Provable Security of (Tweakable) Block Ciphers Based on Substitution-Permutation Networks

Benoît Cogliati[1], Yevgeniy Dodis[2], Jonathan Katz[3], Jooyoung Lee[4], John Steinberger[5], Aishwarya Thiruvengadam[6], and Zhe Zhang[7]

[1] University of Luxembourg, Luxembourg, benoitcogliati@hotmail.fr
[2] New York University, USA, dodis@cs.nyu.edu
[3] University of Maryland, USA, jkatz@cs.umd.edu
[4] KAIST, Korea, hicalf@kaist.ac.kr
[5] jpstein@gmail.com
[6] University of California, Santa Barbara, aish@cs.ucsb.edu
[7] Tsinghua University, Beijing, leomwzz@gmail.com

Abstract. *Substitution-Permutation Networks* (SPNs) refer to a family of constructions which build a wn-bit block cipher from n-bit public permutations (often called S-boxes), which alternate keyless and "local" substitution steps utilizing such S-boxes, with keyed and "global" permutation steps which are non-cryptographic. Many widely deployed block ciphers are constructed based on the SPNs, but there are essentially no provable-security results about SPNs.

In this work, we initiate a comprehensive study of the provable security of SPNs as (possibly tweakable) wn-bit block ciphers, when the underlying n-bit permutation is modeled as a public random permutation. When the permutation step is *linear* (which is the case for most existing designs), we show that 3 SPN rounds are necessary and sufficient for security. On the other hand, even 1-round SPNs can be secure when non-linearity is allowed. Moreover, 2-round non-linear SPNs can achieve "beyond-birthday" (up to $2^{2n/3}$ adversarial queries) security, and, as the number of non-linear rounds increases, our bounds are meaningful for the number of queries approaching 2^n. Finally, our non-linear SPNs can be made *tweakable* by incorporating the tweak into the permutation layer, and provide good multi-user security.

As an application, our construction can turn two public n-bit permutations (or fixed-key block ciphers) into a tweakable block cipher working on wn-bit inputs, $6n$-bit key and an n-bit tweak (for any $w \geq 2$); the tweakable block cipher provides security up to $2^{2n/3}$ adversarial queries in the random permutation model, while only requiring w calls to each permutation, and $3w$ field multiplications for each wn-bit input.

Keywords: substitution-permutation networks, tweakable block ciphers, domain extension of block ciphers, beyond-birthday-bound security

1 Introduction

Substitution-Permutation Networks. Modern block ciphers are generally constructed using two main paradigms [KL15]: Feistel networks [Fei73] or

substitution-permutation networks (SPNs) [Sha49, Fei73]. Examples of block ciphers based on Feistel networks include DES, FEAL, MISTY and KASUMI; block ciphers based on SPNs include AES, Serpent, and PRESENT. These two approaches share the same goal: namely, to extend a "pseudorandom object" on a small domain to a (keyed) pseudorandom permutation on a larger domain by repeating a few, relatively simple operations several times across multiple rounds. Simplifying somewhat, Feistel networks begin with a keyed pseudorandom function on n-bit inputs and extend this to give a keyed pseudorandom permutation on $2n$-bit inputs. On the other hand, SPNs start with one or more public "random permutations" on n-bit inputs (called S-boxes) and extend them to give a keyed pseudorandom permutation on wn-bit inputs for some w, by iterating the following steps:

1. *Substitution step*: break down the wn-bit state into w disjoint n-bit blocks, and compute an S-box on each n-bit block;
2. *Permutation step*: apply a non-cryptographic, keyed permutation to the whole wn-bit state (which is also applied to the plaintext before the first round).

Proving the security of a concrete block cipher unconditionally is currently beyond our capabilities. Thus, the usual approach is to prove that the high-level structure is sound in a relevant security model. For Feistel networks, a substantial line of work, starting with Luby and Rackoff's seminal work [LR88], and culminating with Patarin's results [Pat03, Pat04], proves optimal security with a sufficient number of rounds. Numerous other articles [Pat10, HR10, HKT11, Tes14, CHK+16] study the security of (variants of) Feistel networks in various security models. In contrast, it is somewhat surprising that there are almost no results about provable security of SPNs (see below.) Here, we address this gap and explore conditions under which SPNs can be proven secure.

Domain Extension of Block Ciphers. Block ciphers following the SPNs typically rely on very small S-boxes (e.g. AES uses an 8-bit S-box). However, it is also possible to use a larger domain block cipher with a fixed key (which has non-trivial efficiency gains and avoid related key attacks) as "S-box" in order to extend the domain of the underlying block cipher, or to use a larger dedicated permutation (e.g., Keccak permutation [BDPA09] or with Gimli [BKL+17]), in order to directly obtain a "wide" block cipher. From this point of view, the substitution-permutation networks can also be viewed as enciphering modes of operation (of a fixed input length), in which the length n of the S-box is not necessarily small. Such enciphering modes of operations have applications to disk encryption that protects the confidentiality of data stored on a sector-addressable device, such as a hard disk. In this scenario, the disk is divided into several sectors, and each sector, viewed as a wide block, should be encrypted and decrypted independently of each other. Non-linear 1-round SPNs with *secret* S-boxes have already been used to provide domain extension for block ciphers [CS06, Hal07]. These constructions provide the birthday bound security, while this level of security might not be desirable for an environment where stronger security is required. One of our results will address this limitation.

1.1 Our Contribution

We analyze SPNs in the standard sense as a strong pseudorandom permutation [LR88] (i.e., against adaptive chosen-plaintext and chosen-ciphertext attacks).

Linear SPNs. We first characterize the security of *linear* SPNs, where the permutation layer is a linear function (over $GF(2^n)$, where n is the size of the S-box) of the current wn-bit round key and the current wn-bit state. Indeed, most current SPN-based block ciphers (e.g., AES, Serpent, PRESENT, etc.) use linear permutation step, which involves a simple key-mixing step followed by an invertible linear transformation. For this widely used setting we give a general against any 2-round linear SPNs with $w \geq 2$.[8] Complementing this attack, we show that a 3-round linear SPNs *are* secure, for any w, if the keyed linear permutations satisfy some very mild technical requirements. This result critically uses the H-coefficients technique [Pat08, CS14].

Non-Linear SPNs. In an effort to reduce the number of rounds (and get other benefits we explain below), we then turn our attention to non-linear SPNs, where the permutation step does not have to be linear (although must remain efficient and "non-cryptographic"). Here we show that even a *1-round SPN can be secure*, if appropriate keyed permutations are used. We identify a combinatorial property on the permutations — which we term *blockwise universality* — that suffices for security in this case, and then study the efficiency of constructing permutations satisfying this property. Specifically, we show a construction of a satisfactory permutation with n-bit keys (but having high degree), and another construction with longer keys but having degree 3.

We then show that, by using such blockwise independent permutations, the security of resulting SPNs increases when we increase the number of rounds: while 1 round already achieves "birthday security", as our main technical result we show that 2-round non-linear SPNs (with independent S-boxes and keys in different rounds) achieves "beyond-birthday" security (for up to $2^{2n/3}$ queries). This result uses the refinement of the H-coefficient technique due to [HT16]. We also give an asymptotic analysis of non-linear SPNs built from blockwise universal permutations using the coupling technique of [MRS09, HR10]. In particular, for $r = 2s$ we prove that r-round SPNs are secure as long as the number of adversarial queries is well below $2^{sn/(s+1)}$. Thus, as r grows, our bounds tend towards optimal 2^n security.

As an additional benefit of this setting, we show that the blockwise universal permutations can be efficiently tweaked, meaning that our non-linear SPN constructions yield *tweakable block ciphers* [LRW11], which is important for some settings. Finally, we analyze our non-linear SPNs in the multi-user setting using the point-wise proximity technique of [HT16].

[8] Even a 1-round linear SPN can be secure if $w = 1$, since this corresponds to the famous Even-Mansour cipher [EM97].

Application to Wide Tweakable Block Ciphers. Besides providing theoretical insights on SPN-based block ciphers, our results also have a practical interest in the context of domain extension for block ciphers and permutation-based cryptography. For example, if our construction is instantiated with two n-bit permutations and a tweakable permutation TBPE in the permutation layer (as defined in Section 2.2), then we can build a wide tweakable block cipher with key space $\{0,1\}^{6n}$, tweak space $\{0,1\}^n$ and message space $\{0,1\}^{wn}$ for any integer $w \geq 2$. This tweakable block cipher requires w calls to each permutation and $3w$ field multiplications for each encryption/decryption call. The multi-user advantage of any adversary is shown to be small as long as the number of its queries is well below $2^{2n/3}$. This means that a 192 bit (resp. 384 bit) permutation or block cipher is sufficient to get a provably secure mode of operation as long as the number of adversarial queries is small in front of 2^{128} (resp. 2^{256}). As far as we know, this is the first construction for domain extension of a block cipher/permutation that enjoys beyond birthday-bound security.

Of course, to instantiate this construction we would need a good public permutation with large domain size n. As mentioned earlier, we could either use a larger domain block cipher with a fixed key, or use a larger dedicated permutation on larger domain, such as Keccak permutation [BDPA09] or Gimli [BKL+17].

Open Problems. We conjecture that r-round non-linear SPNs should actually be enough to prove security up to $\mathcal{O}(2^{rn/(r+1)})$ adversarial queries. Proving it using combinatorial techniques seems very challenging and we leave it as an interesting open problem. It is also interesting if we can prove beyond-birthday security bounds for linear SPNs (with 3 or more rounds), as these SPNs appear to be the ones used in practice. More generally, it would be great to prove tight security bounds and matching attacks for r-round linear and non-linear SPNs.

Implications for small block size. While our results are directly meaningful when the length n of public S-boxes in at least security parameter (e.g., for building *wide* tweakable block ciphers), our bounds are too weak for regular SPN-based ciphers, such as AES, which use very low values of n for their S-boxes. This "2^n provable barrier" is inherent using our current modeling, where the S-box of size 2^n is providing the only source of cryptographic hardness. More generally, establishing a sound theory of building block ciphers from small S-boxes is one of the biggest and most important open problems in symmetric-key cryptography. We hope that our structural results for reduced-round SPN ciphers will be useful in establishing such theory, despite not crossing the fundamental "2^n barrier" mentioned above.

1.2 Related Work

There are only a few prior papers looking at provable security of SPNs. The vast majority of such work analyzes the case of secret, key-dependent S-boxes (rather than public S-boxes as we consider here), and so we survey that work first.

SPNs WITH SECRET S-BOXES. Naor and Reingold [NR99] prove security for what can be viewed as a non-linear, 1-round SPN. Their ideas were further developed, in the context of domain extension for block ciphers (see further discussion below), by Chakraborty and Sarkar [CS06] and Halevi [Hal07].

Iwata and Kurosawa [IK00] analyze SPNs in which the linear permutation step is based on the specific permutations used in the block cipher Serpent. They show an attack against 2-round SPNs of this form, and prove security for 3-round SPNs against non-adaptive adversaries. In addition to the fact that we consider public S-boxes, our linear SPN model considers generic linear permutations and we prove security against adaptive attackers.

Miles and Viola [MV15] study SPNs from a complexity-theoretic viewpoint. Two of their results are relevant here. First, they analyze the security of linear SPNs using S-boxes that are not necessarily injective (so the resulting keyed functions are not, in general, invertible). They show that r-round SPNs of this type (for $r \geq 2$) are secure against chosen-plaintext attacks. (In contrast, our results show that 2-round, linear SPNs are not secure against a combination of chosen-plaintext and chosen-ciphertext attacks when $w \geq 2$.) They also analyze SPNs based on a concrete set of S-boxes, but in this case they only show security against linear/differential attacks (a form of chosen-plaintext attack), rather than all possible attacks, and only when the number of rounds is $r = \Theta(\log n)$.

SPNs WITH PUBLIC S-BOXES. A difference between our work and all the work discussed above is that we treat the S-boxes as public. We are aware of only one prior work analyzing the provable security of SPNs in this setting. Dodis et al. [DSSL16] recently studied the *indifferentiability* [MRH04] of confusion-diffusion networks, which can be viewed as *unkeyed* SPNs. One could translate their results to the keyed setting, but that would require using multiple, key-dependent S-boxes (rather than a fixed, public S-box) and so would not imply our results. We remark further that they show positive results only for 5 rounds and above.

As observed earlier, the Even-Mansour construction [EM97] of a (keyed) pseudorandom permutation from a public random permutation can be viewed as a 1-round, linear SPN in the degenerate case where $w = 1$ (i.e., no domain extension) and all round permutations are instantiated using simple key mixing. Security of the 1-round Even-Mansour construction against adaptive chosen-plaintext/ciphertext attacks, using independent keys for the initial and final key mixing, was shown in the original paper [EM97]. Our positive results imply security of the 1-round Even-Mansour construction (with similar concrete security bounds) as a special case. The r-round generalization of the Even-Mansour cipher has seen a lot of interest over the years, culminating with [CS14, HT16] where it was proved that the r-round Even-Mansour construction is secure up to roughly $2^{rn/(r+1)}$ adversarial queries, when the public S-boxes are uniformly random and independent permutations and the round keys are independent. Chen et al. [CLL+14] also proved that several minimized variants of the 2-round Even-Mansour construction are also secure up to roughly $2^{2n/3}$ adversarial queries. None of these results extend to the setting $w > 1$ considered in this work.

Cryptanalysis of SPNs. Researchers have also explored cryptanalytic attacks on generic SPNs [BS10, BBK14, DDKL, BK]. These works generally consider a model of SPNs in which round permutations are secret, random (invertible) linear transformations, and S-boxes may be secret as well; this makes the attacks stronger but positive results weaker. In many cases the complexities of the attacks are exponential in n (though still faster than a brute-force search for the key), and hence do not rule out asymptotic security results. On the positive side, Biryukov et al. [BBK14] show that 2-round SPNs (of the stronger form just mentioned) are secure against some specific types of attacks, but other attacks on such schemes have recently been identified [DDKL].

Attacks. Attacks due to Joux [Jou03] and to Halevi and Rogaway [HR04], originally developed in the afore-mentioned context of block cipher domain extension (or more exactly, in the construction of *tweakable* block ciphers with large domains from standard block ciphers with "small" domains) can be translated to the context of linear SPNs as well. Specifically, these attacks imply that linear 2-round SPNs of width $w \geq 2$ are insecure, as long as the underlying field has characteristic 2.[9]

Domain extension of block ciphers. Non-linear, 1-round SPNs with secret S-boxes have been used for domain extension of block ciphers before [CS06, Hal07]. Other approaches for domain extension, not relying on (pure) SPNs, have also been considered [BD99, HR03, HR04, MF07, CDMS10]. To the best of our knowledge, none of these results achieve beyond-birthday security.

Random Permutation Based Tweakable Block Ciphers. Our tweakable SPNs can be viewed as tweakable block ciphers based on public random permutations. It is easy to see that $T : (h, t, x) \mapsto x \oplus h(t)$ is (δ, δ')-blockwise universal (as defined in Section 2) if h is chosen from a δ'-almost uniform and δ-almost XOR-universal hash family. So with this permutation layer (and with $w = 1$), we obtain the security bound for the Tweakable Even-Mansour constructions [CLS15] in the multi-user setting. In this line of research, a number of efficient constructions have been proposed [GJMN16, Men16].

2 Preliminaries

Throughout this work, we fix positive integers w and n; an element x in $\{0,1\}^{wn}$ can be viewed as a concatenation of w blocks, each of which is of length n. The i-th block of this representation will be denoted x_i for $i = 1, \ldots, w$, so we have

$$x = x_1||x_2||\cdots||x_w,$$

sometimes written as $x = (x_1, \ldots, x_w)$.

[9] Indeed, a technical difference with the attack presented here is that our attack does not require a finite field of characteristic 2. Because of this difference, our attack ends up having little (if anything) in common with the attacks of Joux and Halevi-Rogaway.

For a set R and an integer $s \geq 1$, R^{*s} denotes the set of all sequences that consists of s pairwise distinct elements of R. For any integer r such that $r \geq s$, we will write $(r)_s = r!/(r-s)!$. If $|R| = r$, then $(r)_s$ becomes the size of R^{*s}. The sets of non-negative integers and non-negative real numbers are denoted $\mathbb{N}$ and $\mathbb{R}^{\geq 0}$, respectively. The following inequality will be used in our security proof.

Lemma 1. *Let m be an integer and let x be a real number such that $m \geq 2$ and $-1 \leq x < \frac{1}{m-1}$. Then one has*

$$(1+x)^m \leq 1 + \frac{mx}{1-(m-1)x}.$$

2.1 Tweakable Substitution-Permutation Networks

All the notions below are defined for the general tweak set $\mathcal{T}$; however, the standard "non-tweakable" setting is a special case of the definitions below when $|\mathcal{T}| = 1$.

Tweakable Permutations. For an integer $m \geq 1$, the set of all permutations on $\{0,1\}^m$ will be denoted $\mathsf{Perm}(m)$. A tweakable permutation with tweak space $\mathcal{T}$ and message space $\mathcal{X}$ is a mapping $\widetilde{P} : \mathcal{T} \times \mathcal{X} \to \mathcal{X}$ such that, for any tweak $t \in \mathcal{T}$,

$$x \mapsto \widetilde{P}(t,x)$$

is a permutation of $\mathcal{X}$. The set of all tweakable permutations with tweak space $\mathcal{T}$ and message space $\{0,1\}^m$ will be denoted $\widetilde{\mathsf{Perm}}(\mathcal{T}, m)$.

A keyed tweakable permutation with key space $\mathcal{K}$, tweak space $\mathcal{T}$ and message space $\mathcal{X}$ is a mapping $T : \mathcal{K} \times \mathcal{T} \times \mathcal{X} \to \mathcal{X}$ such that, for any key $k \in \mathcal{K}$,

$$(t,x) \mapsto T(k,t,x)$$

is a tweakable permutation with tweak space $\mathcal{T}$ and message space $\mathcal{X}$. We will sometimes write $T(k,t,x)$ as $T_k(t,x)$ or $T_{k,t}(x)$. For an integer $s \geq 1$, let $\mathbf{t} = (t_1, \ldots, t_s) \in \mathcal{T}^s$, and let $\mathbf{x} = (x_1, \ldots, x_s) \in (\mathcal{X})^{*s}$. We will write $(T(k,t_i,x_i))_{1 \leq i \leq s}$ as $T_k(\mathbf{t}, \mathbf{x})$ or $T_{k,\mathbf{t}}(\mathbf{x})$.

Tweakable SPNs. For fixed parameters w and n, let

$$T : \mathcal{K} \times \mathcal{T} \times \{0,1\}^{wn} \longrightarrow \{0,1\}^{wn}$$

be a keyed tweakable permutation with key space $\mathcal{K}$, tweak space $\mathcal{T}$ and message space $\{0,1\}^{wn}$.

For a fixed number of rounds r, an r-round substitution-permutation network (SPN) based on T, denoted SP^T, takes as input a set of n-bit permutations $\mathcal{S} = (S_1, \ldots, S_r)$, and defines a keyed tweakable permutation $\mathsf{SP}^T[\mathcal{S}]$ operating on wn-bit blocks with key space $\mathcal{K}^{r+1}$ and tweak space $\mathcal{T}$: on input $x \in \{0,1\}^{wn}$, key $\mathbf{k} = (k_0, k_1, \ldots, k_r) \in \mathcal{K}^{r+1}$ and tweak $t \in \mathcal{T}$, the output of $\mathsf{SP}^T[\mathcal{S}]$ is computed as follows (see also Fig. 1).

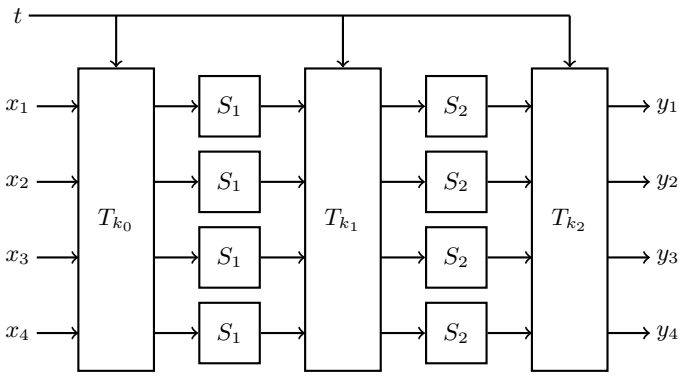

Fig. 1. A 2-round tweakable SPN with $w = 4$. The input and output blocks of the SPN are represented as $x = x_1||x_2||x_3||x_4$ and $y = y_1||y_2||y_3||y_4$, respectively.

$y \leftarrow x$
for $i \leftarrow 1$ to r **do**
  $y \leftarrow T_{k_{i-1},t}(y)$
  Break $y = y_1||\cdots||y_w$ into n-bit blocks
  $y \leftarrow S_i(y_1)||\cdots||S_i(y_w)$
$y \leftarrow T_{k_r,t}(y)$
return y

Remark 1. Both of the permutation layer T and the entire construction SP^T can be viewed as keyed tweakable permutations. However, T will typically be built upon non-cryptographic operations such as filed multiplications, while SP^T are based on S-boxes which are modeled as public random permutations.

Blockwise Universal Tweakable Permutations. A keyed tweakable permutation

$$T : \mathcal{K} \times \mathcal{T} \times \{0,1\}^{wn} \longrightarrow \{0,1\}^{wn}$$

is called (δ, δ')-*blockwise universal* if the following hold.

1. For all distinct (t, x, i), $(t', x', i') \in \mathcal{T} \times \{0,1\}^{wn} \times \{1, \ldots, w\}$, we have
$$\mathsf{Pr}\left[k \xleftarrow{\$} \mathcal{K} : T_{k,t}(x)_i = T_{k,t'}(x')_{i'}\right] \leq \delta.$$
2. For all $(t, x, i, c) \in \mathcal{T} \times \{0,1\}^{wn} \times \{1, \ldots, w\} \times \{0,1\}^n$, we have
$$\mathsf{Pr}\left[k \xleftarrow{\$} \mathcal{K} : T_{k,t}(x)_i = c\right] \leq \delta'.$$

Since each pair of key $k \in \mathcal{K}$ and tweak $t \in \mathcal{T}$ defines a permutation $T_{k,t}$ on $\{0,1\}^{wn}$, one can define a keyed tweakable permutation

$$T^{-1} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^{wn} \longrightarrow \{0,1\}^{wn}$$

such that $T^{-1}(k,t,x) = (T_{k,t})^{-1}(x)$. If T and T^{-1} are both (δ,δ')-blockwise universal, then T is called (δ,δ')-*super blockwise universal*.

2.2 An Efficient Super Blockwise Tweakable Universal Permutation

In this section, we show that an efficient xor-blockwise universal construction, dubbed BPE, proposed by Halevi [Hal07] can be made tweakable with a slight modification. Other constructions of (tweakable) blockwise universal permutations can be found in [DKS+17] some of which support tweaks. We present BPE below and will present the remaining constructions in the full version.

Assuming $2^n \geq w+3$, let $\mathbb{F}$ denote a finite field with 2^n elements. For each $k \in \mathbb{F}$, define a $w \times w$ matrix over $\mathbb{F}$, $M_k =^{\text{def}} A_k + I$, where I is the identity matrix and

$$A_k = \begin{bmatrix} k & k^2 & & k^w \\ k & k^2 & & k^w \\ & & \ddots & \\ k & k^2 & & k^w \end{bmatrix}.$$

Precisely, $(A_k)_{i,j} = k^j$ for $1 \leq i,j \leq w$. Let z be a primitive element of $\mathbb{F}$, and let

$$\mathcal{K} = \left\{ k \in \mathbb{F} : \sum_{i=0}^{w} k^i \neq 0 \right\} \times \mathbb{F}.$$

Then BPE is defined as follows.

$$\begin{aligned} \mathsf{BPE} : \mathcal{K} \times \{0,1\}^{wn} &\longrightarrow \{0,1\}^{wn} \\ ((k,k'),x) &\longmapsto M_k x \oplus a_{k'}, \end{aligned}$$

where we identify $x \in \{0,1\}^{wn}$ with a w-dimensional column vector over $\mathbb{F}$, and

$$a_{k'} = \begin{bmatrix} k' \\ zk' \\ \vdots \\ z^{w-1}k' \end{bmatrix}.$$

It is easy to check that M_k is invertible if $\sum_{i=0}^{w} k^i \neq 0$; precisely,

$$M_k^{-1} = I \oplus \frac{A_k}{k^*},$$

where $k^* =^{\text{def}} \sum_{i=0}^{w} k^i$. For any $(k,k') \in \mathcal{K}$, $\mathsf{BPE}_{k,k'}$ is also invertible with

$$\mathsf{BPE}_{k,k'}^{-1}(x) = M_k^{-1}(x \oplus a_{k'})$$

for any $x \in \{0,1\}^{wn}$. Halevi [Hal07] also proved that for any pair of distinct $(x,i), (x',i') \in \{0,1\}^{wn} \times \{1,\ldots,w\}$ and $\Delta \in \{0,1\}^n$,

$$\begin{aligned} \mathsf{Pr}\left[(k,k') \xleftarrow{\$} \mathcal{K} : \mathsf{BPE}_{k,k'}(x)_i \oplus \mathsf{BPE}_{k,k'}(x')_{i'} = \Delta\right] &\leq \frac{w}{2^n - w}, \\ \mathsf{Pr}\left[(k,k') \xleftarrow{\$} \mathcal{K} : \mathsf{BPE}_{k,k'}^{-1}(x)_i \oplus \mathsf{BPE}_{k,k'}^{-1}(x')_{i'} = \Delta\right] &\leq \frac{w}{2^n - w}. \end{aligned} \tag{1}$$

For a fixed $(x, i, c) \in \{0,1\}^{wn} \times \{1, \ldots, w\} \times \{0,1\}^n$, $\mathsf{BPE}_{k,k'}(x)_i = c$ implies that

$$\sum_{j=1}^{w} x_j k^j \oplus x_i \oplus z^{i-1} k' = c,$$

which holds with probability $\frac{1}{2^n}$ over a random choice of $(k, k') \in \mathcal{K}$. On the other hand, $\mathsf{BPE}^{-1}_{k,k'}(x)_i = c$ implies that

$$\left(z^{i-1} \oplus \frac{1}{k^*} \sum_{j=1}^{w} z^{j-1} k^j \right) k' \oplus \left(c \oplus x_i \oplus \frac{1}{k^*} \sum_{j=1}^{w} x_j k^j \right) = 0.$$

This equation holds with probability at most $\frac{w}{2^n - w} + \frac{1}{2^n}$. To summarize, we have

$$\begin{aligned} \mathsf{Pr}\left[(k, k') \xleftarrow{\$} \mathcal{K} : \mathsf{BPE}_{k,k'}(x)_i = c\right] &\leq \frac{1}{2^n}, \\ \mathsf{Pr}\left[(k, k') \xleftarrow{\$} \mathcal{K} : \mathsf{BPE}^{-1}_{k,k'}(x)_i = c\right] &\leq \frac{w+1}{2^n - w}. \end{aligned} \tag{2}$$

Now we define a tweakable variant of BPE, dubbed TBPE (for Tweakable Blockwise Polynomial-Evaluation), with tweak space $\mathcal{T} = \{0,1\}^n$ as follows.

$$\begin{aligned} \mathsf{TBPE} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^{wn} &\longrightarrow \{0,1\}^{wn} \\ ((k, k'), t, x) &\longmapsto M_k(x \oplus b_t) \oplus a_{k'} \oplus b_t, \end{aligned}$$

where b_t is the column vector whose entries are all t, namely,

$$b_t = \begin{bmatrix} t \\ t \\ \vdots \\ t \end{bmatrix}.$$

Since each pair of key $(k, k') \in \mathcal{K}$ and tweak $t \in \mathcal{T}$ defines a permutation $\mathsf{TBPE}_{k,k',t}$ on $\{0,1\}^{wn}$, one can define a keyed tweakable permutation

$$\mathsf{TBPE}^{-1} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^{wn} \longrightarrow \{0,1\}^{wn}.$$

Then we can prove the following lemma.

Lemma 2. *Let* TBPE *be the keyed tweakable permutation as defined above, and let* TBPE^{-1} *be its inverse.*

1. *For all distinct* (t, x, i), $(t', x', i') \in \mathcal{T} \times \{0,1\}^{wn} \times \{1, \ldots, w\}$, *we have*

$$\mathsf{Pr}\left[(k, k') \xleftarrow{\$} \mathcal{K} : \mathsf{TBPE}_{k,k',t}(x)_i = \mathsf{TBPE}_{k,k',t'}(x')_{i'}\right] \leq \frac{w}{2^n - w}.$$

2. *For all* $(t, x, i, c) \in \mathcal{T} \times \{0,1\}^{wn} \times \{1, \ldots, w\} \times \{0,1\}^n$, *we have*

$$\mathsf{Pr}\left[(k, k') \xleftarrow{\$} \mathcal{K} : \mathsf{TBPE}_{k,k',t}(x)_i = c\right] \leq \frac{1}{2^n}.$$

3. *For all distinct* (t, x, i), $(t', x', i') \in \mathcal{T} \times \{0,1\}^{wn} \times \{1, \ldots, w\}$, *we have*

$$\mathsf{Pr}\left[(k, k') \xleftarrow{\$} \mathcal{K} : \mathsf{TBPE}^{-1}_{k,k',t}(x)_i = \mathsf{TBPE}^{-1}_{k,k',t'}(x')_{i'}\right] \leq \frac{w}{2^n - w}.$$

4. *For all* $(t, x, i, c) \in \mathcal{T} \times \{0,1\}^{wn} \times \{1, \ldots, w\} \times \{0,1\}^n$, *we have*

$$\mathsf{Pr}\left[(k, k') \xleftarrow{\$} \mathcal{K} : \mathsf{TBPE}^{-1}_{k,k',t}(x)_i = c\right] \leq \frac{w+1}{2^n - w}.$$

Proof. For distinct (t, x, i) and (t', x', i'), we have

$$\mathsf{TBPE}_{k,k',t}(x)_i \oplus \mathsf{TBPE}_{k,k',t'}(x')_{i'} = \mathsf{BPE}_{k,k'}(x \oplus b_t)_i \oplus \mathsf{BPE}_{k,k'}(x' \oplus b_{t'})_{i'} \oplus t \oplus t'.$$

If $(x \oplus b_t, i) \neq (x' \oplus b_{t'}, i')$, then $\mathsf{BPE}_{k,k'}(x \oplus b_t)_i \oplus \mathsf{BPE}_{k,k'}(x' \oplus b_{t'})_{i'} \oplus t \oplus t' = 0$ with probability at most $\frac{w}{2^n - w}$ by (1). If $(x \oplus b_t, i) = (x' \oplus b_{t'}, i')$, then it implies $t \neq t'$, and hence $\mathsf{BPE}_{k,k'}(x \oplus b_t)_i \oplus \mathsf{BPE}_{k,k'}(x' \oplus b_{t'})_{i'} \oplus t \oplus t' = t \oplus t' \neq 0$.

For a fixed (t, x, i, c), $\mathsf{TBPE}_{k,k',t}(x)_i = c$ if and only if $\mathsf{BPE}_{k,k'}(x \oplus b_t)_i = c \oplus t$, and this equation holds with probability at most $\frac{1}{2^n}$. The remaining properties are proved similarly. □

From Lemma 2, it follows that TBPE is $\left(\frac{w}{2^n - w}, \frac{w+1}{2^n - w}\right)$-super blockwise universal. Except constant multiplications $z^i k'$, $i = 1, \ldots, w-1$, (which also can be precomputed), each evaluation of $\mathsf{TBPE}_{k,k',t}(x)$ requires w field multiplications.

2.3 Indistinguishability in the Multi-user Setting

Let $\mathsf{SP}^T[\mathcal{S}]$ be an r-round SPN based on a set of S-boxes $\mathcal{S} = (S_1, \ldots, S_r)$ and a keyed tweakable permutation T with key space $\mathcal{K}$ and tweak space $\mathcal{T}$. So $\mathsf{SP}^T[\mathcal{S}]$ becomes a keyed tweakable permutation on $\{0,1\}^{wn}$ with key space $\mathcal{K}^{r+1}$ and tweak space $\mathcal{T}$.

In the multi-user setting, let ℓ denote the number of users. In the *real* world, ℓ secret keys $\mathbf{k}_1, \ldots \mathbf{k}_\ell \in \mathcal{K}^{r+1}$ are chosen independently at random. A set of independent S-boxes $\mathcal{S} = (S_1, \ldots, S_r)$ is also randomly chosen from $\mathsf{Perm}(n)^r$. A distinguisher $\mathcal{D}$ is given oracle access to $(\mathsf{SP}^T_{\mathbf{k}_1}[\mathcal{S}], \ldots, \mathsf{SP}^T_{\mathbf{k}_\ell}[\mathcal{S}])$ as well as $\mathcal{S} = (S_1, \ldots, S_r)$. In the *ideal* world, $\mathcal{D}$ is given a set of independent random tweakable permutations $\widetilde{\mathcal{P}} = (\widetilde{P}_1, \ldots, \widetilde{P}_\ell) \in \widetilde{\mathsf{Perm}}(\mathcal{T}, wn)^\ell$ instead of $(\mathsf{SP}^T_{\mathbf{k}_1}[\mathcal{S}], \ldots, \mathsf{SP}^T_{\mathbf{k}_\ell}[\mathcal{S}])$. However, oracle access to $\mathcal{S} = (S_1, \ldots, S_r)$ is still allowed in this world.

The adversarial goal is to tell apart the two worlds $(\mathsf{SP}^T_{\mathbf{k}_1}[\mathcal{S}], \ldots, \mathsf{SP}^T_{\mathbf{k}_\ell}[\mathcal{S}], \mathcal{S})$ and $(\widetilde{P}_1, \ldots, \widetilde{P}_\ell, \mathcal{S})$ by adaptively making forward and backward queries to each of the constructions and the S-boxes. Formally, $\mathcal{D}$'s distinguishing advantage is defined by

$$\begin{aligned}\mathrm{Adv}^{\mathrm{mu}}_{\mathsf{SP}^T}(\mathcal{D}) = \mathsf{Pr}\left[\widetilde{P}_1, \ldots, \widetilde{P}_\ell \xleftarrow{\$} \widetilde{\mathsf{Perm}}(\mathcal{T}, wn), \mathcal{S} \xleftarrow{\$} \mathsf{Perm}(n)^r : 1 \leftarrow \mathcal{D}^{\mathcal{S}, \widetilde{P}_1, \ldots, \widetilde{P}_\ell}\right] \\ - \mathsf{Pr}\left[\mathbf{k}_1, \ldots, \mathbf{k}_\ell \xleftarrow{\$} \mathcal{K}^{r+1}, \mathcal{S} \xleftarrow{\$} \mathsf{Perm}(n)^r : 1 \leftarrow \mathcal{D}^{\mathcal{S}, \mathsf{SP}^T_{\mathbf{k}_1}[\mathcal{S}], \ldots, \mathsf{SP}^T_{\mathbf{k}_\ell}[\mathcal{S}]}\right].\end{aligned}$$

For p, $q > 0$, we define

$$\mathrm{Adv}_{\mathsf{SP}^T}(p,q) = \max_{\mathcal{D}} \mathrm{Adv}_{\mathsf{SP}^T}(\mathcal{D})$$

where the maximum is taken over all adversaries $\mathcal{D}$ making at most p queries to each of the S-boxes and at most q queries to the outer tweakable permutations. In the single-user setting with $\ell = 1$, $\mathrm{Adv}^{\mathrm{mu}}_{\mathsf{SP}^T}(\mathcal{D})$ and $\mathrm{Adv}^{\mathrm{mu}}_{\mathsf{SP}^T}(p,q)$ will also be written as $\mathrm{Adv}^{\mathrm{su}}_{\mathsf{SP}^T}(\mathcal{D})$ and $\mathrm{Adv}^{\mathrm{su}}_{\mathsf{SP}^T}(p,q)$, respectively.

H-coefficient Technique. Suppose that a distinguisher $\mathcal{D}$ makes p queries to each of the S-boxes, and total q queries to the construction oracles. The queries made to the j-th construction oracle, denoted C_j, are recorded in a query history

$$\mathcal{Q}_{C_j} = (j, t_{j,i}, x_{j,i}, y_{j,i})_{1 \le i \le q_j}$$

for $j = 1, \ldots, \ell$, where q_j is the number of queries made to C_j and $(j, t_{j,i}, x_{j,i}, y_{j,i})$ represents the evaluation obtained by the i-th query to C_j.[10] So according to the instantiation, it implies either $\mathsf{SP}^T_{\mathbf{k}_j}[\mathcal{S}](t_{j,i}, x_{j,i}) = y_{j,i}$ or $\widetilde{P}_j(t_{j,i}, x_{j,i}) = y_{j,i}$. Let

$$\mathcal{Q}_C = \mathcal{Q}_{C_1} \cup \cdots \cup \mathcal{Q}_{C_\ell}.$$

For $j = 1, \ldots, r$, the queries made to S_j are recorded in a query history

$$\mathcal{Q}_{S_j} = (j, u_{j,i}, v_{j,i})_{1 \le i \le p},$$

where $(j, u_{j,i}, v_{j,i})$ represents the evaluation $S_j(u_{j,i}) = v_{j,i}$ obtained by the i-th query to S_j. Let

$$\mathcal{Q}_S = \mathcal{Q}_{S_1} \cup \cdots \cup \mathcal{Q}_{S_r}.$$

Then the pair of query histories

$$\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$$

will be called the *transcript* of the attack: it contains all the information that $\mathcal{D}$ has obtained at the end of the attack. In this work, we will only consider information theoretic distinguishers. Therefore we can assume that a distinguisher is deterministic without making any redundant query, and hence the output of $\mathcal{D}$ can be regarded as a function of τ, denoted $\mathcal{D}(\tau)$ or $\mathcal{D}(\mathcal{Q}_C, \mathcal{Q}_S)$.

Fix a transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$, a key $\mathbf{k} \in \mathcal{K}^{r+1}$, a tweakable permutation $\widetilde{P} \in \widetilde{\mathsf{Perm}}(\mathcal{T}, wn)$, a set of S-boxes $\mathcal{S} = (S_1, \ldots, S_r) \in \mathsf{Perm}(n)^r$ and $j \in \{1, \ldots, \ell\}$: if $S_j(u_{j,i}) = v_{j,i}$ for every $i = 1, \ldots, p$, then we will write $S_j \vdash \mathcal{Q}_{S_j}$. We will write $\mathcal{S} \vdash \mathcal{Q}_S$ if $\mathcal{S}_j \vdash \mathcal{Q}_{S_j}$ for every $j = 1, \ldots, r$. Similarly, if $\mathsf{SP}^T_{\mathbf{k}}[\mathcal{S}](t_{j,i}, x_{j,i}) = y_{j,i}$ (resp. $\widetilde{P}(t_{j,i}, x_{j,i}) = y_{j,i}$) for every $i = 1, \ldots, q_j$, then we will write $\mathsf{SP}^T_{\mathbf{k}}[\mathcal{S}] \vdash \mathcal{Q}_{C_j}$ (resp. $\widetilde{P} \vdash \mathcal{Q}_{C_j}$).

Let $\mathbf{k}_1, \ldots, \mathbf{k}_\ell \in \mathcal{K}^{r+1}$ and $\widetilde{\mathcal{P}} = (\widetilde{P}_1, \ldots \widetilde{P}_\ell) \in \widetilde{\mathsf{Perm}}(\mathcal{T}, wn)^\ell$. If $\mathsf{SP}^T_{\mathbf{k}_j}[\mathcal{S}] \vdash \mathcal{Q}_{C_j}$ (resp. $\widetilde{P}_j \vdash \mathcal{Q}_{C_j}$) for every $j = 1, \ldots, \ell$, then we will write $(\mathsf{SP}^T_{\mathbf{k}_j}[\mathcal{S}])_{j=1,\ldots,\ell} \vdash \mathcal{Q}_C$ (resp. $\widetilde{\mathcal{P}} \vdash \mathcal{Q}_C$).

[10] The index j in a construction query can be dropped out in the single-user setting.

If there exist $\widetilde{\mathcal{P}} \in \widetilde{\mathsf{Perm}}(\mathcal{T}, wn)^\ell$ and $\mathcal{S} \in \mathsf{Perm}(n)^w$ that outputs τ at the end of the interaction with $\mathcal{D}$, then we will call the transcript τ *attainable*. So for any attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$, there exist $\widetilde{\mathcal{P}} \in \widetilde{\mathsf{Perm}}(\mathcal{T}, wn)^\ell$ and $\mathcal{S} \in \mathsf{Perm}(n)^w$ such that $\widetilde{\mathcal{P}} \vdash \mathcal{Q}_C$ and $\mathcal{S} \vdash \mathcal{Q}_S$. For an attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$, let

$$\mathsf{p}_1(\mathcal{Q}_C|\mathcal{Q}_S) = \mathsf{Pr}\left[\widetilde{\mathcal{P}} \xleftarrow{\$} \widetilde{\mathsf{Perm}}(\mathcal{T}, wn)^\ell, \mathcal{S} \xleftarrow{\$} \mathsf{Perm}(n)^r : \widetilde{\mathcal{P}} \vdash \mathcal{Q}_C \;\middle|\; \mathcal{S} \vdash \mathcal{Q}_S\right],$$

$$\mathsf{p}_2(\mathcal{Q}_C|\mathcal{Q}_S) = \mathsf{Pr}\left[\mathbf{k}_1, \ldots, \mathbf{k}_\ell \xleftarrow{\$} \mathcal{K}^{r+1}, \mathcal{S} \xleftarrow{\$} \mathsf{Perm}(n)^r : (\mathsf{SP}^T_{\mathbf{k}_j}[\mathcal{S}])_j \vdash \mathcal{Q}_C \;\middle|\; \mathcal{S} \vdash \mathcal{Q}_S\right].$$

With these definitions, the following lemma, the core of the H-coefficients technique (without defining "bad" transcripts), will be also used in our security proof.

Lemma 3. *Let $\varepsilon > 0$. Suppose that for any attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$,*

$$\mathsf{p}_2(\mathcal{Q}_C|\mathcal{Q}_S) \geq (1 - \varepsilon)\mathsf{p}_1(\mathcal{Q}_C|\mathcal{Q}_S). \tag{3}$$

Then one has

$$\mathrm{Adv}^{\mathrm{mu}}_{\mathsf{SP}^T}(\mathcal{D}) \leq \varepsilon.$$

The lower bound (3) is called *ε-point-wise proximity* of the transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$. The point-wise proximity of a transcript in the multi-user setting is guaranteed by the point-wise proximity of $(\mathcal{Q}_{C_j}, \mathcal{Q}_S)$ for each $j = 1, \ldots, \ell$ in the single-user setting. The following lemma is a restatement of Lemma 3 in [HT16].

Lemma 4. *Let $\varepsilon : \mathbb{N} \times \mathbb{N} \to \mathbb{R}^{\geq 0}$ be a function such that*

1. *$\varepsilon(x, y) + \varepsilon(x, z) \leq \varepsilon(x, y + z)$ for every x, y, $z \in \mathbb{N}$,*
2. *$\varepsilon(\cdot, z)$ and $\varepsilon(z, \cdot)$ are non-decreasing functions on $\mathbb{N}$ for every $z \in \mathbb{N}$.*

Suppose that for any distinguisher $\mathcal{D}$ in the single-user setting that makes p primitive queries to each of the underlying S-boxes and makes q construction queries, and for any attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$ obtained by $\mathcal{D}$, one has

$$\mathsf{p}_2(\mathcal{Q}_C|\mathcal{Q}_S) \geq (1 - \varepsilon(p, q))\mathsf{p}_1(\mathcal{Q}_C|\mathcal{Q}_S).$$

Then for any distinguisher $\mathcal{D}$ in the multi-user setting that makes p primitive queries to each of the underlying S-boxes and makes total q construction queries, and for any attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$ obtained by $\mathcal{D}$, one has

$$\mathsf{p}_2(\mathcal{Q}_C|\mathcal{Q}_S) \geq (1 - \varepsilon(p + wq, q))\mathsf{p}_1(\mathcal{Q}_C|\mathcal{Q}_S).$$

2.4 Coupling Technique

Given a finite event space Ω and two probability distributions μ and ν defined on Ω, the *total variation distance* between μ and ν, denoted $\|\mu - \nu\|$, is defined as

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

The following definitions are also all equivalent.

$$\|\mu - \nu\| = \max_{Z \subset \Omega} \{\mu(Z) - \nu(Z)\} = \max_{Z \subset \Omega} \{\nu(Z) - \mu(Z)\} = \max_{Z \subset \Omega} \{|\mu(Z) - \nu(Z)|\}.$$

A *coupling* of μ and ν is a distribution τ on $\Omega \times \Omega$ such that for all $x \in \Omega$, $\sum_{y \in \Omega} \tau(x, y) = \mu(x)$ and for all $y \in \Omega$, $\sum_{x \in \Omega} \tau(x, y) = \nu(x)$. In other words, τ is a joint distribution whose marginal distributions are respectively μ and ν. We will use the following two lemmas in our security proof.

Lemma 5. *Let μ and ν be probability distributions on a finite event space Ω, let τ be a coupling of μ and ν, and let (X, Y) be a random variable sampled according to distribution τ. Then $\|\mu - \nu\| \leq \mathsf{Pr}[X \neq Y]$.*

Lemma 6. *Let Ω be some finite event space and ν be the uniform probability distribution on Ω. Let μ be a probability distribution on Ω such that $\|\mu - \nu\| \leq \varepsilon$. Then there is a set $Z \subset \Omega$ such that*

1. $|Z| \geq (1 - \sqrt{\varepsilon})|\Omega|$,
2. $\mu(x) \geq (1 - \sqrt{\varepsilon})\nu(x)$ *for every* $x \in Z$.

We refer to [LPS12] for the proof of the above two lemmas.

3 Security of Linear SPNs

All the results in this section are for the "non-tweakable" setting ($|\mathcal{T}| = 1$) Hence, we do not explicitly refer to the tweak in the notation. Further, the results in this section hold even when a single n-bit permutation S is used, i.e., even when $S_1 = \ldots = S_r = S$ and are presented as such. We start by defining linear (non-tweakable) SPNs.

Definition 1. *Keyed permutation $T : \mathcal{K} \times \{0,1\}^{wn} \longrightarrow \{0,1\}^{wn}$ is* linear *if*

$$T(k, x) = (T_k \cdot k) + (T_x \cdot x) + \Delta,$$

where $T_k, T_x \in \mathcal{K} \times \{0,1\}^{wn}$ are linear transformations, T_x is invertible, and $\Delta \in \{0,1\}^{wn}$. An SPN is linear *if all its round permutations $\{T_{k_i}\}_{i=0}^{r}$ are linear.*

We present an attack showing that 2-round, linear SPNs cannot be secure for $w \geq 2$. The attack is based on one shown by Halevi and Rogaway [HR04] in a different context (and is a simple application of the boomerang technique [Wag99]); our contribution here is to observe that the attack is applicable to any 2-round, linear SPN. The attack relies on the fact that the field $\mathbb{F} = \mathrm{GF}(2^n)$ is of characteristic 2. This attack and an attack that works for fields of arbitrary characteristic can be found in [DKS+17].

3.1 Security of 3-Round, Linear (non-tweakable) SPNs

We now explore conditions under which 3-round, linear SPNs are secure. Recall that a 3-round SPN has four round permutations $\{T_i\}_{i=0}^{3}$, and without loss of generality we may assume

$$T_i(k_i, x) = \begin{cases} x \oplus k_i & i \in \{0,3\} \\ T'_i \cdot (x \oplus k_i) & i \in \{1,2\} \end{cases}, \tag{4}$$

where $T'_1, T'_2 \in \mathbb{F}^{w \times w}$ are invertible linear transformations. We prove that a 3-round, linear SPN is secure so long as (i) T'_1 and ${T'}_2^{-1}$ contain no zero entries (Miles and Viola [MV15] show that matrices with maximal branch number [Dae95] satisfy this property), and (ii) round keys k_0 and k_3 are (individually) uniform. The proof of this theorem can be found in [DKS+17].

Theorem 1. *Assume $w > 1$. Let SP^T be a 3-round, linear (non-tweakable) SPN with round permutations as in (4) and with distribution $\mathcal{K}$ over keys k_0, k_1, k_2, k_3. If k_0 and k_3 are uniformly distributed and the matrices T'_1, ${T'}_2^{-1}$ contain no zero entries, then*

$$\mathrm{Adv}^{\mathrm{su}}_{\mathsf{SP}^T}(p,q) \leq \frac{5w^2q^2 + 4wpq}{2^n - p - 2w} + \frac{q^2}{2^{wn}}.$$

A MINIMAL SECURE (LINEAR) SPN. We proved that a 3-round, linear SPN is secure if the keys k_0 and k_3 are individually uniform and $T'_1, {T'}_2^{-1}$ contain no 0-entries. No assumptions were made about independence of k_0, k_3, nor were any assumptions made about the distributions of k_1, k_2. So the theorem implies security for the following "minimal" 3-round, linear SPN: Let $k_0 = k_3 = k$, where k is uniform, set $k_1 = k_2 = 0^{wn}$, and let $T'_1 = {T'_2}^{-1} = T$ be invertible with no 0-entries. Define keyed permutations

$$\pi_i(k, x) = \begin{cases} x \oplus k & i \in \{0,3\} \\ Tx & i = 1 \\ {T'}^{-1}x & i = 2. \end{cases} \tag{5}$$

We have:

Corollary 1. *Assume $w > 1$. Let SP^T be a 3-round, linear SPN with round permutations as in (5) and $\mathcal{K}$ choosing uniform $k_0 = k_3$ and $k_1 = k_2 = 0^{wn}$. Then*

$$\mathrm{Adv}^{\mathrm{su}}_{\mathsf{SP}^T}(p,q) \leq \frac{5w^2q^2 + 4wpq}{2^n - p - 2w} + \frac{q^2}{2^{wn}}.$$

REDUCING KEY-LENGTH. It is in fact sufficient for the wn-bit key k $(= k_0 = k_3)$ in Corollary 1 to satisfy the following conditions: informally, for any n-bit constant c and distinct indices i, i', (a) $k[i]$ equals c with negligible probability, and (b) the sum of $k[i]$ and $k[i']$ equals c with negligible probability. This can be achieved by choosing a uniform n-bit key k' and letting $k[i] = a_i \cdot k'$ where a_i are distinct non-zero elements of $\mathbb{F}$. Thus, one can make do with a "master key" of only n bits, while preserving the same security as in Corollary 1.

4 Security of Non-Linear SPNs

In this section, we first show that keyed tweakable blockwise universal permutations help construct (non-linear) tweakable SPNs. As a preliminary step, we show that 1 round is sufficient to obtain this result. However, the security of the SPN is only up to the birthday attack in this case. Towards obtaining a better security bound, we show that 2 rounds suffice to go beyond the birthday bound and in addition, also present multi-user security beyond the birthday bound for the 2-round tweakable SPN. Finally, we we show that if T is a super blockwise tweakable universal permutation, then the security of SP^T converges to 2^n as the number of rounds r increases.

4.1 Birthday Security of 1-Round SPNs

We show that a tweakable blockwise-universal permutation is useful in constructing non-linear tweakable SPNs. The proof of the theorem is a straightforward extension of the non-tweakable version found in [DKS+17]. Consider the 1-round SPN SP^T with $T_{k_1} := T_{k_0}^{-1}$ where T is a keyed blockwise universal tweakable permutation.

Theorem 2. *Let T be a (δ, δ')-blockwise universal tweakable permutation. Then for any integers p and q, one has $\mathbf{Adv}^{\mathrm{su}}_{\mathsf{SP}^T}(p, q) \leq q^2w^2\delta + pqw\delta'$.*

4.2 Beyond-Birthday Security of 2-Round SPNs

In this section, we will prove the following theorem.

Theorem 3. *Let δ, $\delta' > 0$, and let n and w be positive integers such that $w \geq 2$. Let T be a (δ, δ')-super blockwise universal tweakable permutation. Then for any integers p and q such that $wp + 3w^2q < 2^n/2$, one has*

$$\mathbf{Adv}^{\mathrm{su}}_{\mathsf{SP}^T}(p, q) \leq w^2q(\delta'p + \delta wq)(3\delta'p + 3\delta wq + 2\delta' wq) + \frac{q^2}{2^{wn}} + \frac{q(2wp + 6w^2q)^2}{2^{2n}},$$

$$\mathbf{Adv}^{\mathrm{mu}}_{\mathsf{SP}^T}(p, q) \leq w^2q(\delta'p + (\delta + \delta')wq)(3\delta'p + 3\delta wq + 5\delta' wq) + \frac{q^2}{2^{wn}} + \frac{q(2wp + 8w^2q)^2}{2^{2n}}.$$

Remark 2. For the sake of simplicity, we assume that the three keyed layers are actually the same, which is why we require T to be (δ, δ')-super blockwise tweakable universal. However, if one looks closely at the proof, only the middle layer has to be super-blockwise-universal. The first and the last layer only need to be (δ, δ')-blockwise universal.

Remark 3. When the S-boxes are modeled as block ciphers using secret keys, the security bound (in the standard model) is obtained by setting $p = 0$.

The proof of Theorem 3 relies on the following lemma (with the lower bound simplified) and on Lemma 3 and Lemma 4.

Lemma 7. *Let p and q be positive integers such that $wp + 3w^2q < 2^n/2$, and let $\mathcal{D}$ be a distinguisher in the single-user setting that makes p primitive queries to each of S_1 and S_2 and makes q construction queries. Then for any attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$, one has*

$$\frac{\mathsf{p}_2(\mathcal{Q}_C|\mathcal{Q}_S)}{\mathsf{p}_1(\mathcal{Q}_C|\mathcal{Q}_S)} \geq 1 - w^2q(\delta' p + \delta wq)(3\delta' p + 3\delta wq + 2\delta' wq) - \frac{q^2}{2^{wn}} - \frac{q(2wp + 6w^2q)^2}{2^{2n}}.$$

Outline of Proof of Lemma 7. Throughout the proof, we will write a 2-round SP construction as

$$\mathsf{SP}^T[\mathcal{S}]_{\mathbf{k}}(t, x) = T_{k_2,t}\left(S_2^{||}\left(T_{k_1,t}\left(S_1^{||}\left(T_{k_0,t}(x)\right)\right)\right)\right),$$

where $\mathcal{S} = (S_1, S_2)$ is a pair of two public random permutations of $\{0,1\}^n$, $\mathbf{k} = (k_0, k_1, k_2) \in \mathcal{K}^3$ is the key, $x \in \{0,1\}^{wn}$ is the plaintext, and, for $i = 1, 2$,

$$\begin{aligned} S_i^{||} &: \{0,1\}^{wn} \to \{0,1\}^{wn} \\ x = x_1||x_2||\ldots||x_w &\longmapsto S_i(x_1)||S_i(x_2)||\ldots||S_i(x_w). \end{aligned}$$

We also fix a distinguisher $\mathcal{D}$ as described in the statement and fix an attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$ obtained by $\mathcal{D}$. Let

$$\begin{aligned} \mathcal{Q}_{S_1}^{(0)} &= \{(u, v) \in \{0,1\}^n \times \{0,1\}^n : (1, u, v) \in \mathcal{Q}_S\}, \\ \mathcal{Q}_{S_2}^{(0)} &= \{(u, v) \in \{0,1\}^n \times \{0,1\}^n : (2, u, v) \in \mathcal{Q}_S\} \end{aligned}$$

and let

$$\begin{aligned} U_1^{(0)} &= \{u_1 \in \{0,1\}^n : (u_1, v_1) \in \mathcal{Q}_{S_1}^{(0)}\}, \quad V_1^{(0)} = \{v_1 \in \{0,1\}^n : (u_1, v_1) \in \mathcal{Q}_{S_1}^{(0)}\}, \\ U_2^{(0)} &= \{u_2 \in \{0,1\}^n : (u_2, v_2) \in \mathcal{Q}_{S_2}^{(0)}\}, \quad V_2^{(0)} = \{v_2 \in \{0,1\}^n : (u_2, v_2) \in \mathcal{Q}_{S_2}^{(0)}\} \end{aligned}$$

denote the domains and ranges of $\mathcal{Q}_{S_1}^{(0)}$ and $\mathcal{Q}_{S_2}^{(0)}$, respectively.

This type of lemma is usually proved by defining a large enough set of "good" keys, and then, for each choice of a good key, lower bounding the probability of observing this transcript, again by lower bounding the number of possible "intermediate" values. A key is usually said to be good if the adversary cannot use the transcript to follow the path of computation of the encryption/decryption of a query up to a contradiction. However, since the S-boxes are used several times in each round, there will not be enough information in the transcript to allow such a naive definition. Therefore, instead of summing over the choice of the key, we will define an extension of the transcript, that will provide the necessary information, and then sum over every possible good extension.

We will first define what we mean by an extension of the transcript τ. Then we will define bad extensions and explain the link between good extended transcripts and the ratio $\frac{\mathsf{p}_2(\mathcal{Q}_C|\mathcal{Q}_S)}{\mathsf{p}_1(\mathcal{Q}_C|\mathcal{Q}_S)}$. Finally, we will show that the number of bad extended transcripts is small enough in Lemma 8, and then show that the probability to

obtain any good extension in the real world is sufficiently close to the probability to obtain τ the ideal world in Lemma 9. We stress that extended transcripts are completely virtual and are not disclosed to the adversary. They are just an artificial intermediate step to lower bound the probability to observe transcript τ in the real world.

Extension of a transcript. We will extend the transcript τ of the attack via a certain randomized process. We begin with choosing a pair of keys $(k_0, k_2) \in \mathcal{K}^2$ uniformly at random. Once these keys have been chosen, some construction queries will become involved in collisions. A *colliding query* is defined as a construction query $(t, x, y) \in \mathcal{Q}_C$ such that one of the following conditions holds:

1. there exist an S-box query $(1, u, v) \in \mathcal{Q}_S$ and an integer $i \in \{1, \ldots, w\}$ such that $T_{k_0,t}(x)_i = u$;
2. there exist an S-box query $(2, u, v) \in \mathcal{Q}_S$ and an integer $i \in \{1, \ldots, w\}$ such that $T^{-1}_{k_2,t}(y)_i = v$;
3. there exist a construction query $(t', x', y') \in \mathcal{Q}_C$ and integers $i, j \in \{1, \ldots, w\}$ such that $(t, x, y, i) \neq (t', x', y', j)$ and $T_{k_0,t}(x)_i = T_{k_0,t'}(x')_j$;
4. there exist a construction query $(t', x', y') \in \mathcal{Q}_C$ and integers $i, j \in \{1, \ldots, w\}$ such that $(t, x, y, i) \neq (t', x', y', j)$ and $T^{-1}_{k_2,t}(y)_i = T^{-1}_{k_2,t'}(y')_j$.

We are now going to build a new set $\mathcal{Q}'_S$ of S-box evaluations that will play the role of an extension of $\mathcal{Q}_S$. For each *colliding* query $(t, x, y) \in \mathcal{Q}_C$, we will add tuples $(1, T_{k_0}(t, x)_i, v')_{1 \leq i \leq w}$ (if (t, x, y) collides at the input of S_1) or $(2, u', T^{-1}_{k_2,t}(y)_i)_{1 \leq i \leq w}$ (if (t, x, y) collides at the output of S_2) by lazy sampling $v' = S_1(T_{k_0,t}(x)_i)$ or $u' = S_2^{-1}(T^{-1}_{k_2,t}(y)_i)$, as long as it has not been determined by any existing query in $\mathcal{Q}_S$. We finally choose a key k_1 uniformly at random. An extended transcript of τ will be defined as a tuple $\tau' = (\mathcal{Q}_C, \mathcal{Q}_S, \mathcal{Q}'_S, \mathbf{k})$ where $\mathbf{k} = (k_0, k_1, k_2)$. For each collision between a construction query and a primitive query, or between two construction queries, the extended transcript will contain enough information to compute a complete round of the evaluation of the SPN. This will be useful to lower bound the probability to get the transcript τ in the real world.

Definition of Bad Transcript Extensions. Let

$$\mathcal{Q}^{(1)}_{S_1} = \{(u, v) \in \{0,1\}^n \times \{0,1\}^n : (1, u, v) \in \mathcal{Q}_S \cup \mathcal{Q}'_S\}$$
$$\mathcal{Q}^{(1)}_{S_2} = \{(u, v) \in \{0,1\}^n \times \{0,1\}^n : (2, u, v) \in \mathcal{Q}_S \cup \mathcal{Q}'_S\}.$$

In words, $\mathcal{Q}^{(1)}_{S_i}$ summarizes each constraint that is forced on S_i by $\mathcal{Q}_S$ and $\mathcal{Q}'_S$. Let

$$U_1 = \{u_1 \in \{0,1\}^n : (1, u_1, v_1) \in \mathcal{Q}^{(1)}_{S_1}\}, \quad V_1 = \{v_1 \in \{0,1\}^n : (1, u_1, v_1) \in \mathcal{Q}^{(1)}_{S_1}\},$$
$$U_2 = \{u_2 \in \{0,1\}^n : (2, u_2, v_2) \in \mathcal{Q}^{(1)}_{S_2}\}, \quad V_2 = \{v_2 \in \{0,1\}^n : (2, u_2, v_2) \in \mathcal{Q}^{(1)}_{S_2}\}$$

be the domains and ranges of $\mathcal{Q}_{S_1}^{(1)}$ and $\mathcal{Q}_{S_2}^{(1)}$, respectively. We define two quantities characterizing an extended transcript τ', namely

$$\alpha_1 \stackrel{\text{def}}{=} \left|\{(x,y) \in \mathcal{Q}_C : T_{k_0}(x)_i \in U_1 \text{ for some } i \in \{1,\ldots,w\}\}\right|,$$
$$\alpha_2 \stackrel{\text{def}}{=} \left|\{(x,y) \in \mathcal{Q}_C : T_{k_2}^{-1}(y)_i \in V_2 \text{ for some } i \in \{1,\ldots,w\}\}\right|.$$

In words, α_1 (resp. α_2) is the number of queries $(t,x,y) \in \mathcal{Q}_C$ which collide with a query $(u_1,v_1) \in \mathcal{Q}_{S_1}^{(1)}$ (resp. which collide with a query $(u_2,v_2) \in \mathcal{Q}_{S_2}^{(1)}$) in the extended transcript. This corresponds to the number of queries $(t,x,y) \in \mathcal{Q}_C$ which collide with either an original query $(u_1,v_1) \in \mathcal{Q}_{S_1}^{(0)}$ (resp. $(u_2,v_2) \in \mathcal{Q}_{S_2}^{(0)}$) or with a query $(t',x',y') \in \mathcal{Q}_C$ at an input of S_1 (resp. at the output of S_2), once the choice of (k_0,k_2) has been made. We will also denote

$$\beta_i = |\mathcal{Q}_{S_i}^{(1)}| - |\mathcal{Q}_{S_i}^{(0)}| = |\mathcal{Q}_{S_i}^{(1)}| - p$$

for $i = 1,2$, the number of additional queries included in the extended transcript.

We say an extended transcript τ' is *bad* if at least one of the following conditions is fulfilled:

(C-1) there exist $(t,x,y) \in \mathcal{Q}_C$, $i,j \in \{1,\ldots,w\}$, $u_1 \in U_1$, and $v_2 \in V_2$ such that $T_{k_0,t}(x)_i = u_1$ and $T_{k_2,t}^{-1}(y)_j = v_2$;

(C-2) there exist $(t,x,y) \in \mathcal{Q}_C$, $i,j \in \{1,\ldots,w\}$, $u_1 \in U_1$, and $u_2 \in U_2$ such that $T_{k_0,t}(x)_i = u_1$ and $T_{k_1,t}\left(S_1^{||}\left(T_{k_0,t}(x)\right)\right)_j = u_2$[11];

(C-3) there exist $(t,x,y) \in \mathcal{Q}_C$, $i,j \in \{1,\ldots,w\}$, $v_1 \in V_1$, and $v_2 \in V_2$ such that $T_{k_2,t}^{-1}(y)_i = v_2$ and $T_{k_1,t}^{-1}\left(\left(S_2^{-1}\right)^{||}\left(T_{k_2,t}^{-1}(y)\right)\right)_j = v_1$;

(C-4) there exist $(t,x,y),(t',x',y') \in \mathcal{Q}_C$, $i,i',j,j' \in \{1,\ldots,w\}$ with $(t,x,j) \neq (t',x',j')$, $u_1,u_1' \in U_1$ such that $T_{k_0,t}(x)_i = u_1$, $T_{k_0,t'}(x')_{i'} = u_1'$ and

$$T_{k_1,t}\left(S_1^{||}\left(T_{k_0,t}(x)\right)\right)_j = T_{k_1,t'}\left(S_1^{||}\left(T_{k_0,t'}(x')\right)\right)_{j'};$$

(C-5) there exist $(t,x,y),(t',x',y') \in \mathcal{Q}_C$, $i,i',j,j' \in \{1,\ldots,w\}$ with $(y,j) \neq (y',j')$, $v_2,v_2' \in V_2$ such that $T_{k_2,t}^{-1}(y)_i = v_2$, $T_{k_2,t'}^{-1}(y')_{i'} = v_2'$ and

$$T_{k_1,t}^{-1}\left(\left(S_2^{-1}\right)^{||}\left(T_{k_2,t}^{-1}(y)\right)\right)_j = T_{k_1,t'}^{-1}\left(\left(S_2^{-1}\right)^{||}\left(T_{k_2,t'}^{-1}(y')\right)\right)_{j'}.$$

Any extended transcript that is not bad will be called *good*. Given an original transcript τ, we denote $\Theta_{\text{good}}(\tau)$ (resp. $\Theta_{\text{bad}}(\tau)$) the set of good (resp. bad) extended transcripts of τ and $\Theta'(\tau)$ the set of all extended transcripts of τ.

[11] Note that the value $S_1^{||}\left(T_{k_0,t}(x)\right)$ is well-defined thanks to the additional virtual queries from $\mathcal{Q}_S'$.

From Attainable Transcripts to Good Extended Transcripts. We are now going to justify the usefulness of extended transcripts. For any extended transcript $\tau' = (\mathcal{Q}_C, \mathcal{Q}_S, \mathcal{Q}'_S, \mathbf{k})$, let us denote

$$\mathsf{p}_{\mathrm{re}}(\tau') = \mathsf{Pr}\left[(\mathbf{k}', \mathcal{S}) \xleftarrow{\$} \mathcal{K}^3 \times \mathsf{Perm}(n)^2 : (\mathcal{S} \vdash \mathcal{Q}_S \cup \mathcal{Q}'_S) \wedge (\mathsf{SP}^T_{\mathbf{k}}[\mathcal{S}] \vdash \mathcal{Q}_C) \wedge (\mathbf{k}' = \mathbf{k})\right],$$
$$\mathsf{p}(\tau') = \mathsf{Pr}\left[\mathcal{S} \xleftarrow{\$} \mathsf{Perm}(n)^2 : \mathsf{SP}^T[\mathcal{S}]_{\mathbf{k}} \vdash \mathcal{Q}_C \middle| (S_1 \vdash \mathcal{Q}^{(1)}_{S_1}) \wedge (S_2 \vdash \mathcal{Q}^{(1)}_{S_2})\right].$$

Note that one has

$$\mathsf{Pr}\left[(\widetilde{P}, \mathcal{S}) \xleftarrow{\$} \widetilde{\mathsf{Perm}}(\mathcal{T}, wn) \times \mathsf{Perm}(n)^2 : (\mathcal{S} \vdash \mathcal{Q}_S) \wedge (\widetilde{P} \vdash \mathcal{Q}_C)\right] \leq \frac{1}{(2^{wn})_q (2^n)_p (2^n)_p},$$

$$\mathsf{Pr}\left[(\mathbf{k}, \mathcal{S}) \xleftarrow{\$} \mathcal{K}^3 \times \mathsf{Perm}(n)^2 : (\mathcal{S} \vdash \mathcal{Q}_S) \wedge (\mathsf{SP}^T_{\mathbf{k}}[\mathcal{S}] \vdash \mathcal{Q}_C)\right] \geq \sum_{\tau' \in \Theta_{\mathrm{good}}(\tau)} \mathsf{p}_{\mathrm{re}}(\tau') \geq \sum_{\tau' \in \Theta_{\mathrm{good}}(\tau)} \frac{1}{|\mathcal{K}|^3 (2^n)_{p+\beta_1} (2^n)_{p+\beta_2}} \mathsf{p}(\tau'),$$

which gives

$$\mathsf{p}_1(\mathcal{Q}_C | \mathcal{Q}_S) \leq \frac{1}{(2^{wn})_q},$$
$$\mathsf{p}_2(\mathcal{Q}_C | \mathcal{Q}_S) \geq \sum_{\tau' \in \Theta_{\mathrm{good}}(\tau)} \frac{1}{|\mathcal{K}|^3 (2^n - p)_{\beta_1} (2^n - p)_{\beta_2}} \mathsf{p}(\tau').$$

Thus one has

$$\frac{\mathsf{p}_2(\mathcal{Q}_C | \mathcal{Q}_S)}{\mathsf{p}_1(\mathcal{Q}_C | \mathcal{Q}_S)} \geq \sum_{\tau' \in \Theta_{\mathrm{good}}(\tau)} \frac{(2^{wn})_q}{|\mathcal{K}|^3 (2^n - p)_{\beta_1} (2^n - p)_{\beta_2}} \mathsf{p}(\tau')$$
$$\geq \min_{\tau' \in \Theta_{\mathrm{good}}(\tau)} ((2^{wn})_q \mathsf{p}(\tau')) \sum_{\tau' \in \Theta_{\mathrm{good}}(\tau)} \frac{1}{|\mathcal{K}|^3 (2^n - p)_{\beta_1} (2^n - p)_{\beta_2}}.$$

Note that the weighted sum $\sum_{\tau' \in \Theta_{\mathrm{good}}(\tau)} \frac{1}{|\mathcal{K}|^3 (2^n - p)_{\beta_1} (2^n - p)_{\beta_2}}$ corresponds exactly to the probability that a random extended transcript is good when it is sampled as follows:

1. choose keys $k_0, k_2 \in \mathcal{K}$ uniformly and independently at random;
2. choose the partial extension of the S-box queries based on the new collisions $\mathcal{Q}'_S$ uniformly at random (meaning that each possible u or v is chosen uniformly at random in the set of its authorized values);
3. finally choose k_1 uniformly at random, independently from everything else.

Thus, the exact probability of observing the extended transcript τ' is

$$\frac{1}{|\mathcal{K}|^3 (2^n - p)_{\beta_1} (2^n - p)_{\beta_2}},$$

and we have

$$\sum_{\tau' \in \Theta_{\mathrm{good}}(\tau)} \frac{1}{|\mathcal{K}|^3 (2^n - p)_{\beta_1} (2^n - p)_{\beta_2}} = \mathsf{Pr}\left[\tau' \in \Theta_{\mathrm{good}}(\tau)\right].$$

One finally gets

$$\frac{\mathsf{p}_2(\mathcal{Q}_C | \mathcal{Q}_S)}{\mathsf{p}_1(\mathcal{Q}_C | \mathcal{Q}_S)} \geq \mathsf{Pr}\left[\tau' \in \Theta_{\mathrm{good}}(\tau)\right] \cdot \min_{\tau' \in \Theta_{\mathrm{good}}(\tau)} ((2^{wn})_q \mathsf{p}(\tau')). \tag{6}$$

Lemma 8 and Lemma 9 lower bound $\mathsf{Pr}\left[\tau' \in \Theta_{\mathrm{good}}(\tau)\right]$ (by upper bounding $\mathsf{Pr}\left[\tau' \in \Theta_{\mathrm{bad}}(\tau)\right]$) and $\min_{\tau' \in \Theta_{\mathrm{good}}(\tau)}((2^{wn})_q \mathsf{p}(\tau'))$, respectively. Then combining (6) with Lemma 8 and Lemma 9 will complete the proof of Lemma 7.

Lemma 8. *One has*

$$\mathsf{Pr}\left[\tau' \in \Theta_{\mathrm{bad}}(\tau)\right] \leq w^2 q(\delta' p + \delta w q)(3\delta' p + 3\delta w q + 2\delta' w q).$$

Proof. We fix any attainable transcript, denoted $(\mathcal{Q}_C, \mathcal{Q}_{S_1}^{(0)}, \mathcal{Q}_{S_2}^{(0)})$. For any fixed construction query $(t, x, y) \in \mathcal{Q}_C$, define event

$$\mathsf{Coll}_1(t, x, y) \Leftrightarrow \text{there exist } i \in \{1, \ldots, w\} \text{ and } u_1 \in U_1 \text{ such that } T_{k_0,t}(x)_i = u_1.$$

This event can be broken down into the following two subevents:

- there exist $i \in \{1, \ldots, w\}$, $j \in \{1, \ldots, p\}$ such that $T_{k_0,t}(x)_i = u_j$,
- there exist $(t', x', y') \in \mathcal{Q}_C$, $i, j \in \{1, \ldots, w\}$ such that $(t, x, y, i) \neq (t', x', y', j)$ and $T_{k_0,t}(x)_i = T_{k_0,t'}(x')_j$.

Note that these events only involve queries from the original transcript, which means that the choice of the key is actually independent from these values. By the blockwise uniformity of T, one has

$$\mathsf{Pr}\left[k_0 \in \mathcal{K} : \mathsf{Coll}_1(t, x, y)\right] \leq \delta' w p + \delta w^2 q. \tag{7}$$

Similarly, let

$$\mathsf{Coll}_2(t, x, y) \Leftrightarrow \text{there exist } i \in \{1, \ldots, w\} \text{ and } v_2 \in V_2 \text{ such that } T_{k_2,t}^{-1}(y)_i = v_2.$$

Then one has

$$\mathsf{Pr}\left[k_2 \in \mathcal{K} : \mathsf{Coll}_2(x, y)\right] \leq \delta' w p + \delta w^2 q. \tag{8}$$

Also note that one has $|\mathcal{Q}_{S_1}^{(1)}|, |\mathcal{Q}_{S_2}^{(1)}| \leq p + wq$, as additional tuples in $\mathcal{Q}'_S$ come from the completion of partial information about a construction query.

We now upper bound the probabilities of the five conditions in turn. The sets of attainable transcripts fulfilling condition (C-1), (C-2), (C-3), (C-4), (C-5) will be denoted Θ_1, Θ_2, Θ_3, Θ_4, Θ_5, respectively.

Condition (C-1). One has

$$\mathsf{Pr}\left[\tau' \in \Theta_1\right] \leq \sum_{(t,x,y)\in\mathcal{Q}_C} \mathsf{Pr}\left[\mathsf{Coll}_1(t,x,y) \wedge \mathsf{Coll}_2(t,x,y)\right].$$

Since the random choice of k_0 and k_2 are independent, and by (7) and (8), one has

$$\mathsf{Pr}\left[\tau' \in \Theta_1\right] \leq q(\delta' wp + \delta w^2 q)^2.$$

Condition (C-2) and (C-3). Fix any query $(t,x,y) \in \mathcal{Q}_C$. Since the random choice of k_1 is independent from the queries transcript and from the choice of k_0, the probability, over the random choice of k_1, that there exist $i \in \{1,\ldots,w\}$ and $u_2 \in U_2$ such that $T_{k_1,t}\left(S_1^{||}\left(T_{k_0,t}(x)\right)\right)_i = u_2$, conditioned on $\mathsf{Coll}_1(t,x,y)$, is upper bounded by $\delta' w(p+wq)$. Thus, by summing over every construction query and using (7), one has

$$\mathsf{Pr}\left[\tau' \in \Theta_2\right] \leq \delta' wq(p+wq)(\delta' wp + \delta w^2 q).$$

Similarly, one has

$$\mathsf{Pr}\left[\tau' \in \Theta_3\right] \leq \delta' wq(p+wq)(\delta' wp + \delta w^2 q).$$

Conditions (C-4), and (C-5). Given two distinct pairs $(i,(t,x,y)),(i',(t',x',y')) \in \{1,\ldots,w\}\times\mathcal{Q}_C$ such that (t,x,y) and (t',x',y') are both colliding queries, let us define event

$$\mathsf{Coll}(t,x,y,t',x',y')_{i,i'} \Leftrightarrow T_{k_1,t}\left(S_1^{||}\left(T_{k_0,t}(x)\right)\right)_i = T_{k_1,t'}\left(S_1^{||}\left(T_{k_0,t'}(x')\right)\right)_{i'}.$$

Then for any distinct pairs $(i,(t,x,y)),(i',(t',x',y')) \in \{1,\ldots,w\}\times\mathcal{Q}_C$, one has

$$\begin{aligned}
&\mathsf{Pr}\left[\mathsf{Coll}_1(t,x,y) \wedge \mathsf{Coll}_1(t',x',y') \wedge \mathsf{Coll}(t,x,y,t',x',y')_{i,i'}\right] \\
&\quad = \mathsf{Pr}\left[\mathsf{Coll}(t,x,y,t',x',y')_{i,i'} \mid \mathsf{Coll}_1(t,x,y) \wedge \mathsf{Coll}_1(t',x',y')\right] \\
&\qquad \times \mathsf{Pr}\left[\mathsf{Coll}_1(t',x',y') \mid \mathsf{Coll}_1(t,x,y)\right] \\
&\qquad\quad \times \mathsf{Pr}\left[\mathsf{Coll}_1(t,x,y)\right] \leq \delta \cdot 1 \cdot (\delta' wp + \delta w^2 q),
\end{aligned}$$

where, for the last inequality, we used the (δ,δ')-blockwise uniformity of T and the fact that the event $\mathsf{Coll}_1(t,x,y) \wedge \mathsf{Coll}_1(t',x',y')$ only depends on the choice of k_0 whereas $\mathsf{Coll}(t,x,y,t',x',y')_{i,i'}$ involves the choice of k_1. Thus, by summing over every such pair, one obtains

$$\mathsf{Pr}\left[\tau' \in \Theta_4\right] \leq \delta w^2 q^2(\delta' wp + \delta w^2 q).$$

Similarly, one has

$$\mathsf{Pr}\left[\tau' \in \Theta_5\right] \leq \delta w^2 q^2(\delta' wp + \delta w^2 q).$$

The lemma follows by taking a union bound over all the conditions. □

Our next step is to study good extended transcripts.

Lemma 9. *For any good extended transcript τ', one has*

$$(2^{wn})_q \mathsf{p}(\tau') \geq 1 - \frac{q^2}{2^{wn}} - \frac{q(2wp + 6w^2q)^2}{2^{2n}}.$$

Proof. Fix any good extended transcript $\tau' = (\mathcal{Q}_C, \mathcal{Q}_S, \mathcal{Q}'_S, (k_0, k_1, k_2))$. Let us denote $p_1 = |\mathcal{Q}_{S_1}^{(1)}|$ and $p_2 = |\mathcal{Q}_{S_2}^{(1)}|$.

Our goal is then to prove that $\mathsf{p}(\tau')$ is close enough to $1/(2^{wn})_q$. In order to do so, we are going to group the construction queries according to the type of collision they are involved in:

$$\begin{aligned}
\mathcal{Q}_{U_1} &= \{(t,x,y) \in \mathcal{Q}_C : T_{k_0,t}(x)_i \in U_1 \text{ for } i = 1, \ldots, w\} \\
\mathcal{Q}_{V_2} &= \{(t,x,y) \in \mathcal{Q}_C : T^{-1}_{k_2,t}(y)_i \in V_2 \text{ for } i = 1, \ldots, w\} \\
\mathcal{Q}_0 &= \mathcal{Q}_C \setminus (\mathcal{Q}_{U_1} \cup \mathcal{Q}_{V_2}).
\end{aligned}$$

Note that, thanks to the additional queries from $\mathcal{Q}'_S$, there is an equivalence between the events "$T_{k_0,t}(x)_i \in U_1$ for each $i = 1, \ldots, w$" and "there exists $i \in \{1, \ldots, w\}$ such that $T_{k_0,t}(x)_i \in U_1$". Thus, one has by definition $|\mathcal{Q}_{U_1}| = \alpha_1$. Similarly, one has $|\mathcal{Q}_{V_2}| = \alpha_2$. Also note that these sets form a partition of $\mathcal{Q}_C$:

- $\mathcal{Q}_0 \cap \mathcal{Q}_{U_1} = \emptyset$ by definition;
- $\mathcal{Q}_0 \cap \mathcal{Q}_{V_2} = \emptyset$ by definition;
- $\mathcal{Q}_{U_1} \cap \mathcal{Q}_{V_2} = \emptyset$ since otherwise τ' would satisfy (C-1).

If we denote respectively $\mathsf{E}_{U_1}, \mathsf{E}_{V_2}$ and E_0 the event $\mathsf{SP}^T[\mathcal{S}]_{\mathbf{k}} \vdash \mathcal{Q}_{U_1}, \mathcal{Q}_{V_2}, \mathcal{Q}_0$, the event $\mathsf{SP}^T[\mathcal{S}]_{\mathbf{k}} \vdash \mathcal{Q}_C$ is equivalent to $\mathsf{E}_{U_1} \wedge \mathsf{E}_{V_2} \wedge \mathsf{E}_0$. Note that, by definition of $\mathcal{Q}_{U_1}$, each $(t,x,y) \in \mathcal{Q}_{U_1}$ is such that $T_{k_0,t}(x)_i \in U_1$ for each $i = 1, \ldots, w$; this means that the output of S_1 is already fixed by $\mathcal{Q}_{S_1}^{(1)}$ and E_{U_1} actually only involves S_2. A similar reasoning can be made for E_{V_2}. Thus we have

$$\begin{aligned}
\mathsf{p}(\tau') &= \mathsf{Pr}\left[\mathsf{E}_{U_1} \wedge \mathsf{E}_{V_2} \wedge \mathsf{E}_0 \;\middle|\; S_1 \vdash \mathcal{Q}_{S_1}^{(1)} \wedge S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right] \\
&= \mathsf{Pr}\left[\mathsf{E}_{U_1} \wedge \mathsf{E}_{V_2} \;\middle|\; S_1 \vdash \mathcal{Q}_{S_1}^{(1)} \wedge S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right] \\
&\times \mathsf{Pr}\left[\mathsf{E}_0 \;\middle|\; \mathsf{E}_{U_1} \wedge \mathsf{E}_{V_2} \wedge S_1 \vdash \mathcal{Q}_{S_1}^{(1)} \wedge S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right] \\
&= \mathsf{Pr}\left[\mathsf{E}_{U_1} \;\middle|\; S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right] \cdot \mathsf{Pr}\left[\mathsf{E}_{V_2} \;\middle|\; S_1 \vdash \mathcal{Q}_{S_1}^{(1)}\right] \\
&\times \mathsf{Pr}\left[\mathsf{E}_0 \;\middle|\; \mathsf{E}_{U_1} \wedge \mathsf{E}_{V_2} \wedge S_1 \vdash \mathcal{Q}_{S_1}^{(1)} \wedge S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right], \qquad (9)
\end{aligned}$$

where $\mathsf{Pr}\left[\mathsf{E}_{U_1} \middle| S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right]$ (resp. $\mathsf{Pr}\left[\mathsf{E}_{V_2} \middle| S_1 \vdash \mathcal{Q}_{S_1}^{(1)}\right]$) is the probability, over the random choice of permutation S_2 (resp. permutation S_1), that S_2 (resp. S_1) is compatible with the additional equations implied by $\mathcal{Q}_{U_1}$ (resp. by $\mathcal{Q}_{V_2}$), conditioned on the event $S_2 \vdash \mathcal{Q}_{S_2}^{(1)}$ (resp. $S_1 \vdash \mathcal{Q}_{S_1}^{(1)}$).

In order to evaluate $\mathsf{Pr}\left[\mathsf{E}_{U_1}\middle| S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right]$ and $\mathsf{Pr}\left[\mathsf{E}_{V_2}\middle| S_1 \vdash \mathcal{Q}_{S_1}^{(1)}\right]$, we first note that, since we condition on the event $S_2 \vdash \mathcal{Q}_{S_2}^{(1)}$, S_2 is already fixed on p_2 values. Second, remark that this event is actually equivalent to the following equations:

$$S_2\left(T_{k_1,t}\left(S_1^{||}\left(T_{k_0,t}(x)\right)\right)_i\right) = T_{k_2,t}^{-1}(y)_i$$

for every $(t,x,y) \in \mathcal{Q}_{U_1}$ and $i \in \{1,\ldots,w\}$. All the values $T_{k_1,t}\left(S_1^{||}\left(T_{k_0,t}(x)\right)\right)_i$ are actually pairwise distinct and outside U_2 since otherwise (C-2) or (C-4) would be satisfied. Similarly, the values $T_{k_2,t}^{-1}(y)_i$ are pairwise distinct and outside V_2 since otherwise (C-1) would be satisfied. Indeed, if a collision between two values $T_{k_2,t}^{-1}(y)_i$ had occured, then these values would also appear in V_2. Hence the event E_{U_1} is actually equivalent to $w\alpha_1$ new and distinct equations on S_2, so that

$$\mathsf{Pr}\left[\mathsf{E}_{U_1}\middle| S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right] = \frac{1}{(2^n - p_2)_{w\alpha_1}}. \tag{10}$$

By a similar reasoning, one has

$$\mathsf{Pr}\left[\mathsf{E}_{V_2}\middle| S_1 \vdash \mathcal{Q}_{S_1}^{(1)}\right] = \frac{1}{(2^n - p_1)_{w\alpha_2}}. \tag{11}$$

The next step is to lower bound $\mathsf{Pr}\left[\mathsf{E}_0\middle|\mathsf{E}_{U_1} \wedge \mathsf{E}_{V_2} \wedge S_1 \vdash \mathcal{Q}_{S_1}^{(1)} \wedge S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right]$. Conditioned on $\mathsf{E}_{U_1} \wedge \mathsf{E}_{V_2} \wedge S_1 \vdash \mathcal{Q}_{S_1}^{(1)} \wedge S_2 \vdash \mathcal{Q}_{S_2}^{(1)}$, S_1 and S_2 are fixed on respectively $p_1 + w\alpha_2$ and $p_2 + w\alpha_1$ values. Let U_1' (resp. U_2') be the set of values on which S_1 (resp. S_2) is already fixed and $V_1' = \{S_1(u) : u \in U_1'\}$ (resp. $V_2' = \{S_2(u) : u \in U_2'\}$). Let also $q_0 = |\mathcal{Q}_0|$. For clarity, we denote

$$\mathcal{Q}_0 = \{(t_1,x_1,y_1),\ldots,(t_{q_0},x_{q_0},y_{q_0})\},$$

using an arbitrary ordering of the queries.

Our goal is now to compute a lower bound on the number of possible "intermediate values" such that the event E_0 is equivalent to new and distinct equations on S_1 and S_2. First note that the values $T_{k_0,t}(x)_i$ for each $(t,x,y) \in \mathcal{Q}_0, i \in \{1,\ldots,w\}$ are pairwise distinct and outside U_1'. Indeed, if this were not the case, then at least one query in $\mathcal{Q}_0$ would be a colliding query. By definition of our security experiment, this means that this query would either be in E_{U_1} or E_{V_2}, depending on the type of collision it is involved in. Similarly, the values $T_{k_2,t}^{-1}(y)_i$ for each $(t,x,y) \in \mathcal{Q}_0, i \in \{1,\ldots,w\}$ are pairwise distinct and outside V_2'.

Let N_0 be the number of tuples of distinct values $(v_{1,i,j})_{1\le i\le q_0, 1\le j\le w}$ in $\{0,1\}^n \backslash V_1'$ such that the values $(T_{k_1,t_i}(||_{k=1}^w v_{1,i,k})_j)_{1\le i\le q_0,1\le j\le w}$ are also pairwise distinct and outside U_2'. Let $i \in \{1,\ldots,q_0\}$. There are exactly $(2^n - |V_1'| - w(i-1))_w$ possible tuples of distinct values $(v_{1,i,j})_{1\le j\le w}$ in $\{0,1\}^n \setminus V_1'$ that will also be different from the previous values $v_{1,i,j}$ for $i < q_0$ and $j \in \{1,\ldots,w\}$. Similarly, there are exactly $(2^n - |U_2'| - w(i-1))_w$ possible tuples of distinct values for

$(T_{k_1,t_i}(||_{k=1}^{w} v_{1,i,k}))_{1\le j\le w}$ in $\{0,1\}^n \backslash U_2'$ that will also be different from the previous values $T_{k_1,t_i}(||_{k=1}^{w} v_{1,i,k})$ for $i < q_0$ and $j \in \{1,\dots,w\}$. This removes at most $2^{wn} - (2^n - |U_2'| - w(i-1))_w$ tuples of values for $(T_{k_1,t_i}(||_{k=1}^{w} v_{1,i,k}))_{1\le j\le w}$. Since T_{k_1,t_i} is a permutation, we have to remove at most $2^{wn} - (2^n - |U_2'| - w(i-1))_w$ possible tuples of values for $(v_{1,i,j})_{1\le j\le w}$. Thus

$$N_0 \ge \prod_{i=1}^{q_0} \left((2^n - |V_1'| - w(i-1))_w + (2^n - |U_2'| - w(i-1))_w - 2^{wn}\right). \tag{12}$$

For any tuple of values $(v_{1,i,j})$ fulfilling the previous conditions, then, conditioned on S_1 satisfying $S_1(T_{k_0,t_i}(x_i))_j = v_{1,i,j}$, the event E_0 is equivalent to wq_0 distinct and new equations on S_2. Hence, it follows that

$$\mathsf{Pr}\left[\mathsf{E}_0 \;\middle|\; \mathsf{E}_{U_1} \wedge \mathsf{E}_{V_2} \wedge S_1 \vdash \mathcal{Q}_{S_1}^{(1)} \wedge S_2 \vdash \mathcal{Q}_{S_2}^{(1)}\right] \ge \frac{N_0}{(2^n - p_1 - w\alpha_2)_{wq_0}(2^n - p_2 - w\alpha_1)_{wq_0}}. \tag{13}$$

Combining (9), (10), (11), (12) (13), we obtain

$$\begin{aligned}
(2^{wn})_q \mathsf{p}(\tau') &\ge \frac{(2^{wn})_q \prod_{i=0}^{q_0-1} \left(\begin{array}{l} (2^n - p_1 - w(\alpha_2 + i))_w \\ \quad +(2^n - p_2 - w(\alpha_1 + i))_w - 2^{wn} \end{array} \right)}{(2^n - p_1)_{wq_0 + w\alpha_2}(2^n - p_2)_{wq_0 + w\alpha_1}} \\
&= \frac{(2^{wn})_q}{2^{q_0 wn}(2^n - p_1)_{w\alpha_2}(2^n - p_2)_{w\alpha_1}} \\
&\times \prod_{i=0}^{q_0-1} \frac{2^{wn} \left(\begin{array}{l} (2^n - p_1 - w(\alpha_2 + i))_w \\ \quad +(2^n - p_2 - w(\alpha_1 + i))_w - 2^{wn} \end{array} \right)}{(2^n - p_1 - w\alpha_2 - wi)_w (2^n - p_2 - w\alpha_1 - wi)_w} \\
&\ge \frac{(2^{wn})_q}{2^{q_0 wn}(2^n - p_1)_{w\alpha_2}(2^n - p_2)_{w\alpha_1}} \cdot \prod_{i=0}^{q_0-1} \Delta_i
\end{aligned}$$

where

$$\Delta_i = 1 - \left(\frac{2^{wn}}{(2^n - p_2 - w\alpha_1 - wi)_w} - 1\right)\left(\frac{2^{wn}}{(2^n - p_1 - w\alpha_2 - wi)_w} - 1\right)$$

for $i = 0,\dots,q_0 - 1$. We also have $\alpha_1 \le q$ and $p_2 \le p + wq$, which gives

$$\frac{2^{wn}}{(2^n - p_2 - w\alpha_1 - wi)_w} \le \left(\frac{2^n}{2^n - p - 3wq}\right)^w \le \left(1 + \frac{p + 3wq}{2^n - p - 3wq}\right)^w.$$

Then, since $wp + 3w^2q < 2^n/2$, we can apply Lemma 1 and we get

$$\frac{2^{wn}}{(2^n - p_2 - w\alpha_1 - wi)_w} \le 1 + \frac{wp + 3w^2q}{2^n - wp - 3w^2q} \le 1 + \frac{2wp + 6w^2q}{2^n}.$$

Similarly, one has

$$\frac{2^{wn}}{(2^n - p_1 - w\alpha_2 - wi)_w} \leq 1 + \frac{2wp + 6w^2q}{2^n}.$$

Thus one has

$$\Delta_i \geq 1 - \left(\frac{2wp + 6w^2q}{2^n}\right)^2.$$

Moreover, one has

$$\frac{(2^{wn})_q}{2^{q_0 wn}(2^n - p_1)_{w\alpha_2}(2^n - p_2)_{w\alpha_1}} \geq \frac{(2^{wn} - q)^q}{2^{qwn}} \geq \left(1 - \frac{q}{2^{wn}}\right)^q \geq 1 - \frac{q^2}{2^{wn}}.$$

Finally, we get

$$\begin{aligned}(2^{wn})_q \mathsf{p}(\tau') &\geq \left(1 - \frac{q^2}{2^{wn}}\right)\left(1 - \left(\frac{2wp + 6w^2q}{2^n}\right)^2\right)^{q_0} \\ &\geq \left(1 - \frac{q^2}{2^{wn}}\right)\left(1 - \frac{q(2wp + 6w^2q)^2}{2^{2n}}\right) \\ &\geq 1 - \frac{q^2}{2^{wn}} - \frac{q(2wp + 6w^2q)^2}{2^{2n}}.\end{aligned}$$

□

4.3 Asymptotically Optimal Security of SPNs

In this section, we will prove that if T is a super blockwise tweakable universal permutation, then the security of SP^T converges to 2^n (in terms of the threshold number of queries) as the number of rounds r increases.

Theorem 4. *For an even integer r, let SP^T be an r-round substitution-permutation network based on a (δ, δ')-super blockwise tweakable universal permutation T. Then one has*

$$\mathrm{Adv}^{\mathrm{mu}}_{\mathsf{SP}^T}(p, q) \leq 4\sqrt{q}\left(2wp\delta' + 2w^2q(\delta' + \delta) + w^2\delta\right)^{\frac{r}{4}}.$$

Hence, assuming $\delta, \delta' \simeq 2^{-n}$ and $p = q$, an r-round SP^T is secure up to $2^{\frac{rn}{r+2}}$ queries.

Proof of Theorem 4. We assume that $r = 2s$ for a positive integer s. Let $\overline{\mathsf{SP}}^T[\mathcal{S}]$ denote a variant of $\mathsf{SP}^T[\mathcal{S}]$ without the last permutation layer. Then one has

$$\mathsf{SP}^T[\mathcal{S}] = \left(\overline{\mathsf{SP}}^{T^{-1}}[\mathcal{S}^{(2)}]\right)^{-1} \circ T \circ \overline{\mathsf{SP}}^T[\mathcal{S}^{(1)}]$$

for $\mathcal{S}^{(1)} = (S_1, \ldots, S_s)$ and $\mathcal{S}^{(2)} = (S_{2s}^{-1}, \ldots, S_{s+1}^{-1})$. Our proof strategy is to first prove NCPA-security of $\overline{\mathsf{SP}}$ in the multi-user setting and lift it to CCA-security by doubling the number of rounds.

Suppose that a distinguisher $\mathcal{D}$ makes p primitive queries to each of the underlying S-boxes and makes q construction queries in the multi-user setting, obtaining an attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_S)$. We can partition $\mathcal{Q}_C$ and $\mathcal{Q}_S$ as follows.

$$\begin{aligned}\mathcal{Q}_C &= \mathcal{Q}_{C_1} \cup \cdots \cup \mathcal{Q}_{C_\ell},\\ \mathcal{Q}_S &= \mathcal{Q}_{S_1} \cup \cdots \cup \mathcal{Q}_{S_s} \cup \mathcal{Q}_{S_{s+1}} \cup \cdots \cup \mathcal{Q}_{S_{2s}},\end{aligned}$$

where we will write

$$\begin{aligned}\mathcal{Q}_S^{(1)} &= \mathcal{Q}_{S_1} \cup \cdots \cup \mathcal{Q}_{S_s},\\ \mathcal{Q}_S^{(2)} &= \mathcal{Q}_{S_{s+1}} \cup \cdots \cup \mathcal{Q}_{S_{2s}}.\end{aligned}$$

Throughout the proof, we will write $\mathcal{Q}_{C_j} = (t_{j,i}, x_{j,i}, y_{j,i})_{1\le i\le q_j}$ for $j = 1, \ldots, \ell$. So q_j denotes the number of queries made to the j-th construction oracle C_j, and $(t_{j,i}, x_{j,i}, y_{j,i})$ represents the evaluation obtained by the i-th query to C_j. We will also write $\mathbf{t} = (\mathbf{t}_j)_{1\le j\le \ell}$, $\mathbf{x} = (\mathbf{x}_j)_{1\le j\le \ell}$, $\mathbf{y} = (\mathbf{y}_j)_{1\le j\le \ell}$, where

$$\mathbf{t}_j = (t_{j,1}, \ldots, t_{j,q_j}), \qquad \mathbf{x}_j = (x_{j,1}, \ldots, x_{j,q_j}), \qquad \mathbf{y}_j = (y_{j,1}, \ldots, y_{j,q_j}),$$

for $j = 1, \ldots, \ell$. Without loss of generality, we can assume that the indices (j, i) have been grouped by their tweaks $t_{j,i}$; suppose that $\mathbf{t}_j$ consists of d different tweaks, $t_1^*, \ldots, t_d^* \in \mathcal{T}$. Then by dropping j for simplicity (when it will be clear from the context), we can write

$$\mathbf{x}_j = (\mathbf{x}_1^*, \ldots, \mathbf{x}_d^*),$$

so that $\mathbf{x}_i^* = (x_{i,1}^*, \ldots, x_{i,q_i'}^*)$ corresponds to t_i^* for $i = 1, \ldots, d$, where q_i' is the multiplicity of t_i^* in $\mathbf{t}_j$ (satisfying $q_1' + \ldots + q_d' = q_\beta$). Let

$$\begin{aligned}\Omega_{\mathbf{t}_j} &= \left\{(u_1, \ldots, u_{q_j}) \in (\{0,1\}^n)^{q_j} : \forall i \ne i', (t_{j,i}, u_i) \ne (t_{j,i'}, u_{i'})\right\},\\ \Omega_{\mathbf{t}} &= \Omega_{\mathbf{t}_1} \times \ldots \times \Omega_{\mathbf{t}_\ell}.\end{aligned}$$

With these notations, we define probability distributions μ_1 and μ_2 on $\Omega_{\mathbf{t}}$; for each $\mathbf{z} = (\mathbf{z}_1, \ldots, \mathbf{z}_\ell) \in \Omega_{\mathbf{t}}$,

$$\mu_1(\mathbf{z}) \stackrel{\text{def}}{=} \mathsf{Pr}\left[\mathbf{k}_1, \ldots, \mathbf{k}_\ell \xleftarrow{\$} \mathcal{K}^s, \mathcal{S} \xleftarrow{\$} \mathsf{Perm}(n)^s : \forall j, \overline{\mathsf{SP}}_{\mathbf{k}_j}^T[\mathcal{S}] \vdash (t_{j,i}, x_{j,i}, z_{j,i})_{1\le i\le q_j} \;\middle|\; \mathcal{S} \vdash \mathcal{Q}_S^{(1)}\right],$$

$$\mu_2(\mathbf{z}) \stackrel{\text{def}}{=} \mathsf{Pr}\left[\mathbf{k}_1, \ldots, \mathbf{k}_\ell \xleftarrow{\$} \mathcal{K}^s, \mathcal{S} \xleftarrow{\$} \mathsf{Perm}(n)^s : \forall j, \overline{\mathsf{SP}}_{\mathbf{k}_j}^T[\mathcal{S}] \vdash (t_{j,i}, y_{j,i}, z_{j,i})_{1\le i\le q_j} \;\middle|\; \mathcal{S} \vdash \mathcal{Q}_S^{(2)}\right],$$

where we write $\mathbf{z}_j = (z_{j,i})_{1\le i\le q_j}$ for $j = 1, \ldots, \ell$. Using the coupling technique, we can upper bound the statistical distance between μ_c and the uniform probability distribution for $c = 1, 2$. The proof of the following lemma can be found in [CL18].

Lemma 10. *For $c = 1, 2$, let μ_c be the probability distribution defined as above, and let ν be the uniform probability distribution on $\Omega_{\mathbf{t}}$. Then for $c = 1, 2$, one has $\|\mu_c - \nu\| \le \varepsilon$, where*

$$\varepsilon = \varepsilon(p, q) \stackrel{\text{def}}{=} q\left(2wp\delta' + 2w^2q(\delta' + \delta) + w^2\delta\right)^s.$$

By Lemma 6 and Lemma 10, we have a subset $Z_1 \subset \Omega_{\mathbf{t}}$ such that $|Z_1| \geq (1-\sqrt{\varepsilon})|\Omega_{\mathbf{t}}|$ and

$$\mu_1(\mathbf{z}) \geq (1-\sqrt{\varepsilon})\nu(\mathbf{z}) = \frac{1-\sqrt{\varepsilon}}{|\Omega_{\mathbf{t}}|}$$

for every $\mathbf{z} \in Z_1$. Similarly, we also have a subset $Z_2 \subset \Omega_{\mathbf{t}}$ such that $|Z_2| \geq (1-\sqrt{\varepsilon})|\Omega_{\mathbf{t}}|$ and

$$\mu_2(\mathbf{z}) \geq (1-\sqrt{\varepsilon})\nu(\mathbf{z}) = \frac{1-\sqrt{\varepsilon}}{|\Omega_{\mathbf{t}}|}$$

for every $\mathbf{z} \in Z_2$. For a fixed key $(k_1, \ldots, k_\ell) \in \mathcal{K}^\ell$, let

$$Z_2' = \{(T^{-1}_{k_1,\mathbf{t}_1}(\mathbf{z}_1), \ldots, T^{-1}_{,k_\ell,\mathbf{t}_\ell}(\mathbf{z}_\ell)) : (\mathbf{z}_1, \ldots, \mathbf{z}_\ell) \in Z_2\},$$

and let $Z = Z_1 \cap Z_2'$. Then it follows that

$$\begin{aligned}
\mathsf{p}_2(\mathcal{Q}_C|\mathcal{Q}_S) &= \mathsf{Pr}\left[\forall j, \mathsf{SP}^T_{\mathbf{k}_j}[\mathcal{S}] \vdash \mathcal{Q}_{C_j} \;\middle|\; \mathcal{S} \vdash \mathcal{Q}_S\right] \\
&\geq \frac{1}{|\mathcal{K}|^\ell} \sum_{\substack{k_1,\ldots,k_\ell \in \mathcal{K} \\ \mathbf{z}_1,\ldots,\mathbf{z}_\ell \in Z}} \mathsf{Pr}\left[\forall j, \overline{\mathsf{SP}}^T_{\mathbf{k}_j}[\mathcal{S}] \vdash (\mathbf{t}_j, \mathbf{x}_j, \mathbf{z}_j) \;\middle|\; \mathcal{S} \vdash \mathcal{Q}_S^{(1)}\right] \\
&\quad \times \mathsf{Pr}\left[\forall j, \overline{\mathsf{SP}}^{T^{-1}}_{\mathbf{k}_j}[\mathcal{S}] \vdash (\mathbf{t}_j, \mathbf{y}_j, T_{k_j,\mathbf{t}_j}(\mathbf{z}_j)) \;\middle|\; \mathcal{S} \vdash \mathcal{Q}_S^{(2)}\right] \\
&\geq (1-2\sqrt{\varepsilon})|\Omega_{\mathbf{t}}| \cdot \left(\frac{1-\sqrt{\varepsilon}}{|\Omega_{\mathbf{t}}|}\right)^2 \geq (1-4\sqrt{\varepsilon})\mathsf{p}_1(\mathcal{Q}_C|\mathcal{Q}_S)
\end{aligned}$$

since $|Z| \geq (1-2\sqrt{\varepsilon})|\Omega_{\mathbf{t}}|$. By Lemma 3, we complete the proof of Theorem 4.

Acknowledgments

The work of Aishwarya Thiruvengadam was done while at the University of Maryland. Benoît Cogliati was partially supported by the European Union's H2020 Programme under grant agreement number ICT-644209. The work of Yevgeniy Dodis was done in part while visiting the University of Maryland, and was supported by gifts from VMware Labs and Google, as well as NSF grants 1619158, 1319051, and 1314568. The work of Jonathan Katz and Aishwarya Thiruvengadam was performed under financial assistance award 70NANB15H328 from the U.S. Department of Commerce, National Institute of Standards and Technology. Jooyoung Lee was supported by a National Research Foundation of Korea (NRF) grant funded by the Korean government (Ministry of Science and ICT), No. NRF-2017R1E1A1A03070248.

References

[BBK14] Alex Biryukov, Charles Bouillaguet, and Dmitry Khovratovich. Cryptographic schemes based on the ASASA structure: Black-box, white-box, and

public-key (extended abstract). In *Advances in Cryptology—Asiacrypt 2014, Part I*, volume 8873 of *LNCS*, pages 63–84. Springer, 2014.

[BD99] D. Bleichenbacher and A. Desai. A construction of a super-pseudorandom cipher, February 1999. Unpublished manuscript.

[BDPA09] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Keccak Sponge Function Family Main Document. *Submission to NIST (Round 2)*, 2009. Available at http://keccak.noekeon.org/Keccak-main-2.0.pdf.

[BK] Alex Biryukov and Dmitry Khovratovich. Decomposition attack on SASASASAS. Available at http://eprint.iacr.org/2015/646.

[BKL+17] Daniel J. Bernstein, Stefan Kölbl, Stefan Lucks, Pedro Maat Costa Massolino, Florian Mendel, Kashif Nawaz, Tobias Schneider, Peter Schwabe, François-Xavier Standaert, Yusuke Todo, and Benoît Viguier. GIMLI: a cross-platform permutation. In Wieland Fischer and Naofumi Homma, editors, *Cryptographic Hardware and Embedded Systems - CHES 2017*, volume 10529 of *LNCS*. Springer, 2017. To appear. Also available at http://eprint.iacr.org/2017/630.

[BS10] Alex Biryukov and Adi Shamir. Structural cryptanalysis of SASAS. *Journal of Cryptology*, 23(4):505–518, 2010.

[CDMS10] Jean-Sébastien Coron, Yevgeniy Dodis, Avradip Mandal, and Yannick Seurin. A domain extender for the ideal cipher. In *7th Theory of Cryptography Conference—TCC 2010*, volume 5978 of *LNCS*, pages 273–289. Springer, 2010.

[CHK+16] Jean-Sébastien Coron, Thomas Holenstein, Robin Künzler, Jacques Patarin, Yannick Seurin, and Stefano Tessaro. How to Build an Ideal Cipher: The Indifferentiability of the Feistel Construction. *J. Cryptology*, 29(1):61–114, 2016.

[CL18] Benoît Cogliati, and Jooyoung Lee. Wide Tweakable Block Ciphers Based on Substitution-Permutation Networks: Security Beyond the Birthday Bound. IACR Cryptology ePrint Archive, Report 2018/488, 2018. Available at http://eprint.iacr.org/2018/488.

[CLL+14] Shan Chen, Rodolphe Lampe, Jooyoung Lee, Yannick Seurin, and John P. Steinberger. Minimizing the Two-Round Even-Mansour Cipher. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 (Proceedings, Part I)*, volume 8616 of *LNCS*, pages 39–56. Springer, 2014.

[CLS15] Benoît Cogliati, Rodolphe Lampe, and Yannick Seurin. Tweaking Even-Mansour Ciphers. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 (Proceedings, Part I)*, volume 9215 of *LNCS*, pages 189–208. Springer, 2015.

[CS06] Debrup Chakraborty and Palash Sarkar. A New Mode of Encryption Providing a Tweakable Strong Pseudo-random Permutation. In Matthew Robshaw, editor, *Fast Software Encryption - FSE 2006*, volume 4047 of *LNCS*, pages 293–309. Springer, 2006.

[CS14] Shan Chen and John Steinberger. Tight Security Bounds for Key-Alternating Ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, 2014.

[Dae95] Joan Daemen. *Cipher and Hash Function Design Strategies Based on Linear and Differential Cryptanalysis.* PhD thesis, Katholieke Universiteit Leuven, 1995.

[DDKL] Itai Dinur, Orr Dunkelman, Thorsten Kranz, and Gregor Leander. Decomposing the ASASA block cipher construction. Available at http://eprint.iacr.org/2015/507.

[DKS+17] Yevgeniy Dodis, Jonathan Katz, John P. Steinberger, Aishwarya Thiruvengadam, and Zhe Zhang. Provable Security of Substitution-Permutation Networks. IACR Cryptology ePrint Archive, Report 2017/016, 2017. Available at http://eprint.iacr.org/2017/016.

[DSSL16] Yevgeniy Dodis, Martijn Stam, John P. Steinberger, and Tianren Liu. Indifferentiability of confusion-diffusion networks. In *Advances in Cryptology—Eurocrypt 2016, Part II*, volume 9666 of *LNCS*, pages 679–704. Springer, 2016.

[EM97] Shimon Even and Yishay Mansour. A Construction of a Cipher from a Single Pseudorandom Permutation. *J. Cryptology*, 10(3):151–162, 1997.

[Fei73] Horst Feistel. Cryptography and computer privacy. *Scientific American*, 228(5):15–23, 1973.

[GJMN16] Robert Granger, Philipp Jovanovic, Bart Mennink, and Samuel Neves. Improved Masking for Tweakable Blockciphers with Applications to Authenticated Encryption. In Marc Fischlin and Jean-Sebastien Coron, editors, *Advances in Cryptology - Eurocrypt 2016 (Proceedings, Part I)*, volume 9665 of *LNCS*, pages 263–293. Springer, 2016.

[Hal07] Shai Halevi. Invertible Universal Hashing and the TET Encryption Mode. In Alfred Menezes, editor, *Advances in Cryptology - Crypto 2007*, volume 4622 of *LNCS*, pages 412–429. Springer, 2007.

[HKT11] Thomas Holenstein, Robin Künzler, and Stefano Tessaro. The Equivalence of the Random Oracle Model and the Ideal Cipher Model, Revisited. In Lance Fortnow and Salil P. Vadhan, editors, *Symposium on Theory of Computing - STOC 2011*, pages 89–98. ACM, 2011.

[HR03] Shai Halevi and Phillip Rogaway. A tweakable enciphering mode. In *Advances in Cryptology—Crypto 2003*, volume 2729 of *LNCS*, pages 482–499. Springer, 2003.

[HR04] Shai Halevi and Phillip Rogaway. A parallelizable enciphering mode. In *Cryptographers' Track—RSA 2004*, volume 2964 of *LNCS*, pages 292–304. Springer, 2004.

[HR10] Viet Tung Hoang and Phillip Rogaway. On Generalized Feistel Networks. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010*, volume 6223 of *LNCS*, pages 613–630. Springer, 2010.

[HT16] Viet Tung Hoang and Stefano Tessaro. Key-Alternating Ciphers and Key-Length Extension: Exact Bounds and Multi-user Security. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 (Proceedings, Part I)*, volume 9814 of *LNCS*, pages 3–32. Springer, 2016.

[IK00] Tetsu Iwata and Kaoru Kurosawa. On the pseudorandomness of the AES finalists—RC6 and Serpent. In *Fast Software Encryption—FSE 2000*, volume 1978 of *LNCS*, pages 231–243. Springer, 2000.

[Jou03] Antoine Joux. Cryptanalysis of the EMD mode of operation. In *Advances in Cryptology—Eurocrypt 2003*, volume 2656 of *LNCS*, pages 1–16. Springer, 2003.

[KL15] J. Katz and Y. Lindell. *Introduction to Modern Cryptography, 2nd edition.* Chapman & Hall/CRC Press, 2015.

[LPS12] Rodolphe Lampe, Jacques Patarin, and Yannick Seurin. An Asymptotically Tight Security Analysis of the Iterated Even-Mansour Cipher. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 278–295. Springer, 2012.

[LR88] Michael Luby and Charles Rackoff. How to Construct Pseudorandom Permutations from Pseudorandom Functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.

[LRW11] Moses Liskov, Ronald L. Rivest, and David A. Wagner. Tweakable block ciphers. *J. Cryptology*, 24(3):588–613, 2011.

[Men16] Bart Mennink. XPX: Generalized Tweakable Even-Mansour with Improved Security Guarantees. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - Crypto 2016 (Proceedings, Part I)*, volume 9814 of *LNCS*, pages 64–94. Springer, 2016.

[MF07] David A. McGrew and Scott R. Fluhrer. The security of the extended codebook (XCB) mode of operation. In *14th Annual Intl. Workshop on Selected Areas in Cryptography (SAC)*, volume 4876 of *LNCS*, pages 311–327. Springer, 2007.

[MRH04] Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In *1st Theory of Cryptography Conference—TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, 2004.

[MRS09] Ben Morris, Phillip Rogaway, and Till Stegers. How to Encipher Messages on a Small Domain. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009*, volume 5677 of *LNCS*, pages 286–302. Springer, 2009.

[MV15] Eric Miles and Emanuele Viola. Substitution-permutation networks, pseudorandom functions, and natural proofs. *J. ACM*, 62(6):46, 2015.

[NR99] Moni Naor and Omer Reingold. On the construction of pseudorandom permutations: Luby-Rackoff revisited. *Journal of Cryptology*, 12(1):29–66, 1999.

[Pat03] Jacques Patarin. Luby-Rackoff: 7 Rounds Are Enough for $2^{n(1-\epsilon)}$ Security. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, volume 2729 of *LNCS*, pages 513–529. Springer, 2003.

[Pat04] Jacques Patarin. Security of Random Feistel Schemes with 5 or More Rounds. In Matthew K. Franklin, editor, *Advances in Cryptology - CRYPTO 2004*, volume 3152 of *LNCS*, pages 106–122. Springer, 2004.

[Pat08] Jacques Patarin. The "Coefficients H" Technique. In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *Selected Areas in Cryptography - SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, 2008.

[Pat10] Jacques Patarin. Security of Balanced and Unbalanced Feistel Schemes with Linear Non Equalities. IACR Cryptology ePrint Archive, Report 2010/293, 2010. Available at http://eprint.iacr.org/2010/293.

[Sha49] Claude Shannon. Communication Theory of Secrecy Systems. *Bell System Technical Journal*, 28(4):656–715, 1949.

[Tes14] Stefano Tessaro. Optimally Secure Block Ciphers from Ideal Primitives. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2014*, volume 9453 of *LNCS*, pages 437–462. Springer, 2014.

[Wag99] David Wagner. The boomerang attack. In Lars R. Knudsen, editor, *Fast Software Encryption—FSE '99*, volume 1636 of *LNCS*, pages 156–170. Springer, March 1999.