistance of block ciphers towards cryptanalytic attacks have been proposed: for instance, Pornin [25] proposes a general criterion of resistance against the Davies and Murphy attack; for this purpose, he makes use of statistical hypothesis testing. Vaudenay, in a

sequence of papers (*e.g.* [30,28]) proposes the *decorrelation theory* as a generic technique for estimating the strength of block ciphers against various kinds of attacks. In these papers, he notably derives bounds on the best advantage of any linear and differential distinguishers, however without using statistical hypothesis testing concepts.

As pointed out by many authors, statistical hypothesis tests are convenient in the analysis of statistical problems, since, in certain cases, well-known optimality results (like the Neyman-Pearson lemma, for instance) can be applied.

Contributions of This Paper

In this paper, we consider the resistance of block ciphers against linear and differential cryptanalysis as a statistical hypothesis testing problem, which allows us to improve Vaudenay’s asymptotic bounds on the best advantage of any linear and differential distinguishers and to give optimality results on the decision processes involved during these attacks.

For this, we recall some well-known statistical concepts in Section §2. In Section §3, we treat linear distinguishers and we derive a Chernoff-like bound, which gives the right asymptotic behaviour of the best advantage of such distinguishers. In §4, we do the same for differential distinguishers. In §5, we formalize the notion of *sequential distinguisher*; this kind of statistical procedure has been recognized quite early as potentially useful (in [22,9], for instance). We restate this by showing, with help of a toy-example (a linear cryptanalysis of 5-rounds DES), that sequential sampling procedures may divide the needed number of plaintext-ciphertext pairs by a non-negligible factor in certain statistical cryptanalysis. In §6, we discuss potential applications of statistical hypothesis testing concepts in various attacks, and finally, we conclude in §7.

Notation. The following notation will be used throughout this paper. Random variables¹ X, Y, \dots are denoted by capital letters, while *realizations* $x \in \mathcal{X}, y \in \mathcal{Y}, \dots$ of random variables are denoted by small letters; random vectors $\mathbf{X}, \mathbf{Y}, \dots$ and their realizations $\mathbf{x}, \mathbf{y}, \dots$ are denoted in bold characters. The fact for a random variable X to follow a distribution D is denoted $X \leftarrow D$, while its probability function is denoted by $\Pr_X[x]$. Finally, as usual, “iid” means “independent and identically distributed”.

2 Statistical Hypothesis Testing

We recall some well-known facts about statistical hypothesis testing, both in the classical and in the Bayesian approaches; details can be found in any good book on statistics (*e.g.* see [26]).

¹ In this paper, we are only dealing with discrete random variables.

2.1 Classical Approach

Let D_0 and D_1 be two different probability distributions defined on the same finite set \mathcal{X} . In a *binary hypothesis testing problem*, one is given an element $x \in \mathcal{X}$ which was drawn according either to D_0 or to D_1 and one has to decide which is the case. For this purpose, one defines a so-called *decision rule*, which is a function $\delta : \mathcal{X} \rightarrow \{0, 1\}$ taking a sample of X as input and defining what should be the guess for each possible $x \in \mathcal{X}$. Associated to this decision rule are two different types of error probabilities: $\alpha \triangleq \Pr_{X_0}[\delta(x) = 1]$ and $\beta \triangleq \Pr_{X_1}[\delta(x) = 0]$. The decision rule δ defines a partition of \mathcal{X} in two subsets which we denote by \mathcal{A} and $\overline{\mathcal{A}}$, *i.e.* $\mathcal{A} \cup \overline{\mathcal{A}} = \mathcal{X}$; \mathcal{A} is called the *acceptance region* of δ . We recall now the Neyman-Pearson lemma, which derives the shape of the optimal statistical test δ between two simple hypotheses, *i.e.* which gives the optimal decision region \mathcal{A} .

Lemma 1 (Neyman-Pearson). *Let X be a random variable drawn according to a probability distribution D and let be the decision problem corresponding to hypotheses $X \leftarrow D_0$ and $X \leftarrow D_1$. For $\tau \geq 0$, let \mathcal{A} be defined by*

$$\mathcal{A} \triangleq \left\{ x \in \mathcal{X} : \frac{\Pr_{X_0}[x]}{\Pr_{X_1}[x]} \geq \tau \right\} \quad (1)$$

Let $\alpha^ \triangleq \Pr_{X_0}[\overline{\mathcal{A}}]$ and $\beta^* \triangleq \Pr_{X_1}[\mathcal{A}]$. Let \mathcal{B} be any other decision region with associated probabilities of error α and β . If $\alpha \leq \alpha^*$, then $\beta \geq \beta^*$.*

Hence, the Neyman-Pearson lemma indicates that the optimum test (regarding error probabilities) in case of a binary decision problem is the *likelihood-ratio test*. All these considerations are summarized in Definition 1.

Definition 1 (Optimal Binary Hypothesis Test). *To test $X \leftarrow D_0$ against $X \leftarrow D_1$, choose a constant $\tau > 0$ depending on α and β and define the likelihood ratio*

$$\text{lr}(x) \triangleq \frac{\Pr_{X_0}[x]}{\Pr_{X_1}[x]} \quad (2)$$

The optimal decision function is then defined by

$$\delta_{\text{opt}} \triangleq \begin{cases} 0 & (\text{i.e. accept } X \leftarrow D_0) \text{ if } \text{lr}(x) \geq \tau \\ 1 & (\text{i.e. accept } X \leftarrow D_1) \text{ if } \text{lr}(x) < \tau \end{cases} \quad (3)$$

We note that Lemma 1 does not consider any special hypothesis on the observed random variable X . In the following, we will assume that we are interested in taking a decision about the distribution of a random vector $\mathbf{X} \triangleq (X_1, \dots, X_n)$ where X_1, \dots, X_n are iid random variables, *i.e.* $\mathbf{X} \leftarrow D^n$ is a random vector of n independent samples of the random variable X . This is a typical situation during a known-plaintext attack.

When dealing with error probabilities, one usually proceeds as follows in the classical approach: one of the two possible error probabilities is fixed, and one minimizes the other error probability. In this case, Stein's lemma (we refer to [8] for more details) gives the best error probability expression. As this approach lacks symmetry, we won't describe it in more details.

2.2 Bayesian Approach

The other possibility is to follow a Bayesian approach and to assign *prior* probabilities π_0 and π_1 to both hypotheses, respectively, and *costs* $\kappa_{i,j} \geq 0$ to the possible decisions $i \in \{0, 1\}$ and states of nature $j \in \{0, 1\}$. In this case, we would like to minimize the *expected cost*. If we assign $\kappa_{0,0} = \kappa_{1,1} \triangleq 0$ and $\kappa_{0,1} = \kappa_{1,0} \triangleq 1$, *i.e.* correct decisions are not penalized, while incorrect decisions are penalized equally, then the optimal Bayesian decision rule is given by

$$\delta(\mathbf{x}) \triangleq \begin{cases} 0 & \text{if } \pi_0 \Pr_{X_0^n}[\mathbf{x}] \geq \pi_1 \Pr_{X_1^n}[\mathbf{x}] \\ 1 & \text{if } \pi_0 \Pr_{X_0^n}[\mathbf{x}] < \pi_1 \Pr_{X_1^n}[\mathbf{x}] \end{cases} \quad (4)$$

Clearly, the overall error probability $P_e^{(n)} \triangleq \pi_0 \alpha^{(n)} + \pi_1 \beta^{(n)}$ of such an optimal Bayesian distinguisher must decrease towards zero as the number n of samples increases. It turns out that the decrease asymptotically approaches an exponential in the number of samples drawn before the decision, the exponent being given by the so-called *Chernoff bound* (see Theorem 1; in the long version [14] of this paper, we give some information-theoretic results justifying this bound, and we refer to [8] for a detailed and complete treatment).

Theorem 1 (Chernoff). *The best probability of error of the Bayesian decision rule defined in (4) satisfies*

$$\lim_{n \rightarrow +\infty} \frac{1}{n} \log \frac{P_e^{(n)}}{2^{-n\nu}} = 0 \quad (5)$$

where $\nu = C(D_0, D_1)$ is the Chernoff information² between D_0 and D_1 .

Note that the Bayesian error exponent does not depend on the actual value of π_0 and π_1 , as long as they are non-zero: essentially, the effect of the prior is washed out for large sample sizes.

3 Linear Distinguishers

In this section, we consider the classical model of a linear distinguisher and we present several new results derived using tools of statistical hypothesis testing.

3.1 Introduction

A *linear distinguisher* δ_{lin} is a (possibly computationally unbounded) Turing machine which can play with an oracle Ω implementing a permutation C ; δ_{lin} is bounded in the number n of queries to the oracle Ω . Furthermore, it uses a *linear characteristic* (\mathbf{a}, \mathbf{b}) which is a pair of boolean vectors. Algorithm 1 defines the classical modelization of a linear distinguisher (see [30]).

² The Chernoff information between two discrete probability distributions D_0 and D_1 is

$$C(D_0, D_1) \triangleq - \min_{0 \leq \lambda \leq 1} \log \left(\sum_{x \in \mathcal{X}} \Pr_{X_0}[x]^\lambda \Pr_{X_1}[x]^{1-\lambda} \right)$$

```

1: Parameters: a complexity  $n$ , a characteristic  $(\mathbf{a}, \mathbf{b})$ , an acceptance region  $\mathcal{A}^{(n)}$ 
2: Input: an oracle  $\Omega$  which implements a permutation  $C$ 
3: Initialize a counter  $u$  to 0.
4: for  $i = 1 \dots n$  do
5:   Pick uniformly at random  $x$  and query  $C(x)$  to the oracle  $\Omega$ .
6:   if  $\mathbf{a} \cdot x = \mathbf{b} \cdot C(x)$  then
7:     Increment  $u$ 
8:   end if
9: end for
10: if  $u \in \mathcal{A}^{(n)}$  then
11:   Output 0
12: else
13:   Output 1
14: end if

```

Algorithm 1: Modelization of a linear distinguisher δ_{lin} .

The statistical game is the following. One gives an oracle Ω to Algorithm 1, which is with probability $\pi_0 = \frac{1}{2}$ the permutation C or, with probability $\pi_1 = \frac{1}{2}$, a permutation $C^* \in_U \mathcal{C}_m$ drawn uniformly at random from the set \mathcal{C}_m of all permutations over inputs of size m (C^* is often referred as the ‘‘Perfect Cipher’’). The goal of Algorithm 1 is to decide whether Ω implements C or C^* . One measures the performance of a distinguisher δ_{lin} by the expression

$$\text{Adv}_{\delta_{\text{lin}}}^n(C, C^*) \triangleq \left| \Pr_C[\delta_{\text{lin}}(\mathbf{x}) = 1] - \Pr_{C^*}[\delta_{\text{lin}}(\mathbf{x}) = 1] \right| = \left| 2P_e^{(n)} - 1 \right| \quad (6)$$

where $\mathbf{x} = (x_1, \dots, x_n)$ is the vector of the values queried to the oracle. The distinguisher’s core is the *acceptance region* $\mathcal{A}^{(n)}$: it defines the set of values (x_1, \dots, x_n) which lead to output 0 (*i.e.* it decides that the oracle implements C) or 1 (*i.e.* it decides that the oracle implements C^*).

As pointed out by Chabaud and Vaudenay in [6], linear cryptanalysis is based on the quantity $\text{LP}^C(\mathbf{a}, \mathbf{b}) \triangleq (2 \cdot \Pr_X[\mathbf{a} \cdot X = \mathbf{b} \cdot C(X)] - 1)^2$. This value depends of the (fixed) permutation C and of the distribution of plaintext, which is usually defined to be uniform. Actually, most of the time, a cryptanalyst does not possess any information about the permutation (*i.e.* about the key), so one is more interested in the *average* $\text{LP}^C(\mathbf{a}, \mathbf{b})$ over the permutation space \mathcal{C}_m (or, equivalently, over the key space \mathcal{K}); this quantity is denoted $\text{ELP}(\mathbf{a}, \mathbf{b}) \triangleq \mathbb{E}[\text{LP}^C(\mathbf{a}, \mathbf{b})]$, where the expectation is taken over the permutation distribution.

When studying linear distinguishers, one is interested in bounding the advantage of any linear distinguisher in terms of $\text{ELP}(\mathbf{a}, \mathbf{b})$. We review now a known result of Vaudenay (see [28], for instance).

Theorem 2 (Vaudenay). *For any distinguisher in the model of Algorithm 1*

$$\text{BestAdv}_{\delta_{\text{lin}}}^n(C, C^*) \leq 2.78 \sqrt[3]{n \cdot \text{ELP}(\mathbf{a}, \mathbf{b})} + 2.78 \sqrt[3]{\frac{n}{2^m - 1}} \quad (7)$$

where m is the block size of the permutation.

In the case of a practical linear cryptanalysis of DES [18], we have $\text{ELP}(\mathbf{a}, \mathbf{b}) \approx 4 \cdot (1.19 \cdot 2^{-21})^2 \approx 1.288 \cdot 10^{-12}$ and $m = 64$, which means that (7) is useful as long as $n \leq 2^{35}$. Thus, although of great theoretical interest, we note that (7) is not tight for large n , or, in other words, does not capture the asymptotical behavior of the advantage. In the next part, we reconsider this problem in the statistical hypothesis testing framework and we derive an asymptotically tight Chernoff-like bound on the best advantage of any linear distinguisher.

3.2 New Asymptotic Bounds

First, we note that if δ_{lin} is optimal, then $P_e^{(n)} \leq \frac{1}{2}$ for all $n > 0$ (otherwise, we could modify it such that it outputs the opposite decision as defined in Algorithm 1 and get a smaller error probability). Thus, we have $\text{Adv}_{\delta_{\text{lin}}}^n(C, C^*) = 1 - 2P_e^{(n)}$. As outlined before, the crucial part of δ_{lin} is the acceptance region $\mathcal{A}^{(n)}$. The following lemma, which is a direct application of Lemma 1, gives the optimal $\mathcal{A}_{\text{opt}}^{(n)}$, *i.e.* the region producing the smallest overall error probability. Without loss of generality, we assume that $\mathbb{E}[\text{Pr}_X[\mathbf{a} \cdot X = \mathbf{b} \cdot C(X)]] \triangleq \frac{1}{2} + \epsilon$ with $\epsilon > 0$ where the expectation is taken over a uniformly distributed plaintext space \mathcal{X} and the key space \mathcal{K} .

Lemma 2. *The optimal acceptance region for δ_{lin} is*

$$\mathcal{A}_{\text{opt}}^{(n)} = \left\{ u \in \{0, \dots, n\} : u \geq n \cdot \frac{\log_2(1 - 2\epsilon)}{\log_2(1 - 2\epsilon) - \log_2(1 + 2\epsilon)} \right\} \quad (8)$$

where u is defined in Algorithm 1.

See [14] for a detailed proof. Note that, for ϵ small, one can approximate (8) with

$$\mathcal{A}_{\text{opt}}^{(n)} \approx \left\{ u \in \{0, \dots, n\} : u \geq n \cdot \left(\frac{1}{2} + \frac{\epsilon}{2} \right) \right\} \quad (9)$$

Using a precise version of Chernoff's theorem 1, we can bound the advantage of the best linear distinguisher as follows (see [14] for a detailed proof):

Theorem 3. *Let m be the block size of the involved permutations. For any distinguisher in the model of Algorithm 1*

$$1 - \frac{(n+1)}{2^{n\nu-1}} \leq \text{BestAdv}_{\delta_{\text{lin}}}^n(C, C^*) \leq 1 - \frac{1}{(n+1) \cdot 2^{n\nu-1}} \quad (10)$$

where $\nu = C(\mathbb{D}_0, \mathbb{D}_1)$ is the Chernoff information between \mathbb{D}_0 , a binary distribution having a bias equal to $\max\{\frac{1}{2^{m-1}}, \epsilon\}$ such that $\text{ELP}^C(\mathbf{a}, \mathbf{b}) = 4\epsilon^2$ and the uniform binary distribution \mathbb{D}_1 .

Generally, the Chernoff information cannot be expressed explicitly, because one has to solve a transcendental equation. However, in the case which interests us, $\nu(\lambda) \triangleq 2^{-\lambda} \cdot ((\frac{1}{2} + \epsilon)^\lambda + (\frac{1}{2} - \epsilon)^\lambda)$ and

$$C(D_0, D_1) = \nu(\lambda^*) \text{ for } \lambda^* = \frac{\log\left(-\frac{\log(1-2\epsilon)}{\log(1+2\epsilon)}\right)}{\log\left(\frac{1+2\epsilon}{1-2\epsilon}\right)} \tag{11}$$

We give now a numerical illustration: for $\epsilon = 1.19 \cdot 2^{-21}$ (which is the bias of the best linear approximation of 14 rounds of DES), we obtain a useful lower bound only for $n \geq 2^{48.2}$; unfortunately, even if it captures the asymptotic exponential shape of the best advantage curve, it is not practically useful for “interesting” values of n ; for which concerns the upper bound, it is useful for all n but it is not tight: one may give a tighter lower bound using Bernstein’s inequality (see Theorem 4 and [11] for a proof). In the following, we will assume that ϵ is small and thus that one is using (9) as acceptance region.

Theorem 4 (Bernstein’s Inequality). *Let X_i be iid discrete random variables following a Bernoulli law with parameter $0 \leq p \leq 1$ and let $S_n \triangleq \sum_i X_i$. Then*

$$\Pr[S_n \geq n(p + \epsilon)] \leq e^{-\frac{1}{4}n\epsilon^2} \text{ for } \epsilon > 0 \tag{12}$$

This allows to derive in an easy way the following lower bound:

Theorem 5. *Let m be the block size of the involved permutations. For any distinguisher in the model of Algorithm 1*

$$\text{BestAdv}_{\delta_{\text{lin}}}^n(C, C^*) \geq 1 - e^{-\frac{n\bar{\epsilon}^2}{16}} \tag{13}$$

where $\bar{\epsilon} \triangleq \max\{\frac{1}{2^m-1}, \epsilon\}$ such that $\text{ELP}^C(\mathbf{a}, \mathbf{b}) = 4\epsilon^2$.

4 Differential Distinguishers

Similarly, one can study *differential distinguishers* with the same tools. A differential distinguisher δ_{diff} is a (possibly computationally unbounded) Turing machine which is able to submit chosen *pairs* of plaintexts to an oracle Ω , implementing with probability $\pi_0 = \frac{1}{2}$ a fixed permutation C or, with probability $\pi_1 = \frac{1}{2}$, a permutation drawn uniformly at random from the set \mathcal{C}_m of all permutations on m -bit blocks. Although the cryptanalytic settings are quite different (δ_{diff} can submit *chosen pairs* of plaintext), in a statistical point of view, the distinguishing process is very similar to linear distinguishers. In Algorithm 2, the classical modelization of a differential distinguisher [30] is given.

If we look at Algorithm 2, we note that, although the complexity n is given in advance as input and is (implicitly) *fixed*, the effective number of queries to the oracle is merely a random variable. In other words, δ_{diff} does not make use of all

the information that it could exploit. In fact, we can see the class of distinguishers submitting a *random* number of queries to the oracle as a generalization of the class of distinguishers submitting a *fixed* number of queries. We will call this generalization *sequential distinguishers*; this new concept is formalized and studied in Section 5.

```

1: Parameters: a complexity  $n$ , a characteristic  $(a, b)$ 
2: Input: an oracle  $\Omega$  which implements a permutation  $C$ 
3: for  $i = 1 \dots n$  do
4:   Pick uniformly at random  $x$  and query  $C(x)$  and  $C(x + a)$  to the oracle  $\Omega$ .
5:   if  $C(x + a) = C(x) + b$  then
6:     Output 0 and stop.
7:   end if
8: end for
9: Output 1.

```

Algorithm 2: Classical modelization of a differential distinguisher δ_{diff} .

In order to better understand the statistical decision processes, we give in Algorithm 3 an “unorthodox” modelization, denoted δ'_{diff} , which is very similar to the linear one. As for linear distinguishers, it is well-known [23] that differential

```

1: Parameters: a complexity  $n$ , a characteristic  $(a, b)$ , an acceptance region  $\mathcal{A}^{(n)}$ 
2: Input: an oracle  $\Omega$  which implements a permutation  $C$ 
3: Initialize a counter  $u$  to 0.
4: for  $i = 1 \dots n$  do
5:   Pick uniformly at random  $x$  and query  $C(x)$  and  $C(x + a)$  to the oracle  $\Omega$ .
6:   if  $C(x + a) = C(x) + b$  then
7:     Increment  $u$ 
8:   end if
9: end for
10: if  $u \in \mathcal{A}^{(n)}$  then
11:   Output 0
12: else
13:   Output 1
14: end if

```

Algorithm 3: Unorthodox modelization of a differential distinguisher δ'_{diff} .

cryptanalysis depends on the quantity $\text{DP}^C(a, b) \triangleq \Pr_X[C(X + a) = C(X) + b]$, where the plaintext space \mathcal{X} is uniformly distributed. As this value depends on the choice of the cipher (*i.e.* on the key), one defines $\text{EDP}(a, b) \triangleq \mathbb{E}[\text{DP}^C(a, b)]$, where the expectation is taken over the permutation space. We note that Algorithm 2 outputs 1 if and only if no *differential event* occurs. As for linear distinguishers, and considering this time Algorithm 3, one can define the optimal acceptance region using Lemma 1 and which is given by Lemma 3. As

$\text{EDP}(a, b) = \frac{1}{2^{m-1}}$ (where m is the block size of the permutation), and, typically, $\text{DP}^C(a, b) \triangleq \frac{1+\epsilon}{2^{m-1}}$ with $0 < \epsilon \leq 2^m - 2$, we can note that the *optimal* acceptance region will make δ_{diff} output 0 if

$$\binom{n}{u} \left(\frac{1 + \epsilon}{2^m - 1} \right)^u \left(1 - \frac{1 + \epsilon}{2^m - 1} \right)^{n-u} \geq \binom{n}{u} \left(\frac{1}{2^m - 1} \right)^u \left(\frac{2^m - 2}{2^m - 1} \right)^{n-u}$$

which gives the following result.

Lemma 3. *The optimal acceptance region for δ'_{diff} is*

$$\mathcal{A}_{\text{opt}}^{(n)} = \left\{ u \in \{0, \dots, n\} : u \geq n \cdot \frac{\log(2^m - 2) - \log(2^m - 2 - \epsilon)}{\log((2^m - 2)(1 + \epsilon)) - \log(2^m - 2 - \epsilon)} \right\} \quad (14)$$

where u is defined in Algorithm 3.

For small ϵ , (14) may be approximated by

$$\mathcal{A}_{\text{opt}}^{(n)} \approx \left\{ u \in \{0, \dots, n\} : u \geq n \cdot \left(\frac{1}{2^m - 1} + \frac{2^{m-1} - 1}{(2^m - 2)(2^m - 1)} \cdot \epsilon \right) \right\} \quad (15)$$

Thus, we have

Corollary 1. *δ_{diff} is an optimal differential distinguisher submitting n queries to the oracle if and only if (14) is satisfied for all $u \in \mathbb{N}$ with $1 < u \leq n$ and for all $0 < \epsilon \leq 2^m - 2$.*

It is not difficult to build artificially a situation where Algorithm 2 is not optimal: it is sufficient to take a characteristic (a, b) with $\text{DP}^C(a, b)$ having a *very high probability*. In this case, it is not sufficient for δ_{diff} to wait for only *one* differential event and to stop, since if it is unique during the n samplings, it would have been better to output 1. However, if we have a look at (15), we can note that Algorithm 2 captures well real-world situations, where exploited differential probabilities are only slightly greater than ideal ones.

A very similar proof of Theorem 3 leads to

Theorem 6. *For any distinguisher in the model of δ'_{diff} ,*

$$1 - \frac{n + 1}{2^{n\nu-1}} \leq \text{BestAdv}_{\delta'_{\text{diff}}}^n(C, C^*) \leq 1 - \frac{1}{(n + 1) \cdot 2^{n\nu-1}} \quad (16)$$

where $\nu = C(D_0, D_1)$ is the Chernoff information between D_0 , a binary distribution with $\Pr_{X_0}[X_0 = 0] = 1 - \Pr_{X_0}[X_0 = 1] = \text{DP}^C(a, b)$, and D_1 , a binary distribution with $\Pr_{X_1}[X_1 = 0] = 1 - \Pr_{X_0}[X_1 = 1] = \frac{1}{2^m - 1}$.

Usually, in the context of differential cryptanalysis, one encounters the concept of *signal-to-noise ratio*, which was used by Biham and Shamir in the papers defining the differential cryptanalysis [1,2,3]; it is defined as being the ratio of probability of the right (sub-)key being suggested by a right pair and the

probability of a random (sub-)key being suggested by a random pair, given the initial difference. By empirical evidence, they suggested that when this ratio is around 1-2, about 20-40 right pairs are sufficient for a successful attack; clearly, this is a (implicitly defined) likelihood-ratio test which turns out to be optimal in terms of error probabilities.

5 Sequential Distinguishers

In this section, we formalize the concepts of *generic sequential non-adaptive distinguisher (GSNAD)* and of *n-limited generic sequential non-adaptive distinguisher (n-limited GSNAD)*. These kinds of distinguishers use sequential sampling procedures as their statistical core. We note that this idea was used earlier by Davies and Murphy (see Appendix of [9]) in an attempt to decrease the complexity of their attack against DES.

In the Luby-Rackoff model [16], a non-adaptive attacker (which may be modeled by an *n-limited GNAD* as described in Algorithm 4) is an infinitely powerful Turing machine which has access to an oracle Ω . It aims at distinguishing a cipher C from the “Perfect Cipher” C^* by querying Ω , and with a limited number n of inputs. The attacker must finally take a decision; usually, one is interested in measuring the ability (*i.e.* the advantage as defined in (6)) to distinguish C from C^* for a given, *fixed* amount n of queries. Clearly, in this model, one is interested in maximizing the advantage given a fixed number of queries.

In a more “real-life” situation, a cryptanalyst proceeds usually in an inverse manner: given a fixed success probability (*i.e.* a given advantage), she may look for minimizing the amount of queries to Ω , since such queries are typically expensive. With this model in head, we can now define a *n-limited generic sequential non-adaptive distinguisher* (see Algorithm 5), which turns out to be more efficient in terms of the average number of oracle queries than Algorithm 4 given a fixed advantage. In fact, such a distinguisher is adaptive regarding the decision process.

After having received the i -th response from the oracle, the distinguisher compare the i responses it has at disposal towards an acceptance set \mathcal{A}_i and a rejection set \mathcal{B}_i , which depend on the number of queries and on the (fixed in advance) advantage, and can then take *three* different decisions: either it decides

```

1: Parameters: a complexity  $n$ , an acceptance set  $\mathcal{A}$ .
2: Input: an oracle  $\Omega$  implementing a permutation  $C$ 
3: Compute some messages  $\mathbf{x} = (x_1, \dots, x_n)$ .
4: Query  $\mathbf{y} = (C(x_1), \dots, C(x_n))$  to  $\Omega$ .
5: if  $\mathbf{y} \in \mathcal{A}$  then
6:   Output 0
7: else
8:   Output 1
9: end if

```

Algorithm 4: A n -limited generic non-adaptive distinguisher (GNAD)

to output “0” or “1” and to stop, or to query one more question to the oracle and to repeat the decision process, until it has queried n questions. Note that $\mathcal{A}_i \subseteq \mathcal{Y}^i$ and $\mathcal{B}_i \subseteq \mathcal{Y}^i$ are disjoint sets for all $1 \leq i \leq n$ and that $\mathcal{A}_n \cup \mathcal{B}_n = \mathcal{Y}^n$ (*i.e.* the distinguisher *must* stop and take a decision at step n). In statistics, this process is known as a *sequential decision procedure*.

We note that Algorithm 2 can be viewed as a sequential differential distinguisher which does not take explicitly into account a decision region, since it always outputs 0 as soon as it observes a “differential event”.

```

1: Parameters: a complexity  $n$ , acceptance sets  $\mathcal{A}_i, 1 \leq i \leq n$  and rejection
   sets  $\mathcal{B}_i, 1 \leq i \leq n$ .
2: Input: an oracle  $\Omega$  implementing a permutation  $C$ 
3:  $i \leftarrow 1$ 
4: repeat
5:   Select non-adaptively a message  $x_i$  and get  $y_i = C(x_i)$ .
6:   if  $(y_1, \dots, y_i) \in \mathcal{A}_i$  then
7:     Output 0 and stop.
8:   else if  $(y_1, \dots, y_i) \in \mathcal{B}_i$  then
9:     Output 1 and stop.
10:  end if
11:   $i \leftarrow i + 1$ 
12: until  $i = n$ 

```

Algorithm 5: A n -limited sequential generic non-adaptive distinguisher

5.1 Sequential Statistical Inference

We describe now formally the *sequential decision procedure* behind Algorithm 5. Let \mathcal{D} be the set of possible decisions.

Definition 2 (Sequential decision procedure). *Let X_1, X_2, \dots be random variables observed sequentially. A sequential decision procedure consists in:*

1. a stopping rule σ_n which specifies whether a decision must be taken without taking any further observation. If at least one observation is taken, this rule specifies for every set of observed values (x_1, \dots, x_n) , with $n \geq 1$, whether to stop sampling and take a decision out of \mathcal{D} or to take another observation x_{n+1} .
2. a decision rule δ_n which specifies the decision to be taken. If $n \geq 1$ observations have been taken, then one takes an action $\delta_n(x_1, \dots, x_n) \in \mathcal{D}$. Once a decision has been taken, the sampling process is stopped.

If we consider Algorithm 5 at the light of this formalism, $\mathcal{D} = \{0, 1\}$,

$$\delta_n(x_1, \dots, x_n) = \begin{cases} 0 & \text{if } (x_1, \dots, x_n) \in \mathcal{A}_n \\ 1 & \text{if } (x_1, \dots, x_n) \in \mathcal{B}_n \end{cases} \quad (17)$$

and

$$\sigma_n(x_1, \dots, x_n) = \begin{cases} \text{continue sampling if } (x_1, \dots, x_n) \notin \mathcal{A}_n \cup \mathcal{B}_n \\ \text{stop sampling if } (x_1, \dots, x_n) \in \mathcal{A}_n \cup \mathcal{B}_n \end{cases} \quad (18)$$

5.2 Sequential Decision Procedures

We have seen that Lemma 1 defines the shape of the optimal acceptance region for binary hypothesis testing. Theoretically, if one is able to compute the exact joint probability distribution of the oracle's responses when it implements both ciphers, one is able to compute the optimal acceptance region \mathcal{A} for a generic n -limited distinguisher.

A sequential likelihood-ratio test uses exactly the same process to define *two* types of acceptance regions, denoted \mathcal{A} and \mathcal{B} , respectively. So, it is always possible to define a sequential test when one has a classical test at disposal. In few words, a sequential test has three alternatives once it has received a response from the oracle: either it can conclude for one of both hypotheses, or it can decide to query more samples. In its simpler definition, a sequential ratio test has the possibility to query *as many samples as it is needed to take a decision*, given a fixed error probability. The *expected* number of queries required to reach one of the two possible decision turns out to be less than it would need in order to make the same decision on the basis of a single *fixed-size* sample set. Of course it may happen that the sequential procedure will take more queries than the fixed-size one, but sequential sampling is a definitely economical procedure.

One may define Algorithm 5, as a *truncated sequential test*, *i.e.* one fixes an upper-bound n on the number of queries; it is still clear that such a sequential procedure cannot be worse than a fixed-size sampling procedure. In the following, we state some definitions and results about sequential hypothesis tests.

Definition 3 (Sequential Likelihood-Ratio Test). *To test $\mathbf{X} \leftarrow D_0$ against $\mathbf{X} \leftarrow D_1$, define two constants $\tau_{\text{up}} > \tau_{\text{down}} > 0$ depending on α and β , and define the likelihood ratio*

$$\text{lr}(\mathbf{x}) \triangleq \frac{f_{\mathbf{X}_1}(\mathbf{x})}{f_{\mathbf{X}_0}(\mathbf{x})}$$

The decision function at i -th step is

$$\delta_{\text{opt}} \triangleq \begin{cases} 1 \text{ (i.e. accept } \mathbf{X} \leftarrow D_1) & \text{if } \text{lr}(\mathbf{x}^{(i)}) \geq \tau_{\text{up}} \\ 0 \text{ (i.e. accept } \mathbf{X} \leftarrow D_0) & \text{if } \text{lr}(\mathbf{x}^{(i)}) \leq \tau_{\text{down}} \\ \emptyset \text{ query another sample} & \text{otherwise} \end{cases} \quad (19)$$

When the observations are *independent and identically distributed*, then sequential likelihood-ratio tests have the following nice property (we refer to [27] as an excellent treatment of sequential procedures and for the proof of the following three theorems):

Theorem 7. *For testing a simple hypothesis against a simple alternative with independent, identically distributed observations, a sequential probability ratio test is optimal in the sense of minimizing the expected sample size among all tests having no larger error probabilities.*

The following results relate error probabilities α and β to τ_{up} and τ_{down} , and give an approximation of the expected number of samples.

Theorem 8. *Let be a sequential likelihood-ratio test with stopping bounds τ_{up} and τ_{down} , with $\tau_{\text{up}} > \tau_{\text{down}}$ and error probabilities $0 < \alpha < 1$ and $0 < \beta < 1$, then $\tau_{\text{down}} \geq \frac{\beta}{1-\alpha}$ and $\tau_{\text{up}} \leq \frac{1-\beta}{\alpha}$.*

The approximation $\tau_{\text{down}} \triangleq \frac{\beta}{1-\alpha}$ and $\tau_{\text{up}} \triangleq \frac{1-\beta}{\alpha}$ is known as “Wald’s approximation”. The following theorem gives some credit to this approximation.

Theorem 9. *Let us assume we select for given $\alpha, \beta \in]0, 1[$, where $\alpha + \beta \leq 1$, the stopping bounds $\tau'_{\text{down}} \triangleq \frac{\beta}{1-\alpha}$ and $\tau'_{\text{up}} \triangleq \frac{1-\beta}{\alpha}$. Then it holds that the sequential likelihood-ratio test with stopping bounds τ'_{down} and τ'_{up} has error probabilities α' and β' where $\alpha' \leq \frac{\alpha}{1-\beta}$, $\beta' \leq \frac{\beta}{1-\alpha}$ and $\alpha' + \beta' \leq \alpha + \beta$.*

By taking into account Wald’s approximation, we can compute approximations of the expected number of queries:

$$E_{X_0} [N] \approx \frac{\alpha \log \left(\frac{1-\beta}{\alpha} \right) + (1-\alpha) \log \left(\frac{\beta}{1-\alpha} \right)}{E_{X_1} [\log(f_{X_0}(x)) - \log(f_{X_0}(x))]} \tag{20}$$

and

$$E_{X_1} [N] \approx \frac{(1-\beta) \log \left(\frac{1-\beta}{\alpha} \right) + \beta \log \left(\frac{\beta}{1-\alpha} \right)}{E_{X_1} [\log(f_{X_1}(x)) - \log(f_{X_0}(x))]} \tag{21}$$

5.3 A Toy-Example on DES

In order to illustrate advantages of sequential linear distinguishers, we have implemented a linear cryptanalysis of DES reduced to five rounds which uses a sequential distinguisher for deciding the parity of the linear approximation, *i.e.* the parity of the sum of involved key bits.

Using a static test, we needed 2800 known plaintext-ciphertext pairs in order to get a success probability of 97 %. Using a sequential strategy and for the same success probability, only 1218 samples were necessary on average. We give here both the static and the sequential decision rules.

Let S_n denote the number of times that Matsui’s best linear characteristic [17] on 5-rounds DES evaluates to 0, where n is the number of known plaintext-ciphertext pairs at disposal. This linear approximation holds with probability $\frac{1}{2} + 0.01907$. The static decision rule is given by

$$\begin{cases} \text{Output “key parity = 0” if } S_n \geq \frac{n}{2} \\ \text{Output “key parity = 1” if } S_n < \frac{n}{2} \end{cases} \tag{22}$$

With 2800 known pairs at disposal, the static rule is successful in 97% of the cases.

For $\alpha = \beta \triangleq 0.025$, Wald’s approximation gives $\tau_{\text{up}} = 48$ and $\tau_{\text{down}} = \frac{1}{48}$. The sequential rule is then defined by

$$\begin{cases} \text{Output “key parity = 1” if } S_n \leq \frac{n}{2} - \frac{\log \tau_{\text{up}}}{2 \log\left(\frac{1+2\epsilon}{1-2\epsilon}\right)} \\ \text{Output “key parity = 0” if } S_n \geq \frac{n}{2} + \frac{\log \tau_{\text{down}}}{2 \log\left(\frac{1+2\epsilon}{1-2\epsilon}\right)} \\ \text{Query another sample, otherwise.} \end{cases} \quad (23)$$

where $\epsilon = 0.01907$. We repeated this experiment 1’000’000 times for 5 different keys and got the following results:

	Exp. 1	Exp. 2	Exp. 3	Exp. 4	Exp. 5
Pr[static distinguisher successful]	0.9689	0.9687	0.9684	0.9686	0.9688
Pr[sequential distinguisher successful]	0.9686	0.9684	0.9683	0.9682	0.9684
Average number of queries	1218.7	1218.7	1218.3	1219.1	1218.8

6 Links to Other Statistical Attacks

Potential applications in cryptanalysis of sequential distinguishers are numerous. As soon as one is able to derive underlying probability distributions, it is possible to define likelihood-ratios, and thus to use a sequential distinguisher. However, deriving even approximations of probability distributions may not be a trivial task in certain cases.

Furthermore, even if one has the probability distributions in hand, one should not neglect the amount of computations necessary to get the information which will be fed into the likelihood-ratio.

Under the light of the hypothesis testing paradigm, several known statistical attacks can be summarized (for which concerns their *decisional part*), and thus potentially analyzed in a simple way. The χ^2 statistical test, proposed in [29] for the first time and then used in many cryptanalytic contributions (*e.g.* see [12, 15,10,21]), is closely related to *generalized likelihood-ratio tests*.

Indeed, as outlined in Section §2, likelihood ratio tests are optimal for testing a simple versus a simple hypothesis. It is possible to develop a generalization of this test for use in situations in which the hypotheses are not simple (*e.g.* one tests a probability distribution depending of a parameter $\theta \in \omega_0$ against $\theta \in \omega_1$ where ω_0 and ω_1 are disjoint *subsets* of possible parameters). Such tests are not generally optimal, but they are typically non-optimal in situations for which no optimal test exists, and they usually perform reasonably well.

It is well-known (see for instance [26]) that Pearson’s χ^2 statistic and a generalized likelihood-ratio test for a *multinomial distribution* are asymptotically equivalent. Thus, the underlying statistical decision processes in linear, differential, statistical, χ^2 - and stochastic cryptanalysis are all equivalent in a statistical point of view: they try to distinguish two different (families of) probability distributions with help of a generalized likelihood-ratio test.

Another interesting attack is Harpes and Massey’s *partitioning cryptanalysis* [13]. In such an attack, one defines the *imbalance* of a random variable as

being a non-uniformity measure, *i.e.* as measure of distance between a uniform distribution and the distribution obtained through the partitioning process. In [13], two different imbalance measures are considered, namely the *peak imbalance* and the *squared Euclidean imbalance*: one could consider a χ^2 -value or, equivalently, a generalized likelihood-ratio value as well (and maybe slightly improve its performances). Thus, the statistical problem behind this attack remains the same.

7 Conclusion

In this paper, we have used the power of some tools proposed by the theory of statistical tests for considering various situations in well-known cryptanalytic attacks, like linear and differential cryptanalysis; we improve known bounds on the asymptotical behavior of the best advantage of distinguishers implementing these attacks. Furthermore, we formalize the concept of “sequential distinguisher” and we illustrate its potential power in a toy-example. Finally, we discuss the application of the statistical tools in a couple of known attacks; this suggests that statistical hypothesis testing theory may be a mean to unify, to characterize and to analyze most of the known attacks against block ciphers.

Acknowledgments. The author would like to thank Serge Vaudenay for many interesting and enlightning discussions.

References

1. E. Biham and A. Shamir, *Differential cryptanalysis of DES-like cryptosystems (extended abstract)*, Advances in Cryptology – CRYPTO’90, LNCS, vol. 537, Springer-Verlag, 1990, pp. 2–21.
2. ———, *Differential cryptanalysis of DES-like cryptosystems*, Journal of Cryptology 4 (1991), no. 1, 3–72.
3. ———, *Differential cryptanalysis of the Data Encryption Standard*, Springer-Verlag, 1993.
4. C. Cachin, *An information-theoretic model for steganography*, Information Hiding, 2nd International Workshop, LNCS, vol. 1525, Springer-Verlag, 1998, pp. 306–318.
5. ———, *An information-theoretic model for steganography*, Available on <http://eprint.iacr.org/2000/028/>, 2000.
6. F. Chabaud and S. Vaudenay, *Links between differential and linear cryptanalysis*, Advances in Cryptology - EUROCRYPT’94, LNCS, vol. 950, Springer-Verlag, 1995, pp. 356–365.
7. D. Coppersmith, S. Halevi, and C. Jutla, *Cryptanalysis of stream ciphers with linear masking*, Advances in Cryptology - CRYPTO’02, LNCS, vol. 2442, Springer-Verlag, 2002, pp. 515–532.
8. T. M. Cover and J. A. Thomas, *Information theory*, Wiley Series in Telecommunications, Wiley, 1991.
9. D. Davies and S. Murphy, *Pairs and triples of DES S-boxes*, Journal of Cryptology 8 (1995), no. 1, 1–25.

10. H. Gilbert, H. Handschuh, A. Joux, and S. Vaudenay, *A statistical attack on RC6*, Fast Software Encryption FSE'00, LNCS, vol. 1978, Springer-Verlag, 2000, pp. 65–74.
11. G.R. Grimmett and D.R. Stirzaker, *Probability and random processes*, Oxford University Press, 2001, 3rd edition.
12. H. Handschuh and H. Gilbert, χ^2 *cryptanalysis of the SEAL encryption algorithm*, Fast Software Encryption FSE'97, LNCS, vol. 1267, Springer-Verlag, 1997, pp. 1–12.
13. C. Harpes and J. Massey, *Partitioning cryptanalysis*, Fast Software Encryption FSE'97, LNCS, vol. 1267, Springer-Verlag, 1997, pp. 13–27.
14. P. Junod, *On the optimality of linear, differential and sequential distinguishers (full version)*, Available on <http://eprint.iacr.org> and on <http://crypto.junod.info>, 2003.
15. L. Knudsen and W. Meier, *Correlations in RC6 with a reduced number of rounds*, Fast Software Encryption FSE'00, LNCS, vol. 1978, Springer-Verlag, 2000, pp. 94–108.
16. M. Luby and C. Rackoff, *How to construct pseudorandom permutations from pseudorandom functions*, SIAM Journal on Computing **17** (1988), no. 2, 373–386.
17. M. Matsui, *Linear cryptanalysis method for DES cipher*, Advances in Cryptology – EUROCRYPT'93, LNCS, vol. 765, Springer-Verlag, 1993, pp. 386–397.
18. ———, *The first experimental cryptanalysis of the Data Encryption Standard*, Advances in Cryptology – CRYPTO'94, LNCS, vol. 839, Springer-Verlag, 1994, pp. 1–11.
19. U. Maurer, *A unified and generalized treatment of authentication theory*, Proc. 13th Symp. on Theoretical Aspects of Computer Science (STACS'96), LNCS, vol. 1046, Springer-Verlag, 1996, pp. 387–398.
20. ———, *Authentication theory and hypothesis testing*, IEEE Transactions on Information Theory **46** (2000), no. 4, 1350–1356.
21. M. Minier and H. Gilbert, *Stochastic cryptanalysis of Crypton*, Fast Software Encryption FSE'00, LNCS, vol. 1978, Springer-Verlag, 2000, pp. 121–133.
22. S. Murphy, F. Piper, M. Walker, and P. Wild, *Likelihood estimation for block cipher keys*, Technical report, Information Security Group, University of London, England, 1995.
23. K. Nyberg, *Perfect nonlinear S-boxes*, Advances in Cryptology – EUROCRYPT'91, LNCS, vol. 547, Springer-Verlag, 1991, pp. 378–386.
24. National Bureau of Standards, *Data Encryption Standard*, U.S. Department of Commerce, 1977.
25. T. Pornin, *Optimal resistance against the Davies and Murphy attack*, Advances in Cryptology – ASIACRYPT'98, LNCS, vol. 1514, Springer-Verlag, 2000, pp. 148–159.
26. J. A. Rice, *Mathematical statistics and data analysis*, Duxbury Press, 1995.
27. D. Siegmund, *Sequential analysis – tests and confidence intervals*, Springer-Verlag, 1985.
28. S. Vaudenay, *Decorrelation: a theory for block cipher security*, to appear in the Journal of Cryptology, Available on <http://lasecwww.epfl.ch>.
29. ———, *An experiment on DES statistical cryptanalysis*, 3rd ACM Conference on Computer and Communications Security, ACM Press, 1996, pp. 139–147.
30. ———, *Resistance against general iterated attacks*, Advances in Cryptology – EUROCRYPT'99, LNCS, vol. 1592, Springer-Verlag, 1999, pp. 255–271.