The t-wise Independence of Substitution-Permutation Networks

Tianren Liu[1], Stefano Tessaro[1], and Vinod Vaikuntanathan[2]

[1] University of Washington, Seattle, WA, USA
tianrenl@uw.edu,tessaro@cs.washington.edu
[2] MIT, Cambridge, MA, USA vinodv@mit.edu

Abstract. Block ciphers such as the Advanced Encryption Standard (Rijndael) are used extensively in practice, yet our understanding of their security continues to be highly incomplete. This paper promotes and continues a research program aimed at *proving* the security of block ciphers against important and well-studied classes of attacks. In particular, we initiate the study of (almost) t-wise independence of concrete block-cipher construction paradigms such as substitution-permutation networks and key-alternating ciphers. Sufficiently strong (almost) pairwise independence already suffices to resist (truncated) differential attacks and linear cryptanalysis, and hence this is a relevant and meaningful target. Our results are two-fold.

Our first result concerns substitution-permutation networks (SPNs) that model ciphers such as AES. We prove the almost pairwise-independence of an SPN instantiated with concrete S-boxes together with an appropriate linear mixing layer, given sufficiently many rounds and independent sub-keys. Our proof relies on a *characterization* of S-box computation on input differences in terms of sampling output differences from certain subspaces, and a new randomness extraction lemma (which we prove with Fourier-analytic techniques) that establishes when such sampling yields uniformity. We use our techniques in particular to prove almost pairwise-independence for sufficiently many rounds of both the AES block cipher (which uses a variant of the patched inverse function $x \mapsto x^{-1}$ as the S-box) and the MiMC block cipher (which uses the cubing function $x \mapsto x^3$ as the S-box), assuming independent sub-keys.

Secondly, we show that instantiating a key-alternating cipher (which can be thought of as a degenerate case of SPNs) with most permutations gives us (almost) t-wise independence in $t + o(t)$ rounds. In order to do this, we use the probabilistic method to develop two new lemmas, an *independence-amplification lemma* and a *distance amplification lemma*, that allow us to reason about the evolution of key-alternating ciphers.

1 Introduction

Block ciphers are among the most fundamental building blocks in cryptography, and applications demand strong pseudorandomness properties from them. However, the simplicity of widely adopted designs, such as Substitution-Permutation

Networks (SPNs), which underlie AES, is inherently at odds with the reductionist approach of provable security, as there are no clear underlying hard mathematical problems upon which security can be based. Instead, the security validation of block ciphers has gone through cryptanalysis, and considered a number of different techniques, including *linear* [41] and *differential* [5] cryptanalysis, higher-order [36] and truncated [34] differential attacks, impossible differential attacks [33], algebraic attacks [25], integral cryptanalysis [35], biclique attacks [7], and so on.

Lacking full proofs of security, the next best thing is to prove that certain relevant *classes* of attacks *cannot* possibly succeed. The more "concrete" and less "asymptotic" such a proof is, the better, and the class of attacks should be as large as possible. The most successful such effort has developed provable bounds for linear and differential cryptanalysis, starting with the seminal work of Nyberg and Knudsen [46], and culminating with fairly precise estimates for concrete block ciphers like AES (see e.g. [29–32, 48, 49]).

t-wise independence. In this paper, we move one step forward and study the (almost) *t-wise independence* of concrete block ciphers – namely, for a block cipher $E : \{0,1\}^s \times \{0,1\}^n \to \{0,1\}^n$, we demand that for any distinct t inputs $x_1, \ldots, x_t$ and a random key S, the distribution of

$$E(S, x_1), \ldots, E(S, x_k)$$

is statistically close to that of t uniform, but distinct, n-bit strings.

This property is attractive for two reasons. First and foremost, it is potentially achievable unconditionally by a concrete design, as long as $s \geq t \cdot n$. For example, a variant of AES-128 with 11 independent round keys[3] can (potentially) be 11-wise independent. Second, t-wise independence already implies resilience against a large class of attacks that have been previously studied. Indeed, the case $t = 2$ (i.e., almost pairwise independence) already implies resilience to linear and differential cryptanalysis but also to truncated differential attacks and any other attack that exploits statistical deviations of pairs of outputs. Similarly, t-wise independence implies resilience to order $\log_2(t)$ differential attacks. One caveat with this view point is that actual cipher instances typically have fixed-length keys which do not grow with t – however, similar to prior works on analyzing simpler properties of block ciphers, and in particular expected differential probabilities, we promote the heuristic angle that properties which are true for independent keys (possibly, unconditionally) remain true (computationally) when these keys are derived via a suitable key-scheduling algorithms from a short, single key.[4]

We note that existing bounds on differential probabilities for ciphers such as AES *could* imply pairwise independence, if good enough, but unfortunately, the

[3] Such an "independence assumption" is common across block cipher analyses. For more, see Section 1.2.

[4] We note that the impact of key-schedules on cryptographic attacks is mostly not well understood.

current state of the art (cf. e.g. [49]) proves *upper* bounds of the order 2^{-111} for 128-bit outputs which does not imply anything about (almost) pairwise independence. Without a finer grained understanding of the difference distribution, this *could* well imply a large distance of pairs of outputs from the uniform distribution.

Scope: Substitution-Permutation Networks. Our focus in this paper is on concrete block cipher designs (which likely benefit from other security properties, such as resilience to algebraic attacks), and in particular *Substitution-Permutation Networks* (SPNs), a class for which AES is a special instance, and a generalization thereof called *Key-Alternating Ciphers* (KACs). SPNs alternate *computationally simple* rounds as follows, starting from the state being equal to the block cipher input:

1. A key-mixing step which consists of XORing the keys bit-wise with the current state;
2. A *local* non-linear step where each bit of the output depends only on a few bits of the input; Concretely, this proceeds by partitioning the n-bit state into k b-bit blocks, and applying a non-linear permutation $S : \{0,1\}^b \to \{0,1\}^b$ (a so-called "S-box") to each block in parallel;
3. A linear *mixing step* is then applied to the state.

We will refer to k as the *width* and to the important special case where $b = n$ (i.e., $k = 1$) as a *Key-Alternating Cipher* (or KAC, for short). (For this case, we can omit the mixing step without loss of generality.) Most modern ciphers are SPNs (or KACs). For example, AES uses an S-box obtained from the patched inverse $x \mapsto x^{2^b-2}$ and a mixing layer alternating two simple operations (ShiftRows and MixColumns). The MiMC cipher [1] is a KAC applying the permutation $x \mapsto x^3$ to its state.

A similar viewpoint to ours was already taken by Vaudenay's *decorrelation theory* [51], but we are unaware of any application of decorrelation to SPNs with concrete S-boxes. (In fact, this was left as an open problem.) Similarly, Hoory *et al.* [24] also suggested the use and analysis of t-wise independence, but the resulting constructions, while very elegant and simple, are far from existing practical designs, and better fit in the general *theoretical* pursuit of building t-wise independent permutations [2, 9, 28].

Our Program. This raises the following questions: If we take t-wise independence as our security goal, what are good choices for the non-linear (resp. linear) step? Which choices *provably* work and which do not? Again, we stress that our goal is to find concrete, fixed choices of these layers, without modeling the S-box as a random permutation oracle.

Our results come in two forms:

1. Results about concrete SPN instantiations of SPNs with S-boxes such as the patched inversion function, and where we prove *pairwise* independence of the resulting construction. In particular, one of our results applies to the round structure of AES, without any simplifications or idealized assumptions.

2. Existential results, which hold for most choices of P, where we prove almost t-wise independence for KACs with a number of rounds that grows with t.

Next, we provide a detailed overview of our results, and the underlying techniques. Then, we give an overview of the most relevant related work.

1.1 Our Results and Techniques

This section gives an overview of our results, and the underlying techniques.

Pairwise independence of SPNs. Our first result deals with SPNs of width k with a concrete S-box $S : \mathbb{F}_{2^b} \to \mathbb{F}_{2^b}$ (thus, $n = b \cdot k$ is the block size here). In particular we focus on the case where the S-box is $S(x) = x^{-1}$ (patched so that $0^{-1} = 0$), though the results extend to other S-boxes. Our main theorem here can be cast as follows.

> **Theorem (Informal).** For a suitably instantiated mixing layer,[5] and as long as $\frac{2k+8}{2^b} + \sqrt{k/2^b} < \frac{1}{2}$, the r-round SPN with S-box $S(x) = x^{-1}$ of width k is δ-close to pairwise independent for sufficiently large $r = r(\delta)$. In particular, if $\frac{2k+8}{2^b} + \sqrt{k/2^b} = C/2$, then $r = O(\frac{\log(1/\delta)}{\log(1/C)})$.

We briefly highlight the main ideas behind the proof and note that we will focus in particular on showing that a *three* round SPN is $O(\sqrt{k/2^b})$-close to pairwise independent – this result will rely on a new extraction lemma, which we explain below. We then resort to an amplification result by Maurer, Pietrzak, and Renner [42] to conclude that the $(r/3)$-fold sequential composition of the SPN is δ-close to pairwise independent, as desired.

Our analysis of the output distribution of a three-round SPN for any two distinct inputs $x \neq x'$ will take the standard (and essentially equivalent) approach of studying the distribution of the *difference* of the outputs of the two evaluation. To this end, we start with a (fixed) input difference $\Delta = x \oplus x' \neq 0^n$. Then, our first step is to show, using (mostly) algebraic properties of the field $\mathbb{F}_{2^b}$, that after ignoring some corner cases that happen with probability no more than $O(k/2^b)$, the input differences to the third round – denoted by $V_1, \ldots, V_k$ – satisfy jointly a very strong distributional property, namely:

> any subset of them of size $k' \leq k$ has (jointly) min-entropy at least $k'(b-1)$.

For this to true, we only need mild assumptions on the linear mixing layer. We merely require it to be described by a full-rank $k \times k$ matrix whose entries are all non-zero.

To understand the effect of the third round, at last, we resort to our extraction lemma – we want to show in particular that the distribution of the differences $Z_1, \ldots, Z_k$, which we obtain after applying the final round of S-boxes with input

[5] Our requirement is very mild, and is in particular implied by having maximum branch number, as it is in the case in many SPN analysis.

differences $V_1, \ldots, V_k$, is very close to uniform.[6] Imagine first that the differences Z_i are not sampled via the S-box, but rather each Z_i is sampled independently from the $(n-1)$-dimensional sub-space orthogonal to $\{0^b, V_i\}$. (We interpret the latter as a linear subspace of $\mathbb{F}_2^b$, and V_i as a vector in this space.) Our extraction lemma shows that in this case, the Z_i's are very close to uniform – the proof uses Fourier-analytic techniques.

Of course, the Z_i's are not sampled this way – by applying the S-boxes to inputs with differences V_i – yet, the key insight is that this is almost equivalent to our sub-space representation, in that by applying a lemma of Nyberg [45] we can show that there exist permutations π, π' such that $\pi'(Z_i)$ is $O(k/2^b)$ close to a random vector sampled orthogonal to $\{0^b, \pi(V_i)\}$.

We also give a proof of a weaker bound for a two-round SPN of order $\sqrt{2^{k-b}}$. This bound could be interesting in some parameter regimes.

The AES case. Unfortunately, we cannot apply the above theorem directly to the AES round structure or the AES parameters. First off, the AES S-box combines the inverse with a $\mathbb{F}_2$-affine function – it turns out this is not particularly difficult to handle (the affine function can be cast as part of the mixing). But we encounter other problems, in that the mixing layers does not satisfy the assumptions needed for the theorem to work, and the theorem does not apply when $k = 16$ and $b = 8$. Still, we can adapt our techniques to obtain a refined analysis which tells us that *six* AES rounds (with independent sub-keys) are ϵ-close to pairwise independent, for some $\epsilon < 1/2$. Then, using the MPR result in the iteration, we obtain the following result:

Theorem. $6r$-round AES is $2^{r-1}(0.472)^r$-close to pairwise independence.

The bound is likely far from tight, as we expect much better, but non-trivial further work seems required to obtain a substantial improvement. However, we do stress that barring the use of independent keys (which again, are common in analyses of expected differential probabilities for AES), this theorem applies to the *actual* AES structure.

Existential Results. All of the above results are about pairwise independence. It is interesting to extend them to t-wise independence for $t \geq 3$. While we leave this important question open for SPNs and concrete S-boxes, we investigate the general question whether (almost) t-wise independent constructions exist in the first place.

To this end, we employ the probabilistic method to show that there exist permutations to instantiate a $(t+1)$-round key-alternating cipher so that it is (almost) t-wise independent. We stress that while our probabilistic argument picks such permutations at random to show their existence, these permutations can then be *fixed*.

[6] The last mixing stage does not affect the argument – it will merely preserve uniformity by virtue of being a permutation.

Our probabilistic argument is quite involved and requires the study of martingale sequences and their concentration. Our result follows by showing two new lemmas, and employing a careful alternation between them. The first is an *independence amplification lemma* that shows how to take a KAC that is very close to t-wise independent and by adding an additional round, obtain a KAC that is somewhat close to a $t+1$-wise independent distribution. The second is a *distance amplification lemma* that shows how to get from a somewhat close to t-wise independent KAC to a very close to t-wise independent KAC, again by adding one round.

1.2 Perspectives and Open Problems

On Independent Keys and Other Such "Ideal" Assumptions. We remark that, to date, *all* analyses of block ciphers make ideal assumptions such as the independence of round keys and/or ideal components. For example, analyses of (iterated) Even-Mansour ciphers assume that both the construction and the adversary have *oracle access* to a random permutation P, and that P remains unqueried on an exponential number of points. This is a *highly idealized model*: a random permutation would take exponentially many bits to write down, and indeed, in the real world, P is instantiated with a concrete permutation. The proofs say nothing about what happens to the pseudorandomness of such a cipher when P is instantiated with any concrete permutation. And moreover, analyses of multi-round constructions *all* assume independent keys.

In contrast, our work continues a research program that aims to avoid such "oracle access" assumptions. This line of work, which has its roots in the work of Nyberg in the 1990s, treats the component permutations and mixing functions as concrete functions (indeed, ones that are used in block ciphers such as AES and MiMC). While proving computational pseudorandomness is way out of reach, this line of research aims to understand the security of these constructions against concrete practical attacks.

The "independent round keys" assumption is very common and rooted in the model of Markov Ciphers of Lai, Massey, and Murphy [37], and adopted by Nyberg [45] and follow-up works. The expectation is that t-wise independence becomes t-wise pseudorandomness with an appropriate instantiation of the key schedule; nevertheless, understanding the precise role of key schedules is an important open problem.

On Algebraic and Other Attacks. The research program we undertake is to study several classes of concrete, powerful, attacks against block ciphers. In particular, t-wise independence rules out an important attack vector, but the program does not stop at just t-wise independence. In particular, the two outstanding open problems that come of this work are (a) to prove t-wise independence of multi-round AES with independent round keys, for $t > 2$; and (b) to formalize and prove security against algebraic attacks. We view solving these problems as an important quest that will likely require importing analytic techniques from mathematics and TCS, as well as inventing new ones.

On Differential Attacks vs. Almost Pairwise Independence. We note that meaningful differential probabilities need to be *very* close to 2^{-n}, or else, they do not rule out distinguishers. For example, in the case of AES-128, a 2^{-127} bound on the expected differential probability (see Section 2 for the definition) does not rule out the first bit of the output being always the same as the first bit of the input. In this case, there is a distinguisher that always works!

We note that our analysis can make the statistical distance as small as we want with sufficiently many rounds, and in particular, make the differential probabilities arbitrarily close to the ideal 2^{-128}. We note that ours is the first such result; in particular, our result for AES is the first such optimal bound for the AES design. Showing a tighter tradeoff between the number of rounds and the statistical distance is an interesting open question. Showing a direct bound on the differential probability without going through statistical distance would be interesting as well.

1.3 Related Work

Coppersmith and Grossman [15] and Kaliski, Rivest and Sherman [26] analyzed the groups generated by transition functions of the DES block cipher. [10] show that the group generated by the round functions of a cipher similar to AES is the alternating group. On the other hand, [44] provide a cautionary tale where guarantees on the group generated by the round functions does not guarantee security.

Bounds on Linear and Differential Probabilities. There is an extensive body of literature on *provable* bounds for linear [41] and differential cryptanalysis [5] of block ciphers. We note that while sufficiently strong bounds on the differential probability – say $(1+\epsilon)2^{-n}$ for block size n and $\epsilon = o(1)$ - would imply almost pairwise independence, these works fall short of proving such strong guarantees.

Adopting the formal framework of Lai, Massey, and Murphy [37], Nyberg and Knudsen [46] prove bounds on the differential probability for Feistel ciphers as a function of the underlying non-linear function. Several works have been devoted to studying the differential properties of fixed functions to instantiate these results – relevant to this work, [45] is the first work to show properties of differentials of the inverse permutation $x \mapsto x^{-1}$ in a finite field (these were later revisited by Daemen and Rijmen [17]). We also refer to [6] for a comprehensive survey on the progress in designing non-linear functions suitable for cryptography.

Much effort has also been devoted to provable bounds on linear and differential probabilities for AES and (more abstractly) SPNs. Hong et al. [23] gave the first analysis of two-round SPNs where the mixing layer has optimal branch number. This result was further generalized to arbitrary branch number by Kang et al. [27]. Very concrete bounds for the specific case of AES were then given via refined methods in several works [29–32, 48, 49]. The best known result here shows that the maximum expected differential probability is at most 1.144×2^{-111} for four rounds of AES. Miles and Viola [43] also provide generic bounds (i.e.,

these bounds only depend on the S-box and the number of rounds) for linear and differential attacks against multi-round SPNs – however, the quality of their bounds decreases with a higher number of rounds.

Baignères and Vaudenay [4] proved optimal resilience to differential cryptanalysis whenever the S-boxes are chosen uniformly at random and secret (i.e., their description is part of the key). Later, Miles and Viola [43] improves this result (implicitly) by showing that SPNs with random S-boxes are effectively a pseudorandom function when the number of queries is smaller than the input size of the S-box.

Stronger Differentials. Strong notions of differential attacks have been proposed. For example, Lai [36] introduced the notion of *higher order differentials*, which consider the k-th derivative (as opposed to the simple derivative of a function), whereas Knudsen [34] introduced *truncated differentials*, which only consider a subset of the bits of the output. We note that security against k-th order differential cryptanalysis is implied by the k-wise independence, whereas pairwise independence implies resistance to truncated differential cryptanalysis. Another attack technique introduced by Knudsen is that of "impossible differential attacks" [33], which leverage differences which occur with probability 0 – once again, sufficiently strong pairwise independence implicitly guarantees that differences occur with sufficiently large probability.

Decorrelation theory. Vaudenay [51] takes a similar position to ours, proving properties of block cipher constructions on a bounded number of inputs, and inferring a number of properties from these statements. The work also naturally exploits a natural connection with t-wise independence, like ours. Interestingly, Vaudenay considers a number of different distance measures for the resulting distributions, and use their properties to derive a number of results. However, we are not aware of any use of decorrelation theory about the security of SPNs or KACs with concrete permutations. Still, it would be interesting to considering distance measures from decorrelation theory in the context of our paper to improve tightness.

Analyses with Public Ideal Permutations. A substantial body of works considers analyses in models where the rounds of a KAC are (public) random permutation $P : \{0,1\}^n \to \{0,1\}^n$ given to the adversary. In particular, since the adversary is query-bounded, she cannot obtain the entire truth table of P and therefore, this is an idealized model. (This model is effectively capturing *generic* attacks that treat these components as a black box.) Increasingly tighter bounds for security as a pseudorandom permutation have been developed by several works [8, 12, 22, 38, 50] which assume the permutations *and* the keys are independent. Other works consider identical permutations and/or identical keys [11, 52]. The model was also considered to prove the stronger version of *indifferentiability* for key-alternating ciphers (cf.]e.g. [3, 20, 21]).

The model was then adapted to SPNs by assuming that the individual S-boxes are public random permutations $\{0,1\}^b \to \{0,1\}^b$ [13, 14, 18, 19]. Crucially,

these results assume that the number of queries to the S-box is smaller than 2^b, which is rather unrealistic for small values of b (e.g., $b = 8$ as in AES).

2 Preliminaries

Notational Conventions. When n is a positive integer, let $[n]$ denote the set $\{1, 2, \ldots, n\}$. When p is a prime or prime power, let $\mathbb{F}_p$ denote the finite field of size p. The logarithm function log uses base 2 by default. Probability distributions are typically denoted by calligraphic letters, e.g., $\mathcal{D}$. Sampling an element from $\mathcal{D}$ is denoted by $d \leftarrow \mathcal{D}$. For any finite set S, sampling x uniformly from S is denoted by $x \leftarrow S$.

Definition 1 (Entropy). *For a distribution over domain Ω whose probability mass function is p.*

- *Its Shannon entropy is* $H(p) = -\sum_{x\in\Omega} p(x)\log(p(x))$.
- *Its Min-entropy is* $\mathrm{H}_\infty(p) = -\log\big(\max_{x\in\Omega} p(x)\big)$.
- *Its Rényi entropy of order 2, also known as the collision entropy, is* $\mathrm{H}_2(p) = -\log\big(\sum_{x\in\Omega} p^2(x)\big)$.

2.1 Almost t-wise Independent Permutations and Cryptanalysis

We review notions of almost t-wise independence, and state some connections with standard notions from the cryptanalytic literature.

Definition 2. *The statistical distance (or total variation distance) between two probability distributions p and q with domain Ω is* $\mathrm{d_{TV}}(p, q) := \frac{1}{2} \cdot \sum_{x\in\Omega} |p(x) - q(x)|$. *Moreover,* $\mathrm{d_{TV}}(p, q) := \sum_{x\in\Omega: p(x)>q(x)} p(x) - q(x)$.

For a two argument function $F : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^\ell$ we often write $F_K(x) = F(K, x)$, and refer to F as a *function family*. (Alternatively, we use the set notation $\mathcal{F} = \{F_K\}_{K\in\{0,1\}^m}$ whenever more convenient.) We will be considering mostly *permutation* families, where $\ell = n$, and F_K is one-to-one for each K.

Definition 3 (close to t-wise independence). *We say that a permutation family $F : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^n$ is* ϵ-close to t-wise independent *if for all* distinct $x_1, \ldots, x_t \in \{0,1\}^n$, *and a uniformly random m-bit string K, the distribution of $(F_K(x_1), \ldots, F_K(x_t))$ has statistical distance at most ϵ from that of t uniformly sampled* distinct *n-bit values (i.e., sampled* without *repetition).*

We will use the following amplification lemma, which is due to Maurer, Pietrzak, and Renner [42].

Lemma 1 (MPR Amplification Lemma). *Let F and G be ϵ- and δ-close to t-wise independent permutation families. Then, the permutation family $F \circ G$ such that $(F \circ G)_{K_1||K_2}(x) = F_{K_1}(G_{K_2}(x))$ is $2\epsilon\delta$-close to t-wise independent.*

In particular, this implies that the permutation family F^r obtained by sequential r-fold composition of an ϵ-close to t-wise independent permutation family F is $2^{r-1}\epsilon^r$-close to t-wise independent. We point out that for a meaningful application of this lemma, we require that $\epsilon < 1/2$.

Differential and linear cryptanalysis. For a permutation family $F : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^n$, we define the *expected differential probability* (EDP) for a given pair Δ and Δ' of non-zero input- and output-differences, as

$$\mathrm{EDP}_F(\Delta, \Delta') = \Pr_{K,X}[F_K(X \oplus \Delta) \oplus F_K(X) = \Delta'] \ ,$$

where K and X are independent and uniformly distributed over the m-bit and n-bit strings, respectively. We also define $\mathrm{MEDP}_F = \max_{\Delta,\Delta' \neq 0} \mathrm{EDP}_F(\Delta, \Delta')$. It is easy to see that if F is ϵ-close to pairwise independent, then $\mathrm{MEDP}_F \leq \epsilon + \frac{1}{2^n - 1}$. We note that a similar result extends to any subset of n output bits, and hence to so-called *truncated* differential probabilities.

We note that higher-order differential cryptanalysis [34, 36] generalizes differential cryptanalysis to look at higher order derivatives. It is not hard to see that almost t-wise independence will imply resistance to order-$\log_2 t$ differential cryptanalysis, as the property relies on the evaluation of the cipher on at most t inputs. We note that while (almost) t-wise independence refers to attacks that look at an arbitrary set of t inputs, an order-$\log_2 t$ differential attack looks at all inputs that lie in some $\log_2 t$-dimensional hypercube, so a total of t inputs but they are not arbitrary.

The connection between pairwise independence and linear cryptanalysis is slightly less obvious. For more details, see the final version of our paper [40].

2.2 Key-Alternating Ciphers and Substitution Permutation Networks

A *Key Alternating Cipher* (KAC) (cf. Figure 1) is parameterized by a block size n, number of rounds r, and a fixed permutation $P : \mathbb{F}_{2^n} \to \mathbb{F}_{2^n}$. A KAC is a family of functions indexed by $r+1$ sub-keys $K_0, K_1, \ldots, K_r$, and defined recursively as follows:

$$\begin{aligned} F^{(0)}_P(x) &= x \oplus K_0 \\ F^{(i)}_{P,K_0,\ldots,K_i}(x) &= P(F^{(i-1)}_{P,K_0,\ldots,K_{i-1}}(x)) \oplus K_i \ . \end{aligned}$$

The family of functions is $\mathcal{F}_P := \left\{F^{(r)}_{P,K_0,\ldots,K_r}(x) : K_i \in \mathbb{F}_2^n\right\}$. One can also naturally extend this to have different permutations in each round.

A *Substitution-Permutation Network* (SPN) (cf. Figure 2) can be seen as a special case of a KAC, where $n = k \cdot b$ (we refer to k as the *width*), and the permutation P is obtained from an S-box $S : \mathbb{F}_{2^b} \to \mathbb{F}_{2^b}$ and a linear *mixing layer*, described by a matrix $M \in \mathbb{F}_{2^b}^{k \times k}$. In particular, P splits its input x into k b-bit blocks $x_1, \ldots, x_k$, and computes first $y_i = S(x_i)$ for each i, and finally

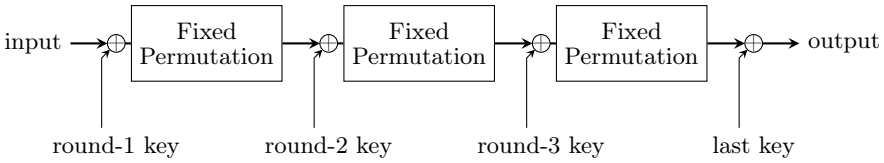

Fig. 1. Illustration of Key Alternating Cipher

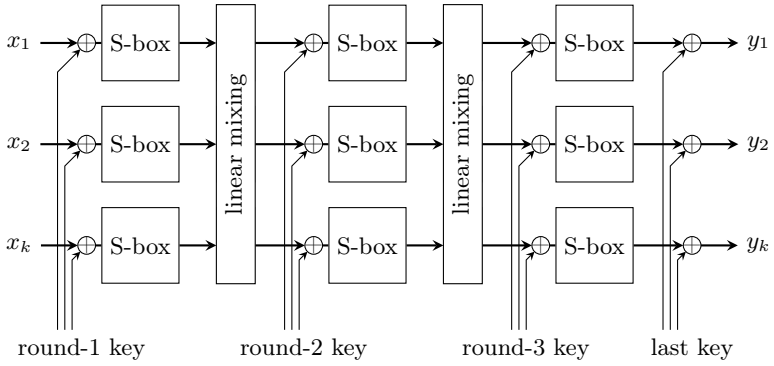

Fig. 2. Illustration of Substitution Permutation Network

outputs $M \cdot (y_1, \ldots, y_k)$. One can of course instead think of a KAC as a special of an SPN with width $k = 1$.

A fact that we will use repeatedly is that in order to bound how close to pairwise independent an SPN or KAC is, it is enough to analyze the distribution of the non-zero difference of outputs of the SPN/KAC, and its distance from the uniform distribution over non-zero strings.

Analyzing Pairwise Independence of KACs and SPNs. We will use the following lemma to reduce the analysis of pairwise independence to analyzing the distribution of differences.

Lemma 2. *Assume that the KAC (resp. SPN) $\mathcal{F}_P$ (resp. $\mathcal{F}_{P,M}$) has the property that for any input difference $\Delta \neq 0$, the distribution of*

$$\Delta' := F_K(x) \oplus F_K(x \oplus \Delta)$$

is ϵ-close to uniform (where the randomness of the distribution is taken over x and K). Then, the KAC (resp. SPN) is ϵ-close to pairwise independent.

The proof is deferred to the full version [40].

Advanced Encryption Standard. The mostly widely used block cipher is the world is Advanced Encryption Standard (AES), which is based on the SPN framework.

The block size is 128 bits, width is 16, i.e. $n = 128, k = 16, b = 8$. AES is a family of ciphers which have 10, 12 or 14 rounds.

The S-box is instantiated by $S(x) = A(x^{2^8-2})$, where $x \mapsto x^{2^8-2}$ is the patched inverse function over $\mathbb{F}_{2^8}$, A is an invertible affine function over $\mathbb{F}_2^8$. The exact form of A is irrelevant for this paper (as shown by Lemma 14).

The linear mixing function is instantiated by the composition of ShiftRows and MixColumns. Their descriptions are deferred to the full version [40].

2.3 Trace in Fields of Characteristic Two

We describe a number of facts related to the finite field $\mathbb{F}_{2^n}$ of characteristic 2 and the trace function over it. For proofs of the claims below, we refer the reader to any standard text on the subject, e.g. [39].

Definition 4. *The trace function* $\mathsf{Tr} : \mathbb{F}_{2^n} \to \mathbb{F}_2$ *is defined as* $\mathsf{Tr}(x) = \sum_{i=0}^{n-1} x^{2^i}$.

Lemma 3. *For every* $x \in \mathbb{F}_{2^n}$, $\mathsf{Tr}(x^2) = \mathsf{Tr}(x)$.

Lemma 4. *For every* $x, y \in \mathbb{F}_{2^n}$, $\mathsf{Tr}(x + y) = \mathsf{Tr}(x) + \mathsf{Tr}(y)$. *In particular, the set of elements* $x \in \mathbb{F}_{2^n}$ *with* $\mathsf{Tr}(x) = 0$ *form an* $\mathbb{F}_2$*-subspace of dimension* $n - 1$.

Lemma 5. *Let* $\alpha \in \mathbb{F}_{2^n}$. *The equation* $y(y \oplus 1) = \alpha$ *over* $\mathbb{F}_{2^n}$ *has two solutions if* $\mathsf{Tr}(\alpha) = 0$ *and no solutions otherwise.*

Corollary 1. *Let* $a, b, c \in \mathbb{F}_{2^n}$ *and* a, b *are non-zero. The equation* $ax^2 + bx + c = 0$ *has two solutions over* $\mathbb{F}_{2^n}$ *if* $\mathsf{Tr}(ac/b^2) = 0$ *and no solutions otherwise.*

Lemma 6. *For every* $x \neq y \in \mathbb{F}_{2^n}$, *let* $S_x := \{z : \mathsf{Tr}(xz) = 0\}$ *and* $S_y := \{z : \mathsf{Tr}(yz) = 0\}$. *Then,* $S_x \neq S_y$. *Indeed, since these are* $(n-1)$*-dimensional subspaces, they intersect at exactly* 2^{n-2} *elements.*

We also need the following Lemma from Nyberg's work [45], which we reprove for completeness.

Lemma 7 ([45]). *Let* $P : \mathbb{F}_{2^n} \to \mathbb{F}_{2^n}$ *be the patched inversion function* $P(x) = x^{2^n-2}$. *For every* $\delta, \gamma \neq 0$, *let* $p_{\delta,\gamma} := \Pr_{x \leftarrow \mathbb{F}_{2^n}}[P(x) \oplus P(x \oplus \delta) = \gamma]$. *Then,*

$$p_{\delta,\gamma} = \begin{cases} 2/2^n, & \text{if } \delta\gamma = 1 \\ 0, & \text{if } \delta\gamma \neq 1 \end{cases} + \begin{cases} 2/2^n, & \text{if } \mathsf{Tr}((\delta\gamma)^{-1}) = 0 \\ 0, & \text{if } \mathsf{Tr}((\delta\gamma)^{-1}) = 1 \end{cases}$$

The following corollary is an immediate consequence.

Corollary 2. *For any non-zero* $\delta \in \mathbb{F}_{2^n}$, *let*

$$p(\gamma) := \Pr_{x \leftarrow \mathbb{F}_{2^n}}[P(x) \oplus P(x \oplus \delta) = \gamma] \ .$$

Let $\mathcal{D}_\delta$ *denote the distribution with probability mass function* p *and let* $\mathcal{D}'_\delta$ *denote the distribution with probability mass function* $p'(\gamma) = p(\gamma^{-1})$, *we have:*

- $\mathcal{D}'_\delta$ *is* $(2/2^b)$*-close to the uniform distribution on a subspace of dimension* $b - 1$.
- $\mathrm{H}_2(\mathcal{D}_\delta) \geq -\log_2\left(\frac{2}{2^b} + \frac{8}{2^{2b}}\right)$.

2.4 Basics of Discrete Fourier Analysis

The characters of the group $\mathbb{F}_2^n$ are functions $\{\chi_{\mathbf{x}} : \mathbb{F}_2^n \to \mathbb{R}\}_{\mathbf{x}\in\mathbb{F}_2^n}$ defined by

$$\chi_{\mathbf{x}}(\mathbf{y}) = (-1)^{\langle \mathbf{x},\mathbf{y}\rangle}$$

The functions $\{\chi_{\mathbf{x}}\}_{\mathbf{x}\in\mathbb{F}_2^n}$ are orthonormal under the inner product[7]

$$\langle \chi_{\mathbf{x}}, \chi_{\mathbf{x}'}\rangle := \frac{1}{2^n}\sum_{\mathbf{y}\in\mathbb{F}_2^n} \chi_{\mathbf{x}}(\mathbf{y})\chi_{\mathbf{x}'}(\mathbf{y}) \ .$$

Let $f : \mathbb{F}_2^n \to \mathbb{R}$ be a real-valued function on $\mathbb{F}_2^n$. Writing $f = \sum_{\mathbf{x}\in\mathbb{F}_2^n} \widehat{f}(\mathbf{x})\chi_{\mathbf{x}}$, we have the Fourier (inversion) formulas

$$f(\mathbf{y}) = \sum_{\mathbf{x}\in\mathbb{F}_2^n} \widehat{f}(\mathbf{x})\chi_{\mathbf{x}}(\mathbf{y}) \quad \text{and} \quad \widehat{f}(\mathbf{x}) = \langle f, \chi_{\mathbf{x}}\rangle = \frac{1}{2^n}\sum_{\mathbf{y}\in\mathbb{F}_2^n} f(\mathbf{y})\chi_{\mathbf{x}}(\mathbf{y})$$

We need the following two facts. For proofs, we refer the reader to [47].

Lemma 8 (Parseval's Theorem). $\frac{1}{2^n}\sum_{\mathbf{y}\in\mathbb{F}_2^n} f(\mathbf{y})^2 = \sum_{\mathbf{x}\in\mathbb{F}_2^n} \widehat{f}(\mathbf{x})^2$.

If S is a subspace of $\mathbb{F}_2^n$, let $S^\perp = \{\mathbf{y} : \langle \mathbf{x}, \mathbf{y}\rangle = 0 \text{ for all } \mathbf{x} \in S\}$ denote its dual subspace. If S is k-dimensional, $S^\perp$ is $(n-k)$-dimensional.

Lemma 9. *Let $S \subseteq \mathbb{F}_2^n$ be a subspace and f_S denote the uniform probability distribution on S. That is, $f_S(\mathbf{y}) = \frac{1}{|S|}$ if $\mathbf{y} \in S$ and 0 otherwise. Then, $\hat{f}_S(\mathbf{x}) = \frac{1}{2^n}$ if $\mathbf{x} \in S^\perp$ and 0 otherwise.*

In particular, let $S \subseteq \mathbb{F}_2^n$ be an $(n-1)$-dimensional subspace which can equivalently be denoted as (the dual subspace) $S = \{0, v\}^\perp$ for some $v \in \mathbb{F}_2^n$. Then,

$$\hat{f}_S(y) = \begin{cases} \frac{1}{2^n}, & \text{if } y \in \{0, v\} \\ 0, & \text{otherwise} \end{cases}$$

Let $f : \mathbb{F}_2^n \to \mathbb{R}, g : \mathbb{F}_2^{n'} \to \mathbb{R}$ be two real-valued functions on $\mathbb{F}_2^n$ and $\mathbb{F}_2^{n'}$ respectively. Their tensor product $f \otimes g : \mathbb{F}_2^{n+n'} \to \mathbb{R}$ is a real-valued function on $\mathbb{F}_2^{n+n'}$ such that

$$(f \otimes g)(x, y) := f(x) \cdot g(y) \text{ for all } x \in \mathbb{F}_2^n, y \in \mathbb{F}_2^{n'}.$$

Assume X, Y are two independent random variables on $\mathbb{F}_2^n$ and $\mathbb{F}_2^{n'}$ respectively, and f, g are the probability mass functions of X, Y. Then $f \otimes g$ is the probability mass function of (X, Y), as

$$\Pr[(X, Y) = (x, y)] = \Pr[X = x] \cdot \Pr[Y = y] = f(x) \cdot g(y) = (f \otimes g)(x, y).$$

The Fourier transform of the tensor equals the tensor of the Fourier transforms.

Lemma 10 (Fourier transform of a Tensor). *For any $f : \mathbb{F}_2^n \to \mathbb{R}, g : \mathbb{F}_2^{n'} \to \mathbb{R}$, $\widehat{f \otimes g} = \hat{f} \otimes \hat{g}$.*

[7] Note that there are two inner products at play here, one over $\mathbb{F}_2^n$ and the other over $\mathbb{R}^{2^n}$, and we are abusing notation by denoting them both as $\langle \cdot, \cdot \rangle$.

3 Pairwise Independence of SPNs

The main result of this section is a proof of pairwise independence of the 3-round substitution-permutation network (see Figure 2) where the non-linear S-box is the patched inverse function over $\mathbb{F}_{2^n}$, used in the AES block cipher. We will show that the 3-round SPN is ϵ-close to pairwise independent for a constant $\epsilon < 1/2$, and note that an application of the MPR amplification lemma (Lemma 1) gives us $2^{-\Omega(r)}$-closeness to pairwise independence in $3r$ rounds.

In Section 3.1, we start with our main technical result, an *S-box extraction lemma*, which says that when the input difference of a single round of SPN has *sufficient Rényi entropy*, the output difference is close to uniformly random. We follow this up by describing mixing functions and their properties in Section 3.2. In Section 3.3, we then use the S-box extraction lemma and properties of mixing functions to show our main result, namely the pairwise independence of 3-round SPN. The reader is encouraged to refer back to Section 2.4 for relevant facts about discrete Fourier analysis as and when necessary.

3.1 The S-box Extraction Lemma

Before we state the S-box extraction lemma, we describe how it will be used to show the pairwise independence of SPNs. As noted in Lemma 2, it is sufficient to show that the distribution of *output differences* on any two inputs is close to uniformly random.

Consider the scenario in the last round of a substitution-permutation network, as illustrated in Figure 3. Before the last round, we will show that the input difference already has high (Rényi) entropy. Indeed, we will show that if there is one round of S-boxes and mixing before the last round, Δ_i has large entropy for any $i \in [k]$; and if there are two rounds of S-boxes and mixing before the last round, the joint distribution of $(\Delta_1, \ldots, \Delta_k)$ has (proportionally) high entropy. The question we ask then is, is the output (difference) vector $(\Delta'_1, \ldots, \Delta'_k)$ close to uniform? The extraction lemma provides an affirmative answer to this question.

Lemma 11 (The S-Box Extraction Lemma). *Let k, b be positive integers and $n = bk$. Let $\mathcal{D}$ be a distribution over $(\mathbb{F}_2^b)^k$ and consider the following probabilistic process called $\mathsf{Samp}_{\mathcal{D}}$.*

1. *Sample $(v_1, \ldots, v_k) \leftarrow \mathcal{D}$. Let $S_1, \ldots, S_k$ be $(b-1)$-dimensional subspaces where each $S_i = \{0, v_i\}^\perp$ is the subspace orthogonal to v_i.*
2. *For each $i \in [k]$, sample $x_i \leftarrow S_i$ independently at random, and output $(x_1, \ldots, x_k)$.*

For any $T \subseteq [k]$, let v_T denote the concatenation of $(v_i)_{i \in T}$, let $\mathcal{D}_T$ denote the distribution of v_T, let $\mathrm{H}_2[\mathcal{D}_T]$ denote its Rényi entropy. Then, the statistical distance between the joint distribution of $(x_1, \ldots, x_k)$ and the uniform distribution over $\mathbb{F}_2^{bk}$ is at most

$$\frac{1}{2}\sqrt{\sum_{T \subseteq [k], T \neq \varnothing} 2^{-\mathrm{H}_2[\mathcal{D}_T]}}\,.$$

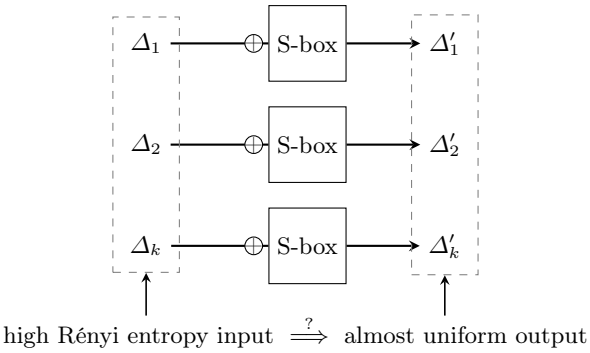

Fig. 3. Application Scenario of the Extraction Lemma

In particular, we have:

- **Weak Extraction:** *Assume that for all $i \in [k]$, $\mathrm{H}_2[v_i] \geq h$ for a fixed real $h \leq b$.Then the statistical distance between the joint distribution of $(x_1, \ldots, x_k)$ and the uniform distribution over $\mathbb{F}_2^{bk}$ is at most $\frac{1}{2} \cdot \sqrt{\frac{2^k-1}{2^h}}$.*
- **Strong Extraction:** *Assume that for any $T \subseteq [k]$, $\mathrm{H}_2[v_T] \geq h \cdot |T|$ where v_T denotes the concatenation of $(v_i)_{i \in T}$. Then the statistical distance between the joint distribution of $(x_1, \ldots, x_k)$ and the uniform distribution over $\mathbb{F}_2^{bk}$ is at most*

$$\frac{1}{2} \cdot \sqrt{\left(1 + \frac{1}{2^h}\right)^k - 1}$$

which, in turn, is at most $\sqrt{\frac{k}{2^{h+1}}}$ assuming $k \leq 2^h$.

Proof. Let f denote the probability mass function of $\mathsf{Samp}_{\mathcal{D}}$. That is, $f(x_1, \ldots, x_k)$ is the probability that $\mathsf{Samp}_{\mathcal{D}}$ outputs $(x_1, \ldots, x_k)$. Let $p(v_1, \ldots, v_k)$ denote the probability assigned by the distribution $\mathcal{D}$ to $(v_1, \ldots, v_k)$ and let ϕ_S denote the probability mass function of the uniform distribution over the subspace $S \subseteq \mathbb{F}_2^b$. Then,

$$f(x_1, \ldots, x_k) = \sum_{v_1, \ldots, v_k \in \mathbb{F}_2^b} p(v_1, \ldots, v_k) \cdot \phi_{S_1}(x_1) \cdot \phi_{S_2}(x_2) \cdot \ldots \cdot \phi_{S_k}(x_k)$$

where $S_i = \{0, v_i\}^\perp$ is an implicit function of v_i, as before. We will write this as

$$f = \sum_{v_1, \ldots, v_k \in \mathbb{F}_2^b} p(v_1, \ldots, v_k) \cdot \left(\phi_{S_1} \otimes \phi_{S_2} \otimes \ldots \otimes \phi_{S_k}\right)$$

We are interested in the statistical distance $\mathrm{d}_{\mathrm{TV}}(f, u) = \frac{1}{2}\|f - u\|_1$, where u is the uniform distribution over $\mathbb{F}_2^{bk}$. It suffices to bound $\|\hat{f} - \hat{u}\|_2^2$ since

$$\|f - u\|_1^2 \leq 2^{kb}\|f - u\|_2^2 = 2^{2kb}\|\hat{f} - \hat{u}\|_2^2. \tag{1}$$

where the inequality comes from Cauchy-Schwartz and the equality comes from Parseval's theorem (Lemma 8).

The Fourier transform of f equals

$$\hat{f}(y_1,\ldots,y_k) = \sum_{v_1,\ldots,v_k\in\mathbb{F}_2^b} p_{(v_1,\ldots,v_k)} \cdot \prod_{i\in[k]} \hat{\phi}_{S_i}(y_i)$$

Observe that by Lemma 9, $\hat{\phi}_{S_i}$ is 0 everywhere except for $\hat{\phi}_{S_i}(v_i) = \hat{\phi}_{S_i}(0) = 1/2^b$. Thus the only inputs $(y_1,\ldots,y_k)$ on which $\hat{f}(y_1,\ldots,y_k)\neq 0$ are those in the set $\{0,v_1\}\times\{0,v_2\}\times\ldots\times\{0,v_k\}$. Thus,

$$\hat{f}(y_1,\ldots,y_k) = \frac{1}{2^{bk}}\cdot \Pr[v_i = y_i \text{ for all } i \text{ s.t. } y_i\neq 0]. \tag{2}$$

The ℓ_2-norm of the Fourier transform of $f-u$ can then be computed as

$$\begin{aligned}\left\|\hat{f}-\hat{u}\right\|_2^2 &= \sum_{\substack{y_1,\ldots,y_k\in\mathbb{F}_2^b\\(y_1,\ldots,y_k)\neq\vec{0}}} \hat{f}^2(y_1,\ldots,y_k)\\ &= \sum_{\substack{T\subseteq[k]\\T\neq\varnothing}}\sum_{\substack{y_1,\ldots,y_k\in\mathbb{F}_2^b\\ y_i\neq 0 \text{ iff } i\in T}} \hat{f}^2(y_1,\ldots,y_k)\\ &= \sum_{\substack{T\subseteq[k]\\T\neq\varnothing}}\sum_{\substack{y_1,\ldots,y_k\in\mathbb{F}_2^b\\ y_i\neq 0 \text{ iff } i\in T}} \frac{1}{2^{2bk}}\cdot \Pr[v_i = y_i \text{ for all } i\in T]^2. \end{aligned}\tag{3}$$

Let $v_T := (v_i)_{i\in T}$ denote the vector v restricted to indices in T, let $\mathcal{D}_T$ denote the distribution of v_T, and let f_T denote the probability mass function of $\mathcal{D}_T$. Then,[8]

$$\left\|\hat{f}-\hat{u}\right\|_2^2 \le \frac{1}{2^{2bk}}\sum_{\substack{T\subseteq[k]\\T\neq\varnothing}} ||f_T||_2^2 = \frac{1}{2^{2kb}}\sum_{\substack{T\subseteq[k]\\T\neq\varnothing}} 2^{-\mathrm{H}_2[\mathcal{D}_T]}. \tag{4}$$

Combining with equation (1) concludes the proof of the general case.

$$\mathrm{d_{TV}}(f,u) \le \frac{1}{2}\cdot 2^{kb}\cdot\|\hat{f}-\hat{u}\|_2 \le \frac{1}{2}\sqrt{\sum_{\substack{T\subseteq[k]\\T\neq\varnothing}} 2^{-\mathrm{H}_2[\mathcal{D}_T]}}. \tag{5}$$

Setting 1: Weak Extraction. Assume for any $i\in[k]$, $\mathrm{H}_2[\mathcal{D}_{\{i\}}]\ge h$. Then, for any non-empty set $T\subseteq[k]$, we have $\mathrm{H}_2[\mathcal{D}_T]\ge h$. Therefore, combining with equation (5),

$$\mathrm{d_{TV}}(f,u) \le \frac{1}{2}\sqrt{\sum_{\substack{T\subseteq[k]\\T\neq\varnothing}} 2^{-\mathrm{H}_2[\mathcal{D}_T]}} \le \frac{1}{2}\cdot\sqrt{\frac{2^k-1}{2^h}}.$$

[8] The first inequality symbol in the equation is tight, if $V_1,\ldots,V_k$ are always non-zero.

Setting 2: Strong Extraction. Assume for any $T \subseteq [k]$, $\mathrm{H}_2[\mathcal{D}_T] \geq h \cdot |T|$. Then

$$\sum_{\substack{T \subseteq [k] \\ T \neq \varnothing}} 2^{-\mathrm{H}_2[\mathcal{D}_T]} \leq \sum_{\substack{T \subseteq [k] \\ T \neq \varnothing}} \left(\frac{1}{2^h}\right)^{|T|} = \left(1 + \frac{1}{2^h}\right)^k - 1$$

using the binomial expansion. Combining with equation (5), we have

$$\mathrm{d}_{\mathrm{TV}}(f, u) \leq \frac{1}{2} \sqrt{\sum_{\substack{T \subseteq [k] \\ T \neq \varnothing}} 2^{-\mathrm{H}_2[\mathcal{D}_T]}} \leq \frac{1}{2} \cdot \sqrt{\left(1 + \frac{1}{2^h}\right)^k - 1}.$$

If we additionally assume that $k \leq 2^h$, then

$$\mathrm{d}_{\mathrm{TV}}(f, u) \leq \frac{1}{2} \cdot \sqrt{\left(1 + \frac{1}{2^h}\right)^k - 1} \leq \frac{1}{2} \sqrt{e^{\frac{k}{2^h}} - 1} \leq \frac{1}{2} \sqrt{\frac{2k}{2^h}}.$$

The last inequality symbol holds only if $\frac{k}{2^h} \leq 1.256\ldots$, which follows from the condition $k \leq 2^h$. □

We remark that Fourier analysis can be bypassed here. The above proof uses Fourier analysis to bound the collision probability. There is an alternative proof of the extraction lemma in the full version [40] that bounds the collision probability using "elementary" non-Fourier methods.

Comparing Figure 3 with the statement of the extraction lemma. The outstanding contrast is that the extraction lemma assumes a very specific linear algebra structure. That is, consider the domain as vector space $\mathbb{F}_2^b$, the output (difference) vector is sampled as a random vector orthogonal to the input (difference) vector. While in each round of SPN, the input is subtracted by the random key and then feed into the S-box. The output difference *is not* sampled uniformly from a subspace.

However, we hope the two can be bridged by change of variables. Say we start with two inputs differing Δ, let Δ' denote the difference after key-subtraction and S-box. We hope there exist 1-to-1 mappings $\pi_{\mathrm{in}}, \pi_{\mathrm{out}} : \mathbb{F}_{2^b} \to \mathbb{F}_2^b$ such that $\pi_{\mathrm{out}}(\Delta')$ is a random vector orthogonal to $\pi_{\mathrm{in}}(\Delta)$.

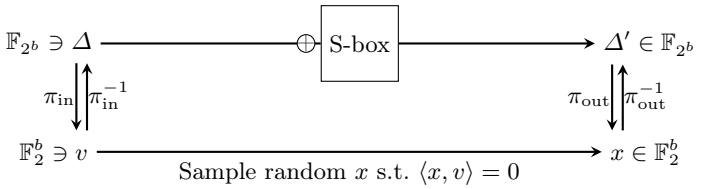

Fig. 4. Subtracting Key followed by S-box ≈ Subspace Sampling, modulo Change of Variables

Figure 4 illustrates the property we are looking for. Although it cannot be exactly satisfied by any S-box — we know $\pi_{\mathrm{out}}(\Delta')$ doesn't equal x by distribution, because $\Delta = 0 \iff \Delta' = 0$ — we show that pragmatic S-boxes almost satisfy the property.

Assuming the S-box is the patched inverse function, the following lemma shows that $\pi_{\mathrm{out}}(\Delta')$ is statistically close to a random vector orthogonal to $\pi_{\mathrm{in}}(\Delta)$, as long as $\Delta \neq 0$.

Lemma 12. *Assume S-box is the patched inverse $P(x) = x^{2^b-2}$. There exist 1-to-1 mappings $\pi_{\mathrm{in}}, \pi_{\mathrm{out}} : \mathbb{F}_{2^b} \to \mathbb{F}_2^b$ such that for any non-zero $\Delta \in \mathbb{F}_{2^b}$, letting Δ' denotes a random variable defined by*

$$\Delta' := P(r) - P(r + \Delta)$$

for a uniformly random $r \in \mathbb{F}_{2^b}$, the statistical distance between $\pi_{\mathrm{out}}(\Delta')$ and the uniform distribution over $\{0, \pi_{\mathrm{in}}(\Delta)\}^\perp$ is no more than $\frac{2}{2^b}$.

Proof. As shown in Lemma 7 (from [45]),

$$\Pr[\Delta' = \delta] = \begin{cases} \frac{2}{2^b}, & \text{if } \delta = \frac{1}{\Delta} \\ 0, & \text{o.w.} \end{cases} + \begin{cases} \frac{2}{2^b}, & \text{if } \mathsf{Tr}(\frac{1}{\delta\Delta}) = 0 \\ 0, & \text{o.w.} \end{cases}$$

Define $\pi_{\mathrm{out}}(x) = x^{2^b-2}$ to be the patched inverse as well. Then

$$\Pr[\pi_{\mathrm{out}}(\Delta') = x] = \begin{cases} \frac{2}{2^b}, & \text{if } x = \Delta \\ 0, & \text{o.w.} \end{cases} + \begin{cases} \frac{2}{2^b}, & \text{if } x \neq 0 \text{ and } \mathsf{Tr}(\frac{x}{\Delta}) = 0 \\ 0, & \text{o.w.} \end{cases}$$

As show in Lemma 4, $x \mapsto \mathsf{Tr}(\frac{x}{\Delta})$ is linear function over $\mathbb{F}_2$. Define $\pi_{\mathrm{in}}(\Delta)$ as the coefficient vector of $x \mapsto \mathsf{Tr}(\frac{x}{\Delta})$. Then

$$\Pr[\pi_{\mathrm{out}}(\Delta') = x] = \begin{cases} \frac{2}{2^b}, & \text{if } x = \Delta \\ 0, & \text{o.w.} \end{cases} + \begin{cases} \frac{2}{2^b}, & \text{if } x \neq 0 \text{ and } \langle \pi_{\mathrm{in}}(\Delta), x \rangle = 0 \\ 0, & \text{o.w.} \end{cases}$$

Apparently, the statistical distance between $\pi_{\mathrm{out}}(\Delta')$ and the uniform distance over $\{0, \pi_{\mathrm{in}}(\Delta)\}^\perp$ is $\frac{2}{2^b}$. □

The following lemma shows the analogous statement for the cube function. The proof is deferred to the full version [40].

Lemma 13. *Assume S-box is the cube function $P(x) = x^3$ over $\mathbb{F}_{2^b}$ where b is even*[9]*. There exist 1-to-1 mappings $\pi_{\mathrm{in}}, \pi_{\mathrm{out}} : \mathbb{F}_{2^b} \to \mathbb{F}_2^b$ such that for any non-zero $\Delta \in \mathbb{F}_{2^b}$, letting Δ' denote a random variable defined by*

$$\Delta' := P(r) - P(r + \Delta)$$

for a uniformly random $r \in \mathbb{F}_{2^b}$, $\pi_{\mathrm{out}}(\Delta')$ is the uniform distribution over $\{0, \pi_{\mathrm{in}}(\Delta)\}^\perp$.

[9] The condition on b being even is necessary to ensure that P is a permutation.

In Section 3.4 we are going to analyze AES. The S-box in AES is called Rijndael S-box, which is not exactly the patched inverse function. Rijndael S-box is the composition of the patched inverse function and an affine transformation. The following lemma shows that the additional affine transformation makes little difference.

Lemma 14. *Assume S-box is $P(x) = A(x^{2^b-2})$, where A is an affine permutation over $\mathbb{F}_2^b$. There exist 1-to-1 mappings $\pi_{\text{in}}, \pi_{\text{out}} : \mathbb{F}_{2^b} \to \mathbb{F}_2^b$. For any non-zero $\Delta \in \mathbb{F}_{2^b}$, let Δ' denote a random variable defined by*

$$\Delta' := P(r) - P(r + \Delta)$$

for a uniformly random $r \in \mathbb{F}_{2^b}$. The statistical distance between $\pi_{\text{out}}(\Delta')$ and the uniform distribution over $\{0, \pi_{\text{in}}(\Delta)\}^\perp$ is no more than $\frac{2}{2^b}$.

Proof. As we are analyzing the differences, any additive constant in the affine function A has no effect. Thus we can safely assume A is a linear permutation.

When input difference is Δ, the output difference is

$$\Delta' = P(r) - P(r+\Delta) = A(r^{2^b-2}) - A((r+\Delta)^{2^b-2}) = A(r^{2^b-2} - (r+\Delta)^{2^b-2}).$$

Define $\Delta^* = r^{2^b-2} - (r+\Delta)^{2^b-2}$, then $\Delta' = A(\Delta^*)$.

Lemma 12 shows that there exists $\pi_{\text{in}}, \pi_{\text{out}}$ such that $\pi_{\text{out}}(\Delta^*)$ is close to uniform distribution over $\{0, \pi_{\text{in}}(\Delta)\}$. Define $\pi'_{\text{out}}(x) := \pi_{\text{out}}(A^{-1}(x))$. Then $\pi'_{\text{out}}(\Delta') = \pi_{\text{out}}(A^{-1}(\Delta')) = \pi_{\text{out}}(\Delta^*)$, which is close to uniform distribution over $\{0, \pi_{\text{in}}(\Delta)\}$. Thus $\pi_{\text{in}}, \pi'_{\text{out}}$ are what we need. □

3.2 Properties of Mixing Functions

Before proceeding to show the almost-pairwise independence of SPN constructions using the extraction lemma, we describe properties that we need the mixing functions to satisfy. We define two such properties below and prove some elementary statements about them.

The first property that we call *diffusion* requires that if one of the input blocks of the (typically linear) function $M : (\mathbb{F}_{2^b})^k \to (\mathbb{F}_{2^b})^k$ has sufficient entropy and the distribution of the k input blocks are independent, then *each output block* has large entropy. It is not hard to see that both the sufficient entropy condition and the independence condition on the input are necessary for such a statement to be true. Looking ahead, this property will turn out to be useful in the first layer (or the first few layers) of the SPN where we wish to propagate differences in one input block to differences in all of them.

Property 1 (Diffusion). Let $M : (\mathbb{F}_{2^b})^k \to (\mathbb{F}_{2^b})^k$ be a function. Let $\mathrm{H}_\alpha \in \{\mathrm{H}_2, \mathrm{H}_\infty\}$ be an entropy function. Let $X_1, \ldots, X_k$ be independent random variables over $\mathbb{F}$ such that there exists an i for which $\mathrm{H}_\alpha(X_i) \geq h$ for a real h, and let $(Y_1, \ldots, Y_k) := M(X_1, \ldots, X_k)$. M is *diffusing* if

$$\text{for all } i \in [k],\ \mathrm{H}_\alpha(Y_i) \geq h.$$

We now show a sufficient condition for a function to be diffusing. The proof is deferred to the full version [40].

Lemma 15. *If $M \in (\mathbb{F}_{2^b})^{k\times k}$ is a matrix with no zero entry, the linear mapping $x \mapsto Mx$ is diffusing* (i.e. satisfies Property 1).

The second property that we call *entropy-preservation* requires that if all of the input blocks of the (typically linear) function $M : (\mathbb{F}_{2^b})^k \to (\mathbb{F}_{2^b})^k$ have sufficient entropy and the distribution of the k blocks are independent, then *each collection of output blocks* have large joint entropy. Looking ahead, this property will turn out to be useful in the subsequent layers of the SPN to ensure that the mixing layers do not *reduce* the entropy. As one might expect, this property comes for free if M is an invertible linear map. The proof is deferred to the full version [40].

Property 2 (Entropy Preservation). A function $M : (\mathbb{F}_{2^b})^k \to (\mathbb{F}_{2^b})^k$ is *entropy preserving* if for any entropy function $\mathrm{H}_\alpha \in \{\mathrm{H}_2, \mathrm{H}_\infty\}$, for any real h, for any independent random variables $X_1, \ldots, X_k$ over $\mathbb{F}_{2^b}$ such that $\mathrm{H}_\alpha(X_i) \geq h$ for all $i \in [k]$, letting $(Y_1, \ldots, Y_k) := M(X_1, \ldots, X_k)$, we have

$$\mathrm{H}_\alpha(Y_{i_1}, \ldots, Y_{i_s}) \geq s \cdot h$$

for any $\{i_1, \ldots, i_s\} \subseteq [k]$.

Lemma 16. *If $M \in (\mathbb{F}_{2^b})^{k\times k}$ is an invertible matrix, the mapping $x \mapsto Mx$ is entropy-preserving* (i.e. satisfies Property 2).

Connection to Branch Number. The branch number of a matrix $M \in (\mathbb{F}_{2^b})^{k\times k}$ is defined to be

$$\mathsf{br}(M) = \mathsf{max}_{\alpha\in(\mathbb{F}_{2^b})^k}(\mathsf{wt}(\alpha) + \mathsf{wt}(M\alpha))$$

where wt denotes the Hamming weight. Having an optimal branch number is considered a desirable feature for mixing functions [16, 27]. An observation by Miles and Viola [43] says that any matrix with the maximal branching number of $k+1$ also satisfies properties 1 and 2, although the converse does not necessarily hold.

3.3 Proofs of Pairwise Independence

In this section, we show several proofs of pairwise independence of SPNs using the patched inverse function $P(x) = x^{2^b-2}$ over the finite field $\mathbb{F}_{2^b}$. The first result (Theorem 1) applies in a regime where $k \leq b$ is relatively small; here, the result says that a 2-round SPN is close to pairwise independent. The second result (Theorem 2) is much more general and applies to large k as long as $k \leq 2^{b-4}$; here, the result says that a 3-round SPN is close to pairwise independent.

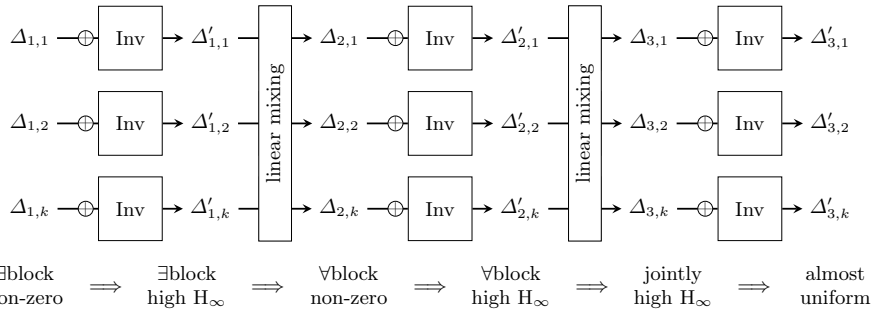

Fig. 5. Illustration of the proof of Theorem 2 and Lemma 17

Theorem 1. *Assume the S-box is $P(x) = x^{2^b-2}$ over $\mathbb{F}_{2^b}$ assume the mixing function is diffusing, that is, it satisfies Property 1. Then a 2-round SPN with k blocks each of which has b bits is ϵ-close to 2-wise independent where*

$$\epsilon \leq \frac{2+4k}{2^b} + \sqrt{\frac{2^k-1}{2^{b+1}}}.$$

Theorem 2. *Assume the S-box is patched inverse $P(x) = x^{2^b-2}$, assume the mixing function satisfies Property 1 and Property 2. Then 3-round SPN is ϵ-close to 2-wise independent where*

$$\epsilon \leq \frac{2+8k}{2^b} + \sqrt{\frac{k}{2^b}}.$$

Proof. Name the variables as in Figure 5, fix any input differences $\Delta_{1,1}, \ldots, \Delta_{1,k}$ which are not all zero. We wish to show that the distribution of $(\Delta'_{3,1}, \ldots, \Delta'_{3,k})$ is ϵ-close to uniform. By Lemma 2, this implies ϵ-closeness to pairwise independence. We proceed via a hybrid argument.

Hybrid 0. Hybrid 0 is the real world hybrid that is illustrated in Figure 5.

Hybrid 1. Pick some j where $\Delta_{1,j} \neq 0$. W.l.o.g., assume $\Delta_{1,1} \neq 0$. Note that the distribution of $\Delta'_{1,1}$ is $(2/2^b)$-close to uniformly random over a subset of size 2^{b-1} (Corollary 2). Call this uniform distribution $\mathcal{D}'_{1,1}$. We have $H_\infty(\mathcal{D}'_{1,1}) = b-1$.

Hybrid 1 is the same as hybrid 0 except that we replace $\Delta'_{1,1}$ by a random sample from the distribution $\mathcal{D}'_{1,1}$. The statistical distance from Hybrid 0 is at most $\frac{2}{2^b}$.

Claim Assume that the mixing function satisfies Property 1. In Hybrid 1, $H_\infty[\Delta_{2,j}] \geq b-1$ for all $j \in [k]$.

Hybrid 2. In this hybrid, we ensure $\Delta_{2,j} \neq 0$ for all $j \in [k]$. Formally, hybrid 2 is the same as hybrid 1 except that we replace $\Delta_{2,j}$ by 1 if $\Delta_{2,j} = 0$ in hybrid 1. The statistical distance from Hybrid 1 is at most $\frac{2k}{2^b}$.

Lemma 17 shows that the joint distribution of $(\Delta'_{3,1}, \ldots, \Delta'_{3,k})$ is $\left(\frac{6k}{2^b} + \sqrt{\frac{k}{2^b}}\right)$ close to uniform in hybrid 2.

Putting everything together, the statistical distance between $(\Delta'_{3,1}, \ldots, \Delta'_{3,k})$ and the uniform distribution is at most $\frac{2+8k}{2^b} + \sqrt{\frac{k}{2^b}}$. □

Lemma 17. *Assume the S-box is patched inverse $P(x) = x^{2^b-2}$, assume the mixing function satisfies Property 2. Starting with a pair of inputs, whose difference is entry-wise-nonzero, after a 2-round SPN, the statistical distance between the output difference and the uniform distribution is no more than $\frac{6k}{2^b} + \sqrt{\frac{k}{2^b}}$.*

Proof. Name the variables as the last two rounds in Figure 5, fix any set of input differences $\Delta_{2,1}, \ldots, \Delta_{2,k}$ which are all non-zero. We wish to show that the distribution of $(\Delta'_{3,1}, \ldots, \Delta'_{3,k})$ is ϵ-close to uniform. We proceed via a hybrid argument.

Hybrid 0. Hybrid 0 is the real world hybrid.

Hybrid 1. Since $\Delta_{2,j} \neq 0$ for all $j \in [k]$, the distribution of $\Delta'_{2,j}$ is $(2/2^b)$-close to uniformly random over a subset of size 2^{b-1} (Corollary 2). Call this uniform distribution $\mathcal{D}'_{2,j}$. We have $\mathrm{H}_\infty(\mathcal{D}'_{2,j}) = b - 1$.

Hybrid 1 is the same as hybrid 0 except that we replace $\Delta'_{2,j}$ by a vector drawn from the distribution $\mathcal{D}'_{2,j}$ for each $j \in [k]$. The statistical distance from Hybrid 0 is at most $\frac{2k}{2^b}$.

Claim Assume that the mixing function satisfies Property 2. In Hybrid 1, $\mathrm{H}_\infty[\Delta_{3,j}] \geq b - 1$ for all $j \in [k]$.

Hybrid 2. In this hybrid, we change the way $\Delta'_{3,j}$ is sampled based on $\Delta_{3,j}$. In particular:

- When $\Delta_{3,j} = \delta \neq 0$, the distribution of $\pi_{\mathrm{out}}(\Delta'_{3,j})$ conditioning on $\Delta_{3,j} = \delta$ is $\frac{2}{2^b}$-close to uniform distribution over $\{0, \pi_{\mathrm{in}}(\delta)\}^\perp$. Let $\pi_{\mathrm{out}}(\Delta'_{3,j})$ sampled uniformly from $\{0, \pi_{\mathrm{in}}(\delta)\}^\perp$ in hybrid 2.
- When $\Delta_{3,j} = 0$, $\Delta'_{3,j}$ is chosen to be uniformly random in hybrid 2.

Let us calculate the statistical distance between hybrids 1 and 2. The first bullet introduces a statistical distance of at most $2k/2^b$. The probability that a fixed coordinate $\Delta_{3,j}$ is 0 is at most $2/2^b$, and therefore, the probability that some coordinate is 0 is at most $2k/2^b$. In total, the statistical distance is at most $\frac{4k}{2^b}$.

By applying our extraction lemma[10] (Lemma 11), we know that, in hybrid 2, the joint distribution of $\Delta'_{3,1}, \ldots, \Delta'_{3,k}$ is at most $\sqrt{\frac{k}{2^b}}$-away from uniform.

Counting them together, the statistical distance between $(\Delta'_{3,1}, \ldots, \Delta'_{3,k})$ and the uniform distribution is at most $\frac{6k}{2^b} + \sqrt{\frac{k}{2^b}}$. □

3.4 AES is Almost Pairwise-Independent

Good asymptotic bounds have been shown in Theorem 1 and 2, but the analysis there is way too loose on AES parameter $(k = 16, b = 8)$. This section emphasizes on better concrete bound. Comparing with Section 3.3, the concrete bound is improved by the following tricks.

- Lemma 11 shows that the statistical distance is less than $\frac{1}{2} \cdot \sqrt{\left(1 + \frac{1}{2^h}\right)^k - 1}$, which is less than $\sqrt{\frac{k}{2^{h+1}}}$. The former is tighter. In particular, when $k = 16$, $b = 8$, $h = -\log_2\left(\frac{2}{2^b} + \frac{8}{2^{2b}}\right)$, the former shows $\mathrm{d_{TV}} \leq 0.18357\ldots \leq \frac{47}{256}$, and the latter shows $\mathrm{d_{TV}} \leq 0.25$.
- Lemma 18 is the strengthening of Lemma 17. Besides using the tighter bound from Lemma 11, it also considers Rényi entropy instead of min-entropy.
- Theorem 3 is the strengthening of Theorem 2. The proof of Theorem 3 (resp. Theorem 2) shows that after two rounds of AES (resp. one round of SPN), all block differences are non-zero with high probability. Then ignoring the rare event, Lemma 18 (resp. Lemma 17) will conclude the proof.
  The proof of Theorem 3 also carefully analyzes the rare event that some block difference is zero after 2 rounds of AES. It observes that, given the rare event happens, after two more rounds, all block differences will be non-zero with high probability.

Lemma 18 (Strengthening of Lemma 17). *Assume the S-box is patched inverse $P(x) = x^{2^b-2}$, assume the mixing function satisfies Property 2. Starting with a pair of inputs, whose difference is entry-wise-nonzero, after a 2-round SPN, the statistical distance between the output difference and the uniform distribution is no more than $\frac{4k}{2^b} + \frac{1}{2}\sqrt{(1+2^{-h})^k - 1}$, where $h = -\log_2\left(\frac{2}{2^b} + \frac{8}{2^{2b}}\right)$.*

In particular, when $k = 16$, $b = 8$, we have $\mathrm{d_{TV}} \leq \frac{64+47}{256}$.

The proof is mostly the same of Lemma 17 and is deferred to the full version [40].

Lemma 19. *Starting with a pair of distinct inputs, after 2-round of AES, including a tailing linear mixing, the output difference has zero entry with probability no more than $\frac{25}{2^7}$.*

Theorem 3. *6-round of AES is 0.472-close to pairwise independence.*

The proof is similar to that of Theorem 2 and is deferred to the full version [40].

[10] Our extraction lemma also requires $k \leq 2^{b-1}$. In the case $k > 2^{b-1}$, Lemma 17 can be trivially proved as $\mathrm{d_{TV}} \leq 1 \leq \frac{6k}{2^b} + \sqrt{\frac{k}{2^b}}$.

3.5 Multi-round SPNs and AES

We now combine the bounds from Theorems 1, 2, and 3 with the MPR amplification lemma (Lemma 1) to obtain the following theorems.

Theorem 4. *Assume the S-box is $P(x) = x^{2^b-2}$ over $\mathbb{F}_{2^b}$ assume the mixing function is diffusing, that is, it satisfies Property 1. Then a $(2r)$-round SPN with k blocks each of which has b bits is ϵ-close to 2-wise independent where*

$$\epsilon \leq 2^{r-1}\left(\frac{2+4k}{2^b} + \sqrt{\frac{2^k-1}{2^{b+1}}}\right)^r .$$

Further, if the mixing function additionally satisfies Property 2, then $(3r)$-round SPN is ϵ-close to 2-wise independent where

$$\epsilon \leq 2^{r-1}\left(\frac{2+8k}{2^b} + \sqrt{\frac{k}{2^b}}\right)^r .$$

Theorem 5. *$6r$-round AES is $2^{r-1}(0.472)^r$-close to pairwise independence.*

4 *t*-wise Independence of KAC

In this section, we consider a key-alternating cipher whose i^{th} round consists of applying a *public, fixed* permutation p_i to the current state followed by adding a (private) round-key s_i. The main result of this section is that for every r, there exist public permutations $p_1, \ldots, p_r$ such that r rounds of KAC using these permutations gets us close to $(r - o(r))$-wise independence. We achieve a strong notion of *pointwise* closeness (see definition 6) much stronger than the statistical distance measures considered in previous sections. Furthermore, it is easy to see that a t-round KAC can at best be (close to) t-wise independence, due to a simple entropy argument, meaning that our result is nearly optimal and entropy-preserving.

We remark that this is an *existential result*: namely, we do not explicitly construct the fixed permutations used by the KAC, but merely show that they exist. Indeed, we show that most permutations work, as is typical of probabilistic arguments. We also remark that the permutations $p_1, \ldots, p_r$ are fixed and known to the adversary, thus the only secret randomness in the construction comes from the round keys s_i.

We start with some new notations. We encourage the reader to consult the full version [40] for tail bounds that are extensively used in our analysis.

4.1 Definitions and Notations

Let $\mathfrak{D}$ denote the domain and let $2^n = N := |\mathfrak{D}|$. Throughout this report, we will consider many distribution of permutations over $\mathfrak{D}$. Permutation distributions will be denoted by calligraphic letters (e.g. $\mathcal{F}, \mathcal{G}, \mathcal{H}$). A random choice of a permutation from such a distribution will act as a key for the KAC. Here is a simple example of permutation distributions:

Example 1 (Shift permutations). Denoted by $\mathcal{S}$, the uniform distribution over

$$\{\sigma_s : x \mapsto x + s \mid s \in \mathfrak{D}\},$$

which consists of all shift permutations σ_s that additively shifts the input by s. The definition assumes $\mathfrak{D}$ to be a group. The support of $\mathcal{S}$ is of size N.

We now define a notation for composition of permutations, the cornerstone of the KAC construction.

Definition 5 (Composition). *Let $\mathcal{F}, \mathcal{G}$ be distributions over permutations, and let p be a permutation over $\mathfrak{D}$. Their compositions are defined as*

$$\begin{aligned}
&\mathcal{F} \circ p \text{ is the distribution of } f \circ p \text{ where } f \leftarrow \mathcal{F},\\
&p \circ \mathcal{G} \text{ is the distribution of } p \circ g \text{ where } g \leftarrow \mathcal{G},\\
&\mathcal{F} \circ \mathcal{G} \text{ is the distribution of } f \circ g \text{ where } f \leftarrow \mathcal{F}, g \leftarrow \mathcal{G} \text{ independently.}
\end{aligned}$$

Key Alternating Cipher. Given the language of permutation distributions from above, we can give an alternative definition of key-alternating ciphers (KACs). A t-round KAC is parametered by fixed permutations $p_1, \ldots, p_{t-1}$, and is the composition

$$\mathcal{S} \circ p_1 \circ \mathcal{S} \circ p_2 \circ \mathcal{S} \circ p_3 \circ \cdots \circ p_{t-1} \circ \mathcal{S}.$$

In words, this means picking t round-keys $s_1, \ldots, s_t \leftarrow \mathfrak{D}$ and letting

$$f_{s_1,\ldots,s_t}(x) = s_t + \underbrace{p_{t-1}(s_{t-1} + p_{t-2}(s_{t-2} + \ldots))}_{\text{repeated } t-1 \text{ times}}$$

as illustrated in Figure 1.

Pointwise Closeness to t-wise Independence. Finally, we define the notion of being pointwise close to t-wise independent which we achieve. It is a stronger notion than being close to t-wise independent (Definition 3), a notion that we worked with in Section 4. This only makes the results of this section stronger.

Definition 6 (pointwise close to t-wise independence). *Let $\mathcal{F}$ be a distribution over permutations. $\mathcal{F}$ is pointwise ϵ-close to t-wise independence if for any distinct $x_1, \ldots, x_t \in \mathfrak{D}$ and any distinct $y_1, \ldots, y_t \in \mathfrak{D}$,*

$$\Pr_{f \leftarrow \mathcal{F}}\Big[f(x_1) = y_1 \wedge f(x_2) = y_2 \wedge \cdots \wedge f(x_t) = y_t\Big] \in \Big(\frac{1-\epsilon}{N^{\underline{t}}}, \frac{1+\epsilon}{N^{\underline{t}}}\Big).$$

4.2 Existential Results for Key Alternating Ciphers

In this section, we will prove our main existential result, that is, for some $r = t + o(t) + s$, there exist permutations $p_1, \ldots, p_r$ such that a r-round KAC using these permutations is $\exp(-s)$-close to t-wise independent.

The result is proved by a careful induction that combines two steps.

- *Independence Amplification:* Lemma 20 shows that if $\mathcal{F}$ is pointwise ε-close to t-wise independent, then $\mathcal{S}\circ p\circ\mathcal{F}$ is pointwise $(c(1+\varepsilon)t^2\log N)$-close to $(t+1)$-wise independent, for most permutations p and for some constant $c>1$. In other words, one more KAC round takes you from *very t-wise independent* to *somewhat $(t+1)$-wise independent*. It is important to note that even though the distance of the resulting permutation is $c(1+\varepsilon)t^2\log N \gg 1$, this is still a non-trivial pointwise guarantee.
  In fact, one can inductively apply Lemma 20 and conclude that t-round KAC is pointwise $((t!)^2(c\log N)^{t-1})$-close to t-wise independence, starting from just 1-wise independence. As mentioned before, although the distance is much larger than 1, this is a non-trivial statement, because it is about pointwise closeness.
- *Distance Amplification:* Lemma 21 will reduce the distance to t-wise independence by adding more rounds. Say $\mathcal{F}$ is pointwise ε-close to t-wise independent *and* is pointwise ε'-close to $(t+1)$-wise independent, where $\varepsilon' \gg \varepsilon$. I.e., $\mathcal{F}$ is very close to t-wise independent and somewhat close to $(t+1)$-wise independent. Lemma 21 shows that adding one more round makes it much closer to $(t+1)$-wise independent. More formally, $\mathcal{S}\circ p\circ\mathcal{F}$ is pointwise $\left(\varepsilon+\tilde{O}(\frac{\varepsilon' t}{\sqrt[3]{N}})\right)$-close to $(t+1)$-wise independent, for most permutations p.

Iterated applications of Lemmas 20 and 21 takes us very close to t-wise independence in $2t$ rounds. Indeed, it is not hard to see that one can do even better: between any two successive applications of distance amplification, one can afford to do a large number ($\approx \log N/\log\log N$ many) of iterations of independence amplification. Therefore, to get to t-wise independence, it suffices to work with a $(t+o(t))$-round KAC.

For example, 1-round KAC is 1-wise independent. Then, 2-round KAC is $O(\log N)$-close to 2-wise independent, due to Lemma 20. By adding one more round, Lemma 21 shows that 3-round KAC is $O(\frac{\log N}{N})$-close to 2-wise independent. Figure 6 illustrates the progression of the inductive argument.

More generally, we show:

Theorem 6 (Main KAC Theorem). *For every t, let $r = t + o(t)$. There exist fixed permutations $p_1,\ldots,p_r$ such that the r-round key-alternating cipher is $1/N^{\Omega(1)}$-close to t-wise independent.*

The theorem follows from Lemma 20 and Lemma 21 below whose proofs are deferred to the full version [40]. Finally, we remark that the proof of the theorem shows more: that an overwhelming fraction of choices of permutations $p_1,\ldots,p_r$ gives us a t-wise independent KAC.

Lemma 20. *Let $\mathcal{F}$ be a distribution which is pointwise ε-close to ℓ-wise independence. At least $1-1/N^{t+1}$ of the possible permutations p satisfy the property that $\mathcal{S}\circ p\circ\mathcal{F}$ is pointwise $O((1+\varepsilon)(t+1)^2\log N)$-close to $(t+1)$-wise independence.*

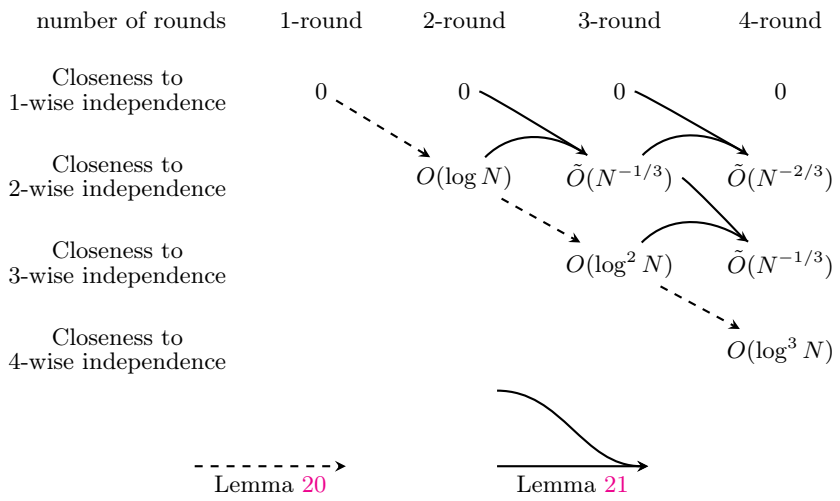

Fig. 6. Illustration of the Inductive Proof using Lemmas 20, 21.

Lemma 21. *Let $\mathcal{F}$ be a permutation distribution that is pointwise ε-close to t-wise independence and is pointwise ε'-close to $(t+1)$-wise independence. At least $1 - 1/N^{t+1}$ of the possible permutations p satisfy the property that $\mathcal{S} \circ p \circ \mathcal{F}$ is pointwise $\left(\varepsilon + 4\varepsilon'(t+1)\sqrt[3]{\ln N/N}\right)$-close to $(t+1)$-wise independence.*

Acknowledgments. The third author thanks Charlie Rackoff for inspiring discussions a decade ago on the topic of what can be proved about the security of block ciphers. He also thanks Orr Dunkelman and Thomas Peyrin for answering his questions about the role of key schedule in the security of AES. We thank the Simons Institute for hosting us at the "Lattices: Algorithms, Complexity and Cryptography" program where the seeds of this work were planted. TL was supported by NSF grants CNS-1528178, CNS-1929901, CNS-1936825 (CAREER), CNS-2026774, a JP Morgan AI research Award, and a Simons Foundation Collaboration Grant on Algorithmic Fairness. ST was supported in part by NSF grants CNS-1930117 (CAREER), CNS-1926324, CNS-2026774, a Sloan Research Fellowship, and a JP Morgan Faculty Award. VV was supported by DARPA under Agreement No. HR00112020023, a grant from the MIT-IBM Watson AI, a grant from Analog Devices, a Microsoft Trustworthy AI grant, and a DARPA Young Faculty Award.

References

1. Martin R. Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. Mimc: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In *ASIACRYPT (1)*, volume 10031 of *Lecture Notes in Computer Science*, pages 191–219, 2016.

2. Noga Alon and Shachar Lovett. Almost k-wise vs. k-wise independent permutations, and uniformity for general group actions. *Theory Comput.*, 9:559–577, 2013.
3. Elena Andreeva, Andrey Bogdanov, Yevgeniy Dodis, Bart Mennink, and John P. Steinberger. On the indifferentiability of key-alternating ciphers. In *CRYPTO 2013*, volume 8042 of *Lecture Notes in Computer Science*, pages 531–550. Springer, 2013.
4. Thomas Baignères and Serge Vaudenay. Proving the security of AES substitution-permutation network. In *Selected Areas in Cryptography*, volume 3897 of *Lecture Notes in Computer Science*, pages 65–81. Springer, 2005.
5. Eli Biham and Adi Shamir. Differential cryptanalysis of des-like cryptosystems. *J. Cryptol.*, 4(1):3–72, 1991.
6. Céline Blondeau and Kaisa Nyberg. Perfect nonlinear functions and cryptography. *Finite Fields Their Appl.*, 32:120–147, 2015.
7. Andrey Bogdanov, Dmitry Khovratovich, and Christian Rechberger. Biclique cryptanalysis of the full AES. In *ASIACRYPT*, volume 7073 of *Lecture Notes in Computer Science*, pages 344–371. Springer, 2011.
8. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, François-Xavier Standaert, John P. Steinberger, and Elmar Tischhauser. Key-alternating ciphers in a provable setting: Encryption using a small number of public permutations - (extended abstract). In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 45–62. Springer, 2012.
9. Alex Brodsky and Shlomo Hoory. Simple permutations mix even better. *Random Struct. Algorithms*, 32(3):274–289, 2008.
10. Andrea Caranti, Francesca Dalla Volta, and Massimiliano Sala. An application of the o'nan-scott theorem to the group generated by the round functions of an aes-like cipher. *Des. Codes Cryptogr.*, 52(3):293–301, 2009.
11. Shan Chen, Rodolphe Lampe, Jooyoung Lee, Yannick Seurin, and John P. Steinberger. Minimizing the two-round even-mansour cipher. In *CRYPTO (1)*, volume 8616 of *Lecture Notes in Computer Science*, pages 39–56. Springer, 2014.
12. Shan Chen and John P. Steinberger. Tight security bounds for key-alternating ciphers. In *EUROCRYPT*, volume 8441 of *Lecture Notes in Computer Science*, pages 327–350. Springer, 2014.
13. Benoît Cogliati, Yevgeniy Dodis, Jonathan Katz, Jooyoung Lee, John P. Steinberger, Aishwarya Thiruvengadam, and Zhe Zhang. Provable security of (tweakable) block ciphers based on substitution-permutation networks. In *CRYPTO 2018*, volume 10991 of *Lecture Notes in Computer Science*, pages 722–753. Springer, 2018.
14. Benoît Cogliati and Jooyoung Lee. Wide tweakable block ciphers based on substitution-permutation networks: Security beyond the birthday bound. *IACR Cryptol. ePrint Arch.*, 2018:488, 2018.
15. Don Coppersmith and Edna Grossman. Generators for certain alternating groups with applications to crytography. *SIAM Journal on Applied Mathematics*, 29(4):624–627, 1975.
16. Joan Daemen. Cipher and hash function design strategies based on linear and differential cryptanalysis, 1995. Ph.D. Thesis, KU Leuven.
17. Joan Daemen and Vincent Rijmen. Understanding two-round differentials in AES. In *SCN 2006*, volume 4116 of *Lecture Notes in Computer Science*, pages 78–94. Springer, 2006.
18. Yevgeniy Dodis, Jonathan Katz, John P. Steinberger, Aishwarya Thiruvengadam, and Zhe Zhang. Provable security of substitution-permutation networks. *IACR Cryptol. ePrint Arch.*, 2017:16, 2017.

19. Yevgeniy Dodis, Martijn Stam, John P. Steinberger, and Tianren Liu. Indifferentiability of confusion-diffusion networks. In *EUROCRYPT 2016*, volume 9666 of *Lecture Notes in Computer Science*, pages 679–704. Springer, 2016.
20. Chun Guo and Dongdai Lin. On the indifferentiability of key-alternating feistel ciphers with no key derivation. In *TCC (1)*, volume 9014 of *Lecture Notes in Computer Science*, pages 110–133. Springer, 2015.
21. Chun Guo and Dongdai Lin. A synthetic indifferentiability analysis of interleaved double-key even-mansour ciphers. In *ASIACRYPT (2)*, volume 9453 of *Lecture Notes in Computer Science*, pages 389–410. Springer, 2015.
22. Viet Tung Hoang and Stefano Tessaro. Key-alternating ciphers and key-length extension: Exact bounds and multi-user security. In *CRYPTO (1)*, volume 9814 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2016.
23. Seokhie Hong, Sangjin Lee, Jongin Lim, Jaechul Sung, Dong Hyeon Cheon, and Inho Cho. Provable security against differential and linear cryptanalysis for the SPN structure. In *FSE*, volume 1978 of *Lecture Notes in Computer Science*, pages 273–283. Springer, 2000.
24. Shlomo Hoory, Avner Magen, Steven A. Myers, and Charles Rackoff. Simple permutations mix well. *Theor. Comput. Sci.*, 348(2-3):251–261, 2005.
25. Thomas Jakobsen and Lars R. Knudsen. The interpolation attack on block ciphers. In *FSE*, volume 1267 of *Lecture Notes in Computer Science*, pages 28–40. Springer, 1997.
26. Burton S. Kaliski Jr., Ronald L. Rivest, and Alan T. Sherman. Is the data encryption standard a group? (results of cycling experiments on DES). *J. Cryptol.*, 1(1):3–36, 1988.
27. Ju-Sung Kang, Seokhie Hong, Sangjin Lee, Okyeon Yi, Choonsik Park, and Jongin Lim. Practical and provable security against differential and linear cryptanalysis for substitution-permutation networks. *Etri Journal*, 23, 02 2002.
28. Eyal Kaplan, Moni Naor, and Omer Reingold. Derandomized constructions of k-wise (almost) independent permutations. In *APPROX-RANDOM*, volume 3624 of *Lecture Notes in Computer Science*, pages 354–365. Springer, 2005.
29. Liam Keliher. Refined analysis of bounds related to linear and differential cryptanalysis for the AES. In *AES 2004*, volume 3373 of *Lecture Notes in Computer Science*, pages 42–57. Springer, 2004.
30. Liam Keliher, Henk Meijer, and Stafford E. Tavares. Improving the upper bound on the maximum average linear hull probability for rijndael. In *Selected Areas in Cryptography*, volume 2259 of *Lecture Notes in Computer Science*, pages 112–128. Springer, 2001.
31. Liam Keliher, Henk Meijer, and Stafford E. Tavares. New method for upper bounding the maximum average linear hull probability for spns. In *EUROCRYPT*, volume 2045 of *Lecture Notes in Computer Science*, pages 420–436. Springer, 2001.
32. Liam Keliher and Jiayuan Sui. Exact maximum expected differential and linear probability for two-round advanced encryption standard. *IET Inf. Secur.*, 1(2):53–57, 2007.
33. Lars Knudsen. Deal - a 128-bit block cipher. In *NIST AES Proposal*, 1998.
34. Lars R. Knudsen. Truncated and higher order differentials. In Bart Preneel, editor, *Fast Software Encryption: Second International Workshop. Leuven, Belgium, 14-16 December 1994, Proceedings*, volume 1008 of *Lecture Notes in Computer Science*, pages 196–211. Springer, 1994.
35. Lars R. Knudsen and David A. Wagner. Integral cryptanalysis. In *FSE*, volume 2365 of *Lecture Notes in Computer Science*, pages 112–127. Springer, 2002.

36. Xuejia Lai. *Higher Order Derivatives and Differential Cryptanalysis*, pages 227–233. Springer US, Boston, MA, 1994.
37. Xuejia Lai, James L. Massey, and Sean Murphy. Markov ciphers and differential cryptanalysis. In *EUROCRYPT*, volume 547 of *Lecture Notes in Computer Science*, pages 17–38. Springer, 1991.
38. Rodolphe Lampe and Yannick Seurin. Security analysis of key-alternating feistel ciphers. In *FSE*, volume 8540 of *Lecture Notes in Computer Science*, pages 243–264. Springer, 2014.
39. Rudolf Lidl and Harald Niederreiter. *Introduction to Finite Fields and Their Applications*. Cambridge University Press, USA, 1986.
40. Tianren Liu, Stefano Tessaro, and Vinod Vaikuntanathan. The t-wise independence of substitution-permutation networks. *IACR Cryptol. ePrint Arch.*, 2021:507, 2021.
41. Mitsuru Matsui and Atsuhiro Yamagishi. A new method for known plaintext attack of FEAL cipher. In *EUROCRYPT*, volume 658 of *Lecture Notes in Computer Science*, pages 81–91. Springer, 1992.
42. Ueli M. Maurer, Krzysztof Pietrzak, and Renato Renner. Indistinguishability amplification. In *CRYPTO 2007*, volume 4622 of *Lecture Notes in Computer Science*, pages 130–149. Springer, 2007.
43. Eric Miles and Emanuele Viola. Substitution-permutation networks, pseudorandom functions, and natural proofs. *J. ACM*, 62(6):46:1–46:29, 2015.
44. Sean Murphy, Kenneth G. Paterson, and Peter R. Wild. A weak cipher that generates the symmetric group. *J. Cryptol.*, 7(1):61–65, 1994.
45. Kaisa Nyberg. Differentially uniform mappings for cryptography. In *EUROCRYPT '93*, volume 765 of *Lecture Notes in Computer Science*, pages 55–64. Springer, 1993.
46. Kaisa Nyberg and Lars R. Knudsen. Provable security against a differential attack. *J. Cryptol.*, 8(1):27–37, 1995.
47. Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, USA, 2014.
48. Sangwoo Park, Soo Hak Sung, Seongtaek Chee, E-Joong Yoon, and Jongin Lim. On the security of rijndael-like structures against differential and linear cryptanalysis. In *ASIACRYPT*, volume 2501 of *Lecture Notes in Computer Science*, pages 176–191. Springer, 2002.
49. Sangwoo Park, Soo Hak Sung, Sangjin Lee, and Jongin Lim. Improving the upper bound on the maximum differential and the maximum linear hull probability for SPN structures and AES. In *FSE*, volume 2887 of *Lecture Notes in Computer Science*, pages 247–260. Springer, 2003.
50. John P. Steinberger. Improved security bounds for key-alternating ciphers via hellinger distance. *IACR Cryptol. ePrint Arch.*, 2012:481, 2012.
51. Serge Vaudenay. Decorrelation: A theory for block cipher security. *J. Cryptol.*, 16(4):249–286, 2003.
52. Yusai Wu, Liqing Yu, Zhenfu Cao, and Xiaolei Dong. Tight security analysis of 3-round key-alternating cipher with a single permutation. In *ASIACRYPT (1)*, volume 12491 of *Lecture Notes in Computer Science*, pages 662–693. Springer, 2020.