Distinguishing Distributions Using Chernoff Information

Thomas Baignères[1], Pouyan Sepehrdad[2], and Serge Vaudenay[2]

[1] CryptoExperts, Paris, France
[2] EPFL, Switzerland
thomas.baigneres@cryptoexperts.com,
{pouyan.sepehrdad,serge.vaudenay}@epfl.ch

Abstract. In this paper, we study the soundness amplification by repetition of cryptographic protocols. As a tool, we use the Chernoff Information. We specify the number of attempts or samples required to distinguish two distributions efficiently in various protocols. This includes weakly verifiable puzzles such as CAPTCHA-like challenge-response protocols, interactive arguments in sequential composition scenario and cryptanalysis of block ciphers. As our main contribution, we revisit computational soundness amplification by sequential repetition in the threshold case, i.e when completeness is not perfect. Moreover, we outline applications to the Leftover Hash Lemma and iterative attacks on block ciphers.

Keywords: distinguishing distributions, Chernoff Information, proof systems, block ciphers

1 Introduction

In many occasions in cryptography we encounter the challenge of distinguishing distributions such as pseudorandom number generators, symmetric key cryptanalysis or challenge-response puzzles. We consider protocols in which one distribution (null) is usually associated with the probability distribution of an adversary winning a game. Similarly, the other distribution (alternate) corresponds to the probability of success of a legitimate party. This concept can be modified depending on the application i.e, the distributions may correspond to a thoroughly uniform distribution and a biased distribution (such as in block ciphers cryptanalysis). For instance, challenge-response puzzles are often deployed to distinguish between a real and a fake solver. Differentiation is obtained by the probability of them solving a randomly chosen challenge. What we focus on in this paper is the application of such distinguishers in weakly verifiable puzzle protocols, sequential repetition of computationally sound protocols and the Leftover Hash lemma.

Initially, we concentrate on interactive protocols, where there always exist a number of false negative and false positive responses by the verifier. They correspond to the completeness and soundness probability of the protocol. One might think of a method to reduce the error associated with the relevant distinguisher. One straightforward strategy to decrease the probability of error in both cases is to provoke the protocol iteratively and output "accept" if all instances accept (non-threshold case). Assuming the passing probability of non-authentic (vs. authentic) parties is b (vs. a), one would like to obtain error probability of b^q after q iterations, but it makes the success probability of authentic parties go down to a^q, which is often not desirable in real life applications. To solve this bottleneck, what we investigate in this paper is the general scenario of threshold repetition, i.e we now accept if the number of accepting repetitions is larger than a given threshold m. We can find an optimal m in which it makes the error probabilities of the protocol arbitrary close to zero. This strategy can be deployed in other similar interactive protocols like weakly verifiable puzzle protocols in which a verifier sends a puzzle to the solver and depends on the solver's response, he outputs *accept* or *reject*. In one section, we principally study CAPTCHA-like protocols as an example of such puzzles. We offer q puzzles to the solver and accept if she replies correctly to at least a threshold m of instances.

The problem of soundness amplification and the previous results. In interactive systems, the soundness probability of the protocol corresponds to upper bounding the probability of success of a malicious party to win the game. We always assume that the verifier is computationally bounded, but depending on computational capability of the prover we can define *argument* or *proof* systems, where the former corresponds to polynomial time provers and the latter to computationally unbounded provers (see section 3.1). We refer to the soundness probability of proof systems as *statistical* soundness versus the *computational* soundness in argument systems. To decrease the soundness error of such protocols, making a problem harder by repetition can be performed using two distinct approaches, namely *sequential* and *parallel* repetition. By sequential repetition we mean repeating the protocol several times, beginning the next run after the previous one terminates. Conversely, in the parallel case, all the instances are yielded to the prover at the same time without waiting for any arbitrary instance to terminate.

Security amplification is a fundamental cryptographic problem and has been studied for a variety of important cryptographic primitives, such as one way functions [37], collision resistance hash functions [10], encryption schemes [18] and weakly verifiable puzzles [9, 22, 23], i.e given a construction C for some primitive P which is "weakly secure", we can construct a "strongly secure" construction C' from C as an example of *direct product theorems*, i.e solving many instances of a hard problem is harder than solving a single instance. It is well-known that sequential and parallel repetition of interactive proof systems reduce the error (statistical soundness) with an exponential rate (see [20]) in the non-threshold case (i.e, when there are no false rejections). In fact, [15] has given the proof that sequential repetition of computationally sound proof systems improves their security with an exponential rate in the "non-uniform model" under non-threshold approach, but it seems there is no explicit proof for the error reduction in the threshold case.

For a long period, it was assumed by the community that there is no distinction between error reduction of interactive arguments (computational soundness) when the protocol is iterated sequentially or in parallel. Finally, Bellare et al. [5] disproved this argument by providing a 4-round protocol in which q iterations does not reduce the computational soundness error probability of the protocol at all. In fact, they showed that there is no "black box" error reduction for such protocols when parallel repetition is concerned. On the other hand, they proved the surprising result that error reduction in parallel case depends fundamentally on the number rounds of the protocol. They proved that error decreases exponentially fast with the increase in the number of iterations if the number of rounds is less than 4. The computation complexity of each instance of their counter-example grows linearly with the number of repetitions and for such protocols the error does not even decrease for some types of interactive proofs. They constructed an artificial oracle to solve this problem. To discard the effect of this oracle, using universal arguments of Barak and Goldreich [4], Pietrzak et al. [33] provided an 8-round protocol in which the q-fold parallel repetition does not decrease the error probability below some constant for any polynomial q (where the communication complexity does not depend on q). This result was extended to Arthur-Merlin games by Pass et al. [32] showing that parallel repetition reduces the soundness-error at an optimal rate (up to a negligible factor) in constant-round public coin arguments and constant-round public-coin proofs of knowledge. As an extension, multi-prover systems were examined in multiple articles such as [19, 34].

In all these cases, we assume that the *secret* is not given to the verifier, otherwise there exist examples that even sequential repetition does not reduces the error at all (see [5, 33]). Recently, Dodis et al. [17] have generalized the previous results to the case of *interactive* cryptographic primitives in the sense that the adversary can query the oracle multiple times before solving the main challenge. They studied the security amplification of MACs, SIGs and PRFs showing how to convert a corresponding weak primitive into a strong primitive. In fact, they proved a direct product theorem and even a Chernoff type theorem for MACs / SIGs with *imperfect completeness* and a regular XOR lemma for PRFs by introducing *Dynamic Weakly Verifiable Puzzles* (DWVPs).

Weakly verifiable puzzles. As another application, we study weakly verifiable puzzles. These are interactive protocols in which the verifier sends a puzzle to the solver and outputs 0 or 1 depending the solver's response. They are *weakly verifiable* in the sense that only the puzzle generator can check the correctness of the responses, either because the challenge may have multiple correct responses and the verifier seeks a particular one of those or because the solver is computationally constrained, for instance in CAPTCHA puzzles [1]. CAPTCHA is a fuzzy challenge response protocol for distinguishing humans from programs (bots) mostly based on a distorted text with extraneous lines [1]. The current vision protocols are not able to pass CAPTCHA efficiently and the probability that a human can pass is much higher than the programs. This is thankful of the non-efficiency of the current image recognition systems not being able to identify distorted texts efficiently, but their passing success rate is still non-negligible. Moreover, many humans (including us) fail a non-negligible fraction of puzzles. This implies that it might not be desirable to consider the non-threshold scenario for such protocols. Previously, Canetti et al. [9] proved that the parallel repetition of weakly verifiable puzzle protocols decreases the error with an exponential rate. In fact, they found a tighter bound than [5]. Their proof is restricted to the non-threshold case which might not be appropriate for CAPTCHAs since their completeness are not perfect. This result can be extended to parallel repetition of interactive arguments. As the pioneers in threshold parallel repetition of such protocols, Impagliazzo et al. [22, 23] have introduced two distinct bounds on the maximum success probability of a malicious algorithm for the parallel repetition of such protocols in the threshold case. The authors observed that the authentic party is on average expect to solve $a.q$ puzzles and if a Chernoff like bound holds, then the probability of fake parties solving $a.q$ puzzles may drop exponentially and they gave an exponential bound. The complication in reducing a single puzzle instance to a direct product puzzle instance originates from the fact that the given single puzzle instance is required to be incorporated in all simulated direct product puzzle instances and thus they are not independent. However, the bound they obtained has a weak constant in the exponent and although their results apply to the parallel composition scenario, they provided values which are irrelevant in practice, CAPTCHA for instance (see section 3.2). This was noticed by the authors themselves motivating to find better bounds as an open problem. Jutla [26] deployed a uniformized parallel solver, who first permutes his given first q-puzzles randomly, solves them as before and permutes the results back. Deploying this strategy, he improved the aforementioned bound and then he plugged it into "trust reduction" strategy in [22] and considered a linearly weighted metric and derived a more optimal bound. In fact, we show by a concrete example that his bound is still not applicable in practice since it asks for solving a huge number of CAPTCHAs in parallel.

Our contribution. The fundamental issue in this area is an approximation on the number of iterations required to effectively tune the probability of false acceptance or false rejection optimally. In fact, we find the optimal threshold m for the best distinguisher in section 3. We show that soundness amplification in the threshold case

- works as expected for statistical soundness.
- works with a small gap for computational soundness when the number of repetitions is logarithmic.

We find a practical bound restricted to sequential repetition of such protocols. Notice that our bounds might not work in the parallel composition scenario but it provides figures which can be deployed in the practice of sequential repetition. It seems more logical for practical applications like CAPTCHAs (see section 3.2).

We also consider the Leftover-Hash lemma. Let assume we have a secret key $\mathcal{K}$ that has t uniform random bits. If ℓ bits of the key are leaked, but it is not clear which one, the Leftover-Hash Lemma [24] tells us that we can produce a key of almost $m = t - \ell$ bits that is ϵ-indistinguishable from uniform distribution over the key space. We define a distinguisher given n samples in Luby-Rackoff model which distinguishes between a universal hash function and a uniform distribution. We derive the same bound as in [13] by deploying

Chernoff Information which turns out to be optimal by introducing the Multi-Session Leftover-Hash Lemma when more than one such key generations are of interest.

In Appendix, we present iterative attacks on block ciphers with applications in linear and differential cryptanalysis and show that we can recover the number of plaintext/ciphertext pairs required to obtain a significant advantage.

Structure of this paper. First, we mention some preliminaries regarding the facts and previous results on hypothesis testing problem and statistical distinguishers. Then, we model our distinguishing games as a challenge of distinguishing two random Boolean sources in section 3 when multiple samples are given to the distinguisher or multiple iterations are of concerned. In section 3.1, we focus on sequential repetition of interactive arguments in the threshold case and derive better bounds to strengthen them. In section 3.2, we investigate the sequential repetition of weakly verifiable puzzles. We compare Impagliazzo et al. [22, 23] bounds and Jutla bound [26] together with the Chernoff-Hoeffding bound of [21] and the asymptotic bound in [3] in a q-sequential-iteration CAPTCHA-like protocol and conclude that the asymptotic estimation is the closest one to the concrete value and appears to be more useful in such specific interactive argument systems. Furthermore, in section 5 we derive a useful bound which we use to investigate the Leftover Hash Lemma when multi sessions of a universal hash function are of concerned. In Appendix, we revisit iterative attack on block ciphers and derive the number of samples required to achieve a significant advantage.

2 Preliminaries

Notations. In this paper, we let $\mathcal{Z}$ denote a finite set and $\mathsf{P}_0, \mathsf{P}_1, \ldots, \mathsf{P}_k$ be $k+1$ probability distributions over $\mathcal{Z}$. The support of a distribution P over $\mathcal{Z}$ is the set $\mathrm{supp}(\mathsf{P}) = \{z \in \mathcal{Z} \ : \ \mathsf{P}[z] > 0\}$. The distribution P is of full-support when $\mathrm{supp}(\mathsf{P}) = \mathcal{Z}$. When considering the two distributions P_0, P_1 we will usually denote $\mathcal{Z}' = \mathrm{supp}(\mathsf{P}_0) \cap \mathrm{supp}(\mathsf{P}_1)$ and have $\mathcal{Z} = \mathrm{supp}(\mathsf{P}_0) \cup \mathrm{supp}(\mathsf{P}_1)$. The natural and base 2 logarithms will respectively be denoted by ln and log. The *Kullback-Leibler divergence* [27] and the *Chernoff Information* [11] between P_0 and P_1 are respectively defined by

$$\mathrm{D}(\mathsf{P}_0 \| \mathsf{P}_1) = \sum_{z \in \mathrm{supp}(\mathsf{P}_0)} \mathsf{P}_0[z] \log \frac{\mathsf{P}_0[z]}{\mathsf{P}_1[z]} \qquad \text{and} \qquad \mathrm{C}(\mathsf{P}_0, \mathsf{P}_1) = - \inf_{0<\lambda<1} \log \sum_{z \in \mathcal{Z}'} \mathsf{P}_0[z]^{1-\lambda} \mathsf{P}_1[z]^{\lambda}$$

When $\mathrm{supp}(\mathsf{P}_0) \nsubseteq \mathrm{supp}(\mathsf{P}_1)$ then $\mathrm{D}(\mathsf{P}_0 \| \mathsf{P}_1) = +\infty$. A sequence of q elements $z_1, \ldots, z_q \in \mathcal{Z}$ and a sequence of random variables $Z_1, \ldots, Z_q \in \mathcal{Z}$ are respectively denoted by $\mathbf{z}^q$ and $\mathbf{Z}^q$. Finally, we say that two functions f and g are asymptotically equivalent when $\lim_{q \to \infty} \frac{1}{q} \ln \frac{f(q)}{g(q)} = 0$ or equivalently when $f(q) = g(q) e^{o(q)}$. This is denoted by $f(q) \doteq g(q)$.

Essential Definitions on Hypothesis Testing. The cryptographic problems we will consider in the following sections can all be formalized as a hypothesis testing problem in which a distinguisher A tries to distinguish between the hypotheses

$$\mathsf{H}_0 : \mathsf{P} = \mathsf{P}_0 \quad \text{and} \quad \mathsf{H}_1 : \mathsf{P} \in \mathcal{D} = \{\mathsf{P}_1, \ldots, \mathsf{P}_k\}$$

on the basis of the knowledge of the P_i's and of $q > 0$ elements $Z_1, \ldots, Z_q \in \mathcal{Z}$ sampled according to the distribution P. It is assumed that one of the hypotheses is true, that the q samples are independent and identically distributed (iid), that the distinguisher A eventually outputs 0 or 1 to indicate its guess and that this distinguisher is computationally unbounded (so that we can assume it is deterministic); for this last reason, A is referred to as a *q-limited distinguisher*. In fact, we are following Luby-Rackoff model of indistinguishability [29] where the only adversarial limitation is the number of queries. In the particular case where $k = 1$, we will refer to the previous problem as a *simple* hypothesis test, whereas when $k > 1$ we call it a *composite* hypothesis test. A q-limited distinguisher A which is given q samples $\mathbf{Z}^q = Z_1, \ldots, Z_q$ is denoted as $\mathsf{A}_q(\mathbf{Z}^q)$. The effectiveness of A is mathematically formulated by its *advantage*.

Definition 1. *The* advantage *of a q-limited distinguisher A_q between the hypotheses H_0 and H_1, based on the q samples $\mathbf{Z}^q = Z_1, \ldots, Z_q$, is defined by*

$$\mathrm{Adv}_{\mathsf{A}_q}(\mathsf{H}_0, \mathsf{H}_1) = \Pr[\mathsf{A}_q(\mathbf{Z}^q) = 1 | \mathsf{H}_0] - \Pr[\mathsf{A}_q(\mathbf{Z}^q) = 1 | \mathsf{H}_1]$$

The hypotheses H_0 and H_1 are (q, ϵ)-indistinguishable if for any q-limited distinguisher A_q we have

$$|\mathrm{Adv}_{\mathsf{A}_q}(\mathsf{H}_0, \mathsf{H}_1)| \leq \epsilon$$

Existence of an Optimal Distinguisher. Since the samples are assumed to be iid, their particular *order* must be irrelevant. What really matters is the number of occurrences of each symbol of $\mathcal{Z}$ in the string $\boldsymbol{Z}^q = Z_1, \ldots, Z_q$ or equivalently the *type* (or *empirical probability distribution*) of this sequence, defined by

$$\mathsf{P}_{\mathbf{Z}^q}[z] = \frac{\#\{i \ : \ Z_i = z\}}{q}$$

Consequently, a distinguisher can be thoroughly specified by the set Π of all types for which it will output 1, i.e.,

$$\mathsf{A}_q(\mathbf{Z}^q) = 1 \ \Leftrightarrow \ \mathsf{P}_{\mathbf{Z}^q} \in \Pi$$

The set Π is called the *acceptance region* of A. Since q is fixed, the number of possible types is finite and thus we can assume wlog that Π is finite. Consequently, there is also a finite number of potential adversaries so that there must be at least one which maximizes the advantage. We call them *best distinguishers* and denote by $\mathrm{BestAdv}_q(\mathsf{H}_0, \mathsf{H}_1)$ (or simply by $\mathrm{BestAdv}_q$) their advantage.

The Optimal Adversary in the Simple Hypothesis Testing Case. We consider the simple case where A must distinguish between

$$\mathsf{H}_0 : \mathsf{P} = \mathsf{P}_0 \quad \text{and} \quad \mathsf{H}_1 : \mathsf{P} = \mathsf{P}_1$$

In that case, we abusively call A a distinguisher between P_0 and P_1 and denote its advantage by $\mathrm{Adv}_{\mathsf{A}_q}(\mathsf{P}_0, \mathsf{P}_1)$. The best possible advantage is obtained by *likelihood ratio test*, where the acceptance region of the distinguisher is such that

$$\mathsf{A}_q(\mathbf{Z}^q) = 1 \ \Leftrightarrow \ \frac{\mathsf{P}_{\mathbf{z}^q|\mathsf{P}_0}}{\mathsf{P}_{\mathbf{z}^q|\mathsf{P}_1}} \leq 1 \tag{1}$$

where $\mathsf{P}_{\mathbf{z}^q|\mathsf{P}_i}$ is the type of the sequence given the distribution P_i has happened. It can be shown [2] that the distinguisher $\mathsf{A}^\star$ defined by the acceptance region

$$\Pi^\star = \{\mathsf{P}_{\mathbf{Z}^q} \ : \ \mathrm{D}(\mathsf{P}_{\mathbf{Z}^q} \| \mathsf{P}_1) \leq \mathrm{D}(\mathsf{P}_{\mathbf{Z}^q} \| \mathsf{P}_0)\} \tag{2}$$

is a best distinguisher.

The following essential theorem allows to relate the advantage of the best distinguisher between P_0 and P_1 to the Chernoff Information[3] [11].

[3] A proof of this result can be found in [14] asymptotically, where it is implicitly assumed that $\mathrm{supp}(\mathsf{P}_0) = \mathrm{supp}(\mathsf{P}_1)$. The general case is treated in [2].

Theorem 1. *Let $\mathsf{P}_0, \mathsf{P}_1$ be two probability distributions. We have*

$$1 - \mathrm{BestAdv}_q(\mathsf{P}_0, \mathsf{P}_1) \doteq 2^{-q\mathrm{C}(\mathsf{P}_0,\mathsf{P}_1)} \tag{3}$$

This result verifies **asymptotically** that having access to $q \approx \frac{1}{\mathrm{C}(\mathsf{P}_0,\mathsf{P}_1)}$ samples we can distinguish P_0 from P_1 with a significant advantage.

3 Application to Boolean Cases

In this paper, we concentrate on applications of distinguishers in scenarios such as soundness amplification and weakly verifiable puzzles. In all these relevant applications, we are trying to differentiate between a legitimate and a malicious party. One strategy is to model this scenario as a distinguishing game between two Boolean random sources. We consider the problem of distinguishing two Boolean random sources with expected values a and b respectively. Suppose $\mathsf{P}_0, \mathsf{P}_1$ be two probability distributions over the set $\mathcal{Z} = \{0, 1\}$. Let

$$\mathsf{P}_0[X] = \begin{cases} a & X = 1 \\ 1 - a & X = 0 \end{cases} \qquad \text{and} \qquad \mathsf{P}_1[X] = \begin{cases} b & X = 1 \\ 1 - b & X = 0 \end{cases}$$

We define a distinguisher which outputs 1 *iff* $n_1 \leq m$, where $bq < m < aq$ and n_1 is the number of occurrences of 1 in the sample set. Intuitively, a refers to the probability that a legitimate user or a program can pass a single challenge successfully and b refers to which of a malicious user or program. As a matter of fact, we mostly investigate the protocols which are distinguishing a legitimate and a malicious user or program offering them q times to try and then if they can pass with a particular minimum threshold, algorithm outputs *accept* otherwise it *rejects*.
It can be shown using (2) that

$$m = \frac{q}{1 - \frac{\ln \frac{b}{a}}{\ln \frac{1-b}{1-a}}} \tag{4}$$

defines the best distinguisher using q samples to distinguish P_0 from P_1 (Note that if $a \approx b$, we have $m \approx q\frac{a+b}{2}$ which is a pretty intuitive threshold).
Then, employing the Chernoff Information, Theorem 1 gives

$$1 - \mathrm{Adv}_q \doteq 2^{-q\mathrm{C}(\mathsf{P}_0,\mathsf{P}_1)}$$

More precisely, having access to q samples and using the binomial distribution

$$\begin{aligned} 1 - \mathrm{Adv}_q &= \sum_{i \leq m} \binom{q}{i} a^i (1-a)^{q-i} + \sum_{i > m} \binom{q}{i} b^i (1-b)^{q-i} \\ &= 1 - \sum_{i \leq m} \binom{q}{i} \left(b^i (1-b)^{q-i} - a^i (1-a)^{q-i} \right) \end{aligned} \tag{5}$$

which is expressed as the *concrete* expression for computing the advantage of the best distinguisher. It might be assumed that this bound only works when the adversary's responses are independent, but we will show in Theorem 3 that it is true even if the adversary's responses are not independent, the only difference is an additive factor of $2^q\epsilon$. In fact, the adversary may decide to answer identically to all challenges or decide to respond to the following challenge as a function of the previous response. The fundamental question is that whether she gains anything by following this approach. What we prove is that she gains an additive factor of $2^q\epsilon$ which can be made arbitrary small for constant values of m and q (see 3.1). The effect of ϵ would be canceled out in the case of statistical soundness when proof systems are of interest since the prover is supposed to be computationally unbounded.

A theorem by *Hoeffding* [21] called *Chernoff-Hoeffding* theorem gives an upper bound on the probability of the addition of q identically independent Boolean random variables referred to as *Chernoff-Hoeffding bound* which can be used as a bound in our distinguishing game.

Theorem 2. ***(Chernoff-Hoeffding Theorem)*** *Let $\{X_1, \ldots, X_q\} \in \{0,1\}^q$ be q identically independent random variables with $\mathrm{E}[X_i] = a$, for $(1 \leq i \leq q)$. Then, for $\forall b > a$, we have*

$$\Pr\left[\frac{1}{q}\sum_{i=1}^{q} X_i \geq b\right] \leq \left(\left(\frac{a}{b}\right)^{b}\left(\frac{1-a}{1-b}\right)^{1-b}\right)^{q} = 2^{-q\mathrm{D}(b\|a)}$$

where $\mathrm{D}(b\|a)$ is the Kullback-Leibler divergence of Boolean random variables of expected values b and a.

As another representation, we can rewrite the Chernoff-Hoeffding bound as

$$\sum_{i=\lceil bq\rceil}^{q} \binom{q}{i} a^i (1-a)^{q-i} \leq 2^{-q\mathrm{D}(b\|a)}$$

Using the above representation of Chernoff-Hoeffding bound, we obtain

$$1 - \mathrm{Adv}_q \leq 2^{-q\mathrm{D}\left(\frac{m}{q}\|a\right)} + 2^{-q\mathrm{D}\left(\frac{m}{q}\|b\right)} \tag{6}$$

We will compare these bounds in section 3.2.

3.1 Soundness Amplification

As an application to the distinguisher in section 3, we consider interactive argument protocols which are methods for one party to prove to another that a statement is true or a string belongs to a language. In fact, we analyze the sequential composition of interactive arguments where the algorithm is expected to repeat q times sequentially and if the number of successes is more than a specific threshold, the protocol outputs *accept* otherwise *reject*. First, we define the notion of proof and argument systems.

Definition 2. *Given a language L over an alphabet Z, an interactive proof system (resp. a computationally proof system or an argument) for L is a pair $(\mathcal{P}, \mathcal{V})$ of interactive machines, where $\mathcal{P}$ is computationally unbounded (resp. $\mathcal{P}$ is computationally bounded) and $\mathcal{V}$ is polynomial-time such that there exist a polynomial P and a, b, where $0 \leq b < a \leq 1$ and*

- ***Termination:*** *for any $x, \omega, r_{\mathcal{P}}, r_{\mathcal{V}}$, the total complexity of $\mathcal{V}$ (until termination) in $\mathcal{P}(\omega; r_{\mathcal{P}}) \overset{x}{\leftrightarrow} \mathcal{V}(r_{\mathcal{V}})$ is bounded by $P(|x|)$, where x is the security parameter.*
- *a-**completeness:** for any $x \in L$, there exists a string ω, such that*

$$\Pr_{r_{\mathcal{P}}, r_{\mathcal{V}}}\left(Out_{\mathcal{V}}(\mathcal{P}(\omega; r_{\mathcal{P}}) \overset{x}{\leftrightarrow} \mathcal{V}(r_{\mathcal{V}})) = accept\right) \geq a(|x|)$$

- *b-**statistical soundness (resp. b-computational soundness)**: for any $x \notin L$ and any computationally unbounded (resp. polynomial-time) interactive machine $\mathcal{P}^\star$*

$$\Pr_{r_{\mathcal{P}}, r_{\mathcal{V}}}\left(Out_{\mathcal{V}}(\mathcal{P}^\star(r_{\mathcal{P}}) \overset{x}{\leftrightarrow} \mathcal{V}(r_{\mathcal{V}})) = accept\right) \leq b(|x|)$$

Given an interactive proof system $(\mathcal{P},\mathcal{V})$ for L which is a-complete and b-sound, we define a new proof system $(\mathcal{P}^q,\mathcal{V}^q_m)$ with threshold m as follows

- *$\mathcal{P}^q$ (resp. $\mathcal{V}^q_m$) simulates $\mathcal{P}$ (resp. $\mathcal{V}$), but have no terminal message until $q(|x|)$ sequential iterations with the same input x are made.*
- *after an iteration completes, they restart the entire protocol with fresh random coins.*
- *$\mathcal{V}^q_m$ accepts if at least $m(|x|)$ iterations of $\mathcal{V}$ are accepted out of $q(|x|)$.*

We use the following Lemma to prove our main theorem.

Lemma 1. *Assume that $(\mathcal{P},\mathcal{V})$ is a b-sound argument for L. Given q and ϵ such that $q\epsilon^{-1}$ is polynomially bounded in terms of $|x|$, we consider $(\mathcal{P}^q,\mathcal{V}^q_m)$ and a polynomially bounded malicious $\mathcal{P}^\star$. For $I \subseteq \{1,\ldots,q\}$ we let p_I be the probability that $\mathcal{P}^\star$ succeeds in every iteration i for $i \in I$. Given $J \subseteq \{1,\ldots,i-1\}$ and $I = J \cup \{i\}$, we have*

$$p_I \leq \max(bp_J, \epsilon)$$

More precisely, if for some I this inequality is not satisfied, then there is a malicious prover for $(\mathcal{P},\mathcal{V})$ with complexity $q\epsilon^{-1}$ times the one by $\mathcal{P}^\star$ to break b-soundness.

Proof. If $p_J \leq \epsilon$, the result is clear since $p_I \leq p_J$. Otherwise, we have $p_J > \epsilon$. We construct a malicious prover for $(\mathcal{P},\mathcal{V})$ who simply simulates $i-1$ iterations for the verifier to $\mathcal{P}^\star$. It repeats the simulation until every iteration j for $j \in J$ succeeds. The number of iterations is expected to be p_J^{-1} which is dominated by ϵ^{-1}. Then it runs an extra simulation with the real verifier in the $(\mathcal{P},\mathcal{V})$ protocol. The complexity of this malicious prover is bounded by $q\epsilon^{-1}$ which is a polynomial. So, it is polynomially bounded and the probability that the last iteration succeeds is bounded by b. Clearly, this is the conditional probability of success given that every iteration j for $j \in J$ succeeds. Hence, $p_I \leq bp_J$.

□

Using the above lemma, we prove that soundness amplification in the threshold case behaves as expected for statistical soundness in proof systems. Furthermore, there is only a small gap between the expected value in statistical soundness and computational soundness when the number of repetitions is logarithmic.

Theorem 3. *For any computationally sound proof system $(\mathcal{P},\mathcal{V})$ and for a language L and any q, m and ϵ such that $q\epsilon^{-1}$ is polynomially bounded in terms of $|x|$, we consider $(\mathcal{P}^q,\mathcal{V}^q_m)$ with threshold m. If $(\mathcal{P},\mathcal{V})$ is a-complete and b-sound, then $(\mathcal{P}^q,\mathcal{V}^q_m)$ is a'-complete and b'-sound where*

$$a' = \sum_{i=m}^{q} \binom{q}{i} a^i (1-a)^{q-i}$$
$$b' = \sum_{i=m}^{q} \binom{q}{i} b^i (1-b)^{q-i} + 2^q\epsilon$$

and the time reduction factor is of $q\epsilon^{-1}$.

Note that if we know $bq < m < aq$ and if we consider the optimal m by equation (4), the above theorem shows that the completeness of the protocol increases and the soundness probability of the protocol declines by q iterations. *Since the reduction factor is $q\epsilon^{-1}$, <u>for constants m and q</u>, the value ϵ can be fixed to an arbitrary low constant*, so we achieve

$$b' = \sum_{i=m}^{q} \binom{q}{i} b^i (1-b)^{q-i}$$

More generally, let $\epsilon = |x|^{-c}$, where c is a constant and set m to equation (4) and q be logarithmic in terms of $|x|$, hence we obtain

$$a' = 1 - O\left(|x|^{-\alpha}\right) \qquad \text{and} \qquad b' = O\left(|x|^{-\beta}\right)$$

with polynomial reduction factor. So, with a logarithmic number of repetitions we can make a', b' tend toward 1 and 0 respectively at a polynomial speed.

Proof. The proof for the a'-completeness is trivial using binomial distribution and considering that repetitions are independent. For b'-soundness the prover may decide to evaluate iterations dependently. In fact, we show that even if the prover does not consider each iteration independently, he may not achieve anything better than responding to each iteration independently except with a gap of $2^q\epsilon$. We define p_I as in the Lemma 1. Let X_j be a 0 or 1 random variable associated with the success of a malicious protocol $\mathcal{P}^\star$ in the j_{th} iteration. We define $p_{x_1 \ldots x_i}$ to be a pattern probability in i iterations as

$$p_{x_1 \ldots x_i} = \Pr\left[\bigwedge_{j=1}^{i} X_j = x_j\right]$$

and T as a random variable enumerating the number of times $\mathcal{P}^\star$ passes the protocol and $P = \Pr(T \geq m)$. Note that p_x can be recursively defined from the set of p_I's, then P can be computed. Due to Lemma 1, p_I's are subject to inequalities. We define an arbitrary $\epsilon > 0$ and we first show that P is lower than a new P called P' defined by a set of p'_I's, where the inequalities in the Lemma 1 are replaced by equalities. Next, we show that for this new set of p_I's we have

$$P \leq \sum_{i \geq m} \binom{q}{i} b^i (1-b)^{q-i} + 2^q \epsilon$$

to obtain b'-soundness.

For the first step, we use a rewriting procedure on the set of p_I's. In the same time we verify that the new set is still consistent with the law of probabilities, with the inequalities from the Lemma 1, and that P only increases. By iterating the rewriting procedure we eventually obtain a new set of p_I's satisfying $p_I = \max(bp_J, \epsilon)$ for all $I = J \cup \{i\}$ with $i > \max J$. The rewriting procedure works as follows.

Initially, we identify $I = J \cup \{i\}$ with $i > \max J$, such that $p_I < \max(bp_J, \epsilon)$, then for any $K \subseteq \{i+1, \ldots, q\}$, we have $p'_{I \cup K} = (1-\lambda)p_{I \cup K} + \lambda p_{J \cup K}$ with λ such that $p'_I = \max(bp'_J, \epsilon)$. Subsequently, we get $\lambda = \frac{\max(bp_J, \epsilon) - p_I}{p_J - p_I}$. All other p_J's are left unchanged. This is equivalent to rewriting $p'_{x0y} = (1-\lambda)p_{x0y}$ and $p'_{x1y} = p_{x1y} + \lambda p_{x0y}$ for $x \in \{0,1\}^{i-1}$ such that $x_j = 1$ for all $j \in I$. It can be shown that p' only updates a subtree starting at position I such that $p'_I = \max(bp'_J, \epsilon)$. Ultimately, all the equalities are reached. To check

$$\sum_{x : x_1 + \cdots + x_q \geq m} p_x \;\leq\; \sum_{x : x_1 + \cdots + x_q \geq m} p'_x$$

we split the sum depending on x:

- for the set of y in which $y_j = 0$ for some $j \in J$, we observe $p'_y = p_y$.
- for the set of y of the form $y = x\beta z$ with the cumulated weight of x and z be at least m and $x_j = 1$ for all $j \in J$, we group by the same x and z, since $p'_{x0z} + p'_{x1z} = p_{x0z} + p_{x1z}$.
- for the set of y of the form $y = x1z$ with the weight m and $x_j = 1$ for all $j \in J$, we observe that $p'_{x1z} \geq p_{x1z}$.

We now assume that the p_I's satisfy $p_I = \max(bp_J, \epsilon)$ for all $I = J \cup \{i\}$ with $i > \max J$ and we want to upper bound P'. Clearly, we have

$$p'_I = \max(b^{\#I}, \epsilon)$$

When turned into p'_x's we have

$$p'_x = \begin{cases} b^{w(x)}(1-b)^{q-w(x)} & \text{if } w(x) \leq \tau \\ \epsilon(1-b)^{q-w(x)} & \text{if } w(x) > \tau \text{ and } x_{q-w(x)+\tau+1} = \cdots = x_q = 1 \\ 0 & \text{otherwise} \end{cases}$$

for all I, where $w(x) = x_1 + \cdots + x_q$ and $\tau = \left\lfloor \frac{\ln \epsilon}{\ln b} \right\rfloor$. We have

$$\begin{aligned} P' &= \sum_{x:w(x)\geq m} p'_x \\ &= \sum_{x:m\leq w(x)\leq\tau} p'_x + \sum_{x:w(x)>\tau} p'_x \\ &\leq \sum_{i\geq m} \binom{q}{i} b^i(1-b)^{q-i} + \epsilon \sum_{x:w(x)>\tau} 1_{x_{q-w(x)+\tau+1}=\cdots=x_q=1}(1-b)^{q-w(x)} \\ &= \sum_{i\geq m} \binom{q}{i} b^i(1-b)^{q-i} + \epsilon \sum_{x:w(x)>\tau} \binom{q-w(x)+\tau}{\tau}(1-b)^{q-w(x)} \\ &\leq \sum_{i\geq m} \binom{q}{i} b^i(1-b)^{q-i} + 2^q\epsilon \end{aligned}$$

□

3.2 Application to Weakly Verifiable Puzzles

A weakly verifiable puzzle protocol is a game $P = (\mathcal{D}, R)$ between a solver and a verifier consisting of a set of distributions $\mathcal{D} = \{\mathcal{D}_1, ..., \mathcal{D}_k\}$ of cardinality k (the security parameter) which are defined on pairs (p_i, c_i) [9]. In fact, p_i is called a *puzzle* which is associated with a challenge from the verifier being sent to the solver and we refer to c_i as the *check string*. The second component is a *relation* $R[(p,c), r]$ where r is a string of a fixed length, which can be assumed as the solver's response. The verifier is aware of p_i and c_i and so he can inspect the response r of the solver. If $R[(p,c), r]$ holds, we say that the solver *passes*, otherwise we say that he *fails*. We define a direct product for P. That is, since q and $m \in [0, q]$, we define $P^q_m = (\mathcal{D}^{\otimes q}, R^q_m)$, where

$$R^q_m[((p_1, \ldots, p_q), (c_1, \ldots, c_q)), r_1, \ldots, r_q] \Leftrightarrow \#\{i \in [0, q]; R[(p_i, c_i), r_i]\} \geq m$$

CAPTCHA is an example of such protocols. Another example is the *Déjà vu* protocol which is used as an authentication method [16]. In fact, weakly verifiable puzzles are essentially 2-round interactive protocols aimed at satisfying b-soundness for a category of malicious provers and a-completeness for a category of honest provers. Clearly, we can apply our previous treatment on sequential iteration to the sequential composition of weakly verifiable puzzles.

Let suppose that honest people pass with probability a and malicious programs pass with probability b. The prominent issue is to find the best method to distinguish a human from a program using q attempts [4]. This can be translated to a hypothesis testing problem, involved is a random variable *accept* with an

[4] Intuitive solution is to ask for many independent challenges.

expected value a (resp. b) associated with hypothesis H_0 (resp. H_1). We can use the results on the previous distinguisher with an application to such puzzles.We use the theorem by *Impagliazzo et al.* [22] to estimate the total probability of error the threshold-based distinguisher attains which can be used for the parallel repetition of such protocols. This was the first bound found on upper bounding the success probability of an adversary in the parallel composition of weakly verifiable puzzles in the threshold case. We consider a pretty good CAPTCHA for which humans pass with probability $a = 90\%$ and such that there exist attacks solving them with probability $b = 33\%$. For instance, we can consider Gimpy. (see [31, 36]).

Theorem 4. ***(Impagliazzo-Jaiswal-Kabanets* 2007**) *If all malicious algorithms can pass a challenge with probability at most b, then the probability that a malicious algorithm passes the challenge at least m times out of q parallel instances is lower than $\beta = 2e^{-\frac{(m-bq)^2}{64q}}$.*

Equivalently, if "pass", b and m are replaced by "fail", $1-a$ and $q-m$ respectively, it leads to the expression that legitimate people succeed less than m times out of q with probability lower than $\alpha = 2e^{-\frac{(m-aq)^2}{64q}}$. Hence, the advantage of a distinguisher which distinguishes the legitimate users from malicious programs using the threshold m can be computed as

$$1 - \mathrm{Adv}_q \leq \alpha + \beta = 2e^{-\frac{(m-bq)^2}{64q}} + 2e^{-\frac{(m-aq)^2}{64q}} \tag{7}$$

Impagliazzo et al. [23] introduced a new bound for the corresponding probability distribution in 2009.

Theorem 5. ***(Impagliazzo-Jaiswal-Kabanets* 2009**) *If all malicious algorithms can pass a challenge with probability at most b, then the probability that a malicious algorithm passes the challenge at least m times out of q parallel instances is lower than $\beta = \frac{100q}{m-bq}e^{-\frac{(m-bq)^2}{40q(1-b)}}$.*

Similarly, using the threshold m

$$1 - \mathrm{Adv}_q \leq \alpha + \beta = \frac{100q}{m-bq}e^{-\frac{(m-bq)^2}{40q(1-b)}} + \frac{100q}{aq-m}e^{-\frac{(m-aq)^2}{40qa}} \tag{8}$$

Recently, Jutla in TCC 2010 [26] and ECCC [25] improved the above bounds by using a uniformized parallel solver who permutes the given q puzzles randomly, solve them and permutes them back. He uses a linearly weighted metric to derive a tighter bound to the Chernoff bound, but as illustrated in the following table, the results are still non-relevant in practice. It is because all three bounds still ask for a huge number of CAPTCHAs which can not be used in real life.

Theorem 6. ***(Jutla* 2010**) *If all malicious algorithms can pass a challenge with probability at most b, then the probability that a malicious algorithm passes the challenge at least m times out of q parallel instances is lower than $\beta = \frac{2(q-bq)^3}{(q-m)^2(m-bq)} \cdot e^{-\left(\frac{q-m}{2}\right)\left(\frac{m-bq}{q-bq}\right)^2}$ if $bq < \min\{m, q-1\}$.*

Similarly, using the threshold m

$$1 - \mathrm{Adv}_q \leq \alpha + \beta = \frac{2(q-bq)^3}{(q-m)^2(m-bq)} \cdot e^{-\left(\frac{q-m}{2}\right)\left(\frac{m-bq}{q-bq}\right)^2} + \frac{2(aq)^3}{m^2(aq-m)} \cdot e^{-\left(\frac{m}{2}\right)\left(\frac{aq-m}{aq}\right)^2} \tag{9}$$

As an improvement, a new bound was derived by Jutla [25], which is still impractical.

$$1 - \mathrm{Adv}_q \leq \alpha + \beta = \frac{4q^2(1-b)^2}{(m-bq)(q-m)} \cdot e^{-\frac{(m-bq)^2}{2q(1-b)}} + \frac{4a^2q^2}{m((1-b)q-m)} \cdot e^{-\frac{((1-b)q-m)^2}{2aq}} \tag{10}$$

		Parallel Repetition				Sequential Repetition		
q	m	IJK07 (7)	IJK09 (8)	J10 (9)	J10$_2$ (10)	asymptotic (3)	concrete (5)	Chernoff (6)
1	0	> 1	> 1	N/A	N/A	0.803	0.430	1.606
3	1	> 1	> 1	> 1	> 1	0.517	0.283	1.035
4	2	> 1	> 1	> 1	> 1	0.415	0.160	0.831
5	3	> 1	> 1	> 1	> 1	0.333	0.125	0.667
7	4	> 1	> 1	> 1	> 1	0.215	0.069	0.430
100	65	> 1	> 1	> 1	> 1	$2^{-31.68}$	$2^{-34.95}$	$2^{-30.68}$
5000	3273	0.019	0.095	≈ 0	> 1	≈ 0	≈ 0	≈ 0

Table 1. $[1 - \mathrm{Adv}_q]$ (total error) comparison for 7 distinct bounds with respect to q for $a = 90\%$ and $b = 33\%$, the exact advantage is given by (5).

We compare the seven distinct bounds already discussed with the concrete value extracted in equation (5). As a summary, the table of advantage bounds we already computed together with the concrete value for the advantage of the distinguisher in section 3 is depicted in Table 1.

As the figures represent, for all the range of q the asymptotic value is the closest one to the concrete value which illustrates a dramatic improvement in the number of samples (less samples for a fixed advantage) or iterations required to run the mentioned protocols compare to the bounds of (7), (8), (6), (9) and (10). Clearly, solving 4 CAPTCHAs in at most 7 sequential attempts provides an error probability below 10% using parameters $a = 90\%$ and $b = 33\%$. "(7), (8) bounds are quite weak when applied to concrete problems such as actual CAPTCHA protocol with reasonable numbers of repetitions" [22, 23], which can be verified by the result in the table above. Although we are comparing sequential with parallel composition, it makes more sense to ask for 7 CAPTCHAs attempts sequentially than requiring to solve 5000 CAPTCHAs (as (7) bound recommends) at the same time. It still remains an open problem to find a better bound which works for the case of parallel repetition, one which provides values which can be implemented in practice. Moreover, as can be observed from the above table the value of the concrete error is always less than the asymptotic value which is the implication of Theorem 7.

4 Useful Bounds

In this section, we derive two bounds (see Appendix for the proof) which we use one in the ongoing section and one which argues that the total error probability in the general case is bounded by its asymptotic value and as was shown in the example in section 3.2, this provides a better bound than (6).

Theorem 7. *Let $\mathcal{Z}$ be a finite set and P_0 and P_1 be two distributions with support of union $\mathcal{Z}$ and intersection $\mathcal{Z}'$. Let $\mathrm{BestAdv}_q(\mathsf{P}_0, \mathsf{P}_1)$ be the best advantage for distinguishing P_0 from P_1 using q samples. We have*

$$1 - \mathrm{BestAdv}_q \leq 2^{-q\mathrm{C}(\mathsf{P}_0,\mathsf{P}_1)}$$

This result yields an upper bound on the probability of error of the best distinguisher. In fact, this result can be verified by the comparison between the concrete value of the error and asymptotic bound derived above.

Theorem 8. *Let P_0 and P_1 be distributions of support $\mathcal{Z}$, We have*

$$\frac{1}{8}\sum_{x\in\mathcal{Z}} \mathsf{P}_0[x]\left(\frac{\mathsf{P}_1[x]-\mathsf{P}_0[x]}{\max(\mathsf{P}_0[x],\mathsf{P}_1[x])}\right)^2 \leq 1 - 2^{-\mathrm{C}(\mathsf{P}_0,\mathsf{P}_1)} \leq \frac{1}{8}\sum_{x\in\mathcal{Z}} \mathsf{P}_0[x]\left(\frac{\mathsf{P}_1[x]-\mathsf{P}_0[x]}{\min(\mathsf{P}_0[x],\mathsf{P}_1[x])}\right)^2$$

As a result, for P_0 be the uniform distribution over a domain of size N, since $\mathsf{P}_0[x] - \|\mathsf{P}_1 - \mathsf{P}_0\|_2 \leq \mathsf{P}_1[x] \leq \mathsf{P}_0[x] + \|\mathsf{P}_1 - \mathsf{P}_0\|_2$, we can rewrite the bound as

$$\frac{1}{8}\frac{N\|\mathsf{P}_1 - \mathsf{P}_0\|_2^2}{(1 + N\|\mathsf{P}_1 - \mathsf{P}_0\|_2)^2} \leq 1 - 2^{-\mathrm{C}(\mathsf{P}_0,\mathsf{P}_1)} \leq \frac{1}{8}\frac{N\|\mathsf{P}_1 - \mathsf{P}_0\|_2^2}{(1 - N\|\mathsf{P}_1 - \mathsf{P}_0\|_2)^2}$$

where $\|\mathsf{P}_0 - \mathsf{P}_1\|_2$ states the Euclidean distance between distribution $\mathsf{P}_0, \mathsf{P}_1$.

5 Multi-Session Leftover-Hash Lemma

Let X be a random variable over a finite set $\mathcal{Z}$, the minimum entropy of X is defined as

$$H_\infty[X] = -\log(\max_z \Pr[X = z])$$

The *Rényi entropy* [35] of order α, where $\alpha \geq 0$, is defined as

$$H_\alpha[X] = \frac{1}{1-\alpha}\log(\sum_z \Pr[X = z]^\alpha)$$

Notice that $2^{-H_2[X]}$ is the collision probability and $2^{-H_2[X]} \leq 2^{-H_\infty[X]}$.

If X is a random variable over a set $\mathcal{Z}$ of order N, the square of Euclidean distance between the distribution of X called $\mathsf{P}_1[X]$ and the uniform distribution $\mathsf{P}_0[X]$ can be expressed as

$$\|\mathsf{P}_1[X] - \mathsf{P}_0[X]\|_2^2 = 2^{-H_2[X]} - \frac{1}{N}$$

Let $d(\mathsf{P}_1, \mathsf{P}_0)$ be the statistical distance between the distribution P_1 and the uniform distribution P_0, the expression

$$d(\mathsf{P}_1[X], \mathsf{P}_0[X]) \leq \sqrt{N}\|\mathsf{P}_1[X], \mathsf{P}_0[X]\|_2$$

shows the link between statistical and Euclidean distance of distributions.

Definition 3. *Let $H = \{H_N\} : D \to \{0,1\}^m$ be a family of functions, where $N \in \mathcal{N}$. H_N is a universal hash function if for any $x, y \in \{0,1\}^m$ such that $x \neq y$, we have*

$$\Pr(H_N[x] = H_N[y]) = 2^{-m}$$

where N is uniformly distributed.

Lemma 2. ***(Leftover Hash Lemma [24]: Impagliazzo-Levin-Luby 1989)*** *If h is a universal hash function with a range of size 2^m and X, N, U are independent random variables where N, U are uniformly distributed and $m \leq H_\infty[X] - 2\log\frac{1}{\epsilon}$, then the distributions of $(h_N[X], N)$ and (U, N) are ϵ-indistinguishable.*

We recall an application of the above Lemma in ElGamal encryption from Boneh [8]. Let $\langle g \rangle$ be a subgroup generated by some g of prime order q in $\mathbb{Z}_p^*$. Consider a scenario in which party A encrypts a message m using the party B's public key e. A picks a random value $r \in \mathbb{Z}_q^*$ and computes the pair $Enc[e, m; r] = (g^r, me^r) = (c_1, c_2)$ and sends it to B. At the other end based on the fact that $e^r = c_1^d$ where d is B's private key (secret key), B decrypts the message by computing $Dec[d, (c_1, c_2)] = m = c_2/(c_1)^d$.

Key recovery in ElGamal encryption is equivalent to the discrete logarithm problem, likewise, the decryption is equivalent to Diffie-Hellman problem [8]. On the other hand, ElGamal is **not** a semantically secure cryptosystem, because $q|\frac{(p-1)}{2}$ and so $g^{\frac{p-1}{2}} = 1$. Let $(\frac{a}{b})$ be the *Legendre symbol* for integers a and b, then

$(\frac{g}{p}) = 1$. We deduce that $(\frac{me^r}{p}) = (\frac{m}{p})$. As a result, if for $b = \{0,1\} : (\frac{m_b}{p}) = (-1)^b$, a distinguisher can distinguish $Enc[e, m_0; r]$ and $Enc[e, m_1; r]$ with advantage 1.

We define a new scheme based on ElGamal encryption which is argued to be $(\epsilon_{\text{DDH}} + \epsilon)$-IND-CPA secure. Let $\langle g \rangle$ be a group generated by some g of prime order q. Following a similar approach as ElGamal, we define the triple $Enc[e, m; N, r] = (g^r, m \oplus h_N[e^r], N) = (c'_1, c'_2, N)$ where $r \in \mathbb{Z}_q^*$ and N is uniformly distributed. Analogously, A sends this triple to B and B decrypts it using $Dec[d, (c'_1, c'_2, N)] = c'_2 \oplus h_N[c'^d_1]$.

Due to the Decisional Diffie-Hellman assumption [8], $(g, g^r, m \oplus h_N[e^r], N)$ is ϵ_{DDH}-indistinguishable from $(g, g^r, m \oplus h_N[g^{r'}], N)$. According to Lemma 2, $(g, g^r, m \oplus h_N[g^{r'}], N)$ is ϵ-indistinguishable from $(g, g^r, m \oplus U, N)$, where U is the uniform distribution. Furthermore, $(g, g^r, m \oplus U, N)$ is perfectly indistinguishable from (g, g^r, U, N). Consequently, $(g, g^r, m \oplus h_N[e^r], N)$ is $(\epsilon_{\text{DDH}} + \epsilon)$-indistinguishable from something independent from m which leads the scheme to be $(\epsilon_{\text{DDH}} + \epsilon)$-IND-CPA secure.

As another application to the Lemma 2, consider the Diffie-Hellman key exchange protocol. Let $\langle g \rangle$ be a group generated by some g of prime order q. In a key exchange between two parties A and B, the party A picks a random $x \in \mathbb{Z}_q^*$ and computes $X \leftarrow g^x$ and sends it to B. The party B aborts if $X \notin \langle g \rangle \backslash \{1\}$, otherwise he picks a random value $y \in \mathbb{Z}_q^*$ and computes $Y \leftarrow g^y$ and sends it to A. The party A aborts if $Y \notin \langle g \rangle \backslash \{1\}$, otherwise $K_{\text{ses}} = g^{xy}$ is computed and is shared between two parties as their session key. Since $\mathbb{Z}_q^*$ is cyclic, K_{ses} is a uniformly distributed non-neutral element of $\langle g \rangle$ (even locally under active attacks). Assume a non-ambiguous representation format for values which may be in $\langle g \rangle$ or not

$$\mathsf{Pr}(K_{\text{ses}} = x) = \begin{cases} \frac{1}{q-1} & x \in \langle g \rangle \backslash \{1\} \\ \\ 0 & otherwise \end{cases}$$

Thus,

$$H_\infty[K_{\text{ses}}] = \log{(q-1)}$$

Consider the protocol that exchanges a random number N and derives the key $K = h_N[K_{\text{ses}}]$. Let $\epsilon = \sqrt{2^m/(q-1)}$ by Leftover Hash Lemma, K is indistinguishable from a random key. Moreover, a protocol using n such key generations is $n\epsilon$-indistinguishable from the same protocol where K is truly random (thanks to the hybrid arguments) implying that it is safe to generate the key n times using the same protocol until n is of order $\sqrt{q.2^{-m}}$. This result is originating from the trivial bound, which can be improved employing a *Multi-Sample Leftover Hash Lemma*.

Lemma 3. ***(Multi-Sample Leftover Hash Lemma)*** *Assume h is a universal hash function with a range of size 2^m and key space $\mathcal{N}$. Let $N \in_U \mathcal{N}$ and $U \in_U \{0,1\}^m$ and X be independent random variables. If $\epsilon = \sqrt{(2^m-1)2^{-H_2[X]}}$ and $\epsilon' = \epsilon\sqrt{2^m \# \mathcal{N}}$, the best advantage for distinguishing $(h_N[X], N)$ from (U, N) using n samples is such that*

$$1 - \text{BestAdv}_n \overset{\bullet}{=} 2^{-n\text{C}}$$

where C *is bounded by*

$$-\log\left(1 - \frac{\epsilon^2}{8(1+\epsilon')^2}\right) \leq \text{C} \leq -\log\left(1 - \frac{\epsilon^2}{8(1-\epsilon')^2}\right)$$

Although this result is not so precise, it already suggests that we could find a better bound. In the above example, we have $H_2(X) = \log(q-1)$, so taking $\epsilon = \sqrt{(2^m-1)/(q-1)}$ and $\#\mathcal{N} \ll q.2^{-2m}$, we obtain that the minimal n for distinguishing is at least within the order of magnitude of ϵ^{-2} which is $q.2^{-m}$.

Proof. Let $\mathsf{P}_0, \mathsf{P}_1$ be two distributions, we proved in Lemma 2 that $\|\mathsf{P}_1 - \mathsf{P}_0\|_2^2 = 2^{-H_2[X]}(1-2^{-m})/\#\mathcal{N}$, where the domain size is $2^m \# \mathcal{N}$. Deploying Theorem 8, we get

$$\begin{aligned} 1 - 2^{-\text{C}(\mathsf{P}_0,\mathsf{P}_1)} &\leq \tfrac{(2^m-1)2^{-H_2[X]}}{8\left(1-\sqrt{(2^m-1)2^{m-H_2[X]}\#\mathcal{N}}\right)^2} \\ &= \tfrac{\epsilon^2}{8(1-\epsilon')^2} \end{aligned}$$

Similar procedure can be shown for the lower bound.

□

It has been shown that the min-entropy $H_\infty(X) = m + 2\log(\frac{1}{\epsilon}) + 2\log n$ suffices for the joint distribution to be ϵ-close to the uniform distribution (see [12, 24, 38]). Furthermore, recently Chung et al. [13] improved the previous bound by reducing $2\log n$ to $\log n$ and they proved that it is optimal for 2-universal hashing by using *Hellinger distance* to evaluate the error accumulation over each hashed instance. In fact they showed that

$$\epsilon = \sqrt{\frac{n}{q.2^{-m}}}$$

Therefore, the minimal n for distinguishing efficiently is of magnitude $q.2^{-m}$ which is the same bound we found by another approach, that is Chernoff Information and asymptotic q-limited distinguisher.

6 Conclusion

We mentioned various applications of distinguishers in cryptography. We evaluated their efficiency using the Chernoff Information. We revisited the interactive argument systems and relying on sequential repetition, we derived new bounds for the soundness property of such protocols (computational soundness) even in the case of dependent responses. Moreover, we compared seven distinct bounds for the error probability of the best distinguisher in weakly verifiable puzzle protocols when q samples are given and concluded that the asymptotic expression is the closest one to the concrete value compared to the bounds of equations (7),(8), (6) and (9). We introduced an application to the Leftover Hash Lemma and by introducing the Multi-Sample Leftover Hash Lemma we derived the same optimal bound as [13] with another approach (Chernoff Information) when the number of iterations is more than unity. We specified the number of samples to obtain a significant advantage in block ciphers cryptanalysis using Chernoff Information approach

References

1. L.V. Ahn, M. Blum, N.J. Hopper, and J. Langford. CAPTCHA: Using Hard AI Problems for Security. In *Advances in Cryptology EUROCRYPT 2003*, volume LNCS 2656, pages 294–311. Springer Berlin / Heidelberg, 2003.
2. T. Baignères. *Quantitative Security of Block Ciphers: Designs and Cryptanalysis Tools*. PhD thesis, EPFL, 2008.
3. T. Baignères and S. Vaudenay. The Complexity of Distinguishing Distributions. In *Information Theoretic Security*, volume LNCS 5155, pages 210–222. Springer Berlin / Heidelberg, 2008.
4. B. Barak and O. Goldreich. Universal arguments and their applications. In *Electronic Colloquium on Computational Complexity*, 2001.
5. M. Bellare, R. Impagliazzo, and M. Naor. Does Parallel Repetition Lower the Error in Computationally Sound Protocols? In *Proceedings of the Thirty-Eighth Annual IEEE Symposium on Foundations of Computer Science*, pages 374–383, 1997.
6. E. Biham and A. Shamir. Differential Cryptanalysis of DES-like Cryptosystems. *Journal of Cryptology*, 4(1):3–72, 1991.
7. C. Blondeau and B. Gérard. On the Data Complexity of Statistical Attacks Against Block Ciphers. In *Cryptology ePrint*, 2009.
8. D. Boneh. The Decision Diffie-Hellman Problem. In *Algorithmic Number Theory*, volume LNCS 1423, pages 48–63. Springer Berlin / Heidelberg, 1998.
9. R. Canetti, S. Halevi, and M. Steiner. Hardness Amplification of Weakly Verifiable Puzzles. In *Theory of Cryptography*, volume LNCS 3378, pages 17–33. Springer Berlin / Heidelberg, 2005.

10. R. Canetti, R. Rivest, M. Sudan, L. Trevisan, S. Vadhan, and H. Wee. Amplifying collision resistance: A complexity-theoretic treatment. In *Advances in Cryptology - CRYPTO 2007*, volume 4622, pages 264–283. Springer - Verlog, 2007.
11. H. Chernoff. *Sequential Analysis and Optimal Design*, volume 8 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, 1972.
12. B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.
13. K. Chung and S. Vadhan. Tight Bounds for Hashing Block Sources. In *Proceedings of the 11th international workshop, APPROX 2008, and 12th international workshop, RANDOM 2008 on Approximation, Randomization and Combinatorial Optimization: Algorithms and Techniques*, volume 5171, pages 357–370. Springer-Verlag, 2008.
14. T.M Cover and J.A. Thomas. *Elements of Information Theory.* Wiley Series in Telecommunications. John Wiley & Sons, 1991.
15. I. Damgård and B. Pfitzmann. Sequential Iteration of Interactive Arguments and an Efficient Zero-knowledge Argument for NP. Technical report, BRICS Report Series, Department of Computer Science, University of Aarhus, 1997.
16. R. Dhamija and A. Perrig. Déjà vu: User Study Using Images for Authentication. In *Proceedings of the 9th conference on USENIX Security Symposium*, volume 9, pages 4–4. USENIX Association, 2000.
17. Y. Dodis, R. Impagliazzo, R. Jaiswal, and V. Kabanets. Security Amplification for Interactive Cryptographic Primitives. In *TCC*, pages 128–145, 2009.
18. C. Dwork, M. Naor, and Reingold. O. Immunizing encryption schemes from decryption errors. In *Advances in Cryptology - EUROCRYPT 2004*, volume 3027, pages 342–360. Springer - Verlog, 2004.
19. U. Feige and O. Verbitsky. Error Reduction by Parallel Repetition - A Negative Result. *Combinatorica*, 22:461–478, 2001.
20. O. Goldreich. *Modern Cryptography, Probabilistic Proofs and Pseudo-randomness.* Algorithms and Combinatorics. Springer-Verlag, 1999.
21. W. Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
22. R. Impagliazzo, R. Jaiswal, and V. Kabanets. Chernoff-Type Direct Product Theorems. In *Advances in Cryptology - CRYPTO 2007*, volume LNCS 4622, pages 500–516. Springer Berlin / Heidelberg, 2007.
23. R. Impagliazzo, R. Jaiswal, and V. Kabanets. Chernoff-Type Direct Product Theorems. *Journal of Cryptology*, 22(1):75–92, 2009.
24. R. Impagliazzo, L.A. Levin, and M. Luby. Pseudo-random Generation from One-way Functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 12–24. ACM Press, 1989.
25. C. S. Juta. Almost Optimal Bounds for Direct Product Threshold Theorem. Technical report, ECCC, 2010.
26. C. S. Jutla. Almost Optimal Bounds for Direct Product Threshold Theorem. In *Theory of Cryptography Conferance*. Springer - Verlog, 2010.
27. S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.
28. X. Lai, J. L. Massey, and S. Murphy. Markov Ciphers and Differential Cryptanalysis. In D. W. Davies, editor, *Advances in Cryptology - EUROCRYPT '91*, volume 547 of *LNCS*, pages 17–38. Springer Berlin / Heidelberg, 1991.
29. M. Luby and C. Rackoff. How to Construct Pseudorandom Permutations from Pseudorandom Functions. *SIAM Journal of Computing*, 17:373–386, 1988.
30. M. Matsui. Linear Cryptanalysis Method for DES Cipher. In *EUROCRYPT 1993*, volume LNCS 765, pages 386–397. Springer Berlin / Heidelberg, 1993.
31. G. Mori and J. Malik. Recognising Objects in Adversarial Clutter: Breaking a Visual CAPTCHA. In *IEEE Conference Compurt Vision and Pattern Recognition*, pages 134–141. IEEE CS Press, 2003.
32. R. Pass and M. Venkitasubramaniam. An Efficient Parallel Repetition Theorem for Arthur-Merlin Games. In *Annual ACM Symposium on Theory of Computing*, pages 420–429, 2007.
33. K. Pietrzak and D. Wikström. Parallel Repetition of Computationally Sound Protocols Revisited. In *Theory of Cryptography*, volume 4392, pages 86–102. Springer Berlin / Heidelberg, 2007.
34. R. Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27:763–803, 1998.

35. A. Rényi. On Measures of Information and Entropy. In *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability*, pages 547–561, 1960.
36. J. Yan and A. Salah. CAPTCHA Security: A Case Study. *Journal of IEEE Security and Privacy*, 7:22–28, 2009.
37. A.C. Yao. Theory and applications of trapdoor functions. In *Proceedings of the Twenty-Third Annual IEEE Symposium on Foundations of Computer Science*, pages 80–91. IEEE Computer Society Press, Los Alamitos, 1982.
38. D. Zuckerman. Simulating BPP using a general weak random source. *Algorithmica*, 16(4-5):367–391, 1996.

A Iterative Attacks on Block Ciphers

We now apply the results regarding simple hypothesis testing to block cipher analysis. We consider a statistical distinguisher who has access to an oracle implementing either an instance c of a block cipher C or an instance c of $\mathsf{C}^\star$, a theoretical ideal scheme (sometimes called the *perfect cipher*) which corresponds to the set of *all* possible permutations over the same text space as C. Viewing both C and $\mathsf{C}^\star$ as *sets* of permutations, the objective of the distinguisher is to choose between the hypotheses[5]

$$\mathsf{H}_0 : \mathsf{c} \in \mathsf{C}^\star \quad \text{and} \quad \mathsf{H}_1 : \mathsf{c} \in \mathsf{C}$$

```
Oracle: a permutation $\mathsf{c}$
for $i$ from 1 to $q$ do
    pick $(X_1, \ldots, X_d)$ according to the distribution $\mathsf{D}$
    for all $1 \leq j \leq d$, query the oracle for $Y_j = \mathsf{c}(X_j)$
    set $Z_i = h(X_1, \ldots, X_d, Y_1, \ldots, Y_d)$
end for
return $\mathsf{A}^\star(Z_1, \ldots Z_q)$
```

Fig. 1. A q-limited iterative h-distinguisher of order d.

Most statistical distinguishers against block ciphers can be seen as q-limited iterative h-distinguishers of order d given some parameters q, h, d. These distinguishers are formalized in Figure 1. At each of the q iterations, the d-tuple $(X_1, \ldots, X_d)$ is chosen according to a certain distribution D. The function h returns at each iteration a value in a finite set $\mathcal{Z}$. Under a hypothesis similar to that of the hypothesis of stochastic equivalence [28], we can assume that the Z_i's follow a distribution P_0 under hypothesis H_0 (when c is an instance of the perfect cipher) or a distribution P_1 under hypothesis H_1 (when c is an instance of the block cipher considered). The two hypotheses can be reformulated as $\mathsf{H}_0 : \mathsf{P} = \mathsf{P}_0$ and $\mathsf{H}_1 : \mathsf{P} = \mathsf{P}_1$, where P is the distribution according to which the Z_i's are sampled. Letting $\mathsf{A}^\star$ be the best distinguisher between P_0 and P_1, the iterated distinguisher finally outputs $\mathsf{A}^\star(\mathbf{Z}^q)$. From Theorem 1 we know that its advantage Adv_q to distinguish H_0 from H_1 (i.e., the block cipher C from the perfect cipher $\mathsf{C}^\star$) verifies $1 - \mathrm{Adv}_q(\mathsf{H}_0, \mathsf{H}_1) = 1 - \mathrm{BestAdv}_q(\mathsf{P}_0, \mathsf{P}_1) \overset{\bullet}{=} 2^{-q\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1)}$. This result verifies asymptotically that having access to

$$q \approx \frac{1}{\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1)} \tag{11}$$

samples derived from the plaintext/ciphertext pairs allows to distinguish C from $\mathsf{C}^\star$ with a significant advantage. As an illustration, we propose to revisit various classical iterated distinguishers, compute their

[5] Note that the fact that the hypotheses are not disjoint is not a problem here, since all our previous results hold in that case too.

complexity based on (11) and derive their strategy from that of $\mathsf{A}^\star$. We focus on the case of differential distinguishers, impossible differentials and linear distinguishers. We attain estimate on q which are similar as in [7]. (see equations (15), (16) and (17)).

In the current application, the two distributions P_0 and P_1 are very close. In that case, it is possible to derive an approximation of the Chernoff Information that is easier to deal with. More formally, considering the case where both distributions are of full support and letting $\epsilon_z = (\mathsf{P}_1[z] - \mathsf{P}_0[z])/\mathsf{P}_0[z]$ be such that $\epsilon_z = o(1)$ for all $z \in \mathcal{Z}$, then it can be shown (see [2, p.50]) that $\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1) = \frac{1}{8 \ln 2} \sum_z \mathsf{P}_0[z]\epsilon_z^2 + o(\|\epsilon\|_2^2)$, where $\epsilon = (\epsilon_z)_{z\in\mathcal{Z}}$. Approximating the Chernoff Information by the right-hand side of the previous equation leads to

$$\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1) \approx \frac{1}{8 \ln 2} \sum_{z\in\mathcal{Z}} \frac{(\mathsf{P}_1[z] - \mathsf{P}_0[z])^2}{\mathsf{P}_0[z]}. \tag{12}$$

A.1 Differential Distinguishers

Differential distinguishers [6] are iterated h-distinguishers of order $d = 2$ where $h(x_1, x_2, y_1, y_2) = y_1 \oplus y_2$ and for which the distribution D is such that X_1 is chosen uniformly at random and $X_2 = X_1 \oplus a$ for some fixed a. Typically, we expect $h(X_1, X_2, Y_1, Y_2) = Y_1 \oplus Y_2$ to be biased under H_1 and uniformly distributed under H_0. Under H_1, we expect in practice for a well chosen b to have $Y_1 \oplus Y_2 = b$ with probability p and $Y_1 \oplus Y_2 = b' \neq b$ with probability $\frac{1-p}{n-1}$, where n is the cardinality of the text space, such that[6] $\frac{1}{n} = o(p)$ and $p = o(1)$. Accordingly, we have that P_0 is the uniform distribution and that

$$\mathsf{P}_1[z] = \begin{cases} p & \text{when } z = b, \\ \frac{1-p}{n-1} = \beta & \text{when } z \neq b \end{cases} \tag{13}$$

Under these notations, we now evaluate $\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1)$ to approximate the number of plaintext/ciphertext pairs required by a differential distinguisher to choose between C and $\mathsf{C}^\star$ with a significant advantage. Letting $\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1) = -\inf_{0<\lambda<1} \log F(\lambda)$ where $F(\lambda) = \sum_z \mathsf{P}_0[z]^{1-\lambda}\mathsf{P}_1[z]^\lambda$, we have

$$F(\lambda) = \frac{p}{(np)^\lambda} + \frac{1-p}{(n\beta)^\lambda}$$

We have $F(0) = F(1) = 1$ and $F'(0) \leq 0$, so that we know that F is minimum for a λ_0 such that $F'(\lambda_0) = 0$. We get

$$\lambda_0 = \frac{\ln\left(\frac{p \ln (np)}{(1-p) \ln \frac{1-1/n}{1-p}}\right)}{\ln\left(np\frac{1-1/n}{1-p}\right)} \sim \frac{\ln \ln (np)}{\ln(np)} \tag{14}$$

Consequently, $(np)^{\lambda_0} \sim \ln(np)$ and $(n\beta)^{\lambda_0} = 1+o(p)$ and thus $F(\lambda_0) = 1-p+o(p)$. The Chernoff Information verifies $\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1) = -\log F(\lambda_0) \sim \frac{p}{\ln 2}$. We conclude from (11) that a differential distinguisher approximately needs

$$q \approx \frac{\ln 2}{p} \tag{15}$$

samples to achieve a significant advantage.

[6] These assumptions simply express the fact that we expect p to be small (otherwise the cipher would be trivial to break), but much larger than $\frac{1}{n}$ (otherwise the cipher would be impossible to break for the chosen a and b).

It is also possible to find the practical (and optimal) strategy of a differential distinguisher. We know that the best distinguisher $\mathsf{A}^\star$ should yield 1 iff $\mathrm{D}(\mathsf{P}_{\mathbf{Z}^q}\|\mathsf{P}_1) \leq \mathrm{D}(\mathsf{P}_{\mathbf{Z}^q}\|\mathsf{P}_0)$ (see (2)). Since this is equivalent to yielding 1 when $2^{q(\mathrm{D}(\mathsf{P}_{\mathbf{Z}^q}\|\mathsf{P}_1)-\mathrm{D}(\mathsf{P}_{\mathbf{Z}^q}\|\mathsf{P}_0))} \leq 1$ and since

$$\mathrm{D}(\mathsf{P}_{\mathbf{Z}^q}\|\mathsf{P}_1) - \mathrm{D}(\mathsf{P}_{\mathbf{Z}^q}\|\mathsf{P}_0) = \sum_z \mathsf{P}_{\mathbf{Z}^q}[z] \log \frac{\mathsf{P}_0[z]}{\mathsf{P}_1[z]} = \frac{1}{q} \log \frac{(\beta/p)^{n_b}}{(n\beta)^q}$$

where n_b denotes the number of times where $Y_1 \oplus Y_2 = b$, then the optimal strategy is to output 1 when

$$\frac{n_b}{q} \geq \frac{\ln(n\beta)}{\ln(\beta/p)} \sim \frac{p}{\ln(np)}$$

Since we take $q \approx \frac{\ln 2}{p}$, this condition is equivalent to $n_b > 0$. Subsequently, we can formalize a differential distinguisher as in Figure 2.

```
Oracle: a permutation c
  for i from 1 to q do
    pick a uniformly distributed random $X$
    query the oracle for $\mathsf{c}(X)$ and $\mathsf{c}(X \oplus a)$
    if $\mathsf{c}(X \oplus a) \oplus \mathsf{c}(X) = b$, output 1 and stop
  end for
  output 0
```

Fig. 2. A differential distinguisher based on the input difference a and output difference b.

A.2 Impossible Differential

The scenario is similar to that considered in the case of differential distinguishers, except that the particular difference b in the ciphertexts can never occur under H_1, i.e., we have $p = 0$. Using the same notations as in Section A.1, we now have $F(\lambda) = (1 - 1/n)^\lambda$ and so $\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1) = -\log(1 - 1/n) \sim \frac{1}{n \ln 2}$. Using (11) we conclude that an iterative distinguisher based on an impossible differential requires

$$q \approx n \ln 2 \tag{16}$$

samples to reach a significant advantage. It is easy to see that this distinguisher should output 1 iff $n_b = 0$.

A.3 Linear Distinguisher

Linear distinguishers [30] are iterated h-distinguishers of order $d = 1$ where $h(x, y) = a \cdot x \oplus b \cdot y \in \{0, 1\}$ (where $\cdot$ denotes the bit-wise xor) for some fixed input mask a and output mask b and for which the distribution D is the uniform distribution. We expect $h(x, y)$ to be biased under H_1 and uniformly distributed under H_0, so that P_0 is assumed to be uniform and $\mathsf{P}_1^{\pm}$ is such that $\mathsf{P}_1^{\pm}[0] = \frac{1}{2}(1 \mp \epsilon)$ and $\mathsf{P}_1^{\pm}[1] = \frac{1}{2}(1 \pm \epsilon)$ for some positive real value ϵ. In this case, we have a composite hypothesis testing problem

$$\mathsf{H}_0 : \mathsf{P} = \mathsf{P}_0 \quad \text{and} \quad \mathsf{H}_1 : \mathsf{P} \in \{\mathsf{P}_1^+, \mathsf{P}_1^-\}$$

In such a case (see [2]), we have a best distinguisher which its acceptance region and advantage can be specified by

$$\Pi^\star = \{\mathsf{P} \ : \ \min_{1 \leq i \leq k} \mathrm{D}(\mathsf{P}\|\mathsf{P}_i) \leq \mathrm{D}(\mathsf{P}\|\mathsf{P}_0)\} \quad \text{and} \quad 1 - \mathrm{BestAdv}_q(\mathsf{P}_0, \mathcal{D}) \doteq \max_{1 \leq i \leq k} 2^{-q\mathrm{C}(\mathsf{P}_0, \mathsf{P}_i)}$$

Assuming that $\epsilon = o(1)$, we have from (12) that $\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1^{\pm}) \approx \frac{\epsilon^2}{8 \ln 2}$ from which we conclude (using (11)) that a linear distinguisher requires

$$q \approx \frac{8 \ln 2}{\epsilon^2} \tag{17}$$

samples to reach a non-negligible advantage. It is easy to see that this linear distinguisher should output 1 iff $\left|2\frac{n_0}{q} - 1\right| \geq \frac{|\epsilon|}{2}$ (where n_0 denotes the number of 0's in the Z_i's), so that we can formalize a linear distinguisher as in Figure 3.

```
Oracle: a permutation $\mathsf{c}$
  initialize a counter $m$ to 0
  for $i$ from 1 to $q$ do
    pick a uniformly distributed random $X$
    query the oracle for $\mathsf{c}(X)$
    if $a \cdot X = b \cdot \mathsf{c}(X)$, increment the counter $m$
  end for
  output 1 if $\left|2\frac{m}{q} - 1\right| \geq \frac{|\epsilon|}{2}$, otherwise output 0.
```

Fig. 3. A linear distinguisher based on the input mask a and output mask b.

B Proof of Theorem 7

Proof. Using (1), we have

$$\begin{aligned} 1 - \mathrm{BestAdv}_q(\mathsf{P}_0, \mathsf{P}_1) &= \sum_{\substack{\mathbf{z}^q \\ \Pr[\mathbf{z}^q|\mathsf{P}_0] > \Pr[\mathbf{z}^q|\mathsf{P}_1]}} \Pr[\mathbf{z}^q|\mathsf{P}_1] + \sum_{\substack{\mathbf{z}^q \\ \Pr[\mathbf{z}^q|\mathsf{P}_0] < \Pr[\mathbf{z}^q|\mathsf{P}_1]}} \Pr[\mathbf{z}^q|\mathsf{P}_0] \\ &= \sum_{\mathbf{z}^q \in \mathcal{Z}'^q} \min\left(\Pr[\mathbf{z}^q|\mathsf{P}_0], \Pr[\mathbf{z}^q|\mathsf{P}_1]\right) \end{aligned}$$

Since for $\forall a, b > 0 : \min(a, b) \leq a^{1-\lambda} b^{\lambda}$ and $0 \leq \lambda \leq 1$, we have

$$\begin{aligned} 1 - \mathrm{BestAdv}_q(\mathsf{P}_0, \mathsf{P}_1) &\leq \inf_{0<\lambda<1} \sum_{\mathbf{z}^q \in \mathcal{Z}'^q} \Pr[\mathbf{z}^q|\mathsf{P}_0]^{1-\lambda} \Pr[\mathbf{z}_q|\mathsf{P}_1]^{\lambda} \\ &= \inf_{0<\lambda<1} \sum_{\mathbf{z}^q \in \mathcal{Z}'^q} \prod_{i=1}^{q} \mathsf{P}_0[z_i]^{1-\lambda} \mathsf{P}_1[z_i]^{\lambda} \\ &= \inf_{0<\lambda<1} \left(\sum_{z \in \mathcal{Z}'} {\mathsf{P}_0}^{1-\lambda}[z] {\mathsf{P}_1}^{\lambda}[z] \right)^q \\ &= 2^{-q\mathrm{C}(\mathsf{P}_0, \mathsf{P}_1)} \end{aligned}$$

□

C Proof of Theorem 8

Proof. Let λ be such that

$$F(\lambda) = \sum_{x \in \mathcal{Z}} \mathsf{P}_0[x]^{1-\lambda} \mathsf{P}_1[x]^{\lambda}$$

and let $\mathsf{P}_1[x] = \mathsf{P}_0[x](1+\epsilon_x)$ with $\epsilon_x \le B_x$, where $B_x = \frac{1}{\mathsf{P}_0[x]} - 1$, We have

$$F(\lambda) = \sum_{x \in \mathcal{Z}} \mathsf{P}_0[x](1+\epsilon_x)^{\lambda}$$

Thanks to the *Taylor Theorem*, for any ϵ there exists $\theta \in [0,1]$, such that

$$(1+\epsilon)^{\lambda} - (1+\lambda\epsilon) = \frac{\lambda(\lambda-1)}{2}\epsilon^2(1+\theta\epsilon)^{\lambda-2}$$

Since $\sum_x \mathsf{P}_0[x](1+\lambda\epsilon_x) = 1$, we obtain

$$\begin{aligned} 1 - F(\lambda) &= \tfrac{\lambda(1-\lambda)}{2} \sum_x \mathsf{P}_0[x]\epsilon_x^2(1+\theta_x\epsilon_x)^{\lambda-2} \\ &= \tfrac{\lambda(1-\lambda)}{2} \sum_x \mathsf{P}_0[x] \frac{(\mathsf{P}_1[x] - \mathsf{P}_0[x])^2}{\mathsf{P}_0[x]^2}(1+\theta_x\epsilon_x)^{\lambda-2} \end{aligned}$$

If $\epsilon_x \ge 0$, then $(1+\theta_x\epsilon_x)^{\lambda-2} \le 1$ and $\mathsf{P}_0[x] \le \mathsf{P}_1[x]$. Otherwise, $(1+\theta_x\epsilon_x)^{\lambda-2} \le \left(\frac{\mathsf{P}_0[x]}{\mathsf{P}_1[x]}\right)^2$ and $\mathsf{P}_1[x] \le \mathsf{P}_0[x]$. Ultimately,

$$1 - \inf_{0<\lambda<1} F(\lambda) \le \frac{1}{8} \sum_{x \in \mathcal{Z}} \mathsf{P}_0[x] \left(\frac{\mathsf{P}_1[x] - \mathsf{P}_0[x]}{\min(\mathsf{P}_0[x], \mathsf{P}_1[x])} \right)^2$$

The other inequality can be shown similarly.

□

D Proof of Lemma 2

Proof. We define P_0 and P_1 as two distributions and compute the Euclidean distance

$$\begin{aligned} \|\mathsf{P}_1 - \mathsf{P}_2\|^2 &= \sum_{k,n} \left(\Pr_{X,N}[h_n[X] = k, N = n] - \frac{1}{2^m \#\mathcal{N}} \right)^2 \\ &= \tfrac{1}{(\#\mathcal{N})^2} \sum_{k,n} \Pr_{X,X'}[h_n[X] = h_n[X'] = k] - \frac{1}{2^m \#\mathcal{N}} \\ &= \tfrac{1}{\#\mathcal{N}} \sum_{x,x'} \Pr[X = x, X' = x', h_N[x] = h_N[x']] - \frac{1}{2^m \#\mathcal{N}} \\ &= \tfrac{1-2^{-m}}{\#\mathcal{N}} \sum_x \Pr[X = x]^2 \\ &\le \tfrac{1-2^{-m}}{\#\mathcal{N}} 2^{-H_\infty[X]} \le \tfrac{1}{2^m \#\mathcal{N}} \epsilon^2 \end{aligned}$$

Applying the link between the statistical distance and Euclidean distance, we obtain $d(\mathsf{P}_1, \mathsf{P}_2) \le \epsilon$.

□