Information-Theoretic Indistinguishability via the Chi-Squared Method

Wei Dai[1], Viet Tung Hoang[2(✉)], and Stefano Tessaro[3(✉)]

[1] Department of Computer Science and Engineering, University of California San Diego, San Diego, USA
weidai@eng.ucsd.edu
[2] Department of Computer Science, Florida State University, Tallahassee, USA
tvhoang@cs.fsu.edu
[3] Department of Computer Science, University of California Santa Barbara, Santa Barbara, USA
tessaro@cs.ucsb.edu

Abstract. Proving tight bounds on information-theoretic indistinguishability is a central problem in symmetric cryptography. This paper introduces a new method for information-theoretic indistinguishability proofs, called "the chi-squared method". At its core, the method requires upper-bounds on the so-called χ^2 divergence (due to Neyman and Pearson) between the output distributions of two systems being queries. The method morally resembles, yet also considerably simplifies, a previous approach proposed by Bellare and Impagliazzo (ePrint, 1999), while at the same time increasing its expressiveness and delivering tighter bounds.

We showcase the chi-squared method on some examples. In particular: (1) We prove an optimal bound of $q/2^n$ for the XOR of two permutations, and our proof considerably simplifies previous approaches using the H-coefficient method, (2) we provide improved bounds for the recently proposed encrypted Davies-Meyer PRF construction by Cogliati and Seurin (CRYPTO '16), and (3) we give a tighter bound for the Swap-or-not cipher by Hoang, Morris, and Rogaway (CRYPTO '12).

Keywords: Symmetric cryptography · Information-theoretic indistinguishability · Provable security

1 Introduction

Information-theoretic indistinguishability proofs are fundamental tools in cryptography, and take a particularly prominent role in symmetric cryptography. In this context, it is imperative to derive bounds which are as precise as possible – a tighter bound yields a better understanding of the actual security of the system at hand, and avoids potential inefficiency provoked by the choice of unnecessarily large parameters, such as the key- and block-lengths, and the number of rounds.

This paper falls within a line of works investigating generic techniques to obtain best-possible information-theoretic bounds. We investigate a new approach to indistinguishability proofs – which we refer to as the *chi-squared method*

J. Katz and H. Shacham (Eds.): CRYPTO 2017, Part III, LNCS 10403, pp. 497–523, 2017.
DOI: 10.1007/978-3-319-63697-9_17

– which will help us tighten (and simplify) proofs for certain examples where proofs so-far have evaded more classical methods, such as the H-coefficient method.

Specifically, we apply our methodology to the analyses of three, a priori seemingly unrelated, constructions – the XOR of permutations (initially studied by Hall, Wagner, Kelsey, and Schneier [12]), the Encrypted Davies-Meyer construction by Cogliati and Seurin [10], and the Swap-or-not construction by Hoang, Morris, and Rogaway [13]. Previously, no connections between these problems have been observed, but we give significantly improved bounds as an application of our framework.

Information-theoretic indistinguishability. Many cryptographic security proofs require showing, for a *distinguisher* A with access to one of two systems, $\mathbf{S}_0$ and $\mathbf{S}_1$,[1] an upper bound on

$$\mathsf{Adv}^{\mathsf{dist}}_{\mathbf{S}_0,\mathbf{S}_1}(A) = \Pr[A(\mathbf{S}_0) = 1] - \Pr[A(\mathbf{S}_1) = 1] \ ,$$

where $A(\mathbf{S}_b)$ denotes the probability that A outputs 1 when interacting with $\mathbf{S}_b$.

While it is customary to only target the case where A is computationally bounded, in many cases, the actual proofs themselves are concerned with the information-theoretic case where the advantage is maximized over *all* distinguishers, only bounded by their number q of *queries*, but with no further restrictions on their time complexities. A first example in this domain is the analysis of Feistel networks in the seminal work of Luby and Rackoff [16], whose main step is a proof that the Feistel construction with truly random round functions is information-theoretically indistinguishable from a random permutation. (This was first pointed out explicitly by Maurer [18].) Another class of inherently information-theoretic analyses – dating back to the analysis of the Even-Mansour [11] block cipher – studies constructions in ideal models (such as the ideal-cipher or random-permutation models), where adversaries are also only bounded in their query-complexity.

In this context, the perhaps most widely-used proof technique is that of bounding the probability of a certain failing condition, where $\mathbf{S}_0$ and $\mathbf{S}_1$ behave *identically*, in some well-defined sense, as long as the condition is not violated. This approach was abstracted e.g. in Maurer's random systems [19] and Bellare-Rogaway game playing [4] frameworks. Unfortunately, such methods are fairly crude, and often fall short of providing tight bounds, especially for so-called beyond-birthday security.[2]

More sophisticated approaches [5,23,25] directly bound the *statistical distance* $\|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_0,A}(\cdot)\|$, where $\mathsf{p}_{\mathbf{S}_1,A}$ and $\mathsf{p}_{\mathbf{S}_0,A}$ are the respective probability distributions of the answers obtained by A, which is assumed to be deterministic. This is an upper bound on $\mathsf{Adv}^{\mathsf{dist}}_{\mathbf{S}_0,\mathbf{S}_1}(A)$. In particular, Patarin's H-coefficient

[1] For now, it suffices to understand such systems informally as interactive objects, or "oracles.".

[2] A not-so-widely known fact is that Maurer, Renner, and Pietrzak [20] show that this method *is* actually optimal, the caveat being however that describing the suitable tight condition may be infeasible, and the result is merely existential.

method [25] has recently re-gained substantial popularity, mostly thanks to Chen and Steinberger's exposition [6]. The technique was further refined by Hoang and Tessaro [14], who provided a "smoothed" version of the H-coefficient method, called the "expectation method."

A DIFFERENT AVENUE. Techniques such as the H-coefficient method heavily exploit computing the probabilities $\mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z})$ and $\mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z})$ that a *full* sequence of q outputs $\boldsymbol{Z} = (Z_1, \ldots, Z_q)$ occur. Often, these probabilities are easy to compute and compare under the condition that the sequence of outputs belongs to a set of good transcripts. One case where such methods however do not yield a good bound is where we are only given local information, e.g., the distance between $\mathsf{p}_{\mathbf{S}_1,A}(\cdot \mid \boldsymbol{Z}_{i-1})$ and $\mathsf{p}_{\mathbf{S}_0,A}(\cdot \mid \boldsymbol{Z}_{i-1})$ for all sequences $\boldsymbol{Z}_{i-1}$ and all $i \geq 1$, where $\boldsymbol{Z}_{i-1}$ is the sequence of the first $i-1$ outputs. Here, the naïve approach is to use a so-called hybrid argument, and bound the distance as

$$\|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_0,A}(\cdot)\| \leq \sum_{i=1}^{q} \mathbf{E}\left[\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot \mid \boldsymbol{X}_{i-1}) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot \mid \boldsymbol{X}_{i-1})\|\right] , \quad (1)$$

where $\boldsymbol{X}_{i-1}$ is the vector of answers to A's first $i-1$ queries, according to $\mathsf{p}_{\mathbf{S}_0,A}(\cdot)$. (Symmetrically, they can be all sampled according to $\mathsf{p}_{\mathbf{S}_1,A}(\cdot)$.) If all summands are upper bounded by ϵ, we obtain a bound of $q\epsilon$. This is rarely tight, and often sub-optimal. A different avenue was explored by Bellare and Impagliazzo (BI) [2], in an unpublished note. They consider the sequence of random variables $U_1, \ldots, U_q$, where

$$U_i = \frac{\mathsf{p}_{\mathbf{S}_1,A}(X_i|\boldsymbol{X}_{i-1})}{\mathsf{p}_{\mathbf{S}_0,A}(X_i|\boldsymbol{X}_{i-1})} ,$$

and $\boldsymbol{X}_{i-1}$ and X_i are sampled from A's interaction with $\mathbf{S}_0$. Roughly, they show that if $|U_i - 1|$ is sufficiently concentrated, say $|U_i - 1| \leq \epsilon$ for all i, except with probability δ, then the bound becomes

$$\|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_0,A}(\cdot)\| \leq O(\sqrt{q} \cdot \epsilon\lambda) + e^{-\lambda^2/2} + \delta .$$

Unfortunately, the BI method is rather complex to use – it requires a careful balancing act in order to assess the trade-off between ϵ and δ, and the additional slackness due to the λ term is also problematic and appear to be an artifact of the proof technique.[3] To the best of our knowledge, the BI method was never used elsewhere.

OUR METHOD: THE CHI-SQUARED METHOD. In this work, we consider a different version of the above method. In particular, we revisit the setting of (1), and change our metric to measure distance between $\mu(\cdot) = \mathsf{p}_{\mathbf{S}_0,A}(\cdot \mid \boldsymbol{Z}_{i-1})$ and

[3] Indeed, it is known [29] that in many cases, if X and Y have statistical distance ϵ, then if one takes vectors $(X_1, \ldots, X_q)$ and $(Y_1, \ldots, Y_q)$ of q independent copies of X and Y, respectively, the statistical distance increases as $\sqrt{q}\epsilon$, this seemingly showing that the BI bound is far from tight.

$\nu(\cdot) = \mathsf{p}_{\mathbf{S}_1,A}(\cdot \mid \boldsymbol{Z}_{i-1})$. Instead of statistical distance, we will use the so-called χ^2-divergence, as proposed by Neyman and Pearson,[4]

$$\chi^2(\mu,\nu) = \sum_x \frac{(\mu(x)-\nu(x))^2}{\nu(x)} \ .$$

where the sum is over all x such that $\nu(x) > 0$, and we assume that if $\mu(x) > 0$, then $\nu(x) > 0$, too. In particular, let $\chi^2(\boldsymbol{Z}_{i-1}) = \chi^2(\mu;\nu)$ as above, then, we show that

$$\|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_0,A}(\cdot)\| \le \sqrt{\frac{1}{2}\sum_{i=1}^{q} \mathbf{E}\left[\chi^2(\boldsymbol{X}_{i-1}))\right]} \ ,$$

where for all $i = 1,\ldots,q$, $\boldsymbol{X}_{i-1}$ is sampled according to $\mathsf{p}_{\mathbf{S}_1,A}(\cdot)$. We refer to the method of obtaining a bound by upper bounding the q expectations $\mathbf{E}\left[\chi^2(\boldsymbol{X}_{i-1}))\right]$ as the *chi-squared method*. A crucial property that will make calculations manageable and elegant is that the distribution of $\boldsymbol{X}_{i-1}$ and the distribution in the denominator of the χ^2-divergence are with respect to different systems. In many case, we will be able to show that $\mathbf{E}\left[\chi^2(\boldsymbol{X}_{i-1}))\right]$ is much smaller than the statistical distance ϵ – even *quadratically*, i.e., $O(\epsilon^2)$ – and thus the method gives a very good bound of the order $O(\sqrt{q}\epsilon)$.

In contrast to the proof behind BI's method, which relies on somewhat heavy machinery, such as Azuma's inequality, the proof behind the chi-squared method is fairly simple, and relies on Pinsker's and Jensen's inequalities. In fact, we are *not* claiming that relations between the statistical distance and χ^2-divergence are novel, but we believe this methodology to be new in the context of cryptography indistinguishability proofs for interactive systems. Our method, as we discuss below in the body of the paper, can also be seen as a generalization of a technique by Chung and Vadhan [8], used in a different context.

We will apply our method to three different problems, improving (or simplifying) existing bounds.

Application: XOR of random permutations. A potential drawback of block ciphers is that their permutation structure makes them unsuitable to be used as good pseudorandom *functions*, as they become distinguishable from a truly random function when reaching $q \approx 2^{n/2}$ queries, where n is the block length. For this reason, Hall, Wagner, Kelsey, and Schneier [12] initiated the study of constructions of good pseudorandom functions from block ciphers with security beyond the so-called Birthday barrier, i.e., above $2^{n/2}$. A particularly simple construction they proposed – which we refer to as the *XOR construction* – transforms a permutation $\pi : \{0,1\}^n \to \{0,1\}^n$ into a function $f : \{0,1\}^{n-1} \to \{0,1\}^n$ by computing $f(x) = \pi(0\|x) \oplus \pi(1\|x)$, where π is meant to be instantiated by a block cipher which is a good pseudorandom permutation, but is treated as a random permutation in the core argument of the proof, which we focus on.

[4] This is in fact Neyman's version—the divergence is not symmetric, and Pearson's version swaps the order of μ and ν.

Lucks [17] proved this construction be secure up to roughly $q = 2^{2n/3}$, whereas Bellare and Impagliazzo [2] gave a better bound of $O(n)q/2^n$, but also only provided a proof sketch. Patarin [24] gave an improved bound of $O(q/2^n)$, but the proof was quite complex. This bound was further improved to $q/2^n$ in an unpublished manuscript [26]. Patarin's tight proof is very involved, using an approach he refers to as "mirror theory",[5] with some claims remaining open or unproved. (Also, as a related problem, Cogliati, Lampe, and Patarin [9] gave weaker bounds for the case of the sum of at least three permutations.) The XOR construction is particularly helpful as a tool for beyond-birthday security, and has been used for example within Iwata's CENC mode of operation [15].

Here, as an application of the chi-squared method, we give a fairly simple proof giving us a bound of $(1.5q + 3\sqrt{q})/2^n$. One can argue that the improvement is small (and in fact, if the bound in [26] is indeed correct, ours is slightly worse). However, we believe the analysis of the XOR construction to be fundamental, and it has evaded simple proofs for nearly two decades. While Patarin's proof deals with precise bounds on number of permutations satisfying a given input-output relationship, our approach is simpler in that it does not require a fine-grained understanding of the underlying distribution, but only requires computing certain expectations.

A related version of the construction is the one computing $f'(x) = \pi_1(x) \oplus \pi_2(x)$ for two independent permutations π_1, π_2. We also analyze this variant in Appendix A, giving a bound of $q^{1.5}/2^{1.5n}$, and in the body focus on the "single-key" variant which is somewhat harder to analyze and more efficient.

Application: The EDM construction. As another application of the chi-squared method, we study the *encrypted Davies-Meyer* (EDM) construction recently introduced by Cogliati and Seurin [10]. The construction depends on two random permutations π and π', and on input x outputs the value $\pi'(\pi(x) \oplus x)$. Again, the goal is to show that this is a good PRF, with security beyond the birthday barrier. In [10], a security bound showing security up to $q = 2^{2n/3}$ queries was shown. Using the chi-squared method, we show that security up to $q = 2^{3n/4}$ is achieved. We note that in concurrent work to ours, Mennink and Neves [21] prove that EDM security approaches 2^n. Their bound uses Patarin's mirror theory, and has a different purpose than ours – we aim for a simpler-to-use framework, and the question of whether our approach yields better bounds remains open for future work.

The EDM construction is the underlying structure of a nonce-based misuse-resistant MAC that CS proposed. CS proved that the MAC construction also achieves $2n/3$-bit of security and conjecture that it actually has n-bit security. While our chi-squared technique seems to be able to handle the MAC construction as well, the combinatorics (also in CS's work) will be very complex, and thus we leave this analysis for future work.

[5] In essence, what mirror theory accounts to is reducing the problem of applying the H-coefficient method to a combinatorial problem counting solutions of a system of linear equations with contraints on their (discrete) solutions.

Application: Swap-or not. As our final application, we consider the swap-or-not block cipher, introduce by Hoang, Morris, and Rogaway [13]. Swap-or-not is a block cipher that supports an arbitrary abelian group $\mathbb{G}$ with size N as its domain, and, for sufficiently many rounds $r = \Omega(\log(N))$, is meant to with stand up to $q < N/c$ queries, for a small constant $c \geq 2$. This makes it particularly suitable as cipher for format-preserving encryption (FPE) [3], both because of its flexibility to support multiple domain formats, as well as for its high security making it suitable to smaller domains. Subsequent work [22,30] focused on boosting its security to $q = N$, at the cost of higher (worst-case) round complexity. The Swap-or-not example is particularly interesting object to analyze, as it uses a very different structure than more usual Feistel-like designs. The original proof in [13] uses a fairly ad-hoc analysis, which however as an intermediate step ends up upper bounding exactly the quantity $\mathbf{E}\left[\chi^2(\boldsymbol{X}_{i-1})\right]$. As a result of this, we end up saving a factor $\sqrt{N}$ on final advantage bound.

For example, for $N = 2^{64}$, $q = 2^{60}$, and r rounds, the original analysis gives a CCA-security advantage $2^{90-0.415r}$ vs one of approximately $2^{62-0.415r}$ for our new analysis. Thus, if we are interested in achieving security 2^{-64}, we would need $r \geq 371$ rounds according to the old analysis, whereas our analysis shows that 293 rounds are sufficient.

A perspective and further related works. We conclude by stressing that with respect to our current state of knowledge, there does not seem to be a universal method to obtain tight bounds on information-theoretic indistinguishability, and ultimately the best method depends on the problem at hand. This situation is not different than what encountered in statistics, where proving bounds on the variational distance require different tools depending on the context.

We are certainly not the first to observe the importance of using different metrics as a tool in cryptographic security proofs and reductions. For example, in symmetric cryptography, Steinberger [31] used the Hellinger distance to sharpen bounds on key-alternating ciphers. The H-coefficient technique itself can be seen as bounding a different distance metric between distributions. Further, cryptographic applications have often relied on using the KL-divergence, e.g., in parallel repetition theorems [7,28], and Renyi divergences, e.g., in lattice-based cryptography [1].

2 Preliminaries

Notation. Let n be a positive integer. We use $[n]$ to denote the set $\{1, \ldots, n\}$. For a finite set S, we let $x \leftarrow\!\!\$\, S$ denote the uniform sampling from S and assigning the value to x. Let $|x|$ denote the length of the string x, and for $1 \leq i < j \leq |x|$, let $x[i, j]$ denote the substring from the ith bit to the jth bit (inclusive) of x. If A is an algorithm, we let $y \leftarrow A(x_1, \ldots; r)$ denote running A with randomness r on inputs $x_1, \ldots$ and assigning the output to y. We let $y \leftarrow\!\!\$\, A(x_1, \ldots)$ be the resulting of picking r at random and letting $y \leftarrow A(x_1, \ldots; r)$.

PRF SECURITY. Let $F : \mathcal{K} \times \{0,1\}^m \to \{0,1\}^n$ be a family of functions. Let $\mathrm{Func}(m,n)$ be the set of all functions $g : \{0,1\}^m \to \{0,1\}^n$. For an adversary A, define

$$\mathsf{Adv}_F^{\mathrm{prf}}(A) = \Pr[K \leftarrow\!\!\$\ \mathcal{K}; A^{F_K(\cdot)} \Rightarrow 1] - \Pr[f \leftarrow\!\!\$\ \mathrm{Func}(m,n); A^{f(\cdot)} \Rightarrow 1]$$

as the PRF advantage of A attacking F.

DISTANCE MEASURES. Let μ and ν be two distributions on a finite event space Ω. The *statistical distance* between μ and ν is defined as

$$\|\mu - \nu\| = \sum_{x \in \Omega} \max\{0, \mu(x) - \nu(x)\}.$$

The *Kullback-Leibler (KL) divergence* between μ and ν is defined as

$$\Delta_{\mathrm{KL}}(\mu,\nu) = \sum_{x \in \Omega} \mu(x) \ln\Big(\frac{\mu(x)}{\nu(x)}\Big).$$

Note that for Δ_{KL} to be well-defined, we need ν to have full support, i.e. Ω. The well-known Pinsker's inequality relates the previous two notions.

Lemma 1 (Pinsker's inequality). Let μ and ν be two distributions on a finite event space Ω such that ν has full support. Then

$$(\|\mu - \nu\|)^2 \le \frac{1}{2}\Delta_{\mathrm{KL}}(\mu,\nu).$$

Another well-known fact for KL-divergence is that it decomposes nicely for product distributions. The *chi-squared divergence* between μ and ν is defined as

$$\chi^2(\mu,\nu) = \sum_{x \in \Omega} \frac{(\mu(x) - \nu(x))^2}{\nu(x)}.$$

Note that for $\chi^2(\mu,\nu)$ to be well-defined, again ν needs to have full support. We remark that $\chi^2(\mu,\nu)$ is related to the notion of collision probability. To justify this remark, let Ω be some finite set and let $M = |\Omega|$. Let ν be the uniform distribution over Ω and μ be any distribution over Ω. Let X_1, X_2 be two i.i.d. samples from μ. Then

$$\begin{aligned}\chi^2(\mu,\nu) &= \sum_{x \in \Omega} M \cdot (\mu(x) - 1/M)^2 \\ &= M \cdot \Pr[X_1 = X_2] - 1.\end{aligned}$$

The following lemma relates the chi-squared divergence and the KL-divergence.

Lemma 2. Let Ω be a finite set, and let μ and ν be two distribution on Ω such that ν has full support. Then

$$\Delta_{\mathrm{KL}}(\mu,\nu) \le \chi^2(\mu,\nu).$$

Proof. Since function $\ln(x)$ is concave, by using Jensen's inequality,

$$\sum_{x\in\Omega} \mu(x)\ln\Big(\frac{\mu(x)}{\nu(x)}\Big) \le \ln\Big(\sum_{x\in\Omega}\frac{(\mu(x))^2}{\nu(x)}\Big). \tag{2}$$

Next,

$$\sum_{x\in\Omega}\frac{(\mu(x)-\nu(x))^2}{\nu(x)} = \sum_{x\in\Omega}\frac{(\mu(x))^2}{\nu(x)} - \sum_{x\in\Omega}(2\mu(x)-\nu(x)) = \sum_{x\in\Omega}\frac{(\mu(x))^2}{\nu(x)} - 1. \tag{3}$$

Finally, using the inequality that $e^t - 1 \ge t$ for any real number t, we have

$$\sum_{x\in\Omega}\frac{(\mu(x))^2}{\nu(x)} - 1 \ge \ln\Big(\sum_{x\in\Omega}\frac{(\mu(x))^2}{\nu(x)}\Big). \tag{4}$$

From Eqs. (2)–(4), we obtain the claimed result.

3 The Chi-Squared Method

In this section, we describe the chi-squared method, which simplifies previous results by Bellare and Impagliazzo (BI), and Chung and Vadhan (CV) [2,8].

Notational setup. Let A be an adversary that tries to distinguish two stateless systems $\mathbf{S}_1$ and $\mathbf{S}_0$. Since we allow A to be computationally unbounded, without loss of generality, assume that A is deterministic. Assume further that the adversary always makes exactly q queries. Since the adversary is deterministic, for any $i \le q-1$, the answers for the first i queries completely determine the first $i+1$ queries. For a system $\mathbf{S} \in \{\mathbf{S}_1, \mathbf{S}_0\}$ and strings $z_1, \ldots, z_i$, let $\mathsf{p}_{\mathbf{S},A}(z_1, \ldots, z_i)$ denote the probability that when the adversary A interacts with system $\mathbf{S}$, the answers for the first i queries that it receives is $z_1, \ldots, z_i$. If $\mathsf{p}_{\mathbf{S},A}(z_1, \ldots, z_i) > 0$, let $\mathsf{p}_{\mathbf{S},A}(z_{i+1} \mid z_1, \ldots, z_i)$ denote the conditional probability that the answer for the $(i+1)$-th query when the adversary interacts with system $\mathbf{S}$ is z_{i+1}, given that the answers for the first i queries are $z_1, \ldots, z_i$ respectively. For each $\boldsymbol{Z} = (z_1, \ldots, z_q)$, let $\boldsymbol{Z}_i = (z_1, \ldots, z_i)$, and for $\mathbf{S} \in \{\mathbf{S}_1, \mathbf{S}_0\}$, let $\mathsf{p}_{\mathbf{S},A}(\cdot \mid \boldsymbol{Z}_i)$ denote $\mathsf{p}_{\mathbf{S},A}(\cdot \mid z_1, \ldots, z_i)$. We let $\boldsymbol{Z}_0$ be the empty vector, and $\mathsf{p}_{\mathbf{S},A}(\cdot \mid \boldsymbol{Z}_0)$ is understood as $\mathsf{p}_{\mathbf{S},A}(\cdot)$.

The technique. We first give a brief intuition regarding our technique. On the high level, the chi-squared method relates the statistical distance of a product distribution to the expected chi-squared divergence of the components, via Kullback-Leibler divergence. The advantage of this approach is that the term that depends on the number of components, say q, is "under the square-root", because of Pinsker's inequality. The details follow.

For each $i \le q$ and each vector $\boldsymbol{Z}_{i-1} = (z_1, \ldots, z_{i-1})$, define (with sligh abuse of notation)

$$\begin{aligned}\chi^2(\boldsymbol{Z}_{i-1}) &= \chi^2(\mathsf{p}_{\mathbf{S}_1,A}(\cdot \mid \boldsymbol{Z}_{i-1}),\ \mathsf{p}_{\mathbf{S}_0,A}(\cdot \mid \boldsymbol{Z}_{i-1})) \\ &= \sum_{z_i} \frac{\Big(\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) - \mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})\Big)^2}{\mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})},\end{aligned}$$

where the sum is taken over all z_i in the support of the distribution $\mathsf{p}_{\mathbf{S}_0,A}(\cdot \mid \boldsymbol{Z}_{i-1})$. We require that if $\mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}_i) > 0$, then so is $\mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z}_i)$. Thus, $\chi^2(\boldsymbol{Z}_{i-1})$ is well-defined. Typically, in applications, $\mathbf{S}_0$ is the "ideal" system, and this technical constraint is always met.

The following lemma bounds the distinguishing advantage of A.

Lemma 3. Suppose whenever $\mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}_i) > 0$ then $\mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z}_i) > 0$. Then,

$$\|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_0,A}(\cdot)\| \le \Big(\frac{1}{2}\sum_{i=1}^{q} \mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})]\Big)^{1/2},$$

where the expectation is taken over vectors $\boldsymbol{X}_{i-1}$ of the $i-1$ first answers sampled according to the interaction with $\mathbf{S}_1$.

Discussion. To illustrate the power of the chi-squared method, suppose that

$$\Big|\frac{\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1})}{\mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})} - 1\Big| \le \varepsilon$$

for every i and every $\boldsymbol{Z}_i$. If one uses the H-coefficient technique, the first step is to give a lower bound for the ratio $\mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z})/\mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z})$, which is

$$\prod_{i=1}^{q} \frac{\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1})}{\mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})} \ge (1-\varepsilon)^q \ge 1 - \varepsilon q.$$

Thus the distinguishing advantage is at most the statistical distance between $\mathsf{p}_{\mathbf{S}_0,A}(\cdot)$ and $\mathsf{p}_{\mathbf{S}_1,A}(\cdot)$, which is

$$\sum_{\boldsymbol{Z}} \max\{0, \mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z}) - \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z})\} \le \sum_{\boldsymbol{Z}} \varepsilon q \cdot \mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z}) \le \varepsilon q.$$

In contrast, from Lemma 3, the distinguishing advantage is at most $\varepsilon\sqrt{q/2}$, because

$$\begin{aligned}\chi^2(\boldsymbol{Z}_{i-1}) &= \sum_{z_i} \mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})\Big(\frac{\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1})}{\mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})} - 1\Big)^2 \\ &\le \sum_{z_i} \mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1}) \cdot \varepsilon^2 = \varepsilon^2.\end{aligned}$$

This is why the chi-square method can substantially improve the security bound in many settings, as we'll demonstrate in subsequent sections.

Proof (of Lemma 3). Recall that the adversary's distinguishing advantage is at most the statistical distance between $\mathsf{p}_{\mathbf{S}_0,A}(\cdot)$ and $\mathsf{p}_{\mathbf{S}_1,A}(\cdot)$. On the other hand, from Pinsker's inequality,

$$\begin{aligned}
2\Big(\|\ \mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_0,A}(\cdot)\|\Big)^2 &\le \sum_{\boldsymbol{Z}} \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}) \ln\Big(\frac{\mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z})}{\mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z})}\Big) \\
&= \sum_{\boldsymbol{Z}=(z_1,\ldots,z_q)} \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}) \ln\Big(\prod_{i=1}^{q} \frac{\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1})}{\mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})}\Big) \\
&= \sum_{\boldsymbol{Z}=(z_1,\ldots,z_q)} \sum_{i=1}^{q} \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}) \ln\Big(\frac{\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1})}{\mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})}\Big) \\
&= \sum_{i=1}^{q} \sum_{\boldsymbol{Z}_i=(z_1,\ldots,z_i)} \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}_{i-1}) \cdot \mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) \cdot \ln\Big(\frac{\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1})}{\mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})}\Big) \quad (5)
\end{aligned}$$

Fix $i \le q$ and $\boldsymbol{Z}_{i-1}$. Let μ and ν be the distributions $\mathsf{p}_{\mathbf{S}_1,A}(\cdot \mid \boldsymbol{Z}_{i-1})$ and $\mathsf{p}_{\mathbf{S}_0,A}(\cdot \mid \boldsymbol{Z}_{i-1})$ respectively. Let S be the support of ν, and recall that the support of μ is a subset of S. Notice that from Lemma 2, we have

$$\sum_{x \in S} \mu(x) \ln\Big(\frac{\mu(x)}{\nu(x)}\Big) \le \sum_{x \in S} \frac{(\mu(x) - \nu(x))^2}{\nu(x)}. \quad (6)$$

From Eqs. (5) and (6),

$$\begin{aligned}
&2\Big(\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\|\Big)^2 \\
&\le \sum_{i=1}^{q} \sum_{\boldsymbol{Z}_i=(z_1,\ldots,z_i)} \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}_{i-1}) \frac{\Big(\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) - \mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})\Big)^2}{\mathsf{p}_{\mathbf{S}_0,A}(z_i \mid \boldsymbol{Z}_{i-1})} \\
&= \sum_{i=1}^{q} \sum_{\boldsymbol{Z}_i=(z_1,\ldots,z_i)} \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}_{i-1}) \cdot \chi^2(\boldsymbol{Z}_{i-1}) = \sum_{i=1}^{q} \mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})].
\end{aligned}$$

This concludes the proof. □

Comparison with CV's framework. Underneath CV's work is, in essence, a specialized treatment of our framework for the case that the ideal system $\mathbf{S}_0$ implements an ideal random function. Thus their method can be used to justify the security of the xor of two permutations (Sect. 4) and Encrypted Davies-Meyer PRF (Sect. 5), but it does not work for the Swap-or-Not shuffle (Sect. 6). CV however do not realize these potential applications, and focus only on the Generalized Leftover Hash Lemma (GLHL) of block sources. To the best of our knowledge, CV's method is never used for any other application, perhaps because it is written in a specific language for the context of GLHL.

Comparison with BI's framework. Compared to BI's framework, ours is better in both usability and tightness.

– In BI's method, the bound is a formula of two user-provided parameters. Consequently, to use BI's method, one has to fine-tune the parameters to optimize the bound. Moreover, since BI's method requires strong concentration bounds, in applications such as the xor of two permutations, one has to make non-trivial use of martingales and Azuma's inequality.[6] In contrast, under the chi-squared method, in Sect. 4, when we handle the xor of two permutations, we only compute an expectation and there's no need to use advanced probabilistic tools.
– Due to BI's requirement of strong concentration bounds, in some settings the results that BI's method obtains can be sub-optimal. The looseness in BI's method varies greatly among different settings. For example, in the xor of two permutations, BI's bound is about $nq/2^n$, whereas ours is just $q/2^n$. For Encrypted Davies-Meyer PRF, BI's method only gives $\frac{2n}{3}$-bit security, which is on par with the result of Cogliati and Seurin via the H-Coefficient technique, but our method yields $\frac{3n}{4}$-bit security. Finally, for the Swap-or-Not shuffle, BI's framework doesn't mesh with the analysis in [13], whereas our method can easily make use of the analysis in [13] to improve their result.

4 The XOR Construction

In this section, we consider the so called xor-construction, which was initially proposed in [12], and which is used to obtain, efficiently, a good pseudorandom function from a block cipher. Here, in particular, we consider a version which only involved *one* permutation (at the price of a slightly smaller domain). We analyze a two-permutation version in Appendix A.

Setup and main theorem. Let $\mathrm{Perm}(n)$ be the set of permutations $\pi : \{0,1\}^n \to \{0,1\}^n$. Define $\mathsf{XOR}[n] : \mathrm{Perm}(n) \times \{0,1\}^{n-1} \to \{0,1\}^n$ to be the construction that takes a permutation $\pi \in \mathrm{Perm}(n)$ as a key, and on input x it returns $\pi(x \,\|\, 0) \oplus \pi(x \,\|\, 1)$. Theorem 1 below gives the PRF security of $\mathsf{XOR}[n]$.

Theorem 1. Fix an integer $n \geq 8$. For any adversary A that makes $q \leq 2^{n-5}$ queries we have

$$\mathsf{Adv}^{\mathrm{prf}}_{\mathsf{XOR}[n]}(A) \leq \frac{1.5q + 3\sqrt{q}}{2^n}.$$

Discussion. Before we proceed into the proof, we have a few remarks. First, the bound in Theorem 1 is tight, since in the real system (the one implementing $\mathsf{XOR}[n]$), no answer can be 0^n. Hence if one simply looks for a 0^n-answer among q queries, one can distinguish the two systems with advantage

[6] This fact was not explicit. Indeed, BI provided only a proof sketch, claiming a bound $O(n)q^{1.5}/2^{1.5n}$ for the xor of two permutations, and their proof relies on the Chernoff bound. However, in their application, the resulting Bernoulli random variables are dependent, and thus a correct proof would need to use Azuma's inequality. We made non-trivial attempts to fix their proof using Azuma inequality, but could only recover a bound around $20nq/2^n$.

$1-(1-1/2^n)^q \approx q/2^n$. Next, if we blindly use the chi-squared method, with $\mathbf{S}_1$ being the real system, and $\mathbf{S}_0$ the ideal one (the one implementing a uniformly random function), then the bound is weak, around $\sqrt{q/2^n}$. The reason is that, for each $i \le q$ and $\boldsymbol{Z}_{i-1} = (z_1, \ldots, z_{i-1})$ that the real system can produce for its first $i-1$ answers,

$$\chi^2(\boldsymbol{Z}_{i-1}) \ge \frac{\left(\mathsf{p}_{\mathbf{S}_1}(0^n \mid \boldsymbol{Z}_{i-1}) - \mathsf{p}_{\mathbf{S}_0}(0^n \mid \boldsymbol{Z}_{i-1})\right)^2}{\mathsf{p}_{\mathbf{S}_0}(0^n \mid \boldsymbol{Z}_{i-1})} = \frac{1}{2^n}.$$

Hence when we sample $\boldsymbol{X}_{i-1}$ according to the interaction with $\mathbf{S}_1$, it holds that $\mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})] \ge 1/2^n$, and consequently we end up with an inferior bound $\sqrt{q/2^n}$. To avoid this issue, the system $\mathbf{S}_0$ in our proof is instead a "normalized" version of the ideal system. It only outputs uniformly random answers in $\{0,1\}^n \backslash \{0^n\}$. This normalization introduces a term $q/2^n$ in the bound, but the important point is that this term won't be under the square-root. We will use the chi-squared method with $\mathbf{S}_1$ being the real system, and $\mathbf{S}_0$ being the normalized ideal system.

*Proof (*Theorem 1*).* Let $\mathbf{S}_1$ be the real system, and let $\mathbf{S}_2$ be the ideal system. To obtain a good advantage, as explained above, we'll first "normalize" $\mathbf{S}_2$ to obtain another system $\mathbf{S}_0$. Let $\mathbf{S}_0$ be the system that implements an ideal random function mapping $\{0,1\}^{n-1}$ to $\{0,1\}^n \backslash \{0^n\}$. Let $\Gamma_{\text{good}} = (\{0,1\}^n \backslash \{0^n\})^q$, and $\Gamma_{\text{bad}} = (\{0,1\}^n)^q \backslash \Gamma_{\text{good}}$. Recall that $\mathsf{Adv}^{\text{xor}}(A,n)$ is at most the statistical distance between $\mathsf{p}_{\mathbf{S}_1,A}$ and $\mathsf{p}_{\mathbf{S}_2,A}$. From triangle inequality,

$$\|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_2,A}(\cdot)\| \le \|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_0,A}(\cdot)\| + \|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_2,A}(\cdot)\|.$$

Let T be the random variable for the q answers in $\mathbf{S}_2$. Then

$$\begin{aligned}\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_2,A}(\cdot)\| &= \sum_{\boldsymbol{Z}} \max\{0, \mathsf{p}_{\mathbf{S}_2,A}(\boldsymbol{Z}) - \mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z})\} \\ &= \sum_{\boldsymbol{Z} \in \Gamma_{\text{bad}}} \mathsf{p}_{\mathbf{S}_2,A}(\boldsymbol{Z}) = \Pr[T \in \Gamma_{\text{bad}}]\end{aligned}$$

where the second equality is due to the fact that $\mathsf{p}_{\mathbf{S}_2,A}(\boldsymbol{Z}) > \mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z})$ if and only if $\boldsymbol{Z} \in \Gamma_{\text{bad}}$, and $\mathsf{p}_{\mathbf{S}_0,A}(\boldsymbol{Z}) = 0$ for every $\boldsymbol{Z} \in \Gamma_{\text{bad}}$. Note that $\Pr[T \in \Gamma_{\text{bad}}]$ is the probability that among q answers in $\mathbf{S}_2$ (the system implementing a uniformly random function), there is at least a 0^n-answer, which happens with probability at most $q/2^n$.

What is left is to bound $\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\|$. We shall use the chi-squared method. Let $\boldsymbol{X} = (X_1, \ldots, X_q)$ be the random variable for the q answers in $\mathbf{S}_1$, and let $\boldsymbol{X}_i = (X_1, \ldots, X_i)$ for every $i \le q$. Fix $i \le q$ and fix $x \in \{0,1\}^n \backslash \{0^n\}$. Let $Y_{i,x}$ be the following random variable. If $\boldsymbol{X}_{i-1}$ takes values $(z_1, \ldots, z_{i-1})$ then $Y_{i,x}$ takes the value $\mathsf{p}_{\mathbf{S}_1,A}(x \mid z_1, \ldots, z_{i-1})$. Recall that

$$\begin{aligned}\chi^2(\boldsymbol{X}_{i-1}) &= \sum_{x \in \{0,1\}^n \backslash \{0^n\}} \frac{(Y_{i,x} - 1/(2^n-1))^2}{1/(2^n-1)} \\ &\le \sum_{x \in \{0,1\}^n \backslash \{0^n\}} 2^n \cdot (Y_{i,x} - 1/(2^n-1))^2. \qquad (7)\end{aligned}$$

We now expand $Y_{i,x}$ into a more expressive and convenient formula to work with. Let $\pi \in \mathrm{Perm}(n)$ be the secret key of $\mathsf{XOR}[n]$. Let $m_1, \ldots, m_i$ be the first i queries of the adversary. Let $V_1 = \pi(m_1 \,\|\, 0), V_2 = \pi(m_1 \,\|\, 1), \ldots, V_{2i-3} = \pi(m_{i-1} \,\|\, 0)$, and $V_{2i-2} = \pi(m_{i-1} \,\|\, 1)$. Regardless of how the adversary chooses its queries, marginally, these $V_1, \ldots, V_{2i-2}$ are simply random variables sampled uniformly without replacement from $\{0,1\}^n$. Let $S = \{V_1, \ldots, V_{2i-2}\}$. Let $D_{i,x}$ be the number of pairs $(u, u \oplus x)$ such that both u and $u \oplus x$ belongs to S. Note that S and $D_{i,x}$ are both random variables, and in fact functions of the random variables $V_1, \ldots, V_{2i-2}$. If $\pi(m_i \,\|\, 0) \oplus \pi(m_i \,\|\, 1) = x$, there are exactly $2^n - 4(i-1) + D_{i,x}$ choices for the pair $(\pi(m_i \,\|\, 0), \pi(m_i \,\|\, 1))$:

- First, $\pi(m_i \| 0)$ must take value in $\{0,1\}^n \backslash (S \cup S^*)$, where $S^* = \{u \oplus x \mid u \in S\}$. There are exactly $2^n - |S \cup S^*| = 2^n - |S| - |S^*| + |S \cap S^*| = 2^n - 4(i-1) + D_{i,x}$ choices for $\pi(m_i \,\|\, 0)$.
- Once $\pi(m_i \,\|\, 0)$ is fixed, the value of $\pi(m_i \,\|\, 1)$ is determined.

Hence

$$Y_{i,x} = \frac{2^n - 4(i-1) + D_{i,x}}{(2^n - 2i + 1)(2^n - 2i)},$$

and thus

$$|Y_{i,x} - 1/(2^n - 1)| = \frac{|(2^n - 1)D_{i,x} - 4(i-1)^2 + 2(2^n - i)|}{(2^n - 2i + 1)(2^n - 2i)(2^n - 1)}.$$

Note that

$$\begin{aligned}
&\frac{|(2^n - 1)D_{i,x} - 4(i-1)^2 + 2(2^n - i)|}{2^n - 1} \\
&= \left| D_{i,x} - \frac{4(i-1)^2}{2^n - 1} + 2 - \frac{2(i-1)}{2^n - 1} \right| \\
&= \left| D_{i,x} - \frac{4(i-1)^2}{2^n} + 2 - \frac{2(i-1)}{2^n - 1} - \frac{4(i-1)^2}{2^n(2^n - 1)} \right| \\
&\leq \left| D_{i,x} - \frac{4(i-1)^2}{2^n} \right| + 2 - \frac{2(i-1)}{2^n - 1} - \frac{4(i-1)^2}{2^n(2^n - 1)} \\
&\leq \left| D_{i,x} - \frac{4(i-1)^2}{2^n} \right| + 2,
\end{aligned}$$

where the first inequality is due to the facts that (i) $|a + b| \leq |a| + |b|$ for any numbers a and b, and (ii) $2 - \frac{2(i-1)}{2^n - 1} - \frac{4(i-1)^2}{2^n(2^n-1)} > 0$, which is in turn due to the hypothesis that $i \leq q \leq 2^{n-5}$, and $n \geq 8$. Dividing both sides by $(2^n - 2i + 1)(2^n - 2i)$ we have

$$
\begin{aligned}
|Y_{i,x} - 1/(2^n-1)| &\le \frac{|D_{i,x} - 4(i-1)^2/2^n| + 2}{(2^n - 2i + 1)(2^n - 2i)} \\
&\le \frac{|D_{i,x} - 4(i-1)^2/2^n| + 2}{\frac{7}{8} \cdot 2^{2n}} \\
&= \frac{\frac{8}{7} \cdot |D_{i,x} - 4(i-1)^2/2^n| + \frac{16}{7}}{2^{2n}} \\
&\le \frac{\frac{8}{7} \cdot |D_{i,x} - 4(i-1)^2/2^n| + 3}{2^{2n}},
\end{aligned}
$$

where the second inequality is also due to the hypothesis that $i \le q \le 2^{n-5}$, and $n \ge 8$. Using the fact that $(a+b)^2 \le 2(a^2+b^2)$ for every real numbers a and b,

$$
\begin{aligned}
(Y_{i,x} - 1/(2^n-1))^2 &\le \frac{\frac{128}{49}(D_{i,x} - 4(i-1)^2/2^n)^2 + 18}{2^{4n}} \\
&\le \frac{3(D_{i,x} - 4(i-1)^2/2^n)^2 + 18}{2^{4n}}.
\end{aligned}
$$

From Eq. (7),

$$
\begin{aligned}
\mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})] &\le \sum_{x \in \{0,1\}^n \setminus \{0^n\}} 2^n \cdot \mathbf{E}\Big[(Y_{i,x} - 1/(2^n-1))^2\Big] \\
&\le \sum_{x \in \{0,1\}^n \setminus \{0^n\}} \frac{18}{2^{3n}} + \frac{3}{2^{3n}} \ \mathbf{E}\Big[\Big(D_{i,x} - \frac{4(i-1)^2}{2^n}\Big)^2\Big].
\end{aligned}
$$

In the last formula, it is helpful to think of each $D_{i,x}$ as a function of $V_1, \ldots, V_{2n-2}$, and the expectation is taken over the choices of $V_1, \ldots, V_{2n-2}$ sampled uniformly without replacement from $\{0,1\}^n$. We will show that for any $x \in \{0,1\}^n \setminus \{0^n\}$,

$$
\mathbf{E}\Big[\Big(D_{i,x} - \frac{4(i-1)^2}{2^n}\Big)^2\Big] \le \frac{4(i-1)^2}{2^n}, \tag{8}
$$

and thus

$$
\mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})] \le \sum_{x \in \{0,1\}^n \setminus \{0^n\}} \Big(\frac{18}{2^{3n}} + \frac{12(i-1)^2}{2^{4n}}\Big) \le \frac{18}{2^{2n}} + \frac{12(i-1)^2}{2^{3n}}.
$$

Summing up, from Lemma 3,

$$
\begin{aligned}
(\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\|)^2 &\le \frac{1}{2}\sum_{i=1}^{q} \mathbf{E}[\chi^2(X_{i-1})] \\
&\le \frac{1}{2}\sum_{i=1}^{q} \frac{18}{2^{2n}} + \frac{12(i-1)^2}{2^{3n}} \\
&\le \frac{1}{2}\Big(\frac{18q}{2^{2n}} + \frac{4q^3}{2^{3n}}\Big) \le \frac{9q + 0.25q^2}{2^{2n}},
\end{aligned}
$$

where the last inequality is due to the hypothesis that $q \le 2^{n-5}$.

We now justify Eq. (8). Fix $x \in \{0,1\}^n \backslash \{0^n\}$. For each $1 \leq j \leq 2i-2$, let B_j be the Bernoulli random variable such that $B_j = 1$ if and only if $V_j \in \{V_1 \oplus x, \ldots, V_{j-1} \oplus x\}$. Then $D_{i,x} = 2(B_1 + \cdots B_{2i-2})$: if $V_j = V_k \oplus x$ for some $k < j$, then these account for two pairs (u, v) such that $v = u \oplus x$, whereas $B_k = 0$ and $B_j = 1$. Let $S_k = B_1 + \cdots + B_k$, and $L_k = S_k - k^2/2^{n+1}$. We will prove by induction that for any $k \leq 2i-2$,

$$\mathbf{E}\Big[(L_k)^2\Big] \leq \frac{2k^2}{2^{n+1}}, \text{ and}$$
$$\mathbf{E}\Big[L_k\Big] \geq \frac{-k}{2^{n+1}}.$$

This subsumes Eq. (8) as the special case for $k = 2i-2$. The base case $k = 1$ is vacuous, since $B_1 = 0$. Suppose this holds for $k-1$; we'll prove that it holds for k as well. Given $B_1, \ldots, B_{k-1}$, the conditional probability that $B_k = 1$ is exactly

$$p = \frac{k - 1 - 2S_{k-1}}{2^n - (k-1)}$$

because it is equally likely for V_k to take any value in $\{0,1\}^n \backslash P$, where $P = \{V_1, \ldots, V_{k-1}\}$ and $2S_{k-1}$ is the number of elements $u \in P$ such that $u \oplus x$ is also in P. Moreover,

$$\frac{k-1-2S_{k-1}}{2^n-(k-1)} = \frac{k-1-2(L_{k-1}+(k-1)^2/2^{n+1})}{2^n-(k-1)} = \frac{k-1}{2^n} - \frac{2L_{k-1}}{2^n-(k-1)}.$$

Hence $p = \frac{k-1}{2^n} - \frac{2L_{k-1}}{2^n-(k-1)}$, and thus

$$\begin{aligned}
\mathbf{E}[L_k] &= \mathbf{E}[L_{k-1} + B_k - (2k-1)/2^{n+1}] = \mathbf{E}[L_{k-1} + p - (2k-1)/2^{n+1}] \\
&= \mathbf{E}\Big[\Big(1 - \frac{2}{2^n-(k-1)}\Big)L_{k-1} - \frac{1}{2^{n+1}}\Big] \\
&= \Big(1 - \frac{2}{2^n-(k-1)}\Big)\mathbf{E}[L_{k-1}] - \frac{1}{2^{n+1}} \\
&\geq \Big(1 - \frac{2}{2^n-(k-1)}\Big)\frac{(1-k)}{2^{n+1}} - \frac{1}{2^{n+1}} \geq \frac{-k}{2^{n+1}},
\end{aligned}$$

where the second last inequality is due to the induction hypothesis. On the other hand,

$$\begin{aligned}
\mathbf{E}[(L_k)^2] &= \mathbf{E}\Big[\Big(L_{k-1} + B_k - (2k-1)/2^{n+1}\Big)^2\Big] \\
&= \mathbf{E}\Big[p\Big(L_{k-1} + 1 - (2k-1)/2^{n+1}\Big)^2 + (1-p)\Big(L_{k-1} - (2k-1)/2^{n+1}\Big)^2\Big]. \quad (9)
\end{aligned}$$

By substituting $p = \frac{k-1}{2^n} - \frac{2L_{k-1}}{2^n-(k-1)}$ and using some simple algebraic manipulations,

$$\begin{aligned}
&p\Big(L_{k-1}+1-(2k-1)/2^{n+1}\Big)^2 + (1-p)\Big(L_{k-1}-(2k-1)/2^{n+1}\Big)^2 \\
&= \Big(1-\frac{4}{2n-k-1}\Big)(L_{k-1})^2 - \Big(\frac{1}{2^n}+\frac{2}{2^n-(k-1)}\Big)L_{k-1} + \frac{(2k-1)^2}{2^{2n+2}} + \frac{(2k-1)}{2^{n+1}} \\
&\le (L_{k-1})^2 - \Big(\frac{1}{2^n}+\frac{2}{2^n-(k-1)}\Big)L_{k-1} + \frac{3(2k-1)}{2^{n+2}}, \qquad (10)
\end{aligned}$$

where the last inequality is due to the fact that $k \le 2q \le 2^{n-4}$. Taking expectation of both sides of Eq. (10), and using the induction hypothesis yield

$$\mathbf{E}\Big[\Big(L_k\Big)^2\Big] \le \frac{2(k-1)^2}{2^{n+1}} + \Big(\frac{1}{2^n}+\frac{2}{2^n-(k-1)}\Big)\frac{k-1}{2^{n+1}} + \frac{3(2k-1)}{2^{n+2}} \le \frac{2k^2}{2^{n+1}},$$

where the last inequality is again due to the fact that $k \le 2q \le 2^{n-4}$. This concludes the proof. □

5 The Encrypted Davies-Meyer Construction

In this section we consider the PRF construction EDM that Cogliati and Seurin (CS) recently propose [10]. They show that EDM achieves $\frac{2n}{3}$-bit security and conjecture that it actually achieves n-bit security. Here we'll give a $\frac{3n}{4}$-bit security proof for EDM. We begin by describing the EDM construction.

Setup and results. The construction $\mathrm{EDM}[n] : (\mathrm{Perm}(n))^2 \times \{0,1\}^n \to \{0,1\}^n$ takes two secret permutations $\pi, \pi' \in \mathrm{Perm}(n)$ as its key, and outputs $\pi'(\pi(x) \oplus x)$ on input x. Theorem 2 below shows that $\mathsf{Adv}^{\mathrm{prf}}_{\mathrm{EDM}[n]}(A) \le \frac{7q^2}{2^{3n/2}}$, namely $\frac{3n}{4}$-bit security, whereas CS's result shows that $\mathsf{Adv}^{\mathrm{prf}}_{\mathrm{EDM}[n]}(A) \le \frac{5q^{3/2}}{2^n}$.

We note that a concurrent work by Mennink and Neves (MN) [21] shows that $\mathsf{Adv}^{\mathrm{prf}}_{\mathrm{EDM}[n]}(A) \le \frac{q}{2^n} + \frac{\binom{q}{t+1}}{2^{nt}}$ for any integer $t \ge 1$ and any $q \le 2^n/67t$. While MN's bound is quite better than ours, their work relies on Patarin's "mirror theory" [26]. Here, our goal is to give a much simpler proof and we leave it as an open question of whether our bound can be tightened without resorting to mirror theory. A graphical comparison of the three bounds is shown in Fig. 1.

Theorem 2. Let $n \ge 16$ be an integer. Then for any adversary A that makes at most q queries,

$$\mathsf{Adv}^{\mathrm{prf}}_{\mathrm{EDM}[n]}(A) \le \frac{7q^2}{2^{1.5n}}.$$

Proof. Without loss of generality, assume that $q \le 2^{n-4}$; otherwise the claimed bound is moot. Assume that the adversary is deterministic and never repeats a

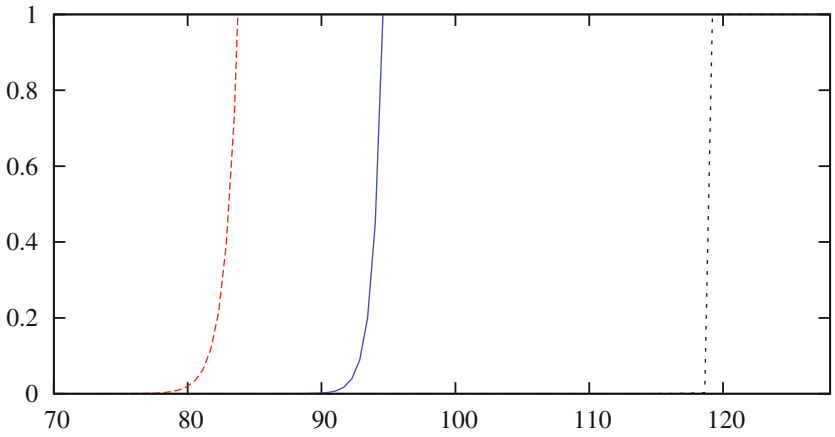

Fig. 1. Comparison among CS's bound (left), ours (middle), and MN's (right) for $n = 128$. The x-axis gives the log (base 2) of q, and the y-axis gives the security bounds. For MN's bound, we use $t = 9$ as suggested by MN.

past query. For convenience of analysis, instead of working directly with the real system (the one implementing EDM), we will "normalize" it to ensure that it has nice behaviors even if the past answers are bad.

Specifically, let $\mathbf{S}_0$ be the ideal system (the one implementing a uniform random function), and $\mathbf{S}_2$ be the real system. We will construct a system $\mathbf{S}_1$ that is the "normalized" version of $\mathbf{S}_2$ as follows. The system $\mathbf{S}_1$ keeps a secret boolean bad that is initially set to false. Initially, it implements $\mathbf{S}_2$, but if among the past queries, there are 4 answers that are the same, then it sets bad to true. Once bad is set, $\mathbf{S}_1$ instead implements $\mathbf{S}_0$. We now show that the advantage $\mathsf{Adv}^{\mathrm{prf}}_{\mathrm{EDM}}(A)$ can be bounded via the statistical distance between $\mathsf{p}_{\mathbf{S}_0,A}(\cdot)$ and $\mathsf{p}_{\mathbf{S}_1,A}(\cdot)$, and then bound the latter via the chi-squared method. First, recall that $\mathsf{Adv}^{\mathrm{prf}}_{\mathrm{EDM}}(A)$ is at most

$$\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_2,A}(\cdot)\| \le \|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\| + \|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_2,A}(\cdot)\|. \tag{11}$$

Let X and X' be the random variables for the q-answers on $\mathbf{S}_0$ and $\mathbf{S}_1$ respectively. Let Γ_{bad} be the subset of $(\{0,1\}^n)^q$ such that for any $\boldsymbol{Z} \in \Gamma_{\mathrm{bad}}$, there are 4 components of $\boldsymbol{Z}$ that are the same. Then $\mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}) = \mathsf{p}_{\mathbf{S}_2,A}(\boldsymbol{Z})$ for every $\boldsymbol{Z} \in (\{0,1\}^n)^q \backslash \Gamma_{\mathrm{bad}}$, and thus

$$\begin{aligned}
\|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_2,A}(\cdot)\| &= \sum_{\boldsymbol{Z} \in (\{0,1\}^n)^q} \max\{0, \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}) - \mathsf{p}_{\mathbf{S}_2,A}(\boldsymbol{Z})\} \\
&= \sum_{\boldsymbol{Z} \in \Gamma_{\mathrm{bad}}} \max\{0, \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}) - \mathsf{p}_{\mathbf{S}_2,A}(\boldsymbol{Z})\} \\
&\le \sum_{\boldsymbol{Z} \in \Gamma_{\mathrm{bad}}} \mathsf{p}_{\mathbf{S}_1,A}(\boldsymbol{Z}) = \Pr[X' \in \Gamma_{\mathrm{bad}}].
\end{aligned}$$

On the other hand, note that $\Pr[X' \in \Gamma_{\mathrm{bad}}] - \Pr[X \in \Gamma_{\mathrm{bad}}]$ can't exceed the statistical distance between X' and X, which is $\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\|$. Hence

$$\begin{aligned}\|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_2,A}(\cdot)\| &\le \Pr[X' \in \Gamma_{\mathrm{bad}}] \\ &\le \Pr[X \in \Gamma_{\mathrm{bad}}] + \|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\|. \end{aligned} \tag{12}$$

From Eqs. (11) and (12),

$$\begin{aligned}\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_2,A}(\cdot)\| &\le 2\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\| + \Pr[X \in \Gamma_{\mathrm{bad}}] \\ &\le 2\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\| + \frac{q^4}{2^{3n}} \\ &\le 2\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\| + \frac{q^2}{2^{1.5n}}.\end{aligned}$$

Hence what's left is to bound $\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\|$. Fix $i \le q$ and $\boldsymbol{Z}_{i-1} = (z_1, \ldots, z_{i-1}) \in (\{0,1\}^n)^{i-1}$. Recall that

$$\chi^2(\boldsymbol{Z}_{i-1}) = \sum_{z_i \in \{0,1\}^n} \frac{(\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) - 1/2^n)^2}{1/2^n}.$$

We claim that if $z_i \in \{z_1, \ldots, z_{i-1}\}$ then

$$\frac{1}{2^n} - \frac{4i}{2^{2n}} \le \mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) \le \frac{1}{2^n} + \frac{2i}{2^{2n}}, \tag{13}$$

and if $z_i \notin \{z_1, \ldots, z_{i-1}\}$

$$\frac{1}{2^n} - \frac{2i^2}{2^{3n}} \le \mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) \le \frac{1}{2^n} + \frac{5i^2}{2^{3n}}. \tag{14}$$

Consequently,

$$\chi^2(\boldsymbol{Z}_{i-1}) \le (i-1)\frac{16i^2}{2^{3n}} + (2^n - i + 1)\frac{25i^4}{2^{5n}} \le \frac{18i^3}{2^{3n}}.$$

Hence from Lemma 3, if one samples vectors $\boldsymbol{X}_{i-1}$ according to interaction with system $\mathbf{S}_1$,

$$(\|\mathsf{p}_{\mathbf{S}_0,A}(\cdot) - \mathsf{p}_{\mathbf{S}_1,A}(\cdot)\|)^2 \le \frac{1}{2}\sum_{i=1}^{q} \mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})] \le \frac{1}{2}\sum_{i=1}^{q} \frac{18i^3}{2^{3n}} \le \frac{9q^4}{2^{3n}}.$$

We now justify the two claims above, namely Eqs. (13) and (14). Note that if there are 4 components of $\boldsymbol{Z}_{i-1}$ that are the same, then the claims are obviously true, as $\mathsf{p}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) = 1/2^n$. Suppose that there are no 4 components of $\boldsymbol{Z}_{i-1}$ that are the same. Let $(m_1, \ldots, m_i)$ be the queries that are uniquely determined from $\boldsymbol{Z}_{i-1}$. Let $v_j = \pi(m_j) \oplus m_j$ for every $j \le i$.

We first justify Eq. (13), namely $z_i \in \{z_1, \ldots, z_{i-1}\}$. First consider the upper bound. Let S be the subset of $\{1, \ldots, i-1\}$ such that $z_i = z_j$, for every $j \in S$. Then $0 < |S| \le 3$. Let ℓ be an arbitrary element of S. Note that $\mathbf{S}_1$ outputs z_i

on query m_i if and only if $\pi(m_i) = v_\ell \oplus m_i$. For each fixed choice of $v_1, \ldots, v_{i-1}$, the conditional probability that $\pi(m_i) = v_\ell \oplus m_i$, given $\pi(m_j) = v_j \oplus m_j$ for every $j \le i-1$, is either 0 or $1/(2^n - i)$. Hence

$$\mathsf{ps}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) \le \frac{1}{2^n - i} \le \frac{1}{2^n} + \frac{2i}{2^{2n}},$$

where the last inequality is due to the hypothesis that $i \le q \le 2^{n-4}$. Next, consider the lower bound in Eq. (13). For each fixed choice of v_j, with $j \in \{1, \ldots, i-1\} \backslash S$, there are at least $2^n - 4i$ choices for v_ℓ, out of at most 2^n possible choices, such that $v_\ell \oplus m_k \neq v_j \oplus m_j$, for every $j \in \{1, \ldots, i-1\} \backslash S$ and every $k \in S \cup \{i\}$. For each such tuple $(v_1, \ldots, v_{i-1})$, the conditional probability that $\pi(m_i) = v_\ell \oplus m_i$, given $\pi(m_j) = v_j \oplus m_j$ for every $j \le i-1$, is exactly $1/(2^n - i)$. Hence

$$\mathsf{ps}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) \ge \frac{2^n - 4i}{2^n(2^n - i)} \ge \frac{1}{2^n} - \frac{4i}{2^{2n}},$$

where the last inequality is due to the hypothesis that $i \le q \le 2^{n-4}$.

We now justify Eq. (14), namely $z_i \notin \{z_1, \ldots, z_{i-1}\}$. First consider the lower bound. Let r be the number of elements in $\{z_1, \ldots, z_{i-1}\}$, and thus $r \le i-1$. The system $\mathbf{S}_1$ will give an answer not in $\{z_1, \ldots, z_{i-1}\}$ if and only if $v_i \notin \{v_1, \ldots, v_{i-1}\}$. Note that for each $x, x' \in \{0,1\}^n \backslash \{z_1, \ldots, z_{i-1}\}$, we have $\mathsf{ps}_{\mathbf{S}_1,A}(x \mid \boldsymbol{Z}_{i-1}) = \mathsf{ps}_{\mathbf{S}_1,A}(x' \mid \boldsymbol{Z}_{i-1})$, since as long as $v_i \notin \{v_1, \ldots, v_{i-1}\}$, $\pi'(v_i)$ is equally likely to take any value in $\{0,1\}^n \backslash \{z_1, \ldots, z_{i-1}\}$. Hence

$$\begin{aligned}
\mathsf{ps}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) &= \frac{1}{2^n - r}\Big(1 - \sum_{x \in \{z_1,\ldots,z_{i-1}\}} \mathsf{ps}_{\mathbf{S}_1,A}(x \mid \boldsymbol{Z}_{i-1})\Big) \\
&\ge \frac{1}{2^n - r}\Big(1 - \sum_{x \in \{z_1,\ldots,z_{i-1}\}} \frac{1}{2^n}(1 + 2i/2^n)\Big) \\
&\ge \frac{1}{2^n - r}\Big(1 - \frac{r}{2^n}(1 + 2i/2^n)\Big) \\
&\ge \frac{1}{2^n} - \frac{2ri}{2^{2n}(2^n - r)} \ge \frac{1}{2^n} - \frac{2i^2}{2^{3n}}.
\end{aligned}$$

For the upper bound of Eq. (14),

$$\begin{aligned}
\mathsf{ps}_{\mathbf{S}_1,A}(z_i \mid \boldsymbol{Z}_{i-1}) &= \frac{1}{2^n - r}\Big(1 - \sum_{x \in \{z_1,\ldots,z_{i-1}\}} \mathsf{ps}_{\mathbf{S}_1,A}(x \mid \boldsymbol{Z}_{i-1})\Big) \\
&\le \frac{1}{2^n - r}\Big(1 - \sum_{x \in \{z_1,\ldots,z_{i-1}\}} \frac{1}{2^n}(1 - 4i/2^n)\Big) \\
&\le \frac{1}{2^n - r}\Big(1 - \frac{r}{2^n}(1 - 4i/2^n)\Big) \\
&\le \frac{1}{2^n} + \frac{4ri}{2^{2n}(2^n - r)} \le \frac{1}{2^n} + \frac{5i^2}{2^{3n}} \quad .
\end{aligned}$$

This concludes the proof. □

6 The Swap-or-Not Construction

As a final application of our framework, we prove a tighter bound on the security of the swap-or-not construction by Hoang, Morris, and Rogaway [13] using the chi-squared method. We start by reviewing the construction, before turning to its analysis.

The swap-or-not construction. Let $r \geq 1$ be a round parameter. Let $\mathbb{G}$ be a finite abelian group, for which we use additive notation to denote the associated operation. Then, the *swap-or-not construction* SN_r uses r functions $f_1, \ldots, f_r : \mathbb{G} \to \{0,1\}$ (to be chosen independently and uniformly at random in the proof), and additionally uses r rounds keys $K = (K_1, \ldots, K_r) \in \mathbb{G}$. Then, on input $X \in \mathbb{G}$, it computes states $X_0, X_1, \ldots, X_r \in \mathbb{G}$, where $X_0 = X$, and for $i \in \{1, \ldots, r\}$, let $V_i = \max\{X_{i-1}, K_i - X_{i-1}\}$,[7]

$$X_i = \begin{cases} X_{i-1} & \text{if } f_i(V_i) = 0\,, \\ K_i - X_{i-1} & \text{else.} \end{cases} \,. \tag{15}$$

Finally, it outputs X_r. The corresponding inversion operation occurs by taking these steps backwards. We denote the resulting construction as $\mathrm{SN}_r[\mathbb{G}]$.

Security notions. For a block cipher $E : \mathcal{K} \times \mathcal{M} \to \mathcal{M}$ and an adversary A, the CCA advantage $\mathsf{Adv}^{\mathrm{cca}}_E(A)$ of A against E is defined as

$$\Pr[K \leftarrow\!\!\$\, \mathcal{K}; A^{E_K(\cdot), E_K^{-1}(\cdot)} \Rightarrow 1] - \Pr[\pi \leftarrow\!\!\$\, \mathrm{Perm}(\mathcal{M}); A^{\pi(\cdot), \pi^{-1}(\cdot)} \Rightarrow 1],$$

where $\mathrm{Perm}(\mathcal{M})$ is the set of all permutations on $\mathcal{M}$. We emphasize that here $\mathcal{M}$ is an arbitrary set. If the adversary only queries its first oracle, and makes only *non-adaptive* queries, then we write $\mathsf{Adv}^{\mathrm{ncpa}}_E(A)$ instead. We write $\mathsf{Adv}^{\mathrm{cca}}_E(q)$ and $\mathsf{Adv}^{\mathrm{ncpa}}_E(q)$ to denote the CCA and NCPA advantage of the best adversaries of q queries against E, respectively.

If we have two block ciphers F and G on the same message space that are just NCPA-secure, one can have a CCA-secure block cipher E by composing $E = F \circ G^{-1}$, meaning that $E_{K,K'}(x) = G^{-1}_{K'}(F_K(x))$. The following well-known theorem by Maurer, Pietrzak, and Renner [20] bounds the CCA security of E based on the NCPA security of F and G.

Lemma 4 ([20]). Let F and G be block ciphers on the same message space, and let $E = F \circ G^{-1}$. Then for any q,

$$\mathsf{Adv}^{\mathrm{cca}}_E(q) \leq \mathsf{Adv}^{\mathrm{ncpa}}_F(q) + \mathsf{Adv}^{\mathrm{ncpa}}_G(q). \qquad \square$$

[7] Here, max is with respect to some encoding. The key point is that $K_i - (K_i - X) = X$, so this will reach a unique representative for this pair of elements of $\mathbb{G}$.

We note that Lemma 4 only holds in the information-theoretic setting where one consider the best possible, computationally unbounded adversaries. Pietrzak shows that this lemma does not hold in the computational setting [27].

NCPA security of Swap-or-Not. Following the route in the analysis of [13], we'll first consider the NCPA security of Swap-or-Not, and then use Lemma 4 to amplify it to CCA security.

Lemma 5. For any adversary A that makes at most q queries and an abelian group $\mathbb{G}$ of N elements,

$$\mathsf{Adv}^{\mathrm{ncpa}}_{\mathrm{SN}_r[\mathbb{G}]}(A) \le \frac{N}{\sqrt{r+1}}\Big(\frac{N+q}{2N}\Big)^{(r+1)/2}.$$

Proof. We assume without loss of generality that A is deterministic, and doesn't make redundant queries. The adversary A interacts with the construction $\mathrm{SN}_r[\mathbb{G}]$ with r secret and randomly chosen functions $f_1, \ldots, f_r : \mathbb{G} \to \{0,1\}$, and r keys $K = (K_1, \ldots, K_r)$. We denote by $\mathbf{S}_1$ the system resulting from $\mathrm{SN}_r[\mathbb{G}]$ and by $\mathbf{S}_0$ the system resulting from interacting with the random permutation π. We will bound

$$\mathsf{Adv}^{\mathrm{ncpa}}_{\mathrm{SN}_r[\mathbb{G}]}(A) \le \|\mathsf{p}_{\mathbf{S}_1,A}(\cdot) - \mathsf{p}_{\mathbf{S}_0,A}(\cdot)\|.$$

For each $i \in \{0, 1, \ldots, q\}$, we define $\boldsymbol{X}_i$ to be the vector of outputs from the first i queries of A to $\mathbf{S}_1$. Let $m_i = N - i + 1$. We will use the following lemma from [13] to bound $\mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})]$.

Lemma 6 ([13]). For any NCPA adversary A making q queries and for any $i \le q$,

$$\mathbf{E}\Big(\sum_{x \in \mathbb{G}\setminus\{x_1,\ldots,x_{i-1}\}} (\mathsf{p}_{\mathbf{S}_1,A}(x \mid \boldsymbol{X}_{i-1}) - 1/m_i)^2\Big) \le \left(\frac{N+i}{2N}\right)^r,$$

where the expectation is taken over a vector $\boldsymbol{X}_{i-1} = (x_1, \ldots, x_{i-1})$ sampled according to interaction with $\mathbf{S}_1$. □

Fix some $\boldsymbol{Z}_{i-1} = (z_1, \ldots, z_{i-1})$ such that $\mathsf{p}_{\mathbf{S}_0}(\boldsymbol{Z}_{i-1}) > 0$. Notice that the i-th output of $\mathbf{S}_0$, given that the first $i-1$ outputs are $\boldsymbol{Z}_{i-1}$, is uniformly distributed over $\mathbb{G}\setminus\{z_1, \ldots, z_{i-1}\}$. In other words, for any $x \in \mathbb{G}\setminus\{z_1, \ldots, z_{i-1}\}$.

$$\mathsf{p}_{\mathbf{S}_0,A}(x \mid \boldsymbol{Z}_{i-1}) = 1/m_i.$$

Hence, from Lemma 6,

$$\begin{aligned}\mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})] &= \mathbf{E}\Big(\sum_{x \in \mathbb{G}\setminus\{x_1,\ldots,x_{i-1}\}} m_i \cdot (\mathsf{p}_{\mathbf{S}_1,A}(x \mid \boldsymbol{X}_{i-1}) - 1/m_i)^2\Big) \\ &\le m_i \left(\frac{N+i}{2N}\right)^r \le N\left(\frac{N+i}{2N}\right)^r. \end{aligned} \tag{16}$$

Using Lemma 3, we obtain,

$$
\begin{aligned}
(\|\mathsf{ps}_{0,A}(\cdot) - \mathsf{ps}_{1,A}(\cdot)\|)^2 &\le \frac{1}{2} \cdot \sum_{i=1}^{q} \mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})] \\
&\le \frac{1}{2} \sum_{i=1}^{q} N \left(\frac{N+i}{2N}\right)^r \\
&\le N^2 \int_0^{q/2N} \left(\frac{1}{2} + x\right)^r dx \le \frac{N^2}{r+1} \left(\frac{N+q}{2N}\right)^{r+1}.
\end{aligned}
$$

CCA security of Swap-or-Not. Note that the inverse of $\mathrm{SN}_r[G]$ is also another $\mathrm{SN}_r[\mathbb{G}]$ (but the round functions and round-keys are bottom up). Hence from Lemmas 4 and 5, we conclude that

Theorem 3. For any $q, r \in \mathbb{N}$ and any abelian group $\mathbb{G}$ of N elements,

$$
\mathsf{Adv}^{\mathrm{cca}}_{\mathrm{SN}_{2r}[\mathbb{G}]}(q) \le \frac{2N}{\sqrt{r+1}} \left(\frac{N+q}{2N}\right)^{(r+1)/2}.
$$

□

Note that in Theorem 3, the number of rounds in the Swap-or-Not shuffle is $2r$. The original bound in [13] is

$$
\mathsf{Adv}^{\mathrm{cca}}_{\mathrm{SN}_{2r}[\mathbb{G}]}(q) \le \frac{4N^{3/2}}{r+2} \left(\frac{N+q}{2N}\right)^{r/2+1}.
$$

Typically one uses $r = \Theta(\log(N))$, and thus our result improves the original analysis by a factor of $\Theta(\sqrt{N/\log(N)})$. We note that our result is probably not tight, meaning that it might be possible to improve the security of Swap-or-Not further.

Acknowledgments. We thank the anonymous reviewers for insightful comments. Wei Dai's work was done in part while at UCSB and he was supported in part by NSF grant CNS-1526801. Viet Tung Hoang was supported in part by the First Year Assistant Professor Award of Florida State University. Stefano Tessaro was supported in part by NSF grants CNS-1423566, CNS-1528178, CNS-1553758 (CAREER), and IIS-152804, and by the Glen and Susanne Culler Chair.

A Another Variant of the Xor of Two Permutations

Let $\mathrm{Perm}(n)$ be the set of permutations $\pi : \{0,1\}^n \to \{0,1\}^n$. In Sect. 4 we show that $\mathsf{XOR}[n]$ is a goof PRF. In this section, we consider the related construction $\mathsf{XOR2}[n] : (\mathrm{Perm}(n))^2 \times \{0,1\}^n \to \{0,1\}^n$ that takes $\pi, \pi' \in \mathrm{Perm}(n)$ as its key, and outputs $\pi(x) \oplus \pi'(x)$ on input x. Theorem 4 below gives a bound on the PRF security of $\mathsf{XOR2}[n]$.

Theorem 4. Fix an integer $n \geq 4$. For any adversary A that makes $q \leq 2^{n-4}$ queries we have

$$\mathsf{Adv}^{\mathrm{prf}}_{\mathsf{XOR2}[n]}(A) \leq \frac{q^{1.5}}{2^{1.5n}}.$$

Proof. Let $\mathbf{S}_1$ be the real system, and let $\mathbf{S}_0$ be the ideal system. We shall use the chi-squared method. Let $\boldsymbol{X} = (X_1, \ldots, X_q)$ be the random variable for the q answers in $\mathbf{S}_1$, and let $\boldsymbol{X}_i = (X_1, \ldots, X_i)$ for every $i \leq q$. Fix $i \leq q$ and fix $x \in \{0,1\}^n$. Let $Y_{i,x}$ be the following random variable. If $\boldsymbol{X}_{i-1}$ takes values $(z_1, \ldots, z_{i-1})$ then $Y_{i,x}$ takes the value $\mathsf{p}_{\mathbf{S}_1,A}(x \mid z_1, \ldots, z_{i-1})$. Recall that

$$\begin{aligned}\chi^2(\boldsymbol{X}_{i-1}) &= \sum_{x\in\{0,1\}^n} \frac{(Y_{i,x} - 1/2^n)^2}{1/2^n} \\ &= \sum_{x\in\{0,1\}^n} 2^n \cdot (Y_{i,x} - 1/2^n)^2. \end{aligned} \tag{17}$$

We now expand $Y_{i,x}$ into a more expressive and convenient formula to work with. Let π and π' be the secret permutations of $\mathsf{XOR2}[n]$. Let $m_1, \ldots, m_i$ be the first i queries of the adversary. Let $V_k = \pi(m_k)$ and $U_k = \pi'(m_k)$ for every $k \leq i$. Regardless of how the adversary chooses its queries, $V_1, V_2, \ldots$ are simply random variables sampled uniformly without replacement from $\{0,1\}^n$. Likewise, $U_1, U_2, \ldots$ are sampled uniformly without replacement from $\{0,1\}^n$ independent of $V_1, V_2, \ldots$. Let $S = \{V_1, \ldots, V_{i-1}\}$ and $S' = \{U_1, \ldots, U_{i-1}\}$. Let $D_{i,x}$ be the number of strings u such that $u \in S$ and $u \oplus x \in S'$. If $\pi(m_i) \oplus \pi'(m_i) = x$, there are exactly $2^n - 2(i-1) + D_{i,x}$ choices for the pair $(\pi(m_i), \pi'(m_i))$:

- First, $\pi(m_i)$ must take value in $\{0,1\}^n \backslash (S \cup S^*)$, where $S^* = \{u \oplus x \mid u \in S'\}$. There are exactly $2^n - |S \cup S^*| = 2^n - |S| - |S^*| + |S \cap S^*| = 2^n - 2(i-1) + D_{i,x}$ choices for $\pi(m_i)$.
- Once $\pi(m_i)$ is fixed, the value of $\pi'(m_i)$ is determined.

Hence

$$Y_{i,x} = \frac{2^n - 2(i-1) + D_{i,x}}{(2^n - i + 1)^2},$$

and thus

$$(Y_{i,x} - 1/2^n)^2 = \frac{(D_{i,x} - (i-1)^2/2^n)^2}{(2^n - 2i + 1)^4} \leq \frac{2(D_{i,x} - (i-1)^2/2^n)^2}{2^{4n}},$$

where the last inequality is due to the fact that $i \leq q \leq 2^{n-4}$. From Eq. (17),

$$\begin{aligned}\mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})] &\leq \sum_{x\in\{0,1\}^n} 2^n \cdot \mathbf{E}\Big[(Y_{i,x} - 1/2^n)^2\Big] \\ &\leq \frac{2}{2^{3n}} \sum_{x\in\{0,1\}^n} \mathbf{E}\Big[\Big(D_{i,x} - \frac{(i-1)^2}{2^n}\Big)^2\Big].\end{aligned}$$

We will show that for any $x \in \{0,1\}^n$,

$$\mathbf{E}\Big[\Big(D_{i,x} - \frac{(i-1)^2}{2^n}\Big)^2\Big] \le \frac{(i-1)^2}{2^n}, \tag{18}$$

and thus

$$\mathbf{E}[\chi^2(\boldsymbol{X}_{i-1})] \le \frac{2(i-1)^2}{2^{3n}} \ .$$

Summing up, from Lemma 3,

$$\begin{aligned}(\|\mathsf{ps}_{0,A}(\cdot) - \mathsf{ps}_{1,A}(\cdot)\|)^2 &\le \frac{1}{2}\sum_{i=1}^{q} \mathbf{E}[\chi^2(X_{i-1})] \\ &\le \sum_{i=1}^{q} \frac{(i-1)^2}{2^{3n}} \le \frac{q^3}{2^{3n}}.\end{aligned}$$

We now justify (18). Fix $x \in \{0,1\}^n$. For each $j \le i-1$, let B_j be the Bernoulli random variable such that $B_j = 1$ if and only if $V_j \oplus x \in S'$. Then $D_{i,x} = B_1 + \cdots + B_{i-1}$. Moreover, for each $j \le i-1$, we have $\mathbf{E}[B_j] = (i-1)/2^n$, because marginally, V_j is uniformly distributed in $\{0,1\}^n$ independent of $U_1, \ldots, U_{i-1}$. Then

$$\mathbf{E}[D_{i,x}] = \sum_{j=1}^{i-1} \mathbf{E}[B_j] = \frac{(i-1)^2}{2^n}.$$

Note that

$$\begin{aligned}\mathbf{E}\Big[\Big(D_{i,x} - \frac{(i-1)^2}{2^n}\Big)^2\Big] &= \mathbf{Var}(D_{i,x}) = \mathbf{E}[(D_{i,x})^2] - (\mathbf{E}[D_{i,x}])^2 \\ &= \mathbf{E}[(D_{i,x})^2] - \frac{(i-1)^4}{2^{2n}}.\end{aligned} \tag{19}$$

On the other hand,

$$\begin{aligned}(D_{i,x})^2 &= \Big(\sum_{j=1}^{i-1} B_j\Big)^2 \\ &= (B_1^2 + \cdots + B_{i-1}^2) + 2\sum_{1 \le j < k \le i-1} B_j B_k \\ &= (B_1 + \cdots + B_{i-1}) + 2\sum_{1 \le j < k \le i-1} B_j B_k,\end{aligned}$$

where the last equality is due to the fact that $R^2 = R$ for any Bernoulli random variable R. Taking expectation of both sides gives us

$$\mathbf{E}[(D_{i,x})^2] = \frac{(i-1)^2}{2^n} + 2\sum_{1 \le j < k \le i-1} \mathbf{E}[B_j B_k].$$

We claim that for any $1 \le j < k \le i$, we have

$$\mathbf{E}[B_j B_k] = \frac{(i-1)(i-2)}{2^n(2^n-1)} \tag{20}$$

and thus

$$\mathbf{E}[(D_{i,x})^2] = \frac{(i-1)^2}{2^n} + \frac{(i-1)^2(i-2)^2}{2^n(2^n-1)}.$$

Combing this with (19) we have

$$\begin{aligned}\mathbf{E}\Big[\Big(D_{i,x} - \frac{(i-1)^2}{2^n}\Big)^2\Big] &= \frac{(i-1)^2}{2^n} + \frac{(i-1)^2(i-2)^2}{2^n(2^n-1)} - \frac{(i-1)^4}{2^{2n}} \\ &\le \frac{(i-1)^2}{2^n}.\end{aligned}$$

What remains is to justify (20). Note that given S' and V_j, we have $V_k \oplus x \in S'$ with condition probability $(i-2)/(2^n-1)$ if $V_j \oplus x \in S'$, and with conditional probability $(i-1)/(2^n-1)$ otherwise. That is, given B_j, the random variable B_k takes value 1 with conditional probability $(i-1-B_j)/(2^n-1)$. Hence

$$\begin{aligned}\mathbf{E}[B_j B_k] = \mathbf{E}\Big[B_j \frac{(i-1-B_j)}{2^n-1}\Big] &= \frac{(i-1)\cdot\mathbf{E}[B_j]}{2^n-1} - \frac{\mathbf{E}[B_j^2]}{2^n-1} \\ &= \frac{(i-1)\cdot\mathbf{E}[B_j]}{2^n-1} - \frac{\mathbf{E}[B_j]}{2^n-1} = \frac{(i-2)(i-1)}{2^n(2^n-1)}.\end{aligned}$$

This completes the proof. □

References

1. Bai, S., Langlois, A., Lepoint, T., Stehlé, D., Steinfeld, R.: Improved security proofs in lattice-based cryptography: using the Rényi divergence rather than the statistical distance. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015. LNCS, vol. 9452, pp. 3–24. Springer, Heidelberg (2015). doi:10.1007/978-3-662-48797-6_1
2. Bellare, M., Impagliazzo, R.: A tool for obtaining tighter security analyses of pseudorandom function based constructions, with applications to PRP to PRF conversion. Cryptology ePrint Archive, Report 1999/024 (1999). http://eprint.iacr.org/1999/024
3. Bellare, M., Ristenpart, T., Rogaway, P., Stegers, T.: Format-preserving encryption. In: Jacobson, M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 295–312. Springer, Heidelberg (2009). doi:10.1007/978-3-642-05445-7_19
4. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006). doi:10.1007/11761679_25
5. Bernstein, D.J.: How to stretch random functions: the security of protected counter sums. J. Cryptol. **12**(3), 185–192 (1999)
6. Chen, S., Steinberger, J.: Tight security bounds for key-alternating ciphers. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 327–350. Springer, Heidelberg (2014). doi:10.1007/978-3-642-55220-5_19

7. Chung, K.-M., Pass, R.: Tight parallel repetition theorems for public-coin arguments using KL-divergence. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015. LNCS, vol. 9015, pp. 229–246. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46497-7_9
8. Chung, K.-M., Vadhan, S.: Tight bounds for hashing block sources. In: Goel, A., Jansen, K., Rolim, J.D.P., Rubinfeld, R. (eds.) APPROX/RANDOM - 2008. LNCS, vol. 5171, pp. 357–370. Springer, Heidelberg (2008). doi:10.1007/978-3-540-85363-3_29
9. Cogliati, B., Lampe, R., Patarin, J.: The indistinguishability of the XOR of k permutations. In: Cid, C., Rechberger, C. (eds.) FSE 2014. LNCS, vol. 8540, pp. 285–302. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46706-0_15
10. Cogliati, B., Seurin, Y.: EWCDM: an efficient, beyond-birthday secure, nonce-misuse resistant MAC. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9814, pp. 121–149. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53018-4_5
11. Even, S., Mansour, Y.: A construction of a cipher from a single pseudorandom permutation. J. Cryptol. **10**(3), 151–162 (1997)
12. Hall, C., Wagner, D., Kelsey, J., Schneier, B.: Building PRFs from PRPs. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 370–389. Springer, Heidelberg (1998). doi:10.1007/BFb0055742
13. Hoang, V.T., Morris, B., Rogaway, P.: An enciphering scheme based on a card shuffle. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 1–13. Springer, Heidelberg (2012). doi:10.1007/978-3-642-32009-5_1
14. Hoang, V.T., Tessaro, S.: Key-alternating ciphers and key-length extension: exact bounds and multi-user security. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9814, pp. 3–32. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53018-4_1
15. Iwata, T.: New blockcipher modes of operation with beyond the birthday bound security. In: Robshaw, M. (ed.) FSE 2006. LNCS, vol. 4047, pp. 310–327. Springer, Heidelberg (2006). doi:10.1007/11799313_20
16. Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. SIAM J. Comput. **17**(2), 373–386 (1988)
17. Lucks, S.: The sum of PRPs is a secure PRF. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 470–484. Springer, Heidelberg (2000). doi:10.1007/3-540-45539-6_34
18. Maurer, U.M.: A simplified and generalized treatment of luby-rackoff pseudorandom permutation generators. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 239–255. Springer, Heidelberg (1993). doi:10.1007/3-540-47555-9_21
19. Maurer, U.: Indistinguishability of random systems. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 110–132. Springer, Heidelberg (2002). doi:10.1007/3-540-46035-7_8
20. Maurer, U., Pietrzak, K., Renner, R.: Indistinguishability amplification. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 130–149. Springer, Heidelberg (2007). doi:10.1007/978-3-540-74143-5_8
21. Mennink, B., Neves, S.: Encrypted Davies-Meyer and its dual: towards optimal security using mirror theory. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, pp. 556–583. Springer, Cham (2017)
22. Morris, B., Rogaway, P.: Sometimes-recurse shuffle. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 311–326. Springer, Heidelberg (2014). doi:10.1007/978-3-642-55220-5_18

23. Nandi, M.: A simple and unified method of proving indistinguishability. In: Barua, R., Lange, T. (eds.) INDOCRYPT 2006. LNCS, vol. 4329, pp. 317–334. Springer, Heidelberg (2006). doi:10.1007/11941378_23
24. Patarin, J.: A proof of security in $O(2^n)$ for the Xor of two random permutations. In: Safavi-Naini, R. (ed.) ICITS 2008. LNCS, vol. 5155, pp. 232–248. Springer, Heidelberg (2008). doi:10.1007/978-3-540-85093-9_22
25. Patarin, J.: The "Coefficients H" technique. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. LNCS, vol. 5381, pp. 328–345. Springer, Heidelberg (2009). doi:10.1007/978-3-642-04159-4_21
26. Patarin, J.: Introduction to mirror theory: analysis of systems of linear equalities and linear non equalities for cryptography. Cryptology ePrint Archive, Report 2010/287 (2010). http://eprint.iacr.org/2010/287
27. Pietrzak, K.: Composition does not imply adaptive security. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 55–65. Springer, Heidelberg (2005). doi:10.1007/11535218_4
28. Raz, R.: A parallel repetition theorem. SIAM J. Comput. **27**(3), 763–803 (1998)
29. Renner, R.: On the variational distance of independently repeated experiments. CoRR, abs/cs/0509013 (2005)
30. Ristenpart, T., Yilek, S.: The mix-and-cut shuffle: small-domain encryption secure against N queries. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8042, pp. 392–409. Springer, Heidelberg (2013). doi:10.1007/978-3-642-40041-4_22
31. Steinberger, J.: Improved security bounds for key-alternating ciphers via hellinger distance. Cryptology ePrint Archive, Report 2012/481 (2012). http://eprint.iacr.org/2012/481