ZCZ – Achieving n-bit SPRP Security with a Minimal Number of Tweakable-block-cipher Calls*

Ritam Bhaumik[1], Eik List[2], and Mridul Nandi[1]**

[1] Indian Statistical Institute, Kolkata, India
[2] Bauhaus-Universität Weimar, Weimar, Germany
bhaumik.ritam@gmail.com, eik.list@uni-weimar.de, mridul.nandi@gmail.com

Abstract. Strong Pseudo-random Permutations (SPRPs) are important for various applications. In general, it is desirable to base an SPRP on a single-keyed primitive for minimizing the implementation costs. For constructions built on classical block ciphers, Nandi showed at ASIACRYPT'15 that at least two calls to the primitive per processed message block are required for SPRP security, assuming that all further operations are linear. The ongoing trend of using tweakable block ciphers as primitive has already led to MACs or encryption modes with high security and efficiency properties. Thus, three interesting research questions are hovering in the domain of SPRPs: (1) if and to which extent the bound of two calls per block can be reduced with a tweakable block cipher, (2) how concrete constructions could be realized, and (3) whether full n-bit security is achievable from primitives with n-bit state size.

The present work addresses all three questions. Inspired by Iwata et al.'s ZHash proposal at CRYPTO'17, we propose the ZCZ (ZHash-Counter-ZHash) construction, a single-key variable-input-length SPRP based on a single tweakable block cipher whose tweak length is at least its state size. ZCZ possesses close to optimal properties with regards to both performance and security: not only does it require only asymptotically $3\ell/2$ calls to the primitive for ℓ-block messages; we show that this figure is close to the minimum by an PRP distinguishing attack on any construction with tweak size of $\tau = n$ bits and fewer than $(3\ell - 1)/2$ calls to the same primitive. Moreover, it provides optimal n-bit security for a primitive with n-bit state and tweak size.

Keywords: Symmetric-key cryptography · provable security · variable-input-length SPRP · tweakable block cipher · encryption

* Full version: https://eprint.iacr.org/2018/819.pdf

** The research by Mridul Nandi has been supported by the Wisekey project at the R. C. Bose Centre for Cryptology and Security, Indian Statistical Institute, Kolkata.

1 Introduction

SPRPs. Strong Pseudo-Random Permutations (or wide-block ciphers), are important symmetric-key schemes for protecting the privacy of variable-length messages. Their tweakable variants (STPRPs) are useful to build strong authenticated encryption [12, 31] or onion AE [30]. During the previous two decades, the symmetric-key community proposed a considerable corpus of SPRPs. From a high-level point of view, existing constructions could be categorized into (1) Generalized Feistel networks, (2) Encrypt-Mix-Encrypt, (3) Hash-ECB-Hash, (4) Hash-Counter-Hash, and (5) miscellaneous designs.

Optimization Goals. The primary goals for optimizations in cryptographic schemes are, in general, low implementation costs, high provable security guarantees, and high performance. For the first criterion, it is desirable to construct higher-level schemes from a single well-analyzed primitive without large internal state and with a single key.

High security is essential in many domains that have to process large amounts of data without the ability of frequent re-keying. In most constructions, however, it comes at the cost of decreased performance. Unsurprisingly, the challenges of combining high security guarantees with high performance have been identified among the hot topics of symmetric-key cryptography at the ESC 2017 workshop [5]. Often, high security is associated with security *beyond the birthday bound*. In the areas of authentication (e.g., [19, 32, 33]), encryption, as well as authenticated encryption (e.g., [13, 14, 28]), beyond-birthday security has undergone a long line of research. In the area of SPRPs, however, the security of the vast majority of existing constructions is still limited by the birthday bound of $n/2$ bits, where n is the state size of the underlying primitive. So, the privacy guarantees are lost if $q \simeq 2^{n/2}$ message blocks have been encrypted under the same key. Assuming the AES as primitive, this would imply that significantly fewer than 2^{64} blocks could safely be encrypted under a single key.

Security of SPRPs: State of the Art. Among the earlier proposals, the LargeBlock1 and LargeBlock2 constructions by Minematsu and Iwata [23] as well as TCT_2 by Shrimpton and Terashima [31] are exceptional for their security guarantees. The LargeBlock designs can achieve optimal n-bit security, whereas TCT_2 is limited by $2n/3$ bits. Both share similarities to the Ψ_2 and Ψ_3 constructions from Coron et al. [9], which use two and three calls to a tweakable block cipher. Both LargeBlock2 and TCT_2 possess a sandwich structure, where an encryption layer is wrapped by two layers of hashing. In the former, the encryption layer is an application of Ψ_2 in ECB-mode; the hashing layers employs two calls to a polynomial hash of $2(\ell - 1)$multiplications each. TCT_2 can be seen as an unbalanced version of Ψ_3, where also $2(\ell-1)$ of ℓ input blocks are hashed in each hashing layer. Both constructions are remarkable for their time. To be comparably efficient, however, they required two primitives, a block cipher and a universal hash function.

A different direction is followed by HHFHFH [2] and its instantiations (e.g., [3]), which is a four-round unbalanced Feistel network, built on a large-state primitive. Instead of providing beyond-birthday security, it possesses large security margins due to a larger birthday bound of their internal primitives. However, the large state size limits its efficiency.

The only approach we are aware of that almost combines both security and performance desiderata is Simpira (v2) [10], a family of Feistel-like constructions built upon the AES round function. Its authors claim 128-bit security and high performance on current processors with support for AES native instructions. However, Simpira's security claim stems purely from heuristics, which will demand intensive further cryptanalysis to increase trust into it.

Tweakable Block Ciphers. One established approach for achieving higher security without considerably sacrificing performance is to use a tweakable block cipher (TBC) [18] as underlying primitive. At the core, tweakable block ciphers employ an additional public input called tweak, which allows to efficiently separate the domains of different calls to the primitive. This fact can reduce the impact of internal collisions on the security of the scheme built around them. For message authentication codes (MACs), a series of recent works pushed the security bounds further [8, 15, 24], but a similar trend is also observable in the domain of encryption modes and authenticated encryption schemes [14, 17, 21, 28, 29]. This approach has also been used earlier for SPRPs [9, 20, 22, 23, 31] – those proposals, however, originate from at least half a decade ago where TBCs used to be constructed in cumbersome fashion from classical block ciphers. Nowadays, we have the option of using efficient dedicated TBCs, such as Deoxys-BC, Joltik-BC [16], or Skinny [1].

The application of TBCs can also boost the efficiency of constructions, as has been demonstrated recently for MACs. At CRYPTO'17, Iwata et al. [15] introduced ZMAC, a TBC-based parallelizable, single-key single-primitive MAC whose internal hash function ZHash processed the message in both the tweak and plaintext simultaneously. The additional message bits per primitive call render ZMAC more efficient than previous MACs and suggest the adoption of the approach to other domains.

Open Research Questions. When abstracting away the details of the primitive, the number of calls to it per input block becomes the main efficiency metric. From Encrypt-Mix-Encrypt-based constructions, it is well-known that the bound is at most two calls per block (plus some minor overhead), assuming all further operations are linear. Thus, it is an interesting question if SPRPs can be built from fewer calls to a single-keyed primitive. Moreover, a strongly related question is that for the minimal number of calls necessary for SPRP security.

From a theoretical perspective, Nandi [26] showed that constructions built from a classical single-keyed block cipher require 2ℓ calls for ℓ-block messages for SPRP security. Though, it seems as though this bound is reducible if one used a TBC instead as the underlying primitive. For Hash-Counter-Hash-based

Table 1: Asymptotic #primitive calls for SPRP paradigms. We assume that hash functions and encryption layers use a single-keyed (tweakable) block cipher with n-bit state and τ-bit tweak size to encrypt an ℓ-block message of σ bits in total. We assume the hashing layers use ZHash (as the most efficient blockcipher-based hash function we are aware of).

	#Block-cipher calls			
Paradigm	Top	Middle	Bottom	Total (asympt.)
LargeBlock2	$2\lceil(\ell-1)/2\rceil$	ℓ	$2\lceil(\ell-1)/2\rceil$	$4\lceil(\ell-1)/2\rceil+\ell$
TCT_2	$2\lceil(\ell-1)/2\rceil$	$2(\ell-1)$	$2\lceil(\ell-1)/2\rceil$	$4\lceil(\ell-1)/2\rceil+2\ell$
Encrypt-Mix-Encrypt	ℓ	$\lceil\ell/n\rceil$	ℓ	$2\ell+\lceil\ell/n\rceil$
Hash-ECB-Hash	ℓ	ℓ	ℓ	3ℓ
Hash-Counter-Hash	$\lceil\sigma/(n+\tau)\rceil$	ℓ	$\lceil\sigma/(n+\tau)\rceil$	$\ell+2\lceil\sigma/(n+\tau)\rceil$
ZCZ	$\ell/2$	$\ell/2+\lceil\ell/2n\rceil$	$\ell/2$	$3\ell/2+\lceil\ell/2n\rceil$

constructions, the most efficient (T)BC-based hash function we are aware of is ZHash. For a TBC with n-bit state and τ-bit tweak length, it would yield a construction of about $\ell + 2\lceil\sigma/(n+\tau)\rceil$ calls for messages of σ bits. For dedicated TBCs, such as Deoxys-BC-128-384 or Skinny-128-384, this figure still implies that approximately $5\ell/3$ calls are necessary. Regarding the other design principles, it is unclear if similar results are applicable to constructions based on the Encrypt-Mix-Encrypt or Hash-ECB-Hash paradigms. We estimate that Hash-ECB-Hash constructions would need about ℓ primitive calls in each hashing layer, plus ℓ calls in the encryption layer. An instantiation of LargeBlock2 with ZHash instead of multiplications would yield $2\lceil(\ell-1)/2\rceil$ calls in each hashing layer, plus ℓ calls in the middle, or 3ℓ calls in sum. TCT_2 could use a ZHash layer each for both top and bottom hashing layer. While further modifications could make it more efficient, its proposal employed $2\ell-2$ calls in the middle. We compare the approaches in Table 1. Altogether, three interesting research questions remain: (1) to which extent can the number of primitive calls be reduced when employing a tweakable block cipher, (2) how can a specific construction be realized, and (3) can it be built with high provable security guarantees.

Contribution. This work tries to answer all three questions above: for the theoretical interest, (1) we show that 1.5ℓ primitive calls per message block is close to minimal by a generic distinguisher on any construction that employs fewer than $(3\ell-1)/2$ calls to a single-keyed primitive per message block, where all further operations are linear. For the practitioner's interest, (2) we propose ZCZ (ZHash-Counter-ZHash), an almost fully parallelizable variable-input-length SPRP based on a single-keyed TBC with n-bit state and n-bit tweak size. ZCZ matches approximately the optimal number of 1.5ℓ calls to the primitive for an ℓ-block message, plus a small overhead. Finally, we show (3) that ZCZ achieves optimal n-bit security, i.e., the SPRP advantage of any adversary that asks at most q queries of σ blocks in total is in $O(\sigma^2/2^{2n})$.

We note that instantiations of Hash-Counter-Hash with ZHash and a TBC with large tweaks of $\tau = 3n$, the number of primitive calls could become equal to that of ZCZ. However, such primitives would introduce a significant slowdown, be it due to the requirements of more rounds in a TWEAKEY-like cipher, or due to the need of calling an additional universal hash function for compressing the tweak. Concerning practical tweak sizes $\tau < 3n$, the number of calls is significantly lower for our construction.

Yet Another Encryption Scheme? It may appear that ZCZ is yet another encryption scheme after all, and with hundreds of encryption schemes already being present in the literature, it is difficult get excited about another one, notwithstanding small improvements in performance and security. We beg to differ on this point primarily for two reasons: (1) very few existing encryption schemes built upon a primitive with an n-bit output provide n-bit security — most in fact are only secure up to the birthday bound. As such, the improvement by ZCZ in terms of security is not a small step, but rather a leap. Since there is a considerable interest in the (still) small group of constructions that achieve this security, we believe that our encryption scheme is an exciting addition to this group. (2) Even more significant is the way that ZCZ exploits the randomness generated by a tweakable blockcipher. While most previous approaches were based on generic replacements of two or more blockcipher calls by a single call to a tweakable block cipher, the approach used by ZCZ is not a corollary of any previous work. Given its efficiency, we believe it can lead to exciting new directions in research on tweakable-blockcipher modes.

Outline. The remainder is structured as follows: first, Section 2 briefly summarizes the necessary preliminaries. Given a primitive with an effective tweak size[3] $\tau = n$, Section 3 illustrates that every PRP with fewer than $3\ell - 1$ primitive calls for 2ℓ-block messages is insecure, which was the core motivation for our search for constructions with about 1.5ℓ calls. Subsequently, Section 4 defines our basic construction, which is first described for messages whose length is a positive multiple of $2n$ bits. Thereupon, Section 5 extends our definition to messages of more general lengths. Section 6 provides the details of our security analysis.

We provide further insights on the starting point of our research in the full version of this work [4]. Therein, we also discuss attacks on insecure preliminary variants that motivated our studies towards the final design of ZCZ.

2 Preliminaries

General Notation. We use lowercase letters x for indices and integers, uppercase letters X, Y for binary strings and functions, and calligraphic uppercase letters $\mathcal{X}, \mathcal{Y}$ for sets. We denote the concatenation of binary strings X and Y by $X \,\|\, Y$; we mostly treat bit strings as representations of elements in the finite field

[3] By effective tweak size, we mean the usable tweak domain without bits that are used for other purposes such as domain separation.

$\mathbb{F}_{2^n}$, which is the Galois Field $\mathbb{GF}(2^n)$ with a fixed irreducible polynomial $p(\mathtt{x})$. There, we interpret a bit string $(x_{n-1} \dots x_1 x_0)$ as polynomial $\sum_{i=0}^{n-1} a_i \cdot \mathtt{x}^i$ in $\mathbb{F}_{2^n}$. Bit x_i represents the coefficient $a_i \in \{0,1\}$, for $0 \leq i \leq n-1$, and the most significant bit is the leftmost, and the least significant bit is the rightmost bit. We denote the result of the addition of two elements as $X + Y$, which is equivalent to the XOR of X and Y. For tuples of bit strings $(X_1, \dots, X_x)$, $(Y_1, \dots, Y_x)$ of equal domain, we denote by $(X_1, \dots, X_x) + (Y_1, \dots, Y_x)$ the element-wise XOR, i.e., $(X_1 + Y_1, \dots, X_x + Y_x)$. Unless stated otherwise, we consider all additions of n-bit values to be in $\mathbb{F}_2^n$. Moreover, we will use $\oplus$ for the XOR of bit strings in illustrations. However, all additions and subtractions in sub- and superscripts that denote indices represent integer additions. We indicate the length of a bit string X in bits by $|X|$, and write X_i for the i-th block. Moreover, we denote by $X \twoheadleftarrow \mathcal{X}$ that X is chosen independently uniformly at random from the set $\mathcal{X}$. We define three sets of particular interest: $\mathsf{Func}(\mathcal{X}, \mathcal{Y})$ be the set of all functions $F : \mathcal{X} \to \mathcal{Y}$, $\mathsf{Perm}(\mathcal{X})$ the set of all permutations over $\mathcal{X}$, and $\widetilde{\mathsf{Perm}}(\mathcal{T}, \mathcal{X})$ for the set of tweaked permutations over $\mathcal{X}$ with associated tweak space $\mathcal{T}$.

$(X_1, \dots, X_x) \xleftarrow{n} X$ denotes that X is split into the minimal number of n-bit blocks possible i.e., $X_1 \,\|\, \dots \,\|\, X_x = X$, and $|X_i| = n$ for $1 \leq i \leq x-1$, and $|X_x| \leq n$. So, when $|X| > 0$, then $|X_x| > 0$. If $|X| = 0$, $Y \xleftarrow{x} X$ sets Y to the empty string. $\langle X \rangle_n$ denotes an encoding of an integer $X \in \mathbb{Z}_n$ as an n-bit string. For two sets $\mathcal{X}$ and $\mathcal{Y}$, a uniform random function $\rho : \mathcal{X} \to \mathcal{Y}$ maps inputs $X \in \mathcal{X}$ independently and uniformly at random to outputs $Y \in \mathcal{Y}$. For an event E, we denote by $\Pr[E]$ the probability of E; ε is the empty string. For a given set $\mathcal{X}$ and integer x, we define $\mathcal{X}^{\leq x} = \bigcup_{i=1}^{x} \mathcal{X}^i$ and $\mathcal{X}^+ = \bigcup_{j=1}^{\infty} \mathcal{X}^j$. For two integers n, k with $n \geq k \geq 1$, we denote the falling factorial as $(n)_k = \prod_{i=0}^{k-1} (n-i)$.

Adversaries. An adversary $\mathbf{A}$ is an efficient Turing machine that interacts with a given set of oracles that appear as black boxes to $\mathbf{A}$. We denote by $\mathbf{A}^{\mathcal{O}}$ the output of $\mathbf{A}$ after interacting with some oracle $\mathcal{O}$. We write $\Delta_{\mathbf{A}}(\mathcal{O}^1; \mathcal{O}^2) := |\Pr[\mathbf{A}^{\mathcal{O}^1} \Rightarrow 1] - \Pr[\mathbf{A}^{\mathcal{O}^2} \Rightarrow 1]|$ for the advantage of $\mathbf{A}$ to distinguish between oracles $\mathcal{O}^1$ and $\mathcal{O}^2$. All probabilities are defined over the random coins of the oracles and those of $\mathbf{A}$, if any. W.l.o.g., we assume that $\mathbf{A}$ never asks queries to which it already knows the answer.

A block cipher E with associated key space $\mathcal{K}$ and message space $\mathcal{M}$ is a mapping $E : \mathcal{K} \times \mathcal{M} \to \mathcal{M}$ such that for every key $K \in \mathcal{K}$, it holds that $E(K, \cdot)$ is a permutation over $\mathcal{M}$. A tweakable block cipher $\widetilde{E}$ with additional tweak space $\mathcal{T}$ is a mapping $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \to \mathcal{M}$ such that for every key $K \in \mathcal{K}$ and tweak $T \in \mathcal{T}$, it holds that $\widetilde{E}(K, T, \cdot)$ is a permutation over $\mathcal{M}$. We also write $\widetilde{E}_K^T(\cdot)$ as short form. In this work, we assume that SPRPs allow variable-length inputs, i.e., there is no single fixed length, but the length of the ciphertext always equals that of the plaintext and vice versa; moreover, over all inputs of equal length, the construction is a permutation. The advantage is defined as follows.

Definition 1 (SPRP Advantage). Let $\mathcal{K}$ be a non-empty set and $\mathcal{M} \subset \{0,1\}^*$. Let $\mathit{\Pi} : \mathcal{K} \times \mathcal{M} \to \mathcal{M}$ be a length-preserving permutation. Let $\pi \twoheadleftarrow$

$\mathsf{Perm}(\mathcal{M})$ be sampled from the set of all length-preserving permutations of $\mathcal{M}$, and $K \twoheadleftarrow \mathcal{K}$. Then, the SPRP advantage of $\mathbf{A}$ with respect to Π is defined as $\mathbf{Adv}^{\mathrm{SPRP}}_{\Pi}(\mathbf{A}) \stackrel{\mathrm{def}}{=} \Delta_{\mathbf{A}}(\Pi_K, \Pi_K^{-1}; \pi, \pi^{-1})$.

Definition 2 (STPRP Advantage). Let $\mathcal{K}$ and $\mathcal{T}$ be non-empty sets and let $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ denote a tweakable block cipher. Let $\widetilde{\pi} \twoheadleftarrow \widetilde{\mathsf{Perm}}(\mathcal{T}, \{0,1\}^n)$ and $K \twoheadleftarrow \mathcal{K}$. Then, the STPRP advantage of $\mathbf{A}$ w.r.t. $\widetilde{E}$ is defined as $\mathbf{Adv}^{\mathrm{STPRP}}_{\widetilde{E}}(\mathbf{A}) \stackrel{\mathrm{def}}{=} \Delta_{\mathbf{A}}(\widetilde{E}_K, \widetilde{E}_K^{-1}; \widetilde{\pi}, \widetilde{\pi}^{-1})$.

Definition 3 (Almost-XOR-Universal Hash Function). Let $\mathcal{K}$, $\mathcal{X}$, and $\mathcal{Y} \subseteq \{0,1\}^*$ be non-empty sets. Let $H : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ be a function keyed by $K \in \mathcal{K}$. We call H ϵ-almost-XOR-universal (ϵ-AXU) if, for all distinct $X, X' \in \mathcal{X}$ and any $\Delta \in \mathcal{Y}$, it holds that $\Pr_{K \twoheadleftarrow \mathcal{K}}[H_K(X) - H_K(X') = \Delta] \leq \epsilon$, where subtraction is in $\mathbb{F}_{2^n}$.

The H-Coefficient Technique. The H-coefficient technique is a proof method by Patarin [27]. It assumes that the results of the interaction of an adversary $\mathbf{A}$ with its oracles are collected in a transcript τ of the attack: $\tau = \langle (M_1, C_1, d_1), \ldots, (M_q, C_q, d_q) \rangle$. (M_i, C_i) denotes the in- and output of the i-th query of $\mathbf{A}$; a Boolean variable d_i denotes the query direction: $d_i = 1$ indicates that C_i was result of an encryption query, and $d_i = 0$that M_i was the result of a decryption query. The task of $\mathbf{A}$ is to distinguish the real world $\mathcal{O}_{\mathrm{real}}$ from the ideal world $\mathcal{O}_{\mathrm{ideal}}$. A transcript τ is called *attainable* if the probability to obtain τ in the ideal world is non-zero. We denote by Θ_{real} and Θ_{ideal} the distribution of transcripts in the real and the ideal world, respectively. Then, the fundamental Lemma of the H-coefficients technique, whose proof is given in [6, 27], states:

Lemma 1 (Fundamental Lemma of the H-coefficient Technique [27]). Assume that the set of attainable transcripts is partitioned into two disjoint sets GoodT and BadT. Further assume that there exist $\epsilon_1, \epsilon_2 \geq 0$ such that for any transcript $\tau \in \mathrm{GoodT}$, it holds that

$$\frac{\Pr[\Theta_{\mathrm{real}} = \tau]}{\Pr[\Theta_{\mathrm{ideal}} = \tau]} \geq 1 - \epsilon_1, \quad \text{and} \quad \Pr[\Theta_{\mathrm{ideal}} \in \mathrm{BadT}] \leq \epsilon_2.$$

Then, for all adversaries $\mathbf{A}$, it holds that $\Delta_{\mathbf{A}}(\mathcal{O}_{\mathrm{real}}; \mathcal{O}_{\mathrm{ideal}}) \leq \epsilon_1 + \epsilon_2$.

3 On the Minimal Number of Required Primitive Calls

This section shows that any PRP with fewer than $3\ell - 1$ calls for messages of 2ℓ blocks to a primitive with n-bit tweak size and n-bit state size is insecure. We follow the approach by [26], who proved that an SPRP based on a single-keyed classical block cipher needs at least 2ℓ calls to the primitive for ℓ-block messages.

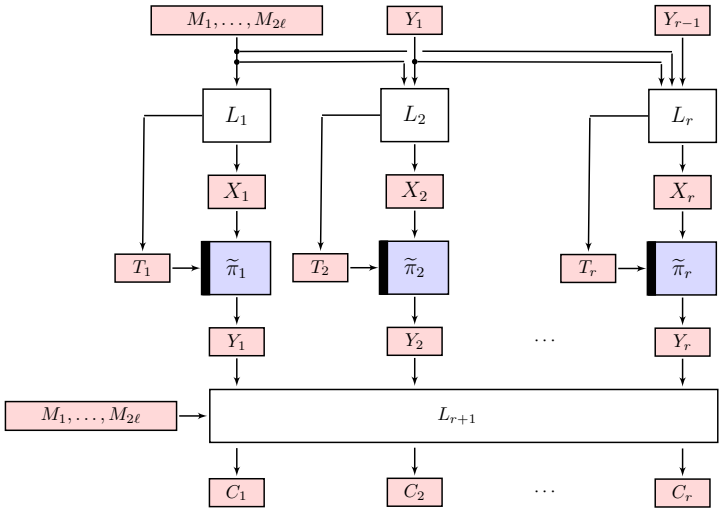

Fig. 1: Generic model of a PRP that consists of at most $r \leq 3\ell - 2$ calls to tweakable block ciphers $\widetilde{\pi}_i$ for messages of 2ℓ blocks.

3.1 Generic Construction

Define positive integers n, τ, and ℓ, and let $\mathcal{M} \subseteq \{0,1\}^*$ denote a space for which $(\{0,1\}^n)^{2\ell} \subseteq \mathcal{M}$. Let $r \leq 3\ell - 2$ and let $\widetilde{\pi}_i : \{0,1\}^\tau \times \{0,1\}^n \to \{0,1\}^n$, for all $1 \leq i \leq r$, denote tweakable permutations with tweak space $\{0,1\}^\tau$ and state size n. Let $\Pi[\widetilde{\pi}_1, \ldots, \widetilde{\pi}_r] : \mathcal{M} \to \mathcal{M}$ be a length-preserving cipher that employs as its only non-linear functions in total r calls to the permutations $\widetilde{\pi}_1, \ldots, \widetilde{\pi}_r$. For simplicity, we also write Π as short form, hereafter. All further components of Π are linear over $\mathbb{F}_{2^n}$. For any such construction, we can formulate this as follows. Let X_i denote the input to π_i, T_i the tweak to π_i, and let $Y_i \leftarrow \pi_i(X_i)$ denote its output. The linear operations in Π must be describable as non-zero linear functions $L_i : \mathcal{M} \times (\{0,1\}^n)^{i-1} \to \{0,1\}^n \times \{0,1\}^\tau$, for $1 \leq i \leq r$, and an additional non-zero linear function $L_{r+1} : \mathcal{M} \times (\{0,1\}^n)^r \to \mathcal{M}$ that, for all given inputs $(M, Y_1, \ldots, Y_r) \in \mathcal{M} \times (\{0,1\}^n)^r$, outputs C s.t. it holds that $|C| = |M|$. Then, we can describe the encryption with $\Pi(M)$ as

$$\begin{aligned} (X_i, T_i) &\leftarrow L_i\left(M, Y_1, \ldots, Y_{i-1}\right), && \text{for all } 1 \leq i \leq r, \\ Y_i &\leftarrow \widetilde{\pi}^{T_i}(X_i), && \text{for all } 1 \leq i \leq r, \text{ and} \\ C &\leftarrow L_{r+1}(M, Y_1, \ldots, Y_r). \end{aligned}$$

Π must be correct for all inputs, i.e., for all $M, C \in \mathcal{M}$, it must hold that $\Pi^{-1}(\Pi(M)) = M$ and $\Pi(\Pi^{-1}(C)) = C$. Figure 1 gives an illustration.

Remark 1. It may not be instantaneously clear why the generic construction above covers all considered schemes. Note that it computes the values X_i and T_i by a non-zero linear function of M, Y_1, Y_2, …, Y_{i-1}. So, the previous values Y_i

can also be used to generate X_i. Indeed, it is generic enough to include all such constructions where the only non-linear components are the permutation calls.

For simplicity, we consider independent permutations with tweak domain $\mathbb{F}_2^\tau$ in this section. For efficiency, our proposal later in this work will employ only a single tweakable primitive with a composite tweak domain $\mathcal{T}_D = \mathcal{D} \times \mathbb{F}_2^\tau$, where $\mathcal{D}$ is a non-empty set of domains. So, this approach achieves the same goal of having independent permutations. We consider that τ is the effectively usable size of the tweaks without domains.

3.2 A PRP Attack on Constructions with At Most $3\ell - 2$ Calls

Case $\tau = n$. Let $\mathbf{A}$ be an adversary with the goal to distinguish the outputs of a variable-input-length PRP Π under a secret key as above from an ideal PRP. First, $\mathbf{A}$ chooses two messages M and M' of 2ℓ blocks each, i.e., $M = (M_1, \ldots, M_{2\ell})$ and $M' = (M'_1, \ldots, M'_{2\ell})$. We define the differences $\Delta M = M - M'$, and analogously the differences ΔX_i, ΔY_i, and ΔC in the obvious manner. Choose M and M' such that it holds that $\Delta X_i = 0$ and $\Delta T_i = 0$, for $1 \leq i \leq \ell - 1$. Note that such a choice of M and M' must be possible since these variables correspond to $2\ell - 2$ equations ($\ell - 1$ equations for adjusting the values ΔX_i and $\ell - 1$ equations for adjusting the values ΔT_i) and there exist 2ℓ blocks ΔM_i. For instance, the adversary can efficiently derive an element N from the null space of $L_1, \ldots, L_{2(\ell-1)}$. It chooses M arbitrarily and derives $M' = M + N$.

From $\Delta X_i = 0^n$ and $\Delta T_i = 0^\tau$ for $1 \leq i \leq \ell - 1$, it follows that $\Delta Y_i = \widetilde{\pi}^{T_i}(X_i) \oplus \widetilde{\pi}^{T'_i}(X'_i) = 0^n$, for all $1 \leq i \leq \ell - 1$. The non-linear layer of calls to the tweakable block cipher maps $(\Delta X_1, \ldots, \Delta X_r)$ to $(\Delta Y_1, \ldots, \Delta Y_r)$. We obtain

$$L_{r+1}(\Delta M, \underbrace{\Delta Y_1, \ldots, \Delta Y_{\ell-1}}_{=(0,\ldots,0)}, \Delta Y_\ell, \ldots, \Delta Y_r) = \Delta C.$$

Since $\mathbf{A}$ fixed ΔM and chose M and M' so that $\Delta X_1 = \ldots = \Delta X_{\ell-1} = 0^n$ and $\Delta T_1 = \ldots = \Delta T_{\ell-1} = 0^\tau$, we obtain $\Delta Y_1, \ldots, \Delta Y_{\ell-1} = 0^n$. So, there are at most $2\ell - 1$ free variables $\Delta Y_\ell, \ldots \Delta Y_r$, and 2ℓ equations for $\Delta C_1, \ldots, \Delta C_{2\ell}$, which implies that 2ℓ blocks of ΔC are a linear combination of $2\ell - 1$ values $\Delta Y_\ell, \ldots, \Delta Y_r$. So, in the real construction, L_{r+1} defines a map from $2\ell - 1$ to 2ℓ n-bit variables, and $\mathbf{A}$ can efficiently derive a solution $\Delta Y_\ell, \ldots, \Delta Y_r$ from the null space of the equation system. This becomes a distinguishing event happening with probability one in the real construction and with probability $1/2^n$ in the ideal world for this example. The distinguishing advantage is hence $1 - 1/2^n$. $\mathbf{A}$ can query it with two messages as above and output real if such a non-zero linear function L exists and random otherwise, as summarized in Algorithm 1.

For general values of τ. A similar attack is applicable for general values of τ. Though, we have to consider linearity over $\mathbb{F}_2$ then. Define

$$s = \left\lfloor \frac{2\ell n}{n + \tau} \right\rfloor - 1.$$

Algorithm 1 PRP attack on generic constructions Π with at most $3\ell - 2$ primitive calls, here for $\tau = n$.

function $\mathbf{A}^{\Pi}$
  Choose M_i for $1 \leq i \leq 2\ell$ arbitrarily
  Choose M'_i for $\ell \leq i \leq 2\ell$ s. t. it holds that
    $L_i(\Delta M_i) = (\Delta X_i, \Delta T_i) = (0^n, 0^\tau)$, for $1 \leq i \leq 2(\ell - 1)$
  Ask for the encryption of $C = \Pi(M)$ and $C' = \Pi(M')$
  Derive $\Delta C = C' - C$
  if there exists $(\Delta Y_\ell, \ldots, \Delta Y_r)$, s. t. $L_{r+1}(\Delta M, \Delta Y) = \Delta C$ **then**
    return "Real"
  return "Random"

The adversary chooses $M \in (\mathbb{F}_2^n)^{2\ell}$ arbitrarily, and $M' \in (\mathbb{F}_2^n)^{2\ell}$ with $M \neq M'$ s. t. $\Delta X_1 = \ldots \Delta X_s = 0^n$ and $\Delta T_1 = \ldots = \Delta T_s = 0^\tau$. Note that we consider the inputs $X_i \in \mathbb{F}_2^n$ and the tweaks $T_i \in \mathbb{F}_2^\tau$ as blocks. Again, such a choice of M' exists for the same reason as above and can be found efficiently from the null space of the linear functions L_1, L_2, ... that are involved in the computation of $\Delta X_1, \ldots, \Delta X_s$ and $\Delta T_1, \ldots, \Delta T_s$. Again, we obtain $\Delta Y_i = 0^n$, for $1 \leq i \leq s$ for the real construction. We obtain the equation

$$L_{r+1}(\Delta M, \underbrace{\Delta Y_1, \ldots, \Delta Y_s}_{=(0,\ldots,0)}, \Delta Y_{s+1}, \ldots, \Delta Y_r) = \Delta C.$$

The blocks $\Delta Y_{s+1}, \ldots, \Delta Y_r$ contain $(r-s)n$ bits, that are mapped through L_{r+1} to $\Delta C_{2\ell n}$ bits. For all schemes Π that use r calls to the primitive with

$$(r - s) \cdot n < 2\ell n, \qquad \text{which leads to} \qquad r < 2\ell \left(1 + \frac{n}{n + \tau}\right) - 1,$$

we obtain a compressing mapping. Then, there exist are more equations than variables, and the distinguisher as before applies. However, the advantage may be smaller and depends on the values of r, n, and τ.

4 Definition of The Basic ZCZ Construction

This section defines the basic ZCZ scheme. First, we consider messages that consist of at most $2n$ blocks, and will extend it thereupon to all messages whose length is a multiple of $2n$ bits. The subsequent section will then further define it for messages whose lengths are not necessarily multiples of $2n$ bits.

Parameters. Let $n, \tau, k, d \geq 1$ be integers with $d \ll n$ and $n = \tau$; we define $N \stackrel{\text{def}}{=} 2^n$ as an alias. Let $\mathcal{B} = \{0,1\}^{2n}$ define a *di-block* (or dual block, double block), i.e., $2n$ bits. We define non-empty sets of tweaks $\mathcal{T} = \{0,1\}^\tau$, keys $\mathcal{K} = \{0,1\}^k$, domains $\mathcal{D} = \{\mathsf{t}, \mathsf{s}, \mathsf{c}, \mathsf{b}, \mathsf{t\$}, \mathsf{s\$}, \mathsf{c\$}, \mathsf{b\$}, \mathsf{xl}, \mathsf{xr}, \mathsf{yl}, \mathsf{yr}, \mathsf{p}, \mathsf{kd}\} \subseteq \{0,1\}^d$, and a set of indices $\mathcal{I} \subseteq \{1, \ldots, 2^n - 1\}$. The purpose of domains and indices is to define an extended tweak set $\mathcal{T}_{D,I} = \mathcal{D} \times \mathcal{I} \times \mathcal{T}$ for a tweakable block cipher $\widetilde{E} : \mathcal{K} \times \mathcal{T}_{D,I} \times \{0,1\}^n \to \{0,1\}^n$.

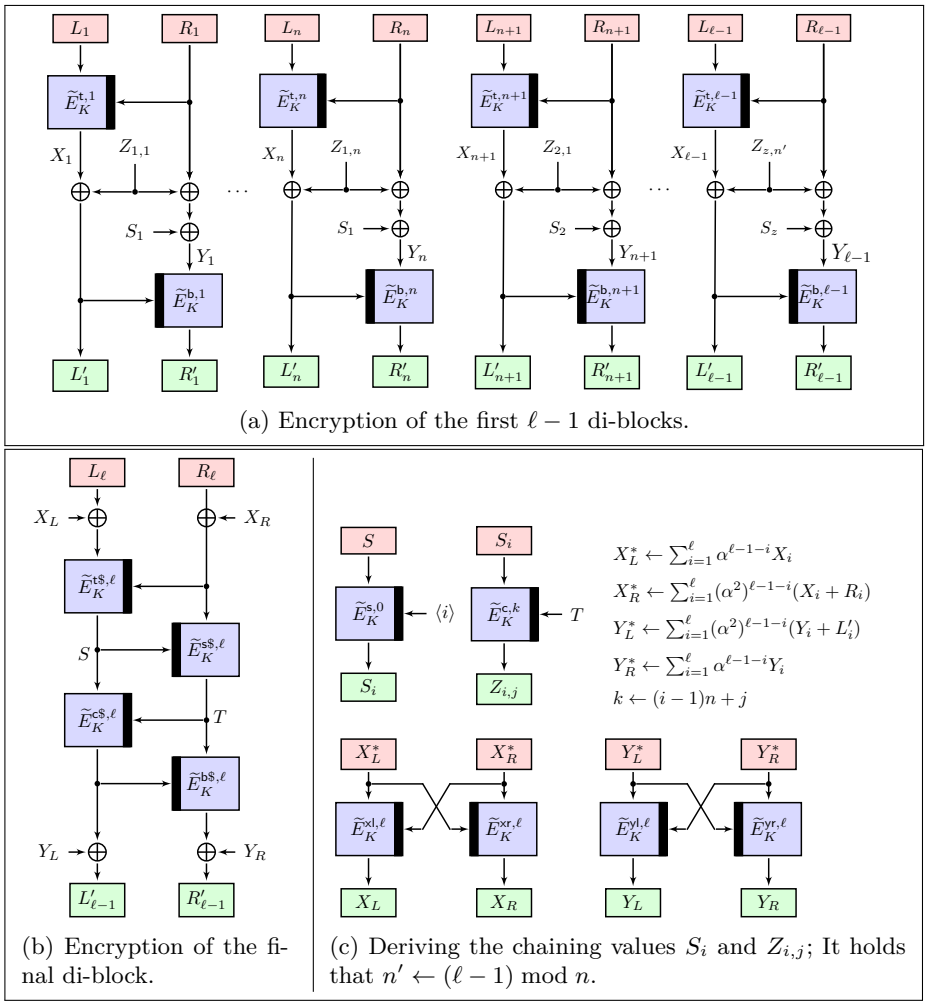

(a) Encryption of the first $\ell - 1$ di-blocks.

(b) Encryption of the final di-block.

(c) Deriving the chaining values S_i and $Z_{i,j}$; It holds that $n' \leftarrow (\ell - 1) \bmod n$.

Fig. 2: Encryption of a message with ℓ complete di-blocks with ZCZ$[\widetilde{E}_K]$.

Overview. The basic ZCZ$[\widetilde{E}_K]$ construction takes as input a secret key $K \in \mathcal{K}$ and a plaintext $M \in \mathcal{B}^{\leq n}$ that is split into $\ell \in [1..n]$ di-blocks. The design can be split into a top, middle, and a bottom layer. In the top layer, the first $\ell - 1$ complete di-blocks (L_i, R_i) are processed similarly as in the $\mathbb{ZHASH}$ construction by Iwata et al. [15]. The TBC outputs X_i are accumulated by an MDS code to two values X_L^* and X_R^* using the Horner rule, which are finally encrypted in a butterfly-like structure [24] to $X_L \leftarrow \widetilde{E}_K^{\mathsf{xl},\ell,X_R^*}(X_L^*)$ and $X_R \leftarrow \widetilde{E}_K^{\mathsf{xr},\ell,X_L^*}(X_R^*)$. X_L and X_R are used to mask the branches of the final di-block, L_ℓ and R_ℓ. The final di-block is processed by a four-round Feistel-like network of four TBC calls in the spirit of the constructions by Coron et al. [9].

This four-round network generates two intermediate values S and T after the first and second call to $\widetilde{E}$. The middle layer derives from S and T a value $S_1 \leftarrow \widetilde{E}_K^{\mathsf{s},0,1}(S)$ and a series of $\ell - 1$ chaining values $Z_{1,j} \leftarrow \widetilde{E}^{\mathsf{c},j,T}(S_1)$. For the j-th di-block, the chaining value $Z_{1,j}$ is added to both branches of the j-th block. Moreover, S_1 is also added to the right branch of each di-block: $L'_j \leftarrow X_j + Z_{i,j}$ and $Y_j \leftarrow R_j + Z_{1,j} + S_1$. So, this middle layer ensures that each di-block depends on all others. Finally, the middle layer generates from the blocks Y_j and L'_j two values Y_L and Y_R symmetrically as X_L and X_R, from the values Y_j.

The bottom layer is then a symmetric version of the top layer. The $\ell - 1$ di-blocks are processed by another $\mathbb{ZHASH}$ layer to compute the ciphertext blocks: $L'_j \leftarrow X_j$ and $R'_j \leftarrow \widetilde{E}_K^{\mathsf{b},i,L'_j}(Y_j)$. The final complete di-block is processed by two further Feistel rounds before Y_L added to the left branch, and Y_R is added to the right branch of the ℓ-th di-block. The resulting values L'_i, R'_i, for $1 \le i \le \ell$, are concatenated and returned as the ciphertext. The details of the encryption with ZCZ$[\widetilde{E}_K]$ is given in Algorithm 2, and is illustrated in parts in Figure 2, already for more than n complete di-blocks.

Rationale. The structure is inspired by ZHash [15] and AEZ [12]. The use of α and α^2 prevents that a collision in X_L would automatically lead to a collision also in X_R and vice versa; considering also the tweak values R_i for X_R renders birthday collisions in X_i from separate tweaks ineffective. Encrypting X_L^*, X_R^*, Y_L^*, and Y_R^* avoids that differences in the masks cancel differences in the final di-block. Finally, adding S_i and $Z_{i,j}$ prevents adversaries from observing differences $\Delta Z_{1,j}$. Using the masks X_L, X_R, Y_L, and Y_R in the final block makes its outputs depend on all blocks; Using S and T for the counter mode in the middle layer creates a dependency of each di-block on all others. We elaborate on attacks on preliminary versions of ZCZ in the full version of this work. We employ pairwise distinct domains for all calls to $\widetilde{E}$ to prevent dependencies between the calls.

Extension to Longer Messages. Messages with more than n di-blocks are partitioned into *chunks*. The i-th (complete) chunk denotes the series of the n consecutive di-blocks $(L_{(i-1)n+1}, R_{(i-1)n+1}, \ldots L_{i \cdot n}, R_{i \cdot n})$, and employs the chaining values S_i and $Z_{i,j}$. We derive all chaining values under distinct domains as before. Furthermore, we derive $\ell - 1$ chaining values $Z_{i,j}$ by a TBC call each from S. For the i-th chunk, S_i is computed as $S_i \leftarrow \widetilde{E}_K^{\mathsf{s},0,i}(S)$. Then, for $j \in [1..n]$, $Z_{i,j}$ for the j-th block of the i-th chunk is generated as $Z_{i,j} \leftarrow \widetilde{E}_K^{\mathsf{c},0,n(i-1)+j}(S_i)$. $Y_{n(i-1)+j}$ is then computed as $Y_{n(i-1)+j} \leftarrow R_{n(i-1)+j} + S_i + Z_{n(i-1)+j}$. The rest of the computations remain unchanged. Letting j take any value in $[1..\ell]$, we can rewrite this as

$$Y_j \leftarrow R_j + S_{\lceil j/n \rceil} + Z_j. \tag{2'}$$

The encryption of ZCZ$[\widetilde{E}_K]$ is defined in Algorithm 2, and illustrated in parts in Figure 2, already for more than n complete di-blocks. The figure employs bold bars in the blocks of $\widetilde{E}$ to indicate the parts of the tweaks that stem from $\mathcal{T}$. The decryption is defined in the obvious way.

Algorithm 2 Definition of the encryption algorithm of $\mathrm{ZCZ}[\widetilde{E}]$ given a tweakable block cipher $\widetilde{E}$. The code in the boxes is only part of $\mathrm{ZCZ}^*[\widetilde{E}]$ in Algorithm 3.

10: **function** $\mathrm{ZCZ}[\widetilde{E}_K](M)$
11: $r \leftarrow |M| \bmod 2n$
12: $\ell \leftarrow (|M| - r)/2n$
13: $z \leftarrow \lceil (\ell - 1)/n \rceil$
14: $L'_* \leftarrow \varepsilon$; $R'_* \leftarrow \varepsilon$
15: $\textsc{Parse}(M, \ell)$
16: $\textsc{TopEnc}[\widetilde{E}_K]()$
17: **if** $r > 0$ **then**
18: $\boxed{\textsc{PartialTopEnc}[\widetilde{E}_K]()}$
19: $\textsc{LastTopEnc}[\widetilde{E}_K](X_L, X_R)$
20: $\textsc{MidLayer}[\widetilde{E}_K](S, T)$
21: $\textsc{BotEnc}[\widetilde{E}_K]()$
22: $\textsc{LastBotEnc}[\widetilde{E}_K](Y_L, Y_R)$
23: **if** $r > 0$ **then**
24: $\boxed{\textsc{PartialBotEnc}[\widetilde{E}_K]()}$
25: $C \leftarrow (L'_1 \| R'_1 \| \cdots \| L'_\ell \| R'_\ell \| L'_* \| R'_*)$
26: **return** C

30: **procedure** $\textsc{TopEnc}[\widetilde{E}_K]$
31: $X^*_L \leftarrow X^*_R \leftarrow 0^n$
32: **for** $i \leftarrow 1 \ldots \ell - 1$ **do**
33: $X_i \leftarrow \widetilde{E}_K^{\mathsf{t},i,R_i}(L_i)$
34: $X^*_L \leftarrow X^*_L + \alpha^{\ell-1-i} X_i$
35: $X^*_R \leftarrow X^*_R + (\alpha^2)^{\ell-1-i}(X_i + R_i)$
36: $X_L \leftarrow \widetilde{E}_K^{\mathsf{xl},\ell,X^*_R}(X^*_L)$
37: $X_R \leftarrow \widetilde{E}_K^{\mathsf{xr},\ell,X^*_L}(X^*_R)$

40: **procedure** $\textsc{MidLayer}[\widetilde{E}_K](S, T)$
41: $S_0 \leftarrow S$
42: **for** $i \leftarrow 1 \ldots z$ **do**
43: $S_i \leftarrow \widetilde{E}_K^{\mathsf{s},0,i}(S_{i-1})$
44: **for** $i \leftarrow 1 \ldots z$ **do**
45: **for** $j \leftarrow 1 \ldots n$ **do**
46: $Z_{i,j} \leftarrow \widetilde{E}_K^{\mathsf{c},(i-1)n+j,T}(S_i)$

50: **procedure** $\textsc{LastTopEnc}[\widetilde{E}_K](X_L, X_R)$
51: $S \leftarrow \widetilde{E}_K^{\mathsf{t\$},\ell,R_\ell+X_R}(L_\ell + X_L)$
52: $T \leftarrow \widetilde{E}_K^{\mathsf{s\$},\ell,S}(R_\ell + X_R)$

60: **procedure** $\textsc{BotEnc}[\widetilde{E}_K]$
61: $Y^*_L \leftarrow 0^n$
62: $Y^*_R \leftarrow 0^n$
63: **for** $i \leftarrow 1 \ldots z - 1$ **do**
64: **for** $j \leftarrow 1 \ldots n$ **do**
65: $k \leftarrow (i-1)n + j$
66: $L'_k \leftarrow X_k + Z_{i,j}$
67: $Y_k \leftarrow R_k + Z_{i,j} + S_i$
68: $R'_k \leftarrow \widetilde{E}_K^{\mathsf{b},k,L'_k}(Y_k)$
69: $Y^*_L \leftarrow Y^*_L + (\alpha^2)^{\ell-1-k}(Y_k + L'_k)$
70: $Y^*_R \leftarrow Y^*_R + (\alpha)^{\ell-1-k} Y_k$
71: **for** $j \leftarrow 1 \ldots \ell - 1 - (z-1)n$ **do**
72: $k \leftarrow (z-1)n + j$
73: $L'_k \leftarrow X_k + Z_{z,j}$
74: $Y_k \leftarrow R_k + Z_{z,j} + S_z$
75: $R'_k \leftarrow \widetilde{E}_K^{\mathsf{b},k,L'_k}(Y_k)$
76: $Y^*_L \leftarrow Y^*_L + (\alpha^2)^{\ell-1-k}(Y_k + L'_k)$
77: $Y^*_R \leftarrow Y^*_R + \alpha^{\ell-1-k} Y_k$
78: $Y_L \leftarrow \widetilde{E}_K^{\mathsf{yl},\ell,Y^*_R}(Y^*_L)$
79: $Y_R \leftarrow \widetilde{E}_K^{\mathsf{yr},\ell,Y^*_L}(Y^*_R)$

80: **procedure** $\textsc{LastBotEnc}[\widetilde{E}_K](Y_L, Y_R)$
81: $L'_\ell \leftarrow \widetilde{E}_K^{\mathsf{c\$},\ell,T}(S) + Y_L$
82: $R'_\ell \leftarrow \widetilde{E}_K^{\mathsf{b\$},\ell,T}(L'_\ell + Y_L) + Y_R$

90: **procedure** $\textsc{Parse}(M, \ell)$
91: $i \leftarrow \ell \cdot 2n$
92: $(L_1, R_1, \cdots, L_\ell, R_\ell) \xleftarrow{n} M[0..i-1]$
93: **if** $r > 0$ **then**
94: $(L_*, R_*) \xleftarrow{n} M[i..|M|]$

5 ZCZ* for Messages with Partial Final Di-block

We extend the definition of ZCZ to messages whose length is not a multiple of $2n$ bits. We denote the last $r \leftarrow |M| \bmod 2n$ bits as *partial di-block*. Our approach for ZCZ* is inspired by the DE domain extender from [25]. Therefore, we briefly recap it.

The Domain Extender $\mathrm{DE}[\Pi, F, H] : \{0,1\}^{\geq n} \to \{0,1\}^{\geq n}$ [25] takes a block-wise-operating length-preserving permutation $\Pi : (\{0,1\}^n)^+ \to (\{0,1\}^n)^+$, a PRF $F : \{0,1\}^n \to \{0,1\}^n$, and an XOR-universal hash function $H : \{0,1\}^n \times$

Algorithm 3 Functions of the encryption algorithm of $\mathrm{ZCZ}^*[\widetilde{E}]$ for messages whose length is not necessarily a multiple of $2n$ bit (but at least $2n$ bit). Recall that $r = |M| \bmod 2n$.

procedure PartialTopEnc$[\widetilde{E}_K]$
  $M_\ell \leftarrow L_\ell \,\|\, R_\ell$
  $M_* \leftarrow \mathsf{pad}_{2n}(L_* \,\|\, R_*)$
  $(\overline{L}_*, \overline{R}_*) \xleftarrow{n} M_*$
  $(U_\ell, V_\ell) \leftarrow \mathcal{H}[\widetilde{E}_K, 0](\overline{L}_*, \overline{R}_*)$
  $L_\ell \leftarrow L_\ell + U_\ell$
  $R_\ell \leftarrow R_\ell + V_\ell$

function $\mathsf{msb}_x(X)$
  return $X[0..x-1]$

function $\mathsf{pad}_x(X)$
  return $X \,\|\, 1 \,\|\, 0^{x-|X|-1}$

procedure PartialBotEnc$[\widetilde{E}_K]$
  $(P, Q) \leftarrow \mathcal{H}[\widetilde{E}_K, 2](L_\ell + L'_\ell, R_\ell + R'_\ell)$
  $W \leftarrow \mathsf{msb}_r(P \,\|\, Q) \,\|\, 0^{2n-r}$
  $(P_*, Q_*) \xleftarrow{n} W$
  $L'_* \leftarrow L_* + P_*$
  $R'_* \leftarrow R_* + Q_*$
  $(\overline{L}'_*, \overline{R}'_*) \xleftarrow{n} \mathsf{pad}_{2n}(L'_* \,\|\, R'_*)$
  $(U'_\ell, V'_\ell) \leftarrow \mathcal{H}[\widetilde{E}_K, 4](\overline{L}'_*, \overline{R}'_*)$
  $L'_\ell \leftarrow L'_\ell + U'_\ell$
  $R'_\ell \leftarrow R'_\ell + V'_\ell$

function $\mathcal{H}[\widetilde{E}_K, i](U, V)$
  $U' \leftarrow \widetilde{E}_K^{\mathsf{p},i,V}(U)$
  $V' \leftarrow \widetilde{E}_K^{\mathsf{p},i+1,V}(U)$
  return (U', V')

$\{0,1\}^{2n} \to \{0,1\}^n$. It produces a length-preserving permutation over bit strings of any length $\geq n$ bits. A message $M \in \{0,1\}^{\geq n}$ is split into blocks $(M_1, \ldots, M_\ell)$; $\mathrm{DE}[\Pi, F, H]$ computes the corresponding ciphertext $C = (C_1, \ldots, C_\ell)$ as: (1) $M^*_{\ell-1} \leftarrow H(M_{\ell-1}, M_\ell)$, (2) $(C_1, \ldots, C_{\ell-2}, C^*_{\ell-1}) \leftarrow \Pi(M_1, \ldots, M_{\ell-2}, M^*_{\ell-1})$, (3) $C_\ell \leftarrow F(M^*_{\ell-1} + C^*_{\ell-1}) +_{|M_\ell|} M_\ell$, and (4) $C_{\ell-1} \leftarrow H(C^*_{\ell-1}, C_\ell)$. where

$$x +_n y \stackrel{\text{def}}{=} \mathsf{msb}_n(x) + y$$

for any $x, y \in \{0,1\}^*$ and integer n. To obtain that DE is a permutation, the hash function H must satisfy $H(H(M_{\ell-1}, M_\ell), M_\ell) = M_{\ell-1}$ for any allowed input $M_{\ell-1}$, M_ℓ (see [25, Remark 2]).

Overview of ZCZ*. Our extension ZCZ^* requires that the message length is still at least $2n$ bits. Let $M_* = (L_*, R_*)$ be the partial message di-block that follows after ℓ complete di-blocks. Further assume that the partial di-block consists of $\geq n$ bits that are split into $|L_*| = n$ and $|R_*| < n$. The right part is padded to n bits by a single 1 and as many zero bits as necessary to extend it to n bits: $\overline{R}_* \leftarrow \mathsf{pad}_n(R_*)$. The values are given as inputs to a hash function $\mathcal{H}[\widetilde{E}_K, i]$, with $i = 0$, that is illustrated on the right side of Figure 3. $\{H\}$ uses one of the two n-bit values as state and the other one as tweak input for two calls to $\widetilde{E}_K$ under distinct tweaks: $U' \leftarrow \widetilde{E}_K^{\mathsf{p},i,V}(U)$ and $V' \leftarrow \widetilde{E}_K^{\mathsf{p},i+1,V}(U)$. The $2n$-bit output (U', V') is added to the final complete di-block. The resulting final di-block (L_ℓ, R_ℓ) is then processed by $\mathrm{ZCZ}[\widetilde{E}_K]$. The sum of $(L_\ell, R_\ell) + (L'_\ell, R'_\ell)$ is then given again into $\mathcal{H}[\widetilde{E}_K, i]$, with $i = 2$ to produce a $2n$-bit value (P'_ℓ, Q'_ℓ). The most significant r bits of it are added to the final partial di-block to obtain the partial ciphertext di-block M'_*. M'_* is again padded to $2n$ bits and given

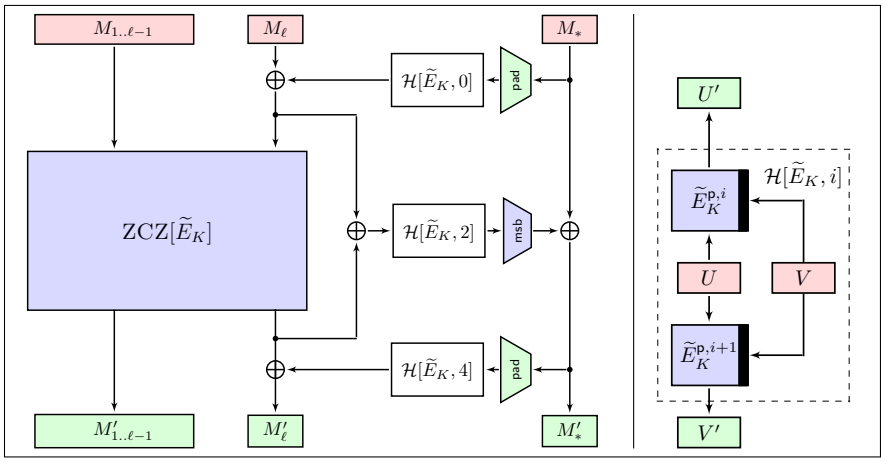

Fig. 3: Encryption of a partial message $M_1, \ldots, M_\ell, M^*$ whose length is not a multiple of $2n$ bit with $\mathrm{ZCZ}^*[\widetilde{E}_K]$. All preceeding di-blocks $M_1, \ldots, M_\ell$ are processed with $\mathrm{ZCZ}[\widetilde{E}_K]$ as before.

as input to a third call to $\mathcal{H}[\widetilde{E}_K, i]$, with $i = 4$. The hash output is added to the final ciphertext di-block to produce M'_ℓ. If the partial di-block consists of less than n bits, it is also padded to $2n$ bits and processed analogously. So, the hash function H from the original definition of $\mathrm{DE}[\Pi, F, H]$ is given by $H(M_\ell, M_*) \stackrel{\text{def}}{=} M_\ell + \mathcal{H}[\widetilde{E}_K, i](\mathsf{pad}_{2n}(M_*))$. One can see that the requirement from above holds for arbitrary M_ℓ and M_*: $H\left(H\left(M_\ell, M_*\right), M_*\right) = M_\ell$.

Remark 2. Note that ZCZ^* still requires messages to consist of at least $2n$ bits. A further minor improvement in future work could be the integration of smaller messages. For instance, the use of the very recent length-doubling construction LDT [7] could reduce the minimal message length to $n + 1$ bits. Though, this step would require an appropriate integration and ZCZ^* is already a variable-input-length SPRP for lengths $\geq 2n$ bit.

6 Security Analysis of ZCZ and ZCZ*

This section studies the SPRP security of ZCZ and ZCZ^*. Figure 4 provides a high-level overview on ZCZ. A given message M is split an input message into (M_L, M_R), where M_R consists of one $2n$-bit di-block, and M_L of the remaining di-blocks; the major part M_L is then processed by a variant of ZHash, that is denoted ZHash* here. It differs from ZHash in two aspects: ZHash* omits the XOR of the TBC output to the tweak input blocks. More prominently, ZHash* does not compress the input to two hash values, but is a permutation over $(n + \tau)^*$. So, the top layer returns the TBC outputs and the tweaks. $\widetilde{V}_1$ and $\widetilde{V}_2$ represent tweakable permutations. Internally, they can use the same primitive

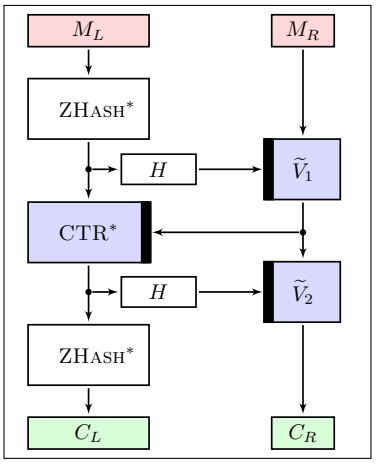

Fig. 4: High-level view on our proposal of ZCZ.

as also for ZHash*, and the tweakable variant of Counter mode, CTR*. H symbolizes an error-correcting code that sums up the inputs to $2n$ bits.

This high-level view allows to give a rationale for a dedicated analysis. A straight-forward use of a rate-1 counter mode would allow to apply a standard generic proof as for HCTR. Though, such an approach would yield 2ℓ calls to the primitive alone in the counter mode. In combination with ZHash*, this approach would need 4ℓ calls to the primitive for messages of 2ℓ blocks. ZCZ considers a special variant of counter mode that uses only ℓ blocks of entropy to mask 2ℓ blocks, similar as has been used in AEZ from version 2 [12]. However, this counter mode disallows to simply adopt the analysis from HCTR-like constructions when the goal is showing n-bit security. So, a dedicated analysis is needed, which is a major contribution of the present work. In the following, we study the security of the basic construction before we consider the extensions for inputs whose length is not necessarily a multiple of $2n$ bits, but at least $2n$ bits. We show the security of the extension ZCZ* at the end of this section.

6.1 Security of The Basic Construction

Theorem 1. Let $\widetilde{\pi} \twoheadleftarrow \widetilde{\mathsf{Perm}}(\mathcal{T}_{D,I}, \{0,1\}^n)$. Let $\mathbf{A}$ be an SPRP adversary on $\mathrm{ZCZ}[\widetilde{\pi}]$, s.t. $\mathbf{A}$ asks at most q queries of domain $\mathcal{B}^{\leq n}$, that sum up to at most σ di-blocks in total. Then

$$\mathbf{Adv}^{\mathrm{SPRP}}_{\mathrm{ZCZ}[\widetilde{\pi}]}(\mathbf{A}) \leq \frac{3\sigma^2 + 9q^2}{2N^2}.$$

Proof. The queries 1 through q by $\mathbf{A}$ are collected in a transcript τ where we define two disjoint sets of indices E and D s.t. $[1..q] = E \sqcup D$, and it holds that E consists of exactly those indices i s.t. the i-th query of $\mathbf{A}$ is an encryption

query; similarly, D consists of exactly those indices i s.t. the i-th query of $\mathbf{A}$ is an decryption query. We define ℓ^i to be the number of di-blocks in the i-query, where $\ell^i \leq n$.

In both worlds, the adversary's queries are answered immediately with the corresponding outputs; certain internal parts of the transcript will be revealed to the adversary after it made all its queries, but before it outputs its decision bit that represents its guess of which world it interacted with. The internal parts consist of $S^i, T^i, S^i_1, X^i_L, X^i_R, Y^i_L, Y^i_R$ for $i \in [1..q]$ and $Z^i_{1,j}$ for $i \in [1..q], j \in [1..\ell^i - 1]$; for ease of notation, we write Z^i_j to refer to $Z^i_{1,j}$.

We will subsequently define certain transcripts to be *good*. More specifically, we describe a mechanism for the ideal oracle to sample the internal values to be given to the adversary at the end of the query phase, and define the event bad as the union of five events $\mathsf{badA}, \mathsf{badB}, \mathsf{badC}, \mathsf{badD}$ and badE. We call a transcript good if it can be obtained by the ideal oracle without encountering the event bad. Now we state two lemmas.

Lemma 2. $\Pr[\mathsf{bad}] \leq \dfrac{3\sigma^2 + 8q^2}{N^2}$.

Lemma 3. For any good transcript τ,

$$\frac{\Pr[\Theta_{\text{real}} = \tau]}{\Pr[\Theta_{\text{ideal}} = \tau]} \geq 1 - \frac{q^2}{N^2}.$$

Then, the proof follows from Lemmas 1, 2, and 3. □

For proving Lemmas 2 and 3, we first define the sampling mechanism of the ideal oracle and the bad events.

Equations. First, we write the internal variables X^i_j, Y^i_j for $i \in [1..q], j \in [1..\ell^i]$ and $U^i_L, U^i_R, V^i_L, V^i_R$ for $i \in [1..q]$ in terms of $S^i, T^i, S^i_1, X^i_L, X^i_R, Y^i_L, Y^i_R, Z^i_j$:

$$X^i_j = L'^i_j + Z^i_j, \tag{1}$$
$$Y^i_j = R^i_j + Z^i_j + S^i_1, . \tag{2}$$

Moreover, we define four auxiliary variables to easier referral:

$$U^i_L = L^i_\ell + X^i_L, \tag{3}$$
$$U^i_R = R^i_\ell + X^i_R, \tag{4}$$
$$V^i_L = L'^i_\ell + Y^i_L, \tag{5}$$
$$V^i_R = R'^i_\ell + Y^i_R. \tag{6}$$

Identifying A Basis. A basis is the set of variables (internal to the constructions) which can be sampled uniformly and independently in the ideal oracles after fixing the inputs and outputs that are known to adversary. By looking at the construction and eliminating the relationships between the internal variables,

plaintexts, and ciphertexts, some internal variables can be chosen almost freely, and still the real construction will behave indistinguishable from the ideal world for the adversary even after observing the plain- and ciphertexts. We call those variables a basis. For $i \in [1..q], j \in [1..\ell^i]$, we define (i, j) to be *fresh* if either of the following is true:

- $i \in E$, and for any $i' \in [1..i-1]$: $(L_j^{i'}, R_j^{i'}) \neq (L_j^i, R_j^i)$;
- $i \in D$, and for any $i' \in [1..i-1]$: $(L_j'^{i'}, R_j'^{i'}) \neq (L_j'^i, R_j'^i)$.

For $i \in [2..q], i' \in [1..i-1]$, we say i is *akin* to i' if either of the following holds:

- $\ell^i = \ell^{i'}$, $i \in E$, and for any $j \in [1..\ell^i - 1]$: $(L_j^{i'}, R_j^{i'}) = (L_j^i, R_j^i)$;
- $\ell^i = \ell^{i'}$, $i \in D$, and for any $j \in [1..\ell^i - 1]$: $(L_j'^{i'}, R_j'^{i'}) = (L_j'^i, R_j'^i)$;

We say i is *new* if it is not akin to any $i' \in [1..i-1]$. Now we define the basis as follows: for $i \in [1..q]$,

- For $j \in [1..\ell^i - 1]$, Z_j^i is in the basis if (i, j) is fresh;
- X_L^i and X_R^i are in the basis if $i \in D$, or if $i \in E$ and i is new;
- Y_L^i and Y_R^i are in the basis if $i \in E$, or if $i \in D$ and i is new;
- S^i, T^i, and S_1^i are in the basis.

Let σ_F represent the total number of fresh pairs in the set $\{(i, j) \mid i \in [q], j \in [\ell^i - 1]\}$. Moreover, let q_ν be the total number of new queries in $[1..q]$. Then, the size of the basis is $\sigma_F + 2q_\nu + 5q$.

Extension from Basis. Now we show how all the internal variables X_j^i, Y_j^i for $i \in [1..q], j \in [1..\ell^i]$ and $U_L^i, U_R^i, V_L^i, V_R^i$ for $i \in [1..q]$ can be written in terms of basis variables. Since we have already seen how to write them in terms of $S^i, T^i, S_1^i, X_L^i, X_R^i, Y_L^i, Y_R^i$ for $i \in [1..q]$ and Z_j^i for $i \in [1..q], j \in [1..\ell^i - 1]$, and S^i, T^i, S_1^i for $i \in [1..q]$ are already in the basis, it suffices to show that Z_j^i for $i \in [1..q], j \in [1..\ell^i - 1]$ and $X_L^i, X_R^i, Y_L^i, Y_R^i$ for $i \in [1..q]$ can be written in terms of basis variables. An expression of an internal variable in terms of basis variables and the oracle inputs and outputs will be called the *extension expression* of the basis variable. Thus, whenever we sample all the basis elements, we can extend this through these equations to assign values to all the internal variables.

For $i \in E, j \in [1..\ell^i]$, let i' be such that (i', j) is fresh, and $(L_j^{i'}, R_j^{i'}) = (L_j^i, R_j^i)$. Then, i' is called the *j-predecessor* of i, denoted $i : j$. Similarly, for $i \in D, j \in [1..\ell^i]$, if for some i' we have (i', j) fresh and $(L_j'^{i'}, R_j'^{i'}) = (L_j'^i, R_j'^i)$, we set $i : j = i'$. (Thus, when (i, j) is fresh, $i : j$ is i itself.) For $i \in E, j \in [1..\ell^i]$ we have from (1) that $X_j^i = X_j^{i:j} = L_j'^{i:j} + Z_j^{i:j}$, so

$$Z_j^i = L_j'^{i:j} + L_j'^i + Z_j^{i:j}; \tag{7}$$

and for $i \in D, j \in [1..\ell^i]$ we have from (2)

$$Y_j^i = Y_j^{i:j} = R_j^{i:j} + Z_j^{i:j} + S_1^{i:j},$$

so

$$Z_j^i = R_j^{i:j} + R_j^i + Z_j^{i:j} + S_1^{i:j} + S_1^i. \tag{8}$$

Now if i and $i : j$ are both in E or both in D, $Z_j^{i:j}$ is a basis element. (In particular, when $i : j = i$, Z_j^i is a basis element.) Otherwise, we can go back one step further to $(i : j) : j$, the j-predecessor of $i : j$, denoted $i : j^2$. We call (1) and (2) the *extension equations*. They will serve useful in the later proofs. Note that it does not hold in general that $(i : j) : j = i : j$. This holds only if $i : j$ and i are both in E or both in D, or when $i : j$ points to a fresh input block.

For $i \in [2..q]$, the smallest query index in $[1..i-1]$ which i is akin to is called the *origin* of i, denoted $\bar{i}$. We also define the origin of 1 to be 1 itself. Thus, for $i \in E$,

$$X_L^i = X_L^{\bar{i}}, \tag{9}$$

$$X_R^i = X_R^{\bar{i}}; \tag{10}$$

and for $i \in D$,

$$Y_L^i = Y_L^{\bar{i}}, \tag{11}$$

$$Y_R^i = Y_R^{\bar{i}}. \tag{12}$$

Since for $i \in E$, $X_L^{\bar{i}}$ and $X_R^{\bar{i}}$ are in the basis, and for $i \in D$, $Y_L^{\bar{i}}$ and $Y_R^{\bar{i}}$ are in the basis, this completes the extensions.

Oracles and Bad Events. The real oracle employs $\mathrm{ZCZ}[\widetilde{\pi}]$ to answer the queries of $\mathbf{A}$. In the ideal world, the encryption oracle samples and returns $L_j'^i, R_j'^i$ for $i \in E, j \in [1..\ell^i]$ uniformly at random; the decryption oracle samples and returns L_j^i, R_j^i for $i \in D$, $j \in [1..\ell^i]$ uniformly at random. Once the interaction phase is over, the ideal world oracle samples and returns each basis element uniformly at random from $\{0,1\}^n$, with two exceptions:

- For $i \in E$, S^i is drawn uniformly from the set $\{0,1\}^n \setminus \left\{ S^{i'} \mid i \text{ is akin to } i', R^i = R^{i'} \right\}$;
- For $i \in D$, T^i is drawn uniformly from the set $\{0,1\}^n \setminus \left\{ T^{i'} \mid i \text{ is akin to } i', L'^i = L'^{i'} \right\}$.

The real world releases the values of the basis variables to the adversary. (Thus, from the extension equations, $\mathbf{A}$ can calculate the values of the inputs, tweaks, and outputs of all internal TBC calls.) $\mathbf{A}$ shall distinguish the real world $\mathcal{O}_{\mathrm{real}}$ from the ideal world $\mathcal{O}_{\mathrm{ideal}}$, given a transcript τ of its interaction with the available oracles. We call a transcript τ bad iff $\tau \in \mathrm{BadT}$, and call it good otherwise. We say that a bad event occurred if at least one of the following occurred:

- badA occurs when:
  - For some $i \in E, j \in [1..\ell^i]$, there exists $i' \in [1..i-1]$ with $\ell^{i'} \geq j$ s.t. $(L_j'^{i'}, R_j'^{i'}) = (L_j'^i, R_j'^i)$;

  - For some $i \in D, j \in [1..\ell^i]$, there exists $i' \in [1..i-1]$ with $\ell^{i'} \geq j$ s.t. $(L_j^{i'}, R_j^{i'}) = (L_j^i, R_j^i)$;
- badB occurs when:
  - For some $i \in [2..q]$ there exists $i' \in [1..i-1]$ with $\ell^i = \ell^{i'}$ s.t. one of the following holds:
    — $(U_L^i, U_R^i) = (U_L^{i'}, U_R^{i'})$;
    — $(S^i, U_R^i) = (S^{i'}, U_R^{i'})$;
    — $(S^i, T^i) = (S^{i'}, T^{i'})$;
    — $(V_L^i, T^i) = (V_L^{i'}, T^{i'})$;
    — $(V_L^i, V_R^i) = (V_L^{i'}, V_R^{i'})$;
- badC occurs when:
  - For some $i \in [1..q]$, there exists $i' \in [1..i-1]$ s.t. $(S_1^i, T^i) = (S_1^{i'}, T^{i'})$;
  - For some $i \in [1..q], j \in [1..\ell^i - 1]$, there exists $i' \in [1..i-1]$ with $\ell^{i'} \geq j+1$ s.t. $(Z_j^i, T^i) = (Z_j^{i'}, T^{i'})$;
- badD occurs when:
  - For some $i \in E, j \in [1..\ell^i - 1]$, there exists $i' \in [1..i-1]$ with $\ell^{i'} \geq j+1$ s.t. $(L_j'^i, Y_j^i) = (L_j'^{i'}, Y_j^{i'})$;
  - For some $i \in D, j \in [1..\ell^i - 1]$, there exists $i' \in [1..i-1]$ with $\ell^{i'} \geq j+1$ s.t. $(R_j^i, X_j^i) = (R_j^{i'}, X_j^{i'})$;
- badE occurs when for some $i \in [2..q]$, there exists $i' \in [1..i-1]$ s.t. i is not akin to i' and yet one of the following holds:
  — $(X_L^{*i}, X_R^{*i}) = (X_L^{*i'}, X_R^{*i'})$;
  — $(Y_L^{*i}, Y_R^{*i}) = (Y_L^{*i'}, Y_R^{*i'})$;

Thus, $\mathsf{bad} \stackrel{\text{def}}{=} \mathsf{badA} \vee \mathsf{badB} \vee \mathsf{badC} \vee \mathsf{badD} \vee \mathsf{badE}$. Clearly,

$$\Pr[\mathsf{bad}] \leq \Pr[\mathsf{badA}] + \Pr[\mathsf{badB}] + \Pr[\mathsf{badC}] + \Pr[\mathsf{badD}] + \Pr[\mathsf{badE}]. \quad (13)$$

Now, we are in a position to prove Lemmas 2 and 3.

Proof of Lemma 2. Below, we show that each of the collision-pairs that would result in one of the bad events has a joint probability of $\leq 1/N^2$. Clearly, we need the assumption that all basis elements are uniformly sampled from $\{0,1\}^n$ for this purpose. Moreover, the values S^i and T^i are sampled without replacement under certain circumstances, their bound is at most $1/N(N-1)$, which can be upper bounded by $1/N(N-1) < 2/N^2$. Thus, for bounding the bad events, we simply need to bound the number of candidate collision-pairs.

For badA, there can be:

- at most $\sigma_E^2/2$ collision events of the form $(L_j'^{i'}, R_j'^{i'}) = (L_j'^i, R_j'^i)$;
- at most $\sigma_D^2/2$ collision events of the form $(L_j^{i'}, R_j^{i'}) = (L_j^i, R_j^i)$;

where σ_E is the total number of encryption query blocks and σ_D is the total number of decryption query blocks, so that $\sigma_E^2 + \sigma_D^2 \leq \sigma^2$. Thus

$$\Pr[\mathsf{badA}] \leq \frac{\sigma^2}{N^2}. \quad (14)$$

For badB, there can be:

– at most $q^2/2$ collision events of the form $(U_L^i, U_R^i) = (U_L^{i'}, U_R^{i'})$;
– at most $q^2/2$ collision events of the form $(S^i, U_R^i) = (S^{i'}, U_R^{i'})$;
– at most $q^2/2$ collision events of the form $(S^i, T^i) = (S^{i'}, T^{i'})$;
– at most $q^2/2$ collision events of the form $(V_L^i, T^i) = (V_L^{i'}, T^{i'})$;
– at most $q^2/2$ collision events of the form $(V_L^i, V_R^i) = (V_L^{i'}, V_R^{i'})$;

Thus

$$\Pr[\mathsf{badB}] \leq \frac{5q^2}{N^2}. \tag{15}$$

For badC, there can be:

– at most $q^2/2$ collision events of the form $(S_1^i, T^i) = (S_1^{i'}, T^{i'})$.
– at most $\sigma^2/2$ collision events of the form $(Z_j^i, T^i) = (Z_j^{i'}, T^{i'})$;

Thus

$$\Pr[\mathsf{badC}] \leq \frac{q^2 + \sigma^2}{N^2}. \tag{16}$$

For badD, there can be:

– at most $\sigma_E^2/2$ collision events of the form $(L_j'^i, Y_j^i) = (L_j'^{i'}, Y_j^{i'})$;
– at most $\sigma_D^2/2$ collision events of the form $(R_j^i, X_j^i) = (R_j^{i'}, X_j^{i'})$.

Thus

$$\Pr[\mathsf{badD}] \leq \frac{\sigma^2}{N^2}. \tag{17}$$

For badE, there can be:

– at most $q^2/2$ collision events of the form $(X_L^{*i}, X_R^{*i}) = (X_L^{*i'}, X_R^{*i'})$;
– at most $q^2/2$ collision events of the form $(Y_L^{*i}, Y_R^{*i}) = (Y_L^{*i'}, Y_R^{*i'})$.

Thus

$$\Pr[\mathsf{badE}] \leq \frac{2q^2}{N^2}. \tag{18}$$

The lemma follows from (13)–(18).

Now, all that is left to do is to establish our claim that each of the collision-pairs that would result in one of the bad events has a joint probability $\leq 1/N^2$. This is to be done by examining each bad event separately. badA, badB and badC are fairly straightforward, and we leave out the proofs. badD is more interesting; we provide below a complete analysis of it. The trickiest case is badE; here, due to space constraints, we only examine two of its main subcases in detail. The complete case-by-case analysis, along with a short analysis of badA, badB and badC, can be found in the Appendix of the full version [4].

Full Analysis of badD. We consider the two cases separately:

– $(L'^i_j, Y^i_j) = (L'^{i'}_j, Y^{i'}_j)$, $i \in E$, $i' < i$: We will show that $Y^i_j = Y^{i'}_j$ always leads to an equation containing at least one basis variable that cannot get canceled out. The required bound follows since the basis variable and L'^i_j are independently sampled. From (2) we have

$$R^i_j + Z^i_j + S^i_1 = R^{i'}_j + Z^{i'}_j + S^{i'}_1 . \tag{19}$$

Note that S^i_1 cannot occur in the expansion of $Z^{i:j}_j$, since $i \in E$. Now we have two options of i':

- $i' \in E$: From (7) and (19) we have

$$R^i_j + L'^{i:j}_j + L'^i_j + Z^{i:j}_j + S^i_1 = R^{i'}_j + L'^{i':j}_j + L'^{i'}_j + Z^{i':j}_j + S^{i'}_1 .$$

  Here the basis element S^i_1 cannot be canceled out, since $i' < i$.
- $i' \in D$: From (7), (8) and (19), we have

$$R^i_j + L'^{i:j}_j + L'^i_j + Z^{i:j}_j + S^i_1 = R^{i'}_j + R^{i':j}_j + R^{i'}_j + Z^{i':j}_j + S^{i':j}_1 .$$

  Again, the basis element S^i_1 cannot be canceled out since $i' : j \leq i' < i$.

– $(R^i_j, X^i_j) = (R^{i'}_j, X^{i'}_j)$, $i \in D$, $i' < i$: As above, we show that $X^i_j = X^{i'}_j$ always leads to an equation containing at least one basis variable that cannot get canceled out, and the required bound follows since the basis variable and R^i_j are independently sampled. From (1), we have

$$L'^i_j + Z^i_j = L'^{i'}_j + Z^{i'}_j . \tag{20}$$

Now, we have two options of i':

- $i' \in E$: From (8), (7) and (20), we have

$$L'^i_j + R^{i:j}_j + R^i_j + Z^{i:j}_j + S^{i:j}_1 + S^i_1 = L'^{i':j}_j + Z^{i':j}_j .$$

  When $i : j < i$, the basis element S^i_1 cannot be canceled out, and when $i = i : j$, we have $i' : j \leq i' < i = i : j$, so the basis element $Z^{i:j}_j = Z^i_j$ cannot be canceled out.
- $i' \in D$: From (8) and (19), we have

$$L'^i_j + R^{i:j}_j + R^i_j + Z^{i:j}_j + S^{i:j}_1 + S^i_1 = L'^{i'}_j + R^{i':j}_j + R^{i'}_j + Z^{i':j}_j + S^{i':j} + S^{i'} ,$$

  Here again, either S^i_1 or the basis element Z^i_j cannot be canceled out, and the argument is identical to the above.

Partial Analysis of badE. This is trickier than the other bad events, and requires some careful case analysis. We examine the two most difficult sub-cases here. Let $i' < i$ and $\ell \stackrel{\text{def}}{=} \ell^{i'} = \ell^{i}$, and let $\alpha_j(\cdot)$ and $\alpha_j^2(\cdot)$ be linear functions defined as

$$\alpha_j(x) \stackrel{\text{def}}{=} \alpha^{\ell-1-j} \cdot x \quad \text{and} \quad \alpha_j^2(x) \stackrel{\text{def}}{=} (\alpha^2)^{\ell-1-j} \cdot x.$$

Both the sub-cases we examine here fall under the case of $(X_L^{*i}, X_R^{*i}) = (X_L^{*i'}, X_R^{*i'})$. We can write this collision as

$$\sum_{j=0}^{\ell-1} \alpha_j(X_j^i + X_j^{i'}) = 0 \quad \text{and} \quad \sum_{j=0}^{\ell-1} \alpha_j^2(X_j^i + X_j^{i'}) = \sum_{j=0}^{\ell-1} \alpha_j^2(R_j^i + R_j^{i'}).$$

Using (1) we can rewrite these as

$$\sum_{j=0}^{\ell-1} \alpha_j(Z_j^i + Z_j^{i'}) = \sum_{j=0}^{\ell-1} \alpha_j(L_j'^i + L_j'^{i'}), \tag{21}$$

$$\sum_{j=0}^{\ell-1} \alpha_j^2(Z_j^i + Z_j^{i'}) = \sum_{j=0}^{\ell-1} \alpha_j^2(L_j'^i + L_j'^{i'} + R_j^i + R_j^{i'}). \tag{22}$$

We first observe that since i is not akin to i', $X_j^i + X_j^{i'}$ cannot trivially disappear for all $j \in [1, .., \ell-1]$. Also, since $\alpha_j(X_j^i + X_j^{i'})$ sum to 0, there must be at least two indices in $[1, .., \ell-1]$ where $X_j^i + X_j^{i'}$ does not trivially disappear; let j_0 and j_1 be the two largest such indices, with $j_0 > j_1$. Now, we first consider the sub-case $i \in E, i' \in E$. From (7), (21) and (22) we have

$$\sum_{j=0}^{\ell-1} \alpha_j(Z_j^{i:j} + Z_j^{i':j}) = \sum_{j=0}^{\ell-1} \alpha_j(L_j'^{i:j} + L_j'^{i':j}), \tag{23}$$

$$\sum_{j=0}^{\ell-1} \alpha_j^2(Z_j^{i:j} + Z_j^{i':j}) = \sum_{j=0}^{\ell-1} \alpha_j^2(L_j'^{i:j} + L_j'^{i':j} + R_j^{i:j} + R_j^{i':j}). \tag{24}$$

By choice of j_0, $i : j_0 \neq i' : j_0$. Suppose $i : j_0 > i' : j_0$. If $i : j_0 \in D$, using (8), we replace $Z_{j_0}^{i:j_0}$ by $R_{j_0}^{i:j_0^2} + R_{j_0}^{i:j_0} + Z_{j_0}^{i:j_0^2} + S_1^{i:j_0^2} + S_1^{i:j_0}$. The basis element $S_1^{i:j_0}$ does not get canceled out; moreover, $R_{j_0}^{i:j_0}$ remains only in the top equation, while it gets canceled out in the bottom equation. Since $i : j = i' : j$ for all $j > j_0$, none of the adversary-queried blocks remaining in either equation came after $R_{j_0}^{i:j_0}$, so it is independent of the rest of the equation; along with the basis element $S_1^{i:j_0}$ (which appears in both equations), this makes the two collisions independent, thus occurring jointly with a probability $1/N^2$.

If $i : j_0 \in E$, $Z_{j_0}^{i:j_0}$ is in the basis, and does not cancel out. On the right hand side of both equations, $L_{j_0}'^{i:j_0}$ remains uncanceled as well, while all later adversary queries get canceled. Thus, the two equations can become dependent with

probability at most $1/N$; then, the common collision can occur with probability at most $1/N$. Thus, in either case, the joint collision can occur with a probability of more than $1/N^2$. The analysis is similar when $i : j_0 < i : j_0$; then we focus on the latter instead.

The other sub-case we consider is $i \in E, i' \in D$. From (7), (8), (21) and (22) we have

$$\sum_{j=0}^{\ell-1} \alpha_j (Z_j^{i:j} + Z_j^{i':j} + S_1^{i'} + S_1^{i':j}) = \sum_{j=0}^{\ell-1} \alpha_j (L_j'^{i:j} + L_j'^{i':j} + R_j^{i'} + R_j^{i':j}), \tag{25}$$

$$\sum_{j=0}^{\ell-1} \alpha_j^2 (Z_j^{i:j} + Z_j^{i':j} + S_1^{i'} + S_1^{i':j}) = \sum_{j=0}^{\ell-1} \alpha_j^2 (L_j'^{i:j} + L_j'^{i':j} + R_j^{i:j} + R_j^{i':j}). \tag{26}$$

By choice of j_0 and j_1, $i : j_0 \neq i'$ and $i : j_1 \neq i'$. Suppose $i : j_0 < i'$. Then $S_1^{i'}$ and $R_{j_0}^{i'}$ remain uncanceled in (25), and no adversary query block queried after $R_{j_0}^{i'}$ remains uncanceled; in (26), $S_1^{i'}$ remains uncanceled again, but there is no $R_{j_0}^{i'}$ and no adversary query block queried after it. Thus these two can occur jointly with a probability at most $1/N^2$.

A symmetric argument can be used when $i : j_0 > i'$ and $i : j_0 \in D$: we replace $Z_{j_0}^{i:j_0}$ by $R_{j_0}^{i:j_0^2} + R_{j_0}^{i:j_0} + Z_{j_0}^{i:j_0^2} + S_1^{i:j_0^2} + S_1^{i:j_0}$ using (8), and observe that $S_1^{i:j_0}$ remains uncanceled in either equation, while $R_{j_0}^{i:j_0}$ remains uncanceled in (25), but gets canceled out in (26), and no adversary query block queried after it remains in either equation.

When $i : j_0 > i'$ and $i : j_0 \in E$, but $i : j_1$ satisfied one of the above two conditions, we can argue as above using $i : j_1$ instead. If we also have $i : j_1 > i'$ and $i : j_1 \in E$, we observe that $Z_{j_0}^{i:j_0}$ and $Z_{j_1}^{i:j_1}$ are basis elements that do not get canceled out in either equation. Their combined contribution to the left-hand side of (25) is $\alpha^{\ell-1-j_0} \cdot Z_{j_0}^{i:j_0} + \alpha^{\ell-1-j_1} \cdot Z_{j_1}^{i:j_1}$ and to the left-hand side of (26) is $(\alpha^2)^{\ell-1-j_0} \cdot Z_{j_0}^{i:j_0} + (\alpha^2)^{\ell-1-j_1} \cdot Z_{j_1}^{i:j_1}$. These two collisions are independent since $\alpha^{\ell-1-j_0} \cdot (\alpha^2)^{\ell-1-j_1} \neq \alpha^{\ell-1-j_1} \cdot (\alpha^2)^{\ell-1-j_0}$, and thus can occur with a probability at most $1/N^2$. The rest of the subcases can be analysed similarly. This completes the proof of Lemma 2. □

Proof of Lemma 3. Let τ be a good transcript, i. e., none of the events badA, badB, badC, badD, or badE occurred. Then, in the ideal world, there are 2σ samplings for generating the query responses and $\sigma_F + 2q_\nu + 5q$ for generating the basis elements. In the ideal world, the basis elements are sampled uniformly at random and independently from each other. Hence, the probability for those is given by $1/N^{\sigma_F + 2q_\nu + 5q}$. The situation differs for the outputs of the scheme. The ideal world is an ideal SPRP; hence, the outputs are sampled without replacement. Since all queries are from the domain $\mathcal{B}^{\leq n}$, we can group encryption and decryption queries into disjoint sets $\mathcal{L}^1, \ldots, \mathcal{L}^n$ s. t. their union contains all queries, and Set $\mathcal{L}^i$ contains exactly the queries of length i di-blocks. We define by $\textsc{Load}\left(\mathcal{L}^i\right)$ the number of queries in Set $\mathcal{L}^i$, for all $1 \leq i \leq n$. The

probability for ciphertext outputs from encryption queries and plaintext outputs from decryption queries is

$$\prod_{i=1}^{n} \frac{1}{(N^{2i})_{\textsc{Load}(\mathcal{L}^i)}}.$$

Since each query has at least $2n$ bits, we can lower bound the probability by

$$\prod_{i=1}^{n} \frac{1}{(N^{2i})_{\textsc{Load}(\mathcal{L}^i)}} \leq \frac{1}{(N^2)_{2q}} \cdot \frac{1}{N^{2\sigma-2q}}.$$

We obtain that

$$\Pr\left[\Theta_{\text{ideal}} = \tau\right] \leq \frac{1}{N^{\sigma_F+2q_\nu+3q+2\sigma}} \cdot \frac{1}{(N^2)_q}. \tag{27}$$

In the real world, the construction employs a permutation $\widetilde{\pi}^{\mathsf{T}}(\cdot)$ for each tweak $\mathsf{T} \in \mathcal{T}_{\mathcal{D}\times\mathcal{I}}$ that was used in the transcript, . We write the set of all occurred tweaks of all di-blocks of all queries in the transcript and write it as $\left\{\mathsf{T}^1, \ldots, \mathsf{T}^\theta\right\}$. We further define by $\textsc{Load}\left(\mathsf{T}\right)$ the load of a tweak T, i.e., the number of distinct inputs used for it over all queries and di-blocks of the transcript. It holds that $\sum_{i=1}^{\theta} \textsc{Load}\left(\mathsf{T}^i\right) = \sigma_F + 2\sigma + 2q_\nu + 5q$. We adopt the notion of transcript-compatible permutations from [6]. We call $\widetilde{\pi}$ *compatible* with τ if for all queries, $\widetilde{\pi}$ produced all intermediate variables as well as all outputs in τ. Let $\mathsf{Comp}(\tau)$ denote the set of tweakable permutations $\widetilde{\pi}$ that are compatible with τ. Thus

$$\Pr\left[\Theta_{\text{real}} = \tau\right] = \Pr\left[\widetilde{\pi} \twoheadleftarrow \widetilde{\mathsf{Perm}}\left(\mathcal{T}_{D,I}, \{0,1\}^n\right) : \widetilde{\pi} \in \mathsf{Comp}(\tau)\right].$$

For a fixed tweak T, the fraction of compatible permutations is

$$\prod_{i=0}^{\textsc{Load}(\mathsf{T})-1} \frac{1}{N-i} = \frac{1}{(N)_{\textsc{Load}(\mathsf{T})}}.$$

Over all tweaks T^i, for $1 \leq i \leq \theta$, the fraction of compatible permutations is given by

$$\prod_{i=1}^{\theta} \frac{1}{(N)_{\textsc{Load}(\mathsf{T}^i)}}$$

It is hard to work with this probability directly. Instead, since we are interested in a bound for the real-world probability of transcripts, we can lower bound the probability of all $\sigma_F + 2q_\nu + 5q$ basis variables by the naive probability that they are all computed from fresh tweaks: $1/N^{\sigma_F+2q_\nu+5q}$. For the ciphertext and plaintext outputs, we can employ similar sets $\mathcal{L}^i$, for $1 \leq i \leq n$, as we had for the ideal world, where Set $\mathcal{L}^i$ again consists of all queries of length i di-blocks. The probability of outputs in the real world can then be lower bounded by

$$\prod_{i=1}^{n} \frac{1}{(N^{2i})_{\textsc{Load}(\mathcal{L}^i)}}.$$

Now, we can upper bound the ratio of the probability of our transcripts by

$$\begin{aligned}
\frac{\Pr\left[\Theta_{\text{real}} = \tau\right]}{\Pr\left[\Theta_{\text{ideal}} = \tau\right]} &\geq \frac{\frac{1}{N^{\sigma_F+2q_\nu+5q}} \cdot \prod_{i=1}^{n} \frac{1}{(N^{2i})_{\text{LOAD}(\mathcal{L}^i)}}}{\frac{1}{N^{\sigma_F+2q_\nu+5q}} \cdot \frac{1}{N^{2\sigma-2q}} \cdot \frac{1}{(N^2)_q}} \\
&\geq \frac{\prod_{i=1}^{n} \frac{1}{(N^{2i})_{\text{LOAD}(\mathcal{L}^i)}}}{\frac{1}{(N^2)_q} \cdot \frac{1}{N^{2\sigma-2q}}} \geq \frac{(N^2)_q \cdot N^{2\sigma-2q}}{N^{2\sigma}} = \frac{(N^2)_q}{(N^2)^q} \\
&= \frac{(N^2)(N^2-1)\cdot\cdots\cdot(N^2-q+1)}{(N^2)^q} \geq \left(\frac{N^2-q+1}{N^2}\right)^q \\
&\geq \left(\frac{N^2-q}{N^2}\right)^q = \left(1-\frac{q}{N^2}\right)^q \geq 1-\frac{q^2}{N^2},
\end{aligned}$$

where the last inequality is Bernoulli's. So, we obtain our claim in Lemma 3. □

6.2 Proof Sketch for Messages with Arbitrary Number of Complete Di-blocks

Theorem 2. Let $\widetilde{\pi} \twoheadleftarrow \widetilde{\mathsf{Perm}}(\mathcal{T}_{D,I}, \{0,1\}^n)$. Let $\mathbf{A}$ be an SPRP adversary on $\mathrm{ZCZ}[\widetilde{\pi}]$ that asks at most q queries queries of domain $\mathcal{B}^+$, whose lengths sum up to at most σ di-blocks in total, and $\mathbf{A}$ runs in time at most TIME. Then

$$\mathbf{Adv}^{\text{SPRP}}_{\mathrm{ZCZ}[\widetilde{\pi}]}(\mathbf{A}) \leq \frac{4\sigma^2+8q^2}{N^2}.$$

Proof Sketch. The proof follows a similar strategy as that of Theorem 1. So, we only consider the equations in the analysis of bad events that differ. We add each S^i_k, $i \in [1..q]$, $k \in \left[1..\left\lceil \ell^i/n \right\rceil\right]$ to the basis. The ideal oracle samples the additional basis elements along with the original basis elements in the second step, and the definitions of the bad cases do not change. From the equations (1)–(6) that we began with, only (2) is now replaced by

$$Y^i_j = R^i_j + Z^i_j + S^i_{\lceil j/n \rceil}. \tag{2'}$$

In the extension equations, this changes only (8), which is replaced by

$$Z^i_j = R^{i:j}_j + R^i_j + Z^{i:j}_j + S^{i:j}_{\lceil j/n \rceil} + S^i_{\lceil j/n \rceil}. \tag{8'}$$

The definitions of the bad cases remain the same except badC, which now occurs when:

- For some $i \in [1..q], k \in \left[1..\left\lceil \ell^i/n \right\rceil\right]$, there exists $i' \in [1..i-1]$ with $\ell^{i'} \geq n(k-1)$ s.t. $(S^i_k, T^i) = (S^{i'}_k, T^{i'})$;
- For some $i \in [1..q], j \in [1..\ell^i - 1]$, there exists $i' \in [1..i-1]$ with $\ell^{i'} \geq j+1$ s.t. $(Z^i_{k,c}, T^i) = (Z^{i'}_{k,c}, T^{i'})$, where $k = \lceil j/n \rceil, c = j - n(k-1)$.

Of these, the counting does not change for the latter; for the former, there are now at most $c_{\max}q^2/2$ possible collision pairs now, where $c_{\max}$ is the maximum number of chunks in one query; we generously bound this by $\sigma^2/2$. This adds $(\sigma^2 - q^2)/2N$ to our earlier bound, to obtain the new bound for the extended version. To ensure that the counting argument for badE still goes through, we only note that for $k \in [1..\lceil \ell/n \rceil]$, S_k^i can only occur in any of the collision equations from badE with coefficients $\beta^{\ell-1-j}$ for $j \in [n(k-1)+1..nk]$, where β is either α or α^2, and for any choice of k, a non-empty subset of these coefficients cannot add to 0.

6.3 Proof Sketch for the Security of ZCZ*

Theorem 3. Let $\widetilde{\pi} \twoheadleftarrow \widetilde{\mathsf{Perm}}(\mathcal{T}_{D,I}, \{0,1\}^n)$. Let $\mathbf{A}$ be an SPRP adversary on $\mathrm{ZCZ}^*[\widetilde{\pi}]$ that asks at most q queries of domain $\{0,1\}^{\geq 2n}$, whose lengths sum up to at most σ di-blocks in total, q' of which contains an incomplete di-block at the end. Then

$$\mathbf{Adv}^{\mathrm{SPRP}}_{\mathrm{ZCZ}^*[\widetilde{\pi}]}(\mathbf{A}) \leq \frac{4\sigma^2 + 8q^2 + 9q'^2}{N^2}.$$

Proof Sketch. The ideal oracle's sampling mechanism for the tweakable blockcipher outputs for the partial di-block messages is slightly trickier. Let $\mathcal{I}$ denote the indices of the queries with incomplete di-blocks. Instead of simulating an ideal permutation, the ideal oracle simulates what [11] calls an $\pm\tilde{\mathbf{rnd}}$ oracle, which always returns random bits, as long as no pointless queries are asked. (It is easy to argue for our construction why not permitting pointless queries does not diminish the adversary's power, so we can confine our attention to the no-pointless-query scenario.)

We use the notation $(U,V), (U_m,V_m), (U',V')$ for outputs of the blockcipher calls in the top, middle, and bottom layers respectively. M_j denotes (L_j, R_j), and $*$ denotes the index of the incomplete di-block.

- For the smallest $i \in \mathcal{I}$, $U_*^i, V_*^i, U_*'^i, V_*^i$ are sampled uniformly from $\{0,1\}^n$;
- For each i in $\mathcal{I}$ such that for no i' in $\mathcal{I}$ with $i' < i$ we have $(L_*^i, R_*^i) \neq (L_*^{i'}, R_*^{i'})$:
  - U_*^i is sampled uniformly from $\{0,1\}^n \setminus \left\{U_*^{i'} \mid i' \in \mathcal{I}, i' < i\right\}$;
  - V_*^i is sampled uniformly from $\{0,1\}^n \setminus \left\{V_*^{i'} \mid i' \in \mathcal{I}, i' < i\right\}$;
- For each i in $\mathcal{I}$ such that for no i' in $\mathcal{I}$ with $i' < i$ we have $(L_*'^i, R_*'^i) \neq (L_*'^{i'}, R_*'^{i'})$:
  - $U_*'^i$ is sampled uniformly from $\{0,1\}^n \setminus \left\{U_*'^{i'} \mid i' \in \mathcal{I}, i' < i\right\}$;
  - $V_*'^i$ is sampled uniformly from $\{0,1\}^n \setminus \left\{V_*'^{i'} \mid i' \in \mathcal{I}, i' < i\right\}$;
- For each $i \in \mathcal{I}$ the $(2n-s)$-bit suffix R^i of (U_{m*}^i, V_{m*}^i) is sampled uniformly from $\{0,1\}^{2n-s}$, and (U_{m*}^i, V_{m*}^i) is set to $(M_*^i + M_*'^i)||R^i$.

The new bad cases are:

– For some distinct i, i' in $\mathcal{I}$ with $\ell^i = \ell^{i'} = \ell$ we have

$$(M^i_{1..\ell-1}, M^i_\ell + (U^i_*, V^i_*)) = (M^{i'}_{1..\ell-1}, M^{i'}_\ell + (U^{i'}_*, V^{i'}_*));$$

– For some distinct i, i' in $\mathcal{I}$ with $\ell^i = \ell^{i'} = \ell$ we have

$$(M'^i_{1..\ell-1}, M'^i_\ell + (U'^i_*, V'^i_*)) = (M'^{i'}_{1..\ell-1}, M'^{i'}_\ell + (U'^{i'}_*, V'^{i'}_*));$$

– For some distinct i, i' in $\mathcal{I}$ with $\ell^i = \ell^{i'} = \ell$ we have

$$\begin{aligned}&(L^i_\ell + L'^i_\ell + U^i_* + U'^i_*, R^i_\ell + R'^i_\ell + V^i_* + V'^i_*)\\ &= (L^{i'}_\ell + L'^{i'}_\ell + U^{i'}_* + U'^{i'}_*, R^{i'}_\ell + R'^{i'}_\ell + V^{i'}_* + V'^{i'}_*);\end{aligned}$$

– For some distinct i, i' in $\mathcal{I}$ with $\ell^i = \ell^{i'} = \ell$ we have

$$(R^i_\ell + R'^i_\ell + V^i_* + V'^i_*, U^i_{m*}) = (R^{i'}_\ell + R'^{i'}_\ell + V^{i'}_* + V'^{i'}_*, U^{i'}_{m*});$$

– For some distinct i, i' in $\mathcal{I}$ with $\ell^i = \ell^{i'} = \ell$ we have

$$(R^i_\ell + R'^i_\ell + V^i_* + V'^i_*, V^i_{m*}) = (R^{i'}_\ell + R'^{i'}_\ell + V^{i'}_* + V'^{i'}_*, V^{i'}_{m*}).$$

The probabilities of these bad cases can be bounded by $q'^2/2N'^2$, $q'^2/2N'^2$, $q'^2/2N'^2$, $q'^2/2NN'$, $q'^2/2NN'$ in that order, where $N' = N - q'$. With the reasonable assumption that $q' \le N/2$, we can replace N' with $N/2$ in these bounds and have them sum to $8q'^2/N^2$, which is our bound for the combined probability of the new bad cases. The theorem follows from Theorem 2 and Lemma 6 of [11].

Our results in Theorems 1 and 3 had considered the instantiation with an ideal random tweaked permutation $\widetilde{\pi} \twoheadleftarrow \widetilde{\mathsf{Perm}}(\mathcal{T}_{I,D}, \{0,1\}^n)$. Corollaries 1 and 3 yield the resulting security bounds when ZCZ and ZCZ* are instantiated with a given tweakable block cipher $\widetilde{E}_K : \mathcal{K} \times \mathcal{T}_{I,D} \times \{0,1\}^n \to \{0,1\}^n$ be a tweakable block cipher with $K \twoheadleftarrow \mathcal{K}$.

Corollary 1. Let $\mathbf{A}$ be an SPRP adversary on ZCZ$[\widetilde{E}_K]$, s.t. $\mathbf{A}$ asks at most q queries of domain $\mathcal{B}^{\le n}$, that sum up to at most σ di-blocks in total, and $\mathbf{A}$ runs in time at most TIME. Then

$$\mathbf{Adv}^{\mathrm{SPRP}}_{\mathrm{ZCZ}[\widetilde{E}_K]}(\mathbf{A}) \le \frac{3\sigma^2 + 10q^2}{2N^2} + \mathbf{Adv}^{\mathrm{STPRP}}_{\widetilde{E}_K, \widetilde{E}_K^{-1}}(\mathbf{A}'),$$

where $\mathbf{A}'$ is an STPRP adversary against $\widetilde{E}_K$ that asks at most $a' = 3\sigma + \lceil \sigma/n \rceil + 6q$ queries and runs in time at most TIME $+ O(a')$.

Corollary 2. Let $\mathbf{A}$ be an SPRP adversary on ZCZ$^*[\widetilde{E}_K]$ that asks at most q queries of domain $\{0,1\}^{\ge 2n}$, whose lengths sum up to at most σ di-blocks in total, q' of which contains an incomplete di-block at the end, and $\mathbf{A}$ runs in time at most TIME. Then

$$\mathbf{Adv}^{\mathrm{SPRP}}_{\mathrm{ZCZ}^*[\widetilde{E}_K]}(\mathbf{A}) \le \frac{4\sigma^2 + 8q^2 + 9q'^2}{N^2} + \mathbf{Adv}^{\mathrm{STPRP}}_{\widetilde{E}_K, \widetilde{E}_K^{-1}}(\mathbf{A}'),$$

where $\mathbf{A}'$ is an STPRP adversary against $\widetilde{E}_K$ that asks at most $a' = 3\sigma + \lceil \sigma/n \rceil + 6q + 6q'$ queries and runs in time at most TIME $+ O(a')$.

Acknowledgments

The authors thank all anonymous reviewers for their fruitful comments that greatly helped to improve this work.

References

1. Beierle, C., Jean, J., Kölbl, S., Leander, G., Moradi, A., Peyrin, T., Sasaki, Y., Sasdrich, P., Sim, S.M.: The SKINNY Family of Block Ciphers and Its Low-Latency Variant MANTIS. In: Robshaw, M., Katz, J. (eds.) CRYPTO II. LNCS, vol. 9815, pp. 123–153. Springer (2016)
2. Bernstein, D.J.: Some Challenges in Heavyweight Cipher Design. Tech. rep. (January 11 2016), https://cr.yp.to/talks/2016.01.15/slides-djb-20160115-a4.pdf
3. Bertoni, G., Daemen, J., Hoffert, S., Peeters, M., Assche, G.V., Keer, R.V.: Farfalle: parallel permutation-based cryptography. IACR Transactions on Symmetric Cryptology **2017**(4), 1–38 (2017)
4. Bhaumik, R., List, E., Nandi, M.: ZCZ – Achieving n-bit SPRP Security with a Minimal Number of Tweakable-block-cipher Calls. Cryptology ePrint Archive, Report 2018/819 (2018), http://eprint.iacr.org/2018/819
5. Biryukov, A., Daemen, J., Lucks, S., Vaudenay, S.: Topics and Research Directions for Symmetric Cryptography. In: Early Symmetric Crypto Workshop. vol. 2017 (2017), https://www.cryptolux.org/mediawiki-esc2017/images/9/9a/ASJS-Topics_SymCrypto-ESC17.pdf
6. Chen, S., Steinberger, J.P.: Tight Security Bounds for Key-Alternating Ciphers. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT. LNCS, vol. 8441, pp. 327–350. Springer (2014)
7. Chen, Y.L., Luykx, A., Mennink, B., Preneel, B.: Efficient Length Doubling From Tweakable Block Ciphers. IACR Trans. Symmetric Cryptol. **2017**(3), 253–270 (2017)
8. Cogliati, B., Lee, J., Seurin, Y.: New Constructions of MACs from (Tweakable) Block Ciphers. In: IACR Transactions on Symmetric Cryptology. vol. 2017, pp. 27–58 (2017)
9. Coron, J., Dodis, Y., Mandal, A., Seurin, Y.: A Domain Extender for the Ideal Cipher. In: Micciancio, D. (ed.) TCC. LNCS, vol. 5978, pp. 273–289. Springer (2010)
10. Gueron, S., Mouha, N.: Simpira v2: A Family of Efficient Permutations Using the AES Round Function. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT, Part I. LNCS, vol. 10031, pp. 95–125 (2016)
11. Halevi, S., Rogaway, P.: A Tweakable Enciphering Mode. In: Boneh, D. (ed.) CRYPTO. LNCS, vol. 2729, pp. 482–499. Springer (2003)
12. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust Authenticated-Encryption AEZ and the Problem That It Solves. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT (1). LNCS, vol. 9056, pp. 15–44. Springer (2015)
13. Iwata, T.: New Blockcipher Modes of Operation with Beyond the Birthday Bound Security. In: Robshaw, M.J.B. (ed.) FSE. LNCS, vol. 4047, pp. 310–327. Springer (2006)
14. Iwata, T., Minematsu, K.: Stronger Security Variants of GCM-SIV. IACR Trans. Symmetric Cryptol. **2016**(1), 134–157 (2016)

15. Iwata, T., Minematsu, K., Peyrin, T., Seurin, Y.: ZMAC: A Fast Tweakable Block Cipher Mode for Highly Secure Message Authentication. In: Katz, J., Shacham, H. (eds.) CRYPTO, Part III. LNCS, vol. 10403, pp. 34–65. Springer (2017)
16. Jean, J., Nikolic, I., Peyrin, T.: Tweaks and Keys for Block Ciphers: The TWEAKEY Framework. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT (2). LNCS, vol. 8874, pp. 274–288 (2014)
17. Jean, J., Nikolić, I., Peyrin, T.: Deoxys v1.41 (2016), third-round submission to the CAESAR competition. https://competitions.cr.yp.to/round3/deoxysv141.pdf
18. Liskov, M., Rivest, R.L., Wagner, D.: Tweakable Block Ciphers. In: Yung, M. (ed.) CRYPTO. LNCS, vol. 2442, pp. 31–46. Springer (2002)
19. Luykx, A., Preneel, B., Tischhauser, E., Yasuda, K.: A MAC Mode for Lightweight Block Ciphers. In: Peyrin, T. (ed.) FSE. LNCS, vol. 9783, pp. 43–59. Springer (2016)
20. Minematsu, K.: Beyond-Birthday-Bound Security Based on Tweakable Block Cipher. In: Dunkelman, O. (ed.) FSE. LNCS, vol. 5665, pp. 308–326. Springer (2009)
21. Minematsu, K.: Parallelizable Rate-1 Authenticated Encryption from Pseudorandom Functions. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT. LNCS, vol. 8441, pp. 275–292. Springer (2014)
22. Minematsu, K.: Building blockcipher from small-block tweakable blockcipher. Designs, Code and Cryptography **74**(3), 645–663 (2015)
23. Minematsu, K., Iwata, T.: Building Blockcipher from Tweakable Blockcipher: Extending FSE 2009 Proposal. In: Chen, L. (ed.) IMACC. pp. 391–412 (2011)
24. Naito, Y.: Full PRF-Secure Message Authentication Code Based on Tweakable Block Cipher. In: Au, M.H., Miyaji, A. (eds.) ProvSec. LNCS, vol. 9451, pp. 167–182. Springer (2015)
25. Nandi, M.: A Generic Method to Extend Message Space of a Strong Pseudorandom Permutation. Computación y Sistemas **12**(3) (2009), http://cys.cic.ipn.mx/ojs/index.php/CyS/article/view/1204
26. Nandi, M.: On the Optimality of Non-Linear Computations of Length-Preserving Encryption Schemes. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT (II). LNCS, vol. 9453, pp. 113–133. Springer (2015)
27. Patarin, J.: The "Coefficients H" Technique. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC. LNCS, vol. 5381, pp. 328–345. Springer (2008)
28. Peyrin, T., Seurin, Y.: Counter-in-Tweak: Authenticated Encryption Modes for Tweakable Block Ciphers. In: Robshaw, M., Katz, J. (eds.) CRYPTO I. LNCS, vol. 9814, pp. 33–63. Springer (2016)
29. Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In: ASIACRYPT. LNCS, vol. 3329, pp. 16–31. Springer (2004)
30. Rogaway, P., Zhang, Y.: Onion-AE: Foundations of Nested Encryption. PoPETs **2018**(2), 85–104 (2018)
31. Shrimpton, T., Terashima, R.S.: A Modular Framework for Building Variable-Input-Length Tweakable Ciphers. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT (1). LNCS, vol. 8269, pp. 405–423. Springer (2013)
32. Yasuda, K.: A New Variant of PMAC: Beyond the Birthday Bound. In: Rogaway, P. (ed.) CRYPTO. LNCS, vol. 6841, pp. 596–609. Springer (2011)
33. Zhang, L., Wu, W., Sui, H., Wang, P.: 3kf9: Enhancing 3GPP-MAC beyond the Birthday Bound. In: Wang, X., Sako, K. (eds.) ASIACRYPT. LNCS, vol. 7658, pp. 296–312. Springer (2012)